

RESEARCH ARTICLE 

Damage localisation using disparate damage states via domain adaptation

Chandula T. Wickramarachchi^{1,2} , Paul Gardner², Jack Poole², Clemens Hübler¹, Clemens Jonscher¹ and Raimund Rolfes¹

¹Institute of Structural Analysis, Leibniz University, Hanover, Germany

²Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Sheffield, UK

Corresponding author: Chandula T. Wickramarachchi; Email: c.t.wickramarachchi@sheffield.ac.uk

Received: 23 June 2023; **Revised:** 16 November 2023; **Accepted:** 03 December 2023


Keywords: damage localisation; distance metrics; domain adaptation; PBSHM; SHM

Abstract

A significant challenge of structural health monitoring (SHM) is the lack of labeled data collected from damage states. Consequently, the collected data can be incomplete, making it difficult to undertake machine learning tasks, to detect or predict the full range of damage states a structure may experience. Transfer learning is a helpful solution, where data from (source) structures containing damage labels can be used to transfer knowledge to (target) structures, for which damage labels do not exist. Machine learning models are then developed that generalize to the target structure. In practical applications, it is unlikely that the source and the target structures contain the same damage states or experience the same environmental and operational conditions, which can significantly impact the collected data. This is the first study to explore the possibility of transfer learning for damage localisation in SHM when the damage states and the environmental variations in the source and target datasets are disparate. Specifically, using several domain adaptation methods, this article localizes severe damage states at a target structure, using labeled information from minor damage states at a source structure. By minimizing the distance between the marginal and conditional distributions between the source and the target structures, this article successfully localizes damage states of disparate severities, under varying environmental and operational conditions. The effect of partial and universal domain adaptation—where the number of damage states in the source and target datasets differ—is also explored in order to mimic realistic industrial applications of these methods.

Impact Statement

One of the main reasons that data-based structural health monitoring systems are not widespread in the industry is because it is challenging to obtain information from damaged structures. One solution here is transfer learning where damage information from related structures is transferred to structures without damage information. However, in realistic industrial scenarios, it is unlikely that the damage states that are transferred from one structure are the same as the damage that is present in another. The situation is further complicated as real structures undergo varying environmental and operational conditions. This article tests the use of transfer learning for damage localisation using disparate types of damage from other structures under different environmental conditions, to evaluate its use in realistic industrial applications.

 This research article was awarded an Open Data badge for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.



1. Introduction

In long-term, online, structural health monitoring (SHM) campaigns, the ability to automatically locate damage within a structure is vital. This increase in information, which aids smarter maintenance decisions, is why localisation is the second level of Rytter's hierarchy of damage identification, directly following damage detection (Rytter, 1993; Farrar and Worden, 2012). In offshore wind farms, for example, significant savings on time and expenses can be made, if the most probable damage location is available prior to physically investigating the site for repair and maintenance. Moreover, knowledge of the location can help in the decision-making process regarding intervention, especially if the damage has occurred in a safety-critical location. In practice, collecting damage labels that describe locations of damage is expensive and often infeasible, making automatic damage localisation using a data-based approach to SHM rather challenging. In SHM, most data-driven damage localisation methods typically require a network of well-placed sensors close to the damage location in the structure of interest (Stubbs et al., 1995; Manson et al., 2003; Chesné and Deraemaeker, 2013; Wernitz et al., 2022). This approach can be expensive and time-consuming to implement. This issue motivates population-based SHM (PBSHM), a field of SHM that aims to increase the available labeled data by leveraging information from a population of structures (Worden et al., 2020; Gardner et al., 2021b).

Traditional machine learning methods rely on the assumption that training and testing data are drawn from the same underlying distribution (Murphy, 2012). However, differences between structures and their environmental and operational variations (EOVs) will lead to discrepancies in data. Thus, when using data from different structures, traditional machine-learning methods will likely have high generalization error. This issue motivates the application of transfer learning (Pan et al., 2010), a field of machine learning that aims to use data from related source domains (structures), with a larger amount of labeled data, to improve performance in a target domain with sparse data. A method for reducing the impact of discrepancies between domains is known as domain adaptation (Pan et al., 2010), which is a branch of transfer learning that aims to map the source and target domains to a shared space. Thus, supervised machine learning methods can be trained using source labels from the source structure and generalized to the target domain in order to undertake diagnostics or prognostics.

In the context of PBSHM, this is the first article to investigate the possibility of damage localisation of operational structures when the source and the target datasets contain different damage severities, and EOVs, via the use of domain adaptation. Here, operational structures refer to real structures in operation that are not under laboratory-controlled environments, and are exposed to the changing environment/weather. Consequently, the procedure investigated in this article can have a number of novel benefits for SHM. If applied across populations of structures, it can locate damage using disparate severities of damage from different structures. For example, one test structure within a population can be used as the source—where large damage states are introduced to explore their effect and collect damage labels—to transfer information to multiple operational target structures with different damage severities. If applied to a single structure, this methodology has the ability to identify and locate damages that have reoccurred in a structure in the same position with a different severity post-repair, for example.

A technique that has been successfully applied in SHM to address domain shift is joint distribution adaptation (JDA) (Long et al., 2013). JDA aims to minimize the distance between both marginal and conditional distributions in the shared latent space. Traditionally, this method uses naïve pseudo-labels to assess the conditional distributions in the unlabelled target domain. However, this process can produce poor target label estimates if the conditional distributions are very different. Recently, metric-informed joint distribution adaptation (M-JDA) was introduced in Gardner et al. (2021a), that uses the Mahalanobis squared distance between the classes in the source and target domains to find better guesses at the initial pseudo-labels. This approach assumes that for each damage class, the source cluster will be closest to the corresponding target cluster. In this article, the M-JDA method is applied to an operational mast structure within a population-based SHM framework, in order to locate severe damage states using information from minor damage states.

This is the first article in the context of PBSHM that explores the effectiveness of the domain adaptation methods for damage localisation under realistic conditions, in the presence of class imbalances—where the marginal distributions differ significantly between source and target domains. The aim is to understand and discuss solutions for class imbalances within domain adaptation, which are likely to occur in realistic damage localisation applications when considering time-varying datasets, EOVs, and so forth. Two types of class imbalances are investigated here. The first is concerned with partial domain adaptation where the available target classes are a subset of available source classes. The second is universal domain adaptation where the target structure contains more classes than the source structure. Class imbalances are a major cause of negative transfer, where the performance in the target domain is negatively affected. Negative transfer can occur if the data from the source and target structures are significantly dissimilar, and not related. Avoiding negative transfer is crucial in realistic transfer learning applications in SHM.

1.1. Original contributions

This article aims to use domain adaptation techniques in transfer learning for continuous damage localisation of operational structures, using a population-based SHM approach. This approach has the potential to extend the application of supervised machine learning methods for SHM to scenarios where labeled data are unavailable in the structure of interest. More specifically, the main contributions of the article are:

- A population-based approach is adopted to localize severe damage states in an unlabelled target structure/domain using knowledge of minor damage states from a labeled source structure/domain. A metric-informed joint distribution adaptation method is applied here.
- The localisation method is applied to time-varying data collected from an operational structure under natural excitation, and environmental and operational variations. The challenges of damage localisation when using a realistic dataset from an operational structure are discussed.
- The effect of, and methods to address class imbalance, in terms of partial and universal domain adaptation are explored.

1.2. Related work

Standard SHM practices require a network of sensors placed in strategic positions in order to locate damage on a structure (Stubbs et al., 1995; Manson et al., 2003; Chesné and Deraemaeker, 2013; Wernitz et al., 2022). Consequently, large amounts of well-placed sensors and collected data are required for damage location. Population-based SHM (Worden et al., 2020; Gardner et al., 2021b; Gosliga et al., 2021) was proposed to address such issues of SHM by transferring information across structures. The idea is to use labeled information from structures that have been studied in the past to make inferences about the health of other structures for which detailed knowledge may not be available. One of the most successful techniques used in PBSHM to aid this task is transfer learning.

Transfer learning (Pan et al., 2010) has been successful in transferring knowledge across different domains (source and target) in a wide variety of applications such as bioinformatics, medicine, transportation, and recommender-system applications (Zhuang et al., 2020). Typically, statistical properties of the data in the source and target are compared (Pan et al., 2010), which can be extended to also comparing the similarities in the geometric structure of the data (Long et al., 2013). For monitoring applications, the majority of work on transfer learning has focused on fine-tuning of neural networks (Cao et al., 2018; Dorafshan et al., 2018; Gao and Mosalam, 2018; Zhu et al., 2020) and domain adaptation (Michau and Fink, 2019; Gardner et al., 2020; Bull et al., 2021; Xu and Noh, 2021; Gardner et al., 2022). Since this article aims to leverage labeled source data where there are no labeled target data, domain adaptation is the most appropriate technology to investigate.

Domain adaptation, which is a branch of techniques within transfer learning, has gained popularity as a useful method to reduce the difference between a labeled source and an unlabelled target domain, in order

to transfer information across. It is used abundantly in visual applications (Ganin et al., 2016; Long et al., 2016; Csurka, 2017) and natural language processing (Ben-David et al., 2006; Blitzer et al., 2007), fault detection/condition monitoring (Li et al., 2019; Jiao et al., 2020; Li et al., 2020; Wang and Liu, 2020; Ding et al., 2021; Zhang and Li, 2022; Li et al., 2023) and PBSHM (Michau and Fink, 2019; Gardner et al., 2020; Bull et al., 2021; Xu and Noh, 2021; Gardner et al., 2022). In PBSHM, domain adaptation—in the form of transfer component analysis—was used to detect damage on a tailplane with incomplete data, by leveraging information from a heterogeneous population of other tailplanes (Bull et al., 2021).

Domain adaptation was applied in PBSHM to localize damage across a heterogeneous population of aircraft wings in Gardner et al. (2022). The authors employed graph matching methods alongside domain adaptation techniques (specifically, balanced distribution adaptation; Wang et al., 2017) in order to identify the most suitable location labels to transfer information across in an unsupervised manner, that is, without using target damage labels. The results showed that the maximum common subgraph (Gosliga et al., 2021) between the damage location of the two structures provided the best candidate features and classification performances.

In Tsialiamanis et al. (2021), damage localisation was performed on the Gnat aircraft wing by fine-tuning neural networks within a PBSHM approach. In Gardner et al. (2021a), the metric-informed JDA method was introduced to address the problem of repair and damage localisation on the Gnat aircraft wing using domain adaptation. The authors found that the domain shift between the healthy state and the repair state was larger than between the healthy state and damage. As a result, the M-JDA method was used to map damage states before and after repair in order to increase classification performance. By doing so, it was possible to locate damage of the post-repair target structure to the pre-repair source structure. This successful implementation of the M-JDA method used repeated measurements of the damage scenarios under laboratory environments, where the effect of class imbalance was not addressed. In the current article, the authors explore the suitability of the M-JDA method for continuous damage localisation when the damage states differ in severity, from a structure under natural excitation, and varying EOVs. This article also explores the effect of class imbalance when localizing damage via transfer learning.

