Statistical alignment in transfer learning to address the repair problem: An experimental case study

Chandula T. Wickramarachchi[1,2], Jack Poole[2], Clemens Hübler[1], Clemens Jonscher[1], Benedikt Hofmeister[1], and Raimund Rolfes[1]

[1]Leibniz University Hannover, Institute of Structural Analysis, Appelstr. 9a, 30167 Hanover, Germany
[2]Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK

January 26, 2024

Abstract

Repair is a critical step in maintenance of civil structures to ensure safe operation. However, repair can pose a problem for data-driven approaches of long-term structural health monitoring, because repairs can change the underlying distributions of the data, which can invalidate models trained on pre-repair data. As a result, models previously trained on pre-repair information fail to generalise to post-repair data, reducing their performances and misrepresenting the actual behaviour of structures. This paper suggests a population-based structural health monitoring approach to address the problem of repair in long-term monitoring of a mast structure, by exploring domain adaptation techniques developed for transfer learning. A combined approach of normal condition alignment and Dirichlet process mixture models is adopted here for damage detection, that can operate unimpeded by post-repair shifts in distributions. The method is able correctly identify 99% of the damage data with a false positive rate of around 1.6%. Moreover, it is able to detect environmental variations such as stiffening due to freezing conditions that can adversely affect the dynamic behaviour of structures.

Keywords: Repair . Domain adaptation . Statistical alignment . SHM . Clustering

1 Introduction

On-line, long-term structural health monitoring campaigns aim to continuously monitor structures to identify and classify damages that may hinder their ability to safely perform, as intended [1]. As time passes, the monitored structures are expected to age and accumulate damage, that can affect their intended performance. Consequently, structures are likely to undergo repairs during their lifetimes. Therefore, an intelligent SHM system should be able to operate continuously, unimpeded by structural repairs [2]. In practice, however, developing data-driven approaches for continuous, long-term SHM can be challenging, especially when taking the process of repair into consideration.

When using conventional machine learning methods within SHM systems, the assumption is that training and testing data are drawn from the same distribution, i.e., the distribution of the data in the healthy state and various damage states are the same during the training and testing periods of the model. However, this assumption is too

strong when considering repairs. Repairs can change the local dynamic properties of the structure, leading to shifts in the distributions of the features used in the SHM models. Subsequently, as structures are repaired, the underlying distribution of the data within the same damage/healthy state can change (i.e., the pre- and post-repair data in the same damage state present different distributions), thus invalidating the assumptions made [3–5]. Consequently, the performance of the trained model can degrade significantly, to the extent that either critical damage is classified as healthy data – potentially leading to unsafe operation, or healthy data could be misidentified as damage – leading to unintended downtimes. Considerable costs to the owners and users of the structures are likely following both aforementioned scenarios.

One expensive and wasteful solution is to discard the data collected before repair and retrain models using new labelled information post-repair. Not only is this approach time consuming and undesirable, but it also nullifies the expensive-to-collect pre-repair damage state data – because the damage state data pre-repair may have a different distribution post-repair.

Another more constructive solution exists in the form of domain adaptation within transfer learning technologies [6]. The purpose of transfer learning is to make inferences about one domain using information from a related domain. Domain adaptation aims to find a mapping that aligns domains, such that labelled data from a source domain can be utilised to learn machine learning models that generalise to a target domain. Within domain adaptation, the change in the underlying distributions of the data pre- and post-repair can be considered a domain shift, where a classifier trained on pre-repair data is unlikely to generalise to post-repair data. Domain adaptation can be used to reduce this domain shift and retain pre-repair information in the model, utilise it to assist future monitoring, increase the pool of available information beyond normal condition data, thus reducing information wastage [7].

Reducing the domain shift introduced by repairs to the structure using domain adaptation can be framed as a population-based SHM (PBSHM) problem; PBSHM aims to make inferences about target structures, using information from other (similar) source structures [8]. Within the PBSHM framework, the structure pre- and post-repair can be treated as source and target structures (respectively) in a homogeneous population (where the structures are nominally identical) with operational and manufacturing variations.

Domain adaptation is prevalent in visual applications [9], in natural language processing [10], fault detection/condition monitoring [11–15] and PBSHM [7, 16–19]. In SHM, domain adaption was successfully applied to address domain shifts during repair of the Gnat aircraft wing [3]. The authors mapped post-repair damages to pre-repair damages by adopting a metric-informed joint distribution adaptation (JDA) approach. JDA aims to map labelled and unlabelled datasets into a shared latent space where the differences between the marginal and conditional distributions are minimised [20], thus allowing transfer across the datasets. In [3] the Mahalanobis squared distance aided pseudo-labelling of the target conditional distribution (the labels are otherwise unknown in JDA, and usually requires a semi-supervised approach to identify). The methods described in [3, 4, 16] are, however, not applicable to this paper because the only class that is shared between the source and target domains in this paper is the normal condition, and as such, it is an open set domain adaptation problem – a challenging scenario where conventional domain adaptation algorithms suffer from serious performance degradation [21].