To demonstrate the use of domain adaptation for damage localisation in SHM with disparate damage states, Section 2 presents the background of transfer learning, domain adaptation, joint distribution adaptation, and metric-informed joint distribution adaptation. In Section 3, the experimental dataset used in this work is introduced. Localizing severe damage states from minor damage states is explored in Section 4. A discussion on the effect of class imbalance on the localisation method, and possible solutions are then presented in Section 5. Finally, in Section 6, conclusions and future work are stated.

2. Transfer learning and domain adaptation

The aim of this article is to explore the use of transfer learning jointly with methods developed for population-based SHM to localize damage on a structure. In particular, metric-informed joint distribution adaptation (M-JDA) (Gardner et al., 2021a) is applied owing to its previous success in SHM. In this section, transfer learning is discussed followed by joint-distribution adaptation and M-JDA.

Transfer learning is a branch of machine learning that seeks to make inferences about data across domains (Pan et al., 2010). In transfer, learning the aim is to improve the predictive function in a target domain using knowledge from a source domain. Or to put simply, infer information about the target by leveraging knowledge from the source domain. A domain is defined as $\mathcal{D} = \{\mathcal{X}, p(X)\}$ containing \mathcal{X} , a feature space and $p(X)$, a marginal probability distribution. Here $X = \{\mathbf{x}_i\}_{i=1}^N \in \mathcal{X}$. The task within the domain is $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$. In this case, \mathcal{Y} is a label space and $f(\cdot)$ is a predictive function or a conditional distribution $p(\mathbf{y}|X)$, learnt from a training set $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$, where $\mathbf{y} \in \mathcal{Y}$. The assumption here is that the source and the target domains are not equal and/or that the source and target tasks are not equal.

In domain adaptation, the aim is to improve the target predictive function by learning a mapping that reduces the discrepancies between the source and target distributions. As such, supervised machine learning methods can be learnt using source information, and generalized to the target domain with limited or, as in this article, no labeled target data. Typically, it is assumed that the source and the target feature spaces and label spaces are equal, that is, $\mathcal{D}_s = \mathcal{D}_t$ and $\mathcal{Y}_s = \mathcal{Y}_t$, but the marginal and conditional distributions are not, that is, $p(X_s) \neq p(X_t)$ and $p(Y_s|X_s) \neq p(Y_t|X_t)$. In this article, homogeneous transfer is considered, where a homogeneous label space is assumed, that is, the set of possible classes in each domain is homogeneous. Within the homogeneous assumption, however, it is possible to have a subset of shared classes from all available classes within a domain.

Next, a method that enables the mapping of the marginal and conditional distributions between domains, named joint-distribution adaptation, is discussed (Long et al., 2013).

2.1. Joint distribution adaptation

First introduced in Long et al. (2013), joint distribution adaptation (JDA) is a method that learns a mapping that projects the source and target data into a shared latent space, where a nonparametric distribution distance metric, the maximum mean discrepancy (MMD) (Gretton et al., 2012), is minimized. The MMD is widely used in domain adaptation (Pan et al., 2010; Wang et al., 2017); it is given by,

$$\text{MMD}(p(X_s), p(X_t)) = \left\| \frac{1}{n_s} \sum_{i=1}^{n_s} \phi(\mathbf{x}_{s,i}) - \frac{1}{n_t} \sum_{i=1}^{n_t} \phi(\mathbf{x}_{t,i}) \right\|^2, \quad (1)$$

where $\phi(\cdot)$ is a kernel mapping. One limitation of typical DA algorithms is that, without knowledge of the labels, minimizing the conditional distribution distance is challenging. JDA aims to alleviate this issue by using pseudo-labels, such that the distance between class-conditional distributions $p(X_s|y_s)$ and $p(X_t|\hat{y}_t)$, can be estimated. As such, the MMD between the marginal distributions (all data) and the class-conditional distributions (the MMD between data corresponding to a particular class), is minimized as a proxy for minimizing the joint distribution distance; this metric is called the joint-MMD (JMMD) and it is given by,

$$\text{JMMD}(p(X_s, y_s), p(X_t, y_t)) = \left\| \frac{1}{n_s^{(c)}} \sum_{\mathbf{x}_{s,i} \in \mathcal{X}_s^{(c)}} \phi(\mathbf{x}_{s,i}) - \frac{1}{n_t^{(c)}} \sum_{\mathbf{x}_{t,i} \in \mathcal{X}_t^{(c)}} \phi(\mathbf{x}_{t,i}) \right\|^2, \quad (2)$$

where $\mathcal{X}_s^{(c)}$ and $\mathcal{X}_t^{(c)}$ denote the samples from class c , with $c \in 1, \dots, C$, for the source and target domain respectively, $n_s^{(c)}$ and $n_t^{(c)}$ are the samples for class c in the source and target, and $c = 0$ represents the data from all classes, giving the MMD between the marginal distributions. For a more in-depth discussion of JDA, the interested reader is directed to Long et al. (2013).

2.2. Metric-informed joint distribution adaptation

The JDA method discussed in the previous section maps not only the marginal distributions, but also the class conditional distributions, making it a powerful tool for transfer learning. However, as the target labels y_t are not usually available for this formulation, the target conditional distribution $p(y_t|X_t)$ is unknown, making it challenging to match the joint distributions between the source and target datasets. A method for obtaining target pseudo-labels \hat{y}_t in a semi-supervised manner has been discussed in Long et al. (2013) where a base classifier trained on source labels is predicted in the unlabelled target domain. Gardner et al. (2021a) built on this method by considering a metric-informed approach; a normalized distance metric—specifically, the Mahalanobis squared distance (MSD)—is used to find the most suitable target pseudo-labels, by assuming that for each class, the source cluster is closest to the target cluster. The MSD has been successfully used for outlier detection in SHM (Worden et al., 2000) where it is used to determine whether a new datapoint is similar or not to normal condition training data.

2.3. Pseudo-labeling using distance metrics

The method described in Gardner et al. (2021a) for identifying target pseudo-labels using the MSD is,

- Calculate the MSD for each class in the labeled source domain $D_c^2(\mathbf{x}_i^s) \forall c \in \{1, 2, \dots, C\}$ by using a sample mean $\hat{\boldsymbol{\mu}}(\mathcal{D}_s^{(c)})$ and covariance $\hat{\boldsymbol{\Sigma}}(\mathcal{D}_s^{(c)})$ per class. The MSD in Worden et al. (2000) is defined as,

$$D_i^2 = (\mathbf{x}_i - \hat{\boldsymbol{\mu}})^T \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}), \quad (3)$$

where \mathbf{x}_i is the current data point, and $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$ are the calculated mean and covariance determined from $X = \{\mathbf{x}_i\}_{i=1}^N$ (X is the source labeled data in this case).

- A threshold T is then calculated for each class using an Monte Carlo approach. The reader is referred to Worden et al. (2000) for more detail.
- The MSD for each class is normalized using the threshold for each class, that is, $\bar{D}_c^2(\mathbf{x}_i) = D_c^2(\mathbf{x}_i)/T_c$. All MSD values below the threshold are subsequently set to zero. The normalization allows for objective comparison between each class.
- A normalized MSD vector is now available for each class in the normalized MSD feature space $D_C^2 \in \mathbb{R}^{(n_s + n_t) \times C} = \{\bar{D}_c^2(\mathbf{x}_i)\} \forall c \in \{1, 2, \dots, C\}$.
- The i th target instance is given the pseudo-label from the class with the minimum normalized MSD, that is, the c where $D_c^2 = \min(D_C^2)$.

For a description of the pseudo-labeling algorithm for metric-informed JDA, the reader is referred to Gardner et al. (2021a).

The assumption for the above method is that the MSD calculated for each source class, when applied to the unlabelled target data $D_c^2(\mathbf{x}_j^t)$ will provide the smallest MSD values for target data from the same corresponding class. In this article, this assumption is used to locate damage on a structure where the severity of damage in source and target clusters may be different, but similar. As a result, it is important to identify damage-sensitive features in both the source and target structures that are also sensitive to the damage location. In the next section, the experimental case study used in this article for damage location is introduced, with a detailed exploration of features used for transfer learning.

3. The experimental dataset from the LUMO structure

The Leibniz University test structure for MONitoring (LUMO) is a steel lattice mast structure that includes reversible damage mechanisms (Wernitz et al., 2021, 2022). Located in Hannover, Germany, the structure provides a benchmark dataset for SHM, containing data pertaining to its dynamic, mechanical, and thermal behavior. Comprehensive documentation of the structure, as well as the open-source measurement data, can be found at <https://data.uni-hannover.de/dataset/lumo>. In this article, the data collected from the structure between September 2020 to July 2021 are explored.

The structure which stands at 9 m in height and weighs 90 kg has three tubular legs, 21 bracing levels, and short connections at each bracing level. Figure 1a presents a photograph of the structure. LUMO is equipped with 18 uni-axial accelerometers located across 9 measurement levels (ML), that measure its orthogonal deflection and the horizontal motion. Pairs of accelerometers are placed at each ML in the x - and y -directions. At the 10th measurement level (ML10), three strain gauges (one on each of the three legs) and a thermocouple also measure the strain and the temperature of the structure, respectively. At six locations throughout the structure, controlled damage states can be introduced using reversible damage mechanisms. Figure 1b presents a schematic that highlights the measurement levels, the reference axes, and the locations that contain the damage mechanisms. Figure 1c shows the damage mechanisms that are

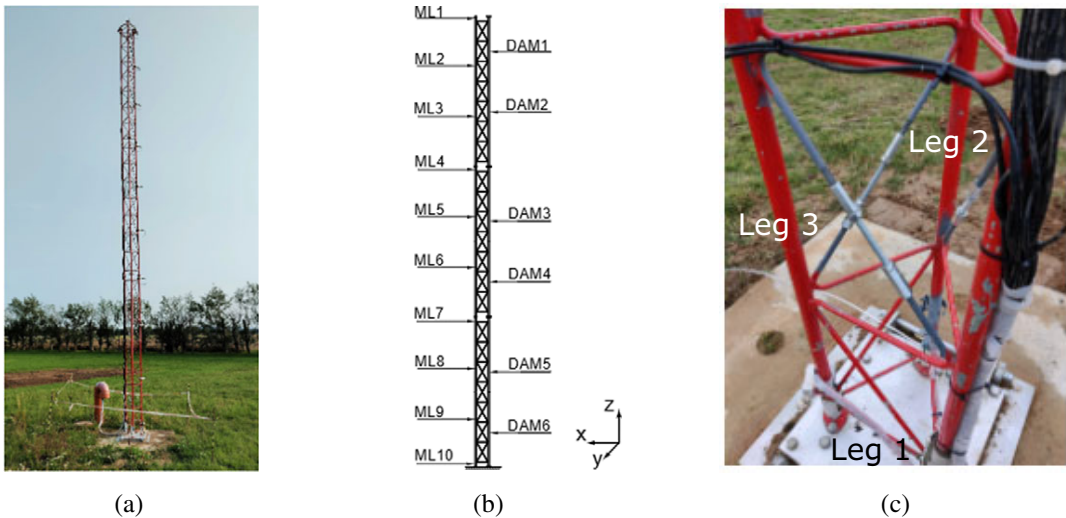


Figure 1. (a) Photograph of the LUMO structure, (b) A schematic of the structure highlighting measurement levels (ML) and damage locations (DAM) with the reference axes, and (c) Damage location 6 (DAM6) displaying the reversible damage mechanisms. Images reconstructed from Wernitz et al. (2022).

installed on the lowest damage level of the structure. An in-depth discussion of the damage states is held in Section 3.1.

As the structure is situated outside, it is excited by natural sources. As a result, the structure represents a realistic case for SHM monitoring. Given its location, the structure is exposed to, and is affected by environmental and operational variations (EOVs). These environmental variations can stem from seasonal and daily changes in temperature, change in wind directions, ice build-up, and freezing temperatures, to name a few. Each time the reversible damage states are repaired, operational variations are also introduced to the structure, which are discussed in detail in Wickramarachchi et al. (2023).

3.1. Damage states of the LUMO dataset

The main feature of the LUMO structure is the ability to introduce and repair damage states in a controlled manner, whilst the structure operates under realistic environments. The damage states essentially introduce stiffness (and possibly mass) changes to the structure by removing the damage mechanisms (Figure 1c) connected to the bracing supports. These mechanisms are present in six locations of the structure (DAM1–DAM6 in Figure 1b). At each of these levels, one, two, or three damage mechanisms can be removed.

The dataset used in this article contains measurements collected during six individual damage scenarios across three locations. At each of these three damage locations, a minor damage and a major/severe damage are introduced. A minor damage is when one damage mechanism is removed, and a major/severe damage is when all three damage mechanisms are removed. Table 1 presents the damage states, their locations, and the corresponding labels assigned in the dataset used in this work.

As the structure has a triangular cross-section, removing just one damage mechanism from a brace level leads to an asymmetrical reduction in structural stiffness. As a result, discrepancies between the data from the accelerometers that capture the dynamic behavior of the structure are to be expected.

It is the aim of this article to locate damage states using domain adaptation when the labels of the target structure are assumed unavailable. Therefore, it is pivotal that damage-sensitive features are obtained from the measured data in order to make inferences about damage locations. In the next section, the feature selection process is discussed and the selected features are provided.

Table 1. The damage states are included in the dataset

| Damage location | Location label | Damage severity | Condition of the structure | Damage label |
|-----------------|----------------|-----------------|------------------------------------|--------------|
| DAM6 | L1 | Severe | All struts removed | D1 |
| DAM4 | L2 | Severe | All struts removed | D2 |
| DAM3 | L3 | Severe | All struts removed | D3 |
| DAM6 | L1 | Minor | Strut between legs 2 and 3 removed | D4 |
| DAM4 | L2 | Minor | Strut between legs 2 and 3 removed | D5 |
| DAM3 | L3 | Minor | Strut between legs 2 and 3 removed | D6 |

Note. Each damage state and each location is assigned a label. The damage states are listed in the order they were introduced to the structure.

3.2. Feature selection

Feature selection is a vitally important step in any monitoring campaign in order to identify damage-sensitive features that provide useful information to the model. In this article, the eigenfrequencies and the mode shapes of structure are considered as features, because they are sensitive to the damage states described in Table 1. The Bayesian operational modal analysis (BAYOMA) and mode tracking methods described in Au et al. (2013) and Jonscher et al. (2023) are used to extract the natural frequencies from the acceleration signals. For a detailed description of the dynamics of the structure, the reader is referred to Wernitz et al. (2022), where 15 modes (first 5 bending modes in x -direction, first 5 bending modes in the y -directions, and first 5 torsional modes) were identified from the current sensor setup.

One of the main benefits of the LUMO dataset is the continuous nature of the monitoring campaign. The structure operates in an environment representative of realistic operational civil structures and is purely excited via natural sources. In order to replicate a realistic application of a continuous SHM campaign, the behavior of the natural frequencies throughout time is considered.

The dynamic behavior of the LUMO structure, given its dimensions and physical composition, is predisposed to producing closely spaced modes, that is, some of the natural frequencies occur very close to one another. In such cases, smallest system changes are particularly noticeable through a change in the alignment of the bending mode shape in the mode subspace (Jonscher et al., 2023). In addition, the alignment is also subject to the greatest uncertainty. Consequently, the mode tracking algorithm used to obtain the natural frequencies can, at times, misclassify the data and introduce large data scatter. This effect is especially prominent during damage (Jonscher et al., 2023; Wickramarachchi et al., 2023). As a result, a number of bending modes that are affected by these misclassifications are excluded from analysis in order to avoid progression of these errors further down the line.

The torsional modes are also excluded from the input features in this work. Although the torsional modes are highly sensitive to damage of the LUMO structure, they are affected by the mode tracking misclassification, as a result of the sensor setup. Moreover, the use of torsional modes is not prevalent in industrial applications as they are not typically highly sensitive to changes within structures. Consequently, the torsional modes are not considered further, to ensure the industrial applicability of the damage localisation method developed here.

The bending modes in the x -direction are also excluded during the feature selection process because of their low sensitivity to the minor damage states. To demonstrate the sensitivities of the x and y -direction eigenfrequencies to the minor damage states, PCA is used. Here, PCA is calculated on the natural frequencies during the time period where minor damages were introduced to the structure (20.04.2021 to 12.07.2021). Table 2 presents the principal component coefficients for the first two principal components (PC1 and PC2) which capture over 95.97% of the total variance. The results show that the bending moments in the y -direction have a much higher contribution to the variance compared to those in the x -direction. As the minor damage states are asymmetric (they are initiated by removing

Table 2. The principle component coefficients calculated on the data pertaining to minor damages. Highlighted are the bending moments in the y -direction showing a much higher contribution to the variance compared to those in the x -direction

| | PC1 | PC2 |
|-------------|--------------|---------------|
| B1-y | 0.092 | 0.018 |
| B1-x | 0.037 | -0.014 |
| B2-y | 0.991 | -0.060 |
| B2-x | 0.061 | -0.030 |
| B4-y | 0.061 | 0.997 |
| B4-x | 0.005 | -0.031 |

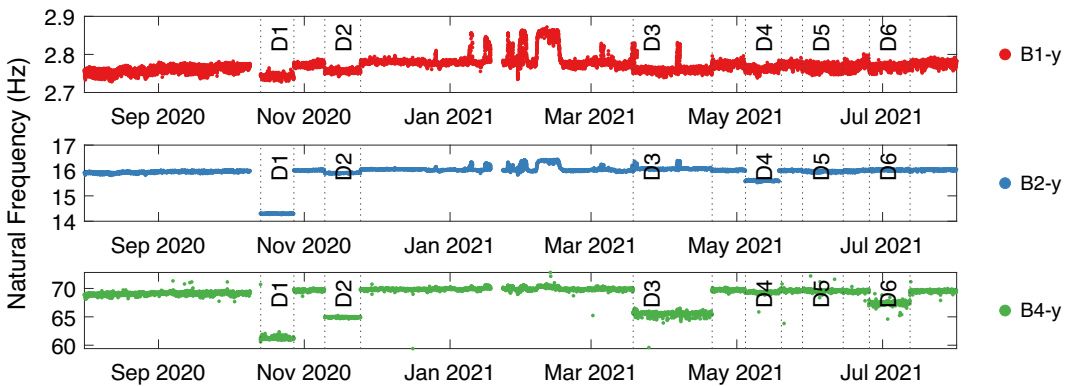


Figure 2. The features selected from the LUMO structure. The behavior of natural frequencies $B1-y$, $B2-y$, and $B4-y$ throughout time. Each datapoint represents 10-minutes of measurements. $D1$ – $D6$ highlight data collected during the damage states.

the damage mechanism between the second and third legs) in this dataset, the accelerometers in the y -direction are more sensitive than those in the x -direction. As a result, x -direction data are omitted from the dataset, purely from a feature-selection standpoint, given their reduced sensitivity to some of the damages.

Finally, considering the above, the y -direction bending modes one ($B1-y$), two ($B2-y$), and four ($B4-y$) are chosen as suitable features that encapsulate damage. These features are presented in Figure 2, where each datapoint represents 10-minutes of measurements. The sensitivity of these features to damage states ($D1$ – $D6$) can be seen in this figure. In this article, the data pertaining to damage states $D1$ – $D6$, are extracted manually from the chosen input features in Figure 2. It is also possible to automatically extract input data collected during damage by using a damage detection method, such as the one suggested in Wickramarachchi et al. (2023).

During the months of January to March, large shifts in the natural frequencies are present. By studying the temperature of the structure and the surrounding area,¹ it is found that these shifts correlate with freezing temperatures. As a result, it is assumed that a stiffening effect (SE) on the structure has caused these shifts, though inspections of the structure during this period were not conducted.

In the next section, domain-adaptation techniques are utilized to localize the severe damage states of the LUMO structure by considering the location of the minor damage states.

¹ The air temperature of the surrounding area is measured by a meteorological mast positioned 20 m away from the structure by the Institute of Meteorology and Climatology of Leibniz University Hannover.

4. Damage localisation using M-JDA

Damage localisation, the second step in Rytter's hierarchy (Rytter, 1993) aims to find the location of the damage present in the system. As discussed previously, damage localisation in SHM through domain adaptation has been attempted in past literature to address the problem of repair, on the Gnat aircraft wing (Gardner et al., 2021a). The methods in question have relied on repeated measurements of the same damage scenarios under laboratory environments, for knowledge transfer. Although these methods have proven to be successful, the scenarios in which they are attempted are not extremely likely in operational structures, that is, it is unlikely that the same damage will repeat in the same position over the life of a structure. A more appropriate situation is that a damage could reoccur in the same position with a different severity. For example, a large damage could follow an inadequate repair of a small damage. Consequently, this section explores damage localisation in the presence of different severity of damage states that appear in the same locations. Here, the metric-informed joint distribution adaptation (M-JDA) method is employed via a population-based approach, where in the first instance, the number of damage classes in the source domain is equal to the number of damage classes in the target domain.

In order to identify source and target domains for M-JDA, the LUMO mast is first viewed as a homogeneous population.

4.1. Viewing the LUMO structure as a population

To localize severe damage states by transferring knowledge from minor damage states, the LUMO structure is considered a homogeneous population with a source and target domains. The source domain includes the data collected during the minor damage states. The target domain includes data collected during the severe damage states. Table 3 presents the way in which the data has been separated. The assumption here is that for the source domain, the class labels and the parameters are known, that is, the shape of each cluster, the data within each cluster, and the location label are known for damage states D4–D6. Within the target domain (D1–D3), the number of classes, as well as the class labels (locations and severity) are assumed to be unknown.