In this paper, the assumption is that, during long-term monitoring, labels are available from only a subset of data within the normal condition pre- and post-repair. This is a sensible assumption as the process of repair requires on-site physical intervention, and post-repair, the structure should be operating in the normal condition. Therefore, the only shared state between the source and target domains are the normal condition, and as a result, the effect of repair is viewed as a partial domain adaptation problem [21] here; in partial domain adaptation, there is also an assumption that the target domain does not contain any classes that are not in the source domain. For this purpose, statistical alignment techniques to aid domain adaptation in a population-based setting are explored [5, 22]. In particular, normal-condition alignment (NCA) [5], is applied in this paper. NCA aims to align the lower order statistics of the data within the normal condition which can also address issues of class imbalance. The method outperformed other domain adaptation techniques such as transfer component analysis and domain adversarial neural networks when classifying damages across domains in [5]. It has also proven useful in a repair scenario previously in [4], where knowledge across two heterogeneous bridges were transferred in order to identify damage classes. This approach is well-suited to the repair problem explored in this paper because the normal condition is the only shared health state across domains. As the typical assumption is that normal condition data have a Gaussian distribution, beyond aligning the correlations, further adaptation cannot be achieved without other classes

in the source domain. Also, by only aligning the normal condition data in the source and target domains, the threat of negative transfer can be minimised, as naively aligning data corresponding to different health-states will cause negative transfer; negative transfer is the process of leveraging information from the source domain that produces undesirable performances in the target domain. Therefore, by including engineering knowledge, it is possible to confidently guarantee that data gathered from similar health states are aligned. Moreover, the alignment method is physically interpretable compared to other dimensionality reduction methods, for example, which is a major benefit to this study.

In this paper, the main aim is to develop a long-term damage detection method that is not invalidated by the domain shift during repair procedures. To that end, statistical alignment techniques developed for domain adaptation in PBSHM are combined with Dirichlet process mixture models (DPMMs) – an online clustering method previously utilised for damage detection in SHM [23]. DPMMs do not require a training period or the expected number of clusters to be set *a-priori*. As a result, the model is able to automatically form new clusters to capture domain shifts that materialise as a result of repairs and other environmental and operational variations. These aforementioned methods are applied to an experimental case study which explores continuous monitoring of a new benchmark mast structure that includes reversible damage states.

The structure of this paper is as follows. Section 2 provides a mathematical introduction to normal condition alignment in domain adaptation. Then, a case study of the long-term monitoring of a mast structure is introduced in Section 3, followed by Section 4 that discusses the domain shift in the mast structure post-repair. Subsequently, Section 5 explored the novel use of statistical alignment methods developed for domain adaptation to address the repair problem. Next, Section 6 demonstrates the novel approach of combined Dirichlet process mixture models and statistical alignment techniques for long-term damage detection of the test structure, within the context of structural repairs. An exploration of the sensitivity of these methods to the amount of data required, as well as the type of input features are also conducted. Finally in Section 7 conclusions and future work are presented.

2 Statistical alignment techniques developed for domain adaptation

Transfer learning, that sits under the umbrella of machine learning, is a method developed for transferring knowledge across domains. It has been successfully used in the field of population-based structural health monitoring (PBSHM) where labelled data from groups of structures are transferred to unlabelled datasets in order to improve the performance of learned models [7] or facilitate more in-depth diagnostics. Domain adaptation is a branch of transfer learning that aims to find a mapping that minimise the distance between distributions from different domains.

Domain adaptation is usually performed across two domains - the source domain and the target domain. A domain $\mathcal{D}$ is defined by its feature space $\mathcal{X}$ and a marginal distribution $P(X)$, where $X = \{\boldsymbol{x}_i\}_{i=1}^{N} \in \mathcal{X}$[24].

The aim here is to improve the performance of the task ($\mathcal{T}_t$) in the target domain $\mathcal{D}_t$ by leveraging the knowledge gained from the source domain $\mathcal{D}_s$ and its task $\mathcal{T}_s$. A task is defined as $\mathcal{T} = \{\mathcal{Y}, f(.)\}$ where $\mathcal{Y}$ is a label space and $f(.)$ is a predictive function or conditional distribution $p(\boldsymbol{y}|X)$ learnt from training data $\{\boldsymbol{x}_i, y_i\}_{i=1}^{N}$ (from the source domain), where $\boldsymbol{y} \in \mathcal{Y}$. The assumption in homogeneous domain adaptation is $\mathcal{X}_s = \mathcal{X}_t$ and $\mathcal{Y}_s = \mathcal{Y}_t$ (in this case, for a shared subset of labels in the source and target domains) but that $p(X_s) \neq p(X_t)$ and $p(Y_s|X_s) \neq p(Y_t|X_t)$ (though generally assumed to be similar). In this paper, to reduce the distance between the source and target domains, statistical alignment will be performed in order to achieve $p(\phi(X_s)) \approx p(\phi(X_t))$ and $p(Y_s|\phi(X_s)) \approx p(Y_t|\phi(X_t))$, where $\phi(.)$ is some mapping.

2.1 Statistical alignment

Statistical alignment techniques investigated for domain adaptation focuses on directly aligning the lower order statistics (the mean and the variance/standard deviation, in this case) of the data. As a result, an affine transformation and scaling will be performed. This method does not require density estimation prior to alignment, reducing the demand for a large number of datapoints from each health state.

In particular, normal-condition alignment described in [5] is explored in this paper, where the the first two statistical moments of data under normal operating conditions (when the structure is healthy and working as

intended) are aligned. This method allows data that corresponds to known shared health-states to be selected, avoiding naive alignment that can lead to negative transfer. Considering the aforementioned PBSHM approach, the procedure to achieve