At this stage, statistical alignment techniques developed for domain adaptation are used on the source and target domains in order to align the lower-order statistics. The aim here is to reduce the distance between the source and target domains—and, therefore, the effects of any EOVs—to aid positive transfer. Specifically, normal condition alignment (NCA) explained in Poole et al. (2022), and used in Wickramarachchi et al. (2023) to address the repair problem of the LUMO structure, is used here. Normal condition alignment is a domain adaptation method that is robust to class imbalance, making it a suitable technique in this work when the number of possible classes in the target domain may be unknown. Normal condition alignment uses affine transformations to align the first two statistical moments of normal condition data in order to reduce distribution distance, and improving the generalization of a source classifier. In this article, the first two statistical moments (extracted from the first 7 days) of the normal condition data—captured prior to each damage state—are used to standardize each normal-condition state

Table 3. The LUMO structure is viewed as a population where the source structure contains the minor damage states and the target structure contains the severe damage states

| Domain | Location label | Damage severity | Condition of the structure | Damage label |
|--------|----------------|-----------------|------------------------------------|--------------|
| Source | L1 | Minor | Strut between legs 2 and 3 removed | D4 |
| | L2 | Minor | Strut between legs 2 and 3 removed | D5 |
| | L3 | Minor | Strut between legs 2 and 3 removed | D6 |
| Target | L1 | Severe | All struts removed | D1 |
| | L2 | Severe | All struts removed | D2 |
| | L3 | Severe | All struts removed | D3 |

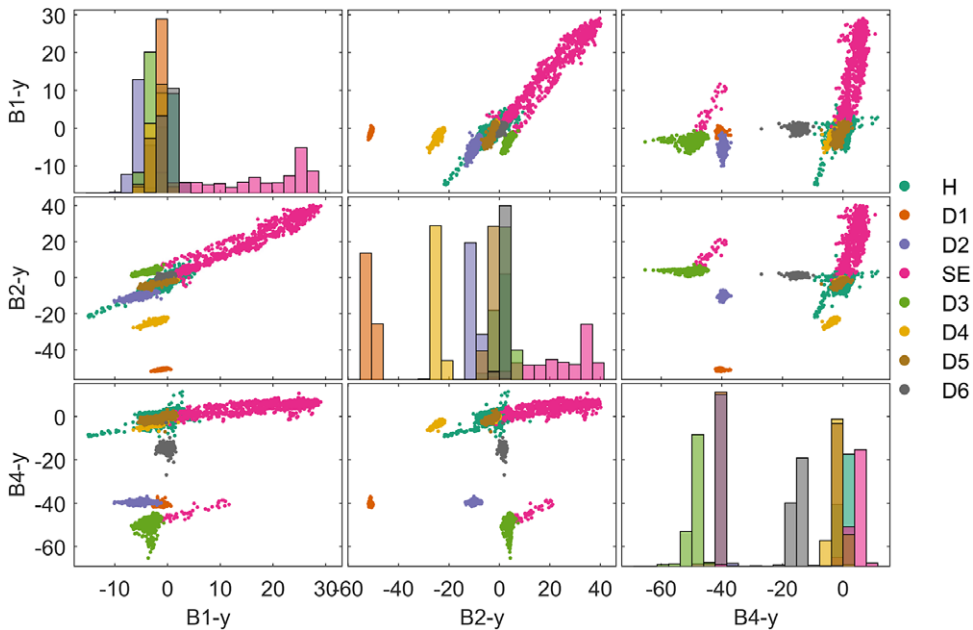


Figure 3. The standardized natural frequencies in the y -direction are plotted against each other. $D1$ – $D6$ represent damage states, and H represents the normal condition data.

and the following damage state. For example, from Figure 2, the statistics from the first 7 days of normal condition after damage state $D1$ are used to standardize the normal condition between $D1$ and $D2$ and damage state $D2$. This type of standardization can be extremely useful in PBSHM to reduce domain shifts initiated from any number of changes—including EOVs—across populations. The resulting features, post-standardization, are presented in Figure 3.

After standardization, the natural frequencies displayed in Figure 3 present interesting observations. Relative to the normal condition/healthy states (H), the clusters that represent damage state data ($D1$ – $D6$) always appear at a different direction to the SE state. For example, when considering $B1$ - y versus $B2$ - y , the damage clusters are formed in the horizontal direction whereas the SE state moves vertically (and diagonally) from the H state. Given the damage states introduce a stiffness reduction to the structure, this apparent change in direction of the SE state suggests that it is a result of a stiffness increase globally across the structure due to freezing weather. The statistical alignment techniques used here have made it possible to find physical interpretation of the clusters and their position in the feature space, which can be extremely useful when attempting to localize damage.

From the features in Figure 3, the damage states described in Table 3 are then extracted for damage localisation. The behavior of these features used in the source and target structures can be studied in Figure 4a, where standardized data are presented. It is clear that there is a common pattern between the source and target structures, as the locations have an influence on the damage state data in the feature space. Physical interpretability of the feature space can be extremely helpful in transfer learning when labels for the target structure are unavailable (as assumed in this article and demonstrated in Figure 4b); the confidence in the mapping may be increased as the feature space makes sense in reality. Some physical knowledge could also be incorporated into the process in order to make decisions about feature choice.

The features used in this work have been carefully engineered in order to pursue a successful localisation technique whilst reducing the risk of negative transfer (Bull et al., 2021). In transfer learning, negative transfer refers to the scenario where the performance of the transfer learner is worse than if a learning model was conducted solely on the target domain (Gardner et al., 2021b). The choice of features is important in any machine learning technique, as the chosen features should reflect the system well, and

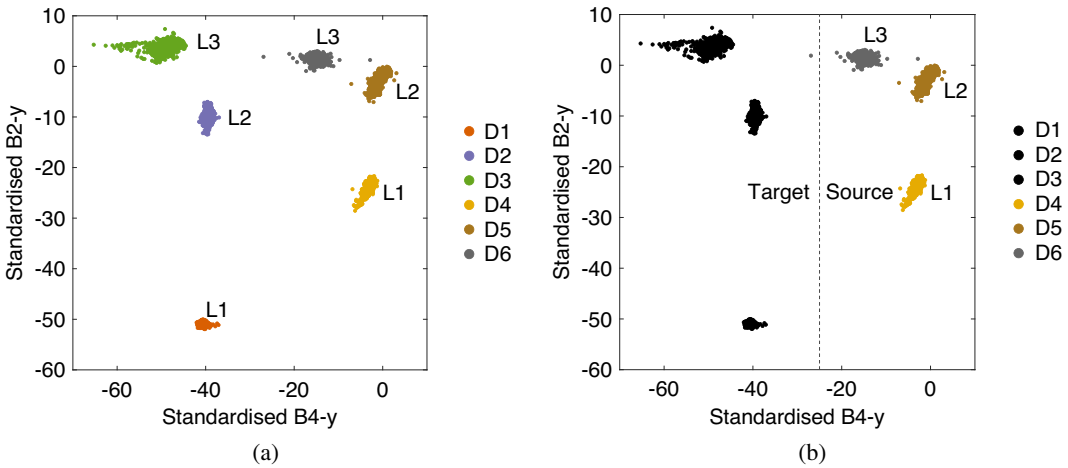


Figure 4. (a) The standardized natural frequencies B2-y versus B4-y. (b) The assumed available labels in the source (in color) and the unavailable labels in the target (in black). The normal condition data as well as the SE state data are excluded here to highlight the damage states and their locations. L1–L3 are the damage location labels.

be sensitive to the attributes that are being discriminated. In the presence of unsuitable features, negative transfer can occur even if the system in question is very similar (Pan et al., 2010; Gardner et al., 2021a, 2021b). The features used in this work are a sensible choice, as they represent the natural frequencies of the structure in the direction of the largest deflection, which are known to be sensitive to the presence of damage. The normalization techniques used here transform the data in a way in which the physical interpretability is preserved.

Figure 4a also shows that the conditional and marginal distributions between the source and target clusters are different, that is, $p(X_s) \neq p(X_t)$ and $p(y_s|X_s) \neq (y_t|X_t)$. As the severity of the damage states are different in the source structure compared to the target structure, this observation is not surprising. However, the structure of the classes within each domain shows similarity. This structural information in the feature space is leveraged to aid transfer. The differences in the direction and distance between each pair of classes in the source and target domains motivate the use of joint distribution adaptation, which aims to match both the conditional and marginal distributions. Additionally, the metric-informed approach is well suited here given the similarity in the within-domain classes, as well as the differences in the conditional distributions; a very large difference in conditional distributions can cause negative transfer when using naïve pseudo-labeling (Gardner et al., 2021a).

4.2. Localizing severe damage states using M-JDA

Once the source and target domains are identified (Table 3) and suitable features are selected, the M-JDA approach can be applied to the LUMO dataset for damage localisation.

First, the location information is inserted into the source domain data by replacing their labels y_s . For example, damage cluster D4 is relabelled as L1. Then, the source and the target datasets in Table 3 are separated further into training and testing sets. The training data are used to obtain the estimates for y_s and W ; the training data from the source and target sets are used to infer the mapping of the domain adaptation, where the target training data are unlabelled (Gardner et al., 2021a). In this work, 300 samples are randomly chosen from each class for the training sets and the rest are used for testing.

The training and testing data from the source and the target domains are then normalized using the first and second moments of the corresponding source and target training sets, respectively. Then, the pseudo-labeling procedure described in Section 2.3 is undertaken on the training data. Figure 5 presents the

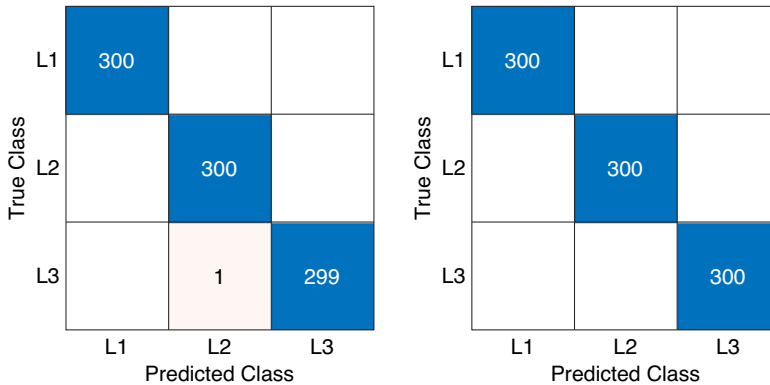


Figure 5. Confusion matrices between the true labels and predicted pseudo-labels were obtained by using the MSD metric for (left) the training datasets in the source domain and (right) training datasets in the target domain. Here, the source domain contains the minor damage states and the target domain contains the severe damage states.

(predicted) pseudo-labels obtained for the training sets. It is clear that the initial labels in the target training set matches the source training set, suggesting that the metric-informed approach has proven helpful here.

Next, the JDA method presented in Section 2.1 is conducted on the training data to calculate W , which is used to find the transformed feature spaces for source and target training and testing data. Here, the Gaussian kernel $K(x, y) = \exp(-\|x - y\|^2 / (2L^2))$ (where L is a scale parameter) is used to provide an embedding of the features into the RKHS to calculate the MMD in equation (1). The scale parameter or the width of the kernel L is calculated here using the median heuristic (Gretton et al., 2012; Garreau et al., 2017), as it has been shown to be robust in numerous studies for PBSHM (Gardner et al., 2020, 2022; Poole et al., 2022). The Gaussian kernel is a characteristic kernel that is continuous and universal in the RKHS, making the MMD a metric on the space of Borel probability measures, that is, $\text{MMD}(p, q) = 0$ if, and only if, $p = q$ (Fukumizu et al., 2007; Gretton et al., 2012).

In order to classify the transformed data within and following the M-JDA approach, a K -nearest neighbor (KNN) classifier with 10-fold cross-validation is employed. As the JDA and M-JDA methods map the data between the source and target domains to reduce their Euclidean distances, the KNN classifier is an appropriate choice here, as its aim is to find the closest class in the training set. These classifiers are only trained on the source training data. They are then applied independently to the source testing, target training, and testing sets.

4.3. Locating severe damage states: results

The results of severe damage localisation in the target domain when using minor damage states in the source domain are presented in Figures 5 and 6. Here, the regularization (or trade-off) parameter μ is set to 0.1 (following the analysis in Pan et al., 2010; Gardner et al., 2021a), and the number of transfer space dimensions k is set to 2 (as dimensionality reduction is expected from a feature space with a dimension of 3). For the Monte Carlo approach to obtain the pseudo-labels, 10,000 simulations are calculated to find each threshold with a 99% confidence bound.

In Figure 5, the confusion matrices between the true labels and predicted labels obtained using the MSD metric shows that the pseudo-labeling regime is robust during training.

In Figure 6, the left panel presents the training data mapping whilst the right panel displays the testing data mapping. It is clear that the M-JDA approach has successfully mapped the severe damage states to the minor damage states, thereby localizing the damage on the target structure, using information from the source. In comparison to classifying the normal-condition aligned features (Figure 4a), the M-JDA approach alongside NCA data has improved the classification accuracy of the target dataset significantly,

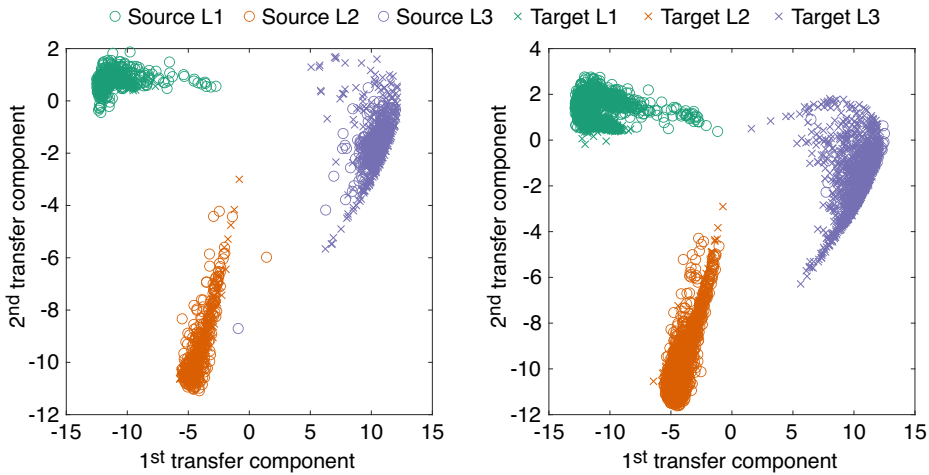


Figure 6. The results of the M-JDA approach where the left panel presents the training data and the right panel presents the testing data. Each class in the target domain has mapped well on to the corresponding class in the source domain suggesting that damage localisation from minor damage states to severe damage states is possible.

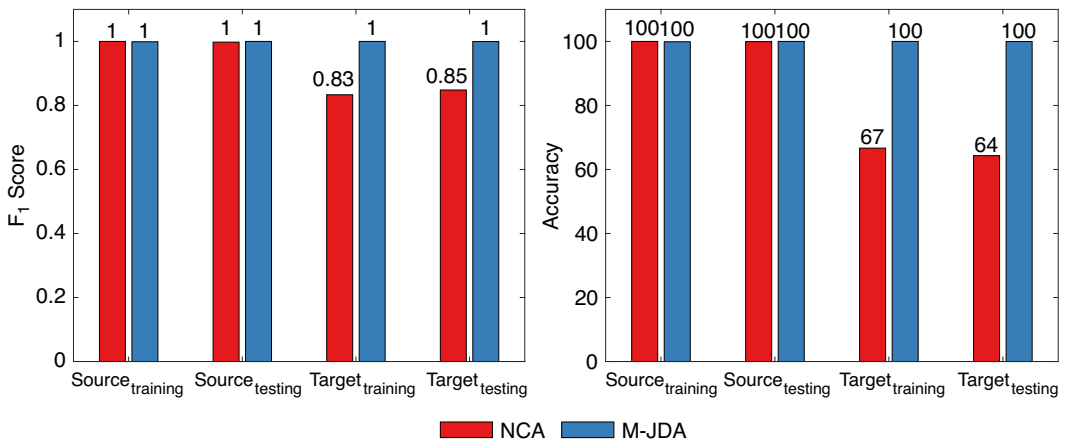


Figure 7. The comparison of KNN model performance between using NCA and NCA + M-JDA approaches. Here, the location of the severe damage states in the target domain is learnt from the labeled minor damage states in the source domain.

as seen in Figure 7. These results demonstrate the advantage of using metric-informed labels over naïve pseudo labels.

In order to evaluate the performance of the classifiers, the F_1 scores and accuracy values are calculated. The F_1 scores are calculated by,

$$F_1 = \frac{TP}{TP + \frac{1}{2}(FP + FN)}, \tag{4}$$

where TP are the true positives, FP are the true negatives and FN are the false negatives. The macro F_1 score is used to demonstrate the average for all classes. The accuracy is calculated by,

Table 4. The performance of the classifiers when severe damage states (target) are localized using information from the minor damage states (source)

| | | F ₁ score | | | | Accuracy (%) |
|--------|----------|----------------------|-------|-------|-------|--------------|
| | | L1 | L2 | L3 | All | |
| Source | Training | 1.000 | 0.998 | 0.998 | 0.999 | 99.89 |
| | Testing | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| Target | Training | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | Testing | 0.999 | 1.000 | 1.000 | 1.000 | 99.97 |

Note. Here, normal condition-aligned data are used in the M-JDA method.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}, \quad (5)$$

where TN are the true negatives. The classification performances of the training data and testing data mapping are presented in Table 4. The results show that the classifier has performed extremely well following the M-JDA approach when using normal-condition aligned data. For comparison, Table 5 presents the KNN classification performances when using non-aligned data in the M-JDA formulation. It is clear that the model is unable to generalize in the target domain because there is a lack of the physical interpretability within the feature space in this scenario. These results indicate that a method such as NCA is required to reduce EOv differences between source and target structures, in order to attempt successful transfer of operational structures. Consequently, the combination of M-JDA and NCA has proven useful for SHM to locate damage using labels from disparate damage states and structures under different EOvs.

These results can be extremely useful for population-based SHM, where information pertaining to the type, location, or severity of damage may not be available from many structures within the population. By transferring knowledge from the source structure where labels are available, it is possible to learn more information about the target structure; in this case, the location of the damage is determined. The results here show that severe damage states can be located using smaller/minor damage states, as long as the data-based features follow a similar pattern. Localizing damage can be extremely useful when deploying engineers to repair structures, as they can immediately concentrate on an optimal location without having to inspect the entire structure.

The success of this method is attributed to the similarity in the damage mechanisms; both severe and minor damage states are caused by a local reduction of stiffness in the structure, which drives the dynamic response measured by the accelerometers. However, features collected from the accelerometers in the y-direction are more sensitive to the damage states compared to those facing the x-direction. This may be because, the severe damage states in the LUMO structure introduce a symmetric reduction in stiffness (as all three bracing supports are affected), whereas, the minor damage states introduce an asymmetric

Table 5. The performance of the classifiers when using the M-JDA method without conducting NCA

| | | F ₁ score | | | | Accuracy (%) |
|--------|----------|----------------------|-------|-------|-------|--------------|
| | | L1 | L2 | L3 | All | |
| Source | Training | 0.933 | 0.929 | 0.913 | 0.925 | 92.56 |
| | Testing | 0.879 | 0.921 | 0.766 | 0.855 | 89.65 |
| Target | Training | 0.411 | 0.555 | 0 | 0.321 | 37.56 |
| | Testing | 0.321 | 0.566 | 0 | 0.296 | 36.27 |

Note. Here, minor damage states are in the source domain and severe damage states are in the target domain.




Figure 8. The influence of the percentage of training data (from each class) on the performance of the classifier.

reduction in stiffness (given only one bracing support is affected in a structure with a triangular cross-section). Therefore, it is possible that the sensitivity of the accelerometers to the damage may vary if a different bracing support (in the same location) is removed instead. This issue can be avoided by placing the accelerometers in the y -direction on a separate leg to the x -direction. Also, in structures with circular cross sections such as wind turbine towers, the issue of asymmetric damage states does not arise.

Given the natural frequencies of the source and target structures have been used here for damage localisation, drawbacks of traditional SHM localisation techniques are avoided for a structure of this scale; specifically, a network of well-placed sensors near the damage location of the target structure is not necessary for damage localisation in a PBSHM setting, when using damage sensitive natural frequencies. Consequently, the number of sensors within the target structure could be reduced when using this method, saving installation times and costs.

4.4. Influence of the training dataset size


As the transfer mappings are ascertained using training data, the influence of the training dataset size on the performance of the classifiers is scrutinized in Figure 8a,b. In these figures, the F₁ scores and classification accuracies are presented for each training and testing phase in both the source and target domains. As mentioned previously, 300 datapoints from each class (each damage state) is used for training in this article up to now. To test the influence of the training dataset size, this value is increased to 306 datapoints—which amounts to 80% of the data within the smallest class (the class with the least number of data points). Therefore, the total number of datapoints in source training set is 918 (80% of data from the three classes), whilst the same number is used in the target training set. The lowest amount of training data used is 10% of the smallest class which is 38 datapoints (or around 6 hr of data because each datapoint represents 10 mins). It is clear that the training dataset size has little influence on the performance of the classifiers when conducting metric-informed joint distribution adaptation; the classification performance values are all in the range of 99.5–100. This result demonstrates the importance of the feature selection process, as well-chosen features can reduce the burden on the methods that follow. This M-JDA method could therefore be implemented in operational structures where very little training data is necessary, as long as the chosen features show a similar behavior in both domains (in order to avoid negative transfer).

4.5. Transferring severe damage information to locate minor damage states

Unsurprisingly, the NCA + M-JDA method can be implemented in the opposite direction where the location of the severe damage states (source) can be used to determine the location of the minor damage states (target). Table 6 presents the results from this case, where a high classification performance is achieved. Figure 9 presents comparison of model performances between classifying the normal-condition aligned data and the M-JDA method in this case. The M-JDA method shows much higher classification performances, yet again. These results show that utilizing transfer learning methods in a

Table 6. The performance of the classifiers when minor damage states (target) are localized using information from the severe damage states (source)

| | | F ₁ score | | | | Accuracy (%) |
|--------|----------|----------------------|-------|-------|-------|--------------|
| | | L1 | L2 | L3 | All | |
| Source | Training | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | Testing | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| Target | Training | 0.998 | 0.997 | 0.998 | 0.998 | 99.78 |
| | Testing | 0.996 | 0.997 | 1.000 | 0.998 | 99.67 |

**Figure 9.** The comparison of KNN model performance between using NCA and NCA + M-JDA approaches. Here the source domain contains the severe damage states and the target domain contains the minor damage states.

PBSHM setting enables successful localisation of minor damage states in the LUMO structure. In comparison, when using standard SHM practices—analyzing the LUMO structure as a single structure—Kalman filter-based methods found it challenging to localize minor damage states (Wernitz, 2022). In practice, the M-JDA method may be suitable if one test structure within a population is used as the source, where large damage states are introduced to explore its behavior and collect damage labels. Then, label information could be mapped onto structures in operation that present minor damage states (unlabelled target structures).

So far, the damage localisation method using domain adaptation techniques has relied heavily on balanced classes in the source and target domain, that is, the number of damage classes in the source domain is equal to the number of damage classes in the target domain. In the next section, the effect of class imbalance on damage localisation of the LUMO structure is investigated.

5. The effect of class imbalance on damage localisation

In the previous section, a successful method for localizing severe damage states from minor damage states is developed via the use of metric-informed joint distribution adaptation combined with normal condition alignment. The method mapped three damage states in the target domain on to three damage states that shared the same location in the source domain. In realistic applications of this method, however, the presence of class imbalances is likely—specifically, the extreme case where no data for a given class is

available, that is, the number of damage classes in the source domain may not be equal to the number of damage classes in the target domain. Therefore, in this section, the effect of class imbalance on M-JDA and the classifier performance is evaluated, and methods for addressing issues with class imbalances are suggested. First, the imbalances of damage classes within the formulation are explored. Then, the class imbalances from environmental and operational variations, specifically, the stiffening effects of the LUMO structure due to freezing conditions are studied.

5.1. Damage class imbalance

In the first example explored, one damage class is withheld from either the source or the target dataset, so the problem is deviating from the homogeneous label space assumption made by conventional domain adaptation. The idea here is to investigate (a) the effect of having a source structure that contains more classes than the target structure—this paradigm is referred to as partial domain adaptation (Cao et al., 2018), where the available target classes consist of subset of available source tasks—and (b) the effect of the target structure containing more classes than the source structure—this idea is also known as universal domain adaptation (Zhuang et al., 2020). The same method of NCA + M-JDA and KNN classification from the previous section is applied here, with the only difference being the subset of classes used in the source and target domains. Table 7 presents the classification performances during a class imbalance. In this table, the results from six scenarios are presented alongside the results from the previous section where the classes are balanced. In scenarios one to three, one class is withheld from the target domain. For example, in the entry Scenario 1 that states Source (All), Target (L2 and L3), the data pertaining to L1 in the target domain is omitted from the mapping. As a result, the source domain contains more classes than the target domain (partial domain adaptation). In scenarios four to six, one location is excluded from the source domain, thereby the target domain containing more classes than the source domain (universal domain adaptation).

Some interesting observations can be made from the classification performances. Firstly, as expected, the performance of the classifiers reduces significantly in the presence of class imbalances. When a target class is excluded, the resulting F_1 score for that class is zero.

When a class is withheld from the target domain, the metric-informed procedure finds pseudo-labels for the target data by considering all the classes in the source that are most similar. This process leads to mislabelling of the existing target domain classes, resulting in poor classification performances for some damage locations. The pseudo-labels assigned for Scenario 2 demonstrate this effect in Figure 10a. In Scenario 2, high classification performances are obtained as the pseudo-labeling procedure is not hindered by the removal of the L2 class in the target domain. The classification performances are high here because the excluded target damage class (L2) appears between L1 and L3 in the feature space (Figure 4a) as well as in the physical locations within the structure (Figure 1b). Consequently, the data in target classes L1 and L3 are more similar and closer in distance to the corresponding source classes. In scenarios one and three the opposite is true, leading to negative transfer.

Given the true target labels are not available to determine the classification performances in reality, it is challenging to ensure that the assigned target labels are correct. This further supports the use of physically interpretable features, such as those used in this article, which hopefully provides some structure to the feature space, that leads to a sensible M-JDA mapping.

In universal domain adaptation scenarios four to six (where there are fewer classes in the source domain than the target), it is not possible to obtain the F_1 scores for a target class if it does not appear in the source domain. Here, the target pseudo-labels are only assigned from existing classes in the source domain, which may not always correspond to the same damage scenario, as seen in Figure 10b. As a result, the classification performances decrease in Table 7, as expected.

In order to investigate the influence of the initial target pseudo-labels from the metrics-informed approach on the JDA method when considering a class imbalance, Figures 11 and 12 are presented. In Figure 11 where partial domain adaptation is concerned, the M-JDA and JDA methods perform relatively similarly (except in Scenario 3 where the performance of the M-JDA is comparatively worse). However,

Table 7. The performance of the KNN classifiers in the presence of a class imbalances within the source and target domains

| Scenario | | | F ₁ score | | | | Accuracy (%) |
|----------|--------------------|----------|----------------------|-------|-------|-------|--------------|
| | | | L1 | L2 | L3 | All | |
| 1 | Source (All) | Training | 1.000 | 0.998 | 0.998 | 0.999 | 99.89 |
| | | Testing | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | Target (All) | Training | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | | Testing | 0.999 | 1.000 | 1.000 | 1.000 | 99.97 |
| | Source (All) | Training | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | | Testing | 0.998 | 0.998 | 0.994 | 0.997 | 99.78 |
| | Target (L2 and L3) | Training | 0 | 0.020 | 0.997 | 0.339 | 50.17 |
| | | Testing | 0 | 0.016 | 1.000 | 0.339 | 53.77 |
| 2 | Source (All) | Training | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | | Testing | 0.999 | 0.999 | 0.994 | 0.997 | 99.89 |
| | Target (L1 and L3) | Training | 1.000 | 0 | 0.995 | 0.665 | 99.50 |
| | | Testing | 1.000 | 0 | 0.998 | 0.666 | 99.69 |
| 3 | Source (All) | Training | 0.925 | 0.955 | 0.897 | 0.926 | 92.56 |
| | | Testing | 0.589 | 0.784 | 0.328 | 0.567 | 67.40 |
| | Target (L1 and L2) | Training | 0.634 | 0.718 | 0 | 0.451 | 62.00 |
| | | Testing | 0.639 | 0.758 | 0 | 0.465 | 64.41 |
| 4 | Source (L2 and L3) | Training | | 0.868 | 0.872 | 0.870 | 87.00 |
| | | Testing | | 0.811 | 0.324 | 0.567 | 70.41 |
| | Target (All) | Training | | 0.623 | 0.885 | 0.754 | 59.78 |
| | | Testing | | 0.689 | 0.880 | 0.784 | 67.83 |
| 5 | Source (L1 and L3) | Training | 0.948 | | 0.945 | 0.947 | 94.67 |
| | | Testing | 0.846 | | 0.364 | 0.605 | 75.23 |
| | Target (All) | Training | 0.393 | | 0.627 | 0.510 | 44.78 |
| | | Testing | 0.394 | | 0.705 | 0.550 | 49.72 |
| 6 | Source (L1 and L2) | Training | 0.846 | 0.807 | | 0.826 | 82.83 |
| | | Testing | 0.807 | 0.782 | | 0.795 | 79.52 |
| | Target (All) | Training | 0.995 | 0.664 | | 0.830 | 66.44 |
| | | Testing | 0.994 | 0.636 | | 0.815 | 59.00 |

Note. The locations (or classes) that are included in each M-JDA formulation are given in the second column. For example, the results for Scenario 1, Source (All), Target (L2 and L3) refer to the case where all the classes (L1–L3) are included in the source domain while only the L2 and L3 classes are included in the target domain. Macro-averaging of the F₁ scores is used here as it treats all classes as equal, and therefore, is insensitive to class imbalances.

in Figure 12 that presented the universal domain adaptation results, the JDA performance is much better than that of the M-JDA. This result suggests that the initial pseudo-labeling procedure using the metric-informed approach hinders the target pseudo-labeling procedure down the line. Figure 13a,b presents the JDA pseudo-labeling results for Scenario 2 and 4 in order to provide a comparison to the pseudo-labeling of the M-JDA process in Figure 10a,b. The assumption “for each class, the source cluster is closest to the target cluster” made in the metric-informed method is not applicable in the class-imbalance problem, leading to the higher degree of negative transfer.

5.2. Methods to address class imbalance in damage localisation

In the presence of class imbalance, the risk of negative transfer can increase, as seen in the previous section. However, analysis in Figure 14 shows that if the number of classes in the source and target domain




Figure 10. Confusion matrices between the true labels and predicted pseudo-labels were obtained by using the MSD metric for (left) the training datasets in the source domain and (right) testing datasets in the target domain. (a) The partial domain adaptation Scenario 2 is tested where all locations are in the source domain and only L1 and L3 are in the target domain. (b) The universal domain adaptation Scenario 5 is tested where all locations are in the target domain and only L1 and L3 are in the source domain.




Figure 11. The comparison of model performances between JDA and M-JDA when considering the partial domain adaptation scenarios one to three. Here, L1–L3 are included in the source domain.




Figure 12. The comparison of model performances between JDA and M-JDA when considering the universal domain adaptation scenarios four to six. Here, L1–L3 are included in the target domain.




Figure 13. Confusion matrices between the true labels and predicted pseudo-labels obtained by using the JDA approach for (left) the training datasets in the source domain and (right) testing datasets in the target domain. (a) The partial domain adaptation Scenario 2 is tested where all locations are in the source domain and only L1 and L3 are in the target domain. (b) The universal domain adaptation Scenario 5 is tested where all locations are in the target domain and only L1 and L3 are in the source domain.




Figure 14. The classification performance when using a leave-one-out method with the M-JDA approach.

are balanced, the highest classification performance is always achieved when using the same classes in the source and target domains. Given the number of classes in the source domain is known a priori, a leave-one-out method such as that used to obtain Figure 14 may be useful here. However, given target labels are assumed to be unavailable, it is not possible to obtain the classification performance metrics in real time. Again, the physical interpretability of the feature space may be extremely useful here to provide more confidence in the results; knowledge relating to the patterns between the clusters (that are driven by physics) in each domain, may be transferred across, leading to a better understanding of the mappings. For example, a Bayesian multi-task approach for fleet analysis described in Bull et al. (2022) encodes the domain expertise and constrains the model assumptions/prior beliefs in order to better share knowledge across domains. A similar method could be helpful here.

From the literature, a number of other methods have been suggested to address class imbalance when the target labels are unavailable. For addressing the class imbalance in the field of image recognition using convolutional neural networks, researchers have developed a number of techniques. For example, feature transfer learning (Yin et al., 2019) seeks to augment the feature space of under-represented subjects using common subjects, large-scale fine-grained categorization (Cui et al., 2018) employs the Earth movers distance to estimate domain similarity. Adding auxiliary weights into the formulation to handle class imbalances is also explored, where class-specific auxiliary weights are placed into the MMD formulation in Yan et al. (2017), and into the samples in the source and target datasets in Al-Stouhi and Reddy (2016). In Khan et al. (2019), a loss function that focuses on classes with low probability is defined, to classify biomedical images with a class imbalance.

When considering the JDA method explored in this article, the most suitable development from previous literature that address class imbalances is balanced distribution adaptation (BDA) (Wang et al., 2017). BDA considers both marginal and conditional distributions just as JDA, but does not give equal importance to both distributions. It does this by adaptively adjusting the importance of the distributions based on specific tasks. A weighted-BDA (W-BDA) method is also proposed in Wang et al. (2017) that adaptively varies the weight of each class to address class imbalance.

In BDA, the equivalent distance between the joint distributions is described as,

$$D(\mathcal{D}_s, \mathcal{D}_t) \approx (1 - \mu)D(p(X_s), p(X_t)) + \mu D(p(\mathbf{y}_s|X_s), p(\mathbf{y}_t|X_t)), \quad (6)$$

where $\mu \in [0, 1]$ is a balance factor that leverages the importance of distributions. The conditional distributions for the class imbalance problem are calculated using,

$$\|p(\mathbf{y}_s|X_s) - p(\mathbf{y}_t|X_t)\|_{\mathcal{H}}^2 = \|\alpha_s p(X_s|\mathbf{y}_s) - \alpha_t p(X_t|\mathbf{y}_t)\|_{\mathcal{H}}^2, \quad (7)$$

where α_s and α_t are weights that are approximated by the class priors of the respective domains. For a detailed explanation on how the marginal and conditional distributions are calculated in BDA and W-BDA, the reader is referred to Wang et al. (2017).

Figure 15 presents the BDA and W-BDA results when using $\mu = 0.5$ for Scenario 1 from Table 7. The left figure shows the M-JDA mapping, the middle is the BDA mapping, and on the right is the W-BDA mapping. To reiterate, in scenario one, the source domain contains data from all classes, whereas the target domain has classes two and three. In these figures, the target classes have been given new labels B and C (equivalent to the source labels L2–L3), as their location labels are assumed to be unavailable during the mapping.

From the classification performance results from Table 7 suggest that in Scenario 1, the source class L3 should be mapped well to target class C. The other classes do not classify well. However, visual inspection of the M-JDA mappings in Figure 15 suggests that source class L1 is also similar to target class B (which should correspond to source class L2). The outlier-driven pseudo-labeling procedure leads to this result. The BDA and W-BDA mappings, however, present source class L3 and target class C together with the other classes separated, providing a helpful solution to the class imbalance problem when target labels are unavailable.

These algorithms typically rely on minimizing marginal distribution distance, which is part of the objective in M-JDA. However, where there is significant class imbalance, accurate estimation of the

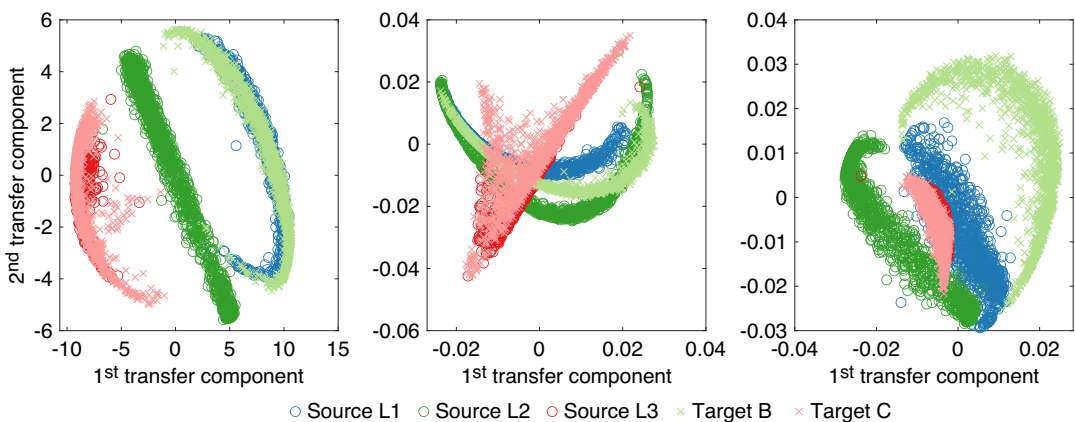


Figure 15. The results of M-JDA (left), BDA (middle), and W-BDA (right) mappings for Scenario 1 are in Table 7. Target class B has been incorrectly mapped to source class L1 in both the M-JDA and BDA results, leading to negative transfer.

underlying marginal distributions is challenging, leading to alignment based on biased estimations. Thus, this objective may lead to suboptimal mappings. An alternative solution is to rely on adapting the conditional distributions, as would be the case with BDA where $\mu = 1$, but this increases the risk of negative transfer because of noisy pseudo-labels (Poole et al., 2022). Weighting schemes have also been proposed to correct for class imbalance (Cao et al., 2018, 2019; Zhang et al., 2018), but these issues may further limit the application of domain adaptation to similar domains, as considered in this article. In Rombach et al. (2023) an approach for generating synthetic data to represent unseen/unavailable target classes is presented to address the partial domain adaptation problem. Although not considered here, this approach may be considered in future work on this dataset.

5.3. Class imbalance from EOVs

During long-term monitoring of structures, it is likely that environmental and operational variations will affect the collected data. Data affected by EOVs could present as abnormalities and should be dealt with appropriately. Automated outlier detection methods within SHM campaigns may identify data affected by EOVs as anomalous. Although detecting these abnormalities can be helpful to owners—as discussed in Wickramarachchi et al. (2023)—methods to distinguish EOVs from damage are necessary. Without such a discrimination, and without labels, it is entirely possible that the assumed SE cluster from Figure 3, for example, is misrepresented as damage, and included in the formulation when conducting damage localisation.

As the stiffening effect of the structure is unlikely to be localized in the same way as the damage states, including the SE state can be detrimental to the M-JDA mapping, and the following classifier performance (Table 8). As the methods used in this article focus on mapping unlabelled target data to labeled source data, the assumption is that information about the source data is known a priori. As a result, it is assumed that SE data would not be included in the source domain when attempting damage localisation. Clearly, it is also important to disregard the SE data from the target domain when attempting to localize damage. Given the unavailability of labels in the target class, the input features (the bending modes) within the target class could be studied as a population during the feature selection process, in order to eliminate the SE data from analysis; by considering a distance metric such as the MMD introduced in Section 2.1, the dissimilarity of the SE cluster can be identified in a principled way, without damage labels. Figure 16 presents the MMD values between the classes in the target training data. A high MMD value (dark blue) across the population suggests dissimilarity. The SE state has a higher MMD compared to other classes in the target training data, as it behaves differently to the damage state data; the stiffening effect increases the stiffness, whereas, the damage states introduce a reduction in stiffness. The effect of these different behaviors is also evident in Figure 3. The information in Figure 16 can, therefore, aid feature selection at the start of the analysis chain to discount SE data from the transfer learning process.

Table 8. The classification performance after M-JDA when considering the SE state data within the target domain

| | | F ₁ score | | | | Accuracy (%) |
|--------|----------|----------------------|-------|-------|-------|--------------|
| | | L1 | L2 | L3 | All | |
| Source | Training | 1.000 | 1.000 | 1.000 | 1.000 | 100.00 |
| | Testing | 1.000 | 0.998 | 0.976 | 0.991 | 99.78 |
| Target | Training | 0.992 | 0.580 | 0.913 | 0.828 | 70.25 |
| | Testing | 0.836 | 0.038 | 0.702 | 0.526 | 53.38 |




Figure 16. The MMD between each class within the training data. Darker shades indicate large MMD values and dissimilarities. It is clear that the SE class has a higher overall MMD compared to other damage classes.

6. Conclusions and future work

This is the first article in the context of SHM that investigated the use of domain adaptation for damage localisation when there exists discrepancies between the damage states in the labeled source structure and unlabelled target structure, under natural excitations. Differences such as severity and type of damage between source and target domains are likely in realistic applications of transfer learning, outside of laboratory environments. This is also the first article in SHM to address the effects of EOVs on real, operational structures that may affect the damage localisation process in transfer learning.

For damage localisation, metric-informed joint distribution adaptation was used to reduce the distance between the source and target domains with the help of informative pseudo-labels. This approach was able to leverage information from a mast structure (under natural excitations) with minor damage states in order to localize severe damage states, by only using the natural frequencies as input features. The M-JDA method was applied alongside normal correlation alignment in order to reduce the effect of disparities in the environments between the source and target datasets. The classification accuracy of the KNN classifier exceeded 99.97% in testing, where an overall F_1 score of 1.0 was achieved. The M-JDA + NCA method is extremely helpful in identifying the location of damage in a target structure (that does not contain any damage labels), which can then be used to optimize any following repair and maintenance procedures. By using a population-based approach, it may be possible to avoid installing a network of strategically-placed sensors for damage localisation, that are essential to traditional SHM methods.

Realistic applications of transfer learning are likely to encounter class imbalance where the source and target structures contain different number of classes. Consequently, this article also tested the suitability of the NCA + M-JDA method when there exists a class imbalance within the source and target domains. Classification performances remained high in some instances and reduced significantly in others. Given the unavailability in target labels in realistic implementations of this method, incorrect mappings can be detrimental to the decision-making process for intervention. Methods to address class imbalance have, therefore, been suggested here. In future work, a comprehensive method to address the issue of class imbalance will be studied.

Data availability statement. The data used for this work from the LUMO structure can be found at <https://data.uni-hannover.de/dataset/lumo>.

Acknowledgments. We gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—SFB-1463—434502799. C.T.W. would also like to thank the Mercator fellowship program as well as The Dynamics Research Group at the University of Sheffield for supporting this work.

Author contribution. Conceptualization: C.T.W.; Formal analysis: C.T.W.; Funding acquisition: R.R.; Investigation: C.T.W.; Methodology: C.T.W., P.G.; Resources: R.R.; Supervision: C.H., R.R.; Visualization: C.T.W.; Writing—original draft: C.T.W.; Writing—review and editing: P.G., J.P., C.J., C.H., R.R. All authors approved the final submitted draft.

Funding statement. This research was supported by grants from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—SFB-1463—434502799.

Competing interest. The authors declare none.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Al-Stouhi S and Reddy CK** (2016) Transfer learning for class imbalance problems with inadequate data. *Knowledge and Information Systems* 48(1), 201–228.
- Au S-K, Zhang F-L and Ni Y-C** (2013) Bayesian operational modal analysis: Theory, computation, practice. *Computers & Structures* 126, 3–14.
- Ben-David S, Blitzer J, Crammer K and Pereira F** (2006) Analysis of representations for domain adaptation. *Advances in Neural Information Processing Systems* 19, 1–8.
- Blitzer J, Dredze M and Pereira F** (2007) Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. Prague: ACL, pp. 440–447.
- Bull L, Di Francesco D, Dhada M, Steinert O, Lindgren T, Parlikad AK, Duncan A and Girolami M** (2022) Hierarchical Bayesian modeling for knowledge transfer across engineering fleets via multitask learning. *Computer-Aided Civil and Infrastructure Engineering* 38, 821.
- Bull L, Gardner P, Dervilis N, Papatheou E, Haywood-Alexander M, Mills R and Worden K** (2021) On the transfer of damage detectors between structures: An experimental case study. *Journal of Sound and Vibration* 501, 116072.
- Cao Z, Ma L, Long M and Wang J** (2018) Partial adversarial domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Berlin: Springer.
- Cao Z, You K, Long M, Wang J and Yang Q** (2019) Learning to transfer examples for partial domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Long Beach, CA: CVPR, pp. 2985–2994.
- Cao P, Zhang S and Tang J** (2018) Preprocessing-free gear fault diagnosis using small datasets with deep convolutional neural network-based transfer learning. *IEEE Access* 6, 26241–26253.
- Chesné S and Deraemaeker A** (2013) Damage localisation using transmissibility functions: A critical review. *Mechanical Systems and Signal Processing* 38(2), 569–584.
- Csurka G** (2017) A comprehensive survey on domain adaptation for visual applications. In *Domain Adaptation in Computer Vision Applications*. Cham: Springer, pp. 1–35.
- Cui Y, Song Y, Sun C, Howard A and Belongie S** (2018) Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT: IEEE, pp. 4109–4118.
- Ding Y, Jia M and Cao Y** (2021) Remaining useful life estimation under multiple operating conditions via deep subdomain adaptation. *IEEE Transactions on Instrumentation and Measurement* 70, 1–11.
- Dorafshan S, Thomas RJ and Maguire M** (2018) Comparison of deep convolutional neural networks and edge detectors for image-based crack detection in concrete. *Construction and Building Materials* 186, 1031–1045.
- Farrar CR and Worden K** (2012) *Structural Health Monitoring: A Machine Learning Perspective*. Hoboken, NJ: John Wiley & Sons.
- Fukumizu K, Gretton A, Sun X and Schölkopf B** (2007) Kernel measures of conditional dependence. *Advances in Neural Information Processing Systems* 20, 1–8.
- Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, Marchand M and Lempitsky V** (2016) Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17(1), 2096–2030.
- Gao Y and Mosalam KM** (2018) Deep transfer learning for image-based structural damage recognition. *Computer-Aided Civil and Infrastructure Engineering* 33(9), 748–768.
- Gardner P, Bull L, Dervilis N and Worden K** (2021a) Overcoming the problem of repair in structural health monitoring: Metric-informed transfer learning. *Journal of Sound and Vibration* 510, 116245.
- Gardner P, Bull L, Gosliga J, Dervilis N and Worden K** (2021b) Foundations of population-based shm, part III: Heterogeneous populations—mapping and transfer. *Mechanical Systems and Signal Processing* 149, 107142.
- Gardner P, Bull L, Gosliga J, Poole J, Dervilis N and Worden K** (2022) A population-based SHM methodology for heterogeneous structures: Transferring damage localisation knowledge between different aircraft wings. *Mechanical Systems and Signal Processing* 172, 108918.
- Gardner P, Liu X and Worden K** (2020) On the application of domain adaptation in structural health monitoring. *Mechanical Systems and Signal Processing* 138, 106550.
- Garreau D, Jitkritum W and Kanagawa M** (2017) Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*.
- Gosliga J, Gardner P, Bull L, Dervilis N and Worden K** (2021) Foundations of population-based SHM, part II: Heterogeneous populations—graphs, networks, and communities. *Mechanical Systems and Signal Processing* 148, 107144.

- Gretton A, Borgwardt KM, Rasch MJ, Schölkopf B and Smola A** (2012) A kernel two-sample test. *The Journal of Machine Learning Research* 13, 723–773.
- Gretton A, Sejdinovic D, Strathmann H, Balakrishnan S, Pontil M, Fukumizu K and Sriperumbudur BK** (2012) Optimal kernel choice for large-scale two-sample tests. *Advances in Neural Information Processing Systems* 25, 1–9.
- Jiao J, Zhao M, Lin J and Liang K** (2020) Residual joint adaptation adversarial network for intelligent transfer fault diagnosis. *Mechanical Systems and Signal Processing* 145, 106962.
- Jonscher C, Hofmeister B, Griebmann T and Rolfes R** (2023) Influence of environmental conditions and damage on closely spaced modes. In *European Workshop on Structural Health Monitoring*, vol. 270. Cham: Springer International Publishing and Imprint Springer, pp. 902–911.
- Jonscher C, Liesecke L, Penner N, Hofmeister B, Griebmann T and Rolfes R** (2023) Influence of system changes on closely spaced modes of a large-scale concrete tower for the application to structural health monitoring. *Journal of Civil Structural Health Monitoring* 29, 328.
- Khan NM, Abraham N, Hon M and Guan L** (2019) Machine learning on biomedical images: Interactive learning, transfer learning, class imbalance, and beyond. In *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*. San Jose, CA: IEEE, pp. 85–90.
- Li Y, Song Y, Jia L, Gao S, Li Q and Qiu M** (2020) Intelligent fault diagnosis by fusing domain adversarial training and maximum mean discrepancy via ensemble learning. *IEEE Transactions on Industrial Informatics* 17, 2833–2841.
- Li X, Yu S, Lei Y, Li N and Yang B** (2023) Intelligent machinery fault diagnosis with event-based camera. *IEEE Transactions on Industrial Informatics* 20, 380–389.
- Li X, Zhang W, Ding Q and Sun J-Q** (2019) Multi-layer domain adaptation method for rolling bearing fault diagnosis. *Signal Processing* 157, 180–197.
- Long M, Wang J, Ding G, Shen D and Yang Q** (2013) Transfer learning with graph co-regularization. *IEEE Transactions on Knowledge and Data Engineering* 26(7), 1805–1818.
- Long M, Wang J, Ding G, Sun J and Yu PS** (2013) Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*. Sydney, NSW: IEEE, pp. 2200–2207.
- Long M, Zhu H, Wang J and Jordan MI** (2016) Unsupervised domain adaptation with residual transfer networks. *Advances in Neural Information Processing Systems* 29, 1–9.
- Manson G, Worden K and Allman D** (2003) Experimental validation of a structural health monitoring methodology: Part III. Damage location on an aircraft wing. *Journal of Sound and Vibration* 259(2), 365–385.
- Michau G and Fink O** (2019) Domain adaptation for one-class classification: Monitoring the health of critical systems under limited information. In *International Journal of Prognostics and Health Management, Deep Learning and Emerging Analytics Special Issue*. State College, PA: PHM Society.
- Murphy KP** (2012) *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT Press.
- Pan SJ, Tsang IW, Kwok JT and Yang Q** (2010) Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks* 22, 199–210.
- Poole J, Gardner P, Dervilis N, Bull L and Worden K** (2022) On statistic alignment for domain adaptation in structural health monitoring. *Structural Health Monitoring* 22, 1–20.
- Poole J, Gardner P, Dervilis N, Bull L and Worden K** (2022) On the application of partial domain adaptation for pbsh. *European Workshop on Structural Health Monitoring: EWSHM 2022* 3, 408–418.
- Rombach K, Michau G and Fink O** (2023) Controlled generation of unseen faults for partial and open-partial domain adaptation. *Reliability Engineering & System Safety* 230, 108857.
- Rytter A** (1993) *Vibrational Based Inspection of Civil Engineering Structures*. PhD Thesis, Aalborg University.
- Stubbs N, Kim J-T and Farrar CR** (1995) Field verification of a nondestructive damage localisation and severity estimation algorithm. In *Proceedings-SPIE the International Society for Optical Engineering*. Nashville, TN: SPIE International Society for Optical, pp. 210.
- Tsialiamanis GP, Wagg DJ, Gardner PA, Dervilis N and Worden K** (2021) On partitioning of an SHM problem and parallels with transfer learning. In *Topics in Modal Analysis & Testing, Volume 8: Proceedings of the 38th IMAC, A Conference and Exposition on Structural Dynamics 2020*. Houston, TX: Springer International Publishing, pp. 41–50.
- Wang J, Chen Y, Hao S, Feng W and Shen Z** (2017) Balanced distribution adaptation for transfer learning. In *2017 IEEE International Conference on Data Mining (ICDM)*. New Orleans, LA: IEEE, pp. 1129–1134.
- Wang X and Liu F** (2020) Triplet loss guided adversarial domain adaptation for bearing fault diagnosis. *Sensors* 20, 320.
- Wernitz S** (2022) *Damage Localisation in Data-Driven Vibration-Based Structural Health Monitoring using Linear Quadratic Estimation Theory*. PhD Thesis, Institut für Statik und Dynamik, Gottfried Wilhelm Leibniz Universität, Hannover.
- Wernitz S, Chatzi E, Hofmeister B, Wolniak M, Shen W and Rolfes R** (2022) On noise covariance estimation for kalman filter-based damage localisation. *Mechanical Systems and Signal Processing* 170, 108808.
- Wernitz S, Hofmeister B, Jonscher C, Griebmann T and Rolfes R** (2021) Dataset: Lumo – Leibniz university test structure for monitoring. <https://doi.org/10.25835/0027803>.
- Wernitz S, Hofmeister B, Jonscher C, Griebmann T and Rolfes R** (2022) A new open-database benchmark structure for vibration-based structural health monitoring. *Structural Control and Health Monitoring* 29, e3077.
- Wickramarachchi CT, Poole J, Hübler C, Jonscher C, Hofmeister B and Rolfes R** (2023) Statistical alignment in transfer learning to address the repair problem: An experimental case study. Preprint submitted to MSSP, pp. 1–25.

- Worden K, Bull LA, Gardner P, Gosliga J, Rogers TJ, Cross EJ, Papatheou E, Lin W and Dervilis N** (2020) A brief introduction to recent developments in population-based structural health monitoring. *Frontiers in Built Environment* 6, 146.
- Worden K, Manson G and Fieller NR** (2000) Damage detection using outlier analysis. *Journal of Sound and Vibration* 229, 647–667.
- Xu S and Noh HY** (2021) Phymdan: Physics-informed knowledge transfer between buildings for seismic damage diagnosis through adversarial learning. *Mechanical Systems and Signal Processing* 151, 107374.
- Yan H, Ding Y, Li P, Wang Q, Xu Y and Zuo W** (2017) Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Honolulu, HI: IEEE, pp. 2272–2281.
- Yin X, Yu X, Sohn K, Liu X and Chandraker M** (2019) Feature transfer learning for face recognition with under-represented data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Princeton, NJ: NEC Laboratories America, pp. 5704–5713.
- Zhang J, Ding Z, Li W and Ogunbona P** (2018) Importance weighted adversarial nets for partial domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT: IEEE, pp. 8156–8164.
- Zhang W and Li X** (2022) Data privacy preserving federated transfer learning in machinery fault diagnostics using prior distributions. *Structural Health Monitoring* 21(4), 1329–1344.
- Zhu J, Zhang C, Qi H and Lu Z** (2020) Vision-based defects detection for bridges using transfer learning and convolutional neural networks. *Structure and Infrastructure Engineering* 16(7), 1037–1049.
- Zhuang F, Qi Z, Duan K, Xi D, Zhu Y, Zhu H, Xiong H and He Q** (2020) A comprehensive survey on transfer learning. *Proceedings of the IEEE* 109(1), 43–76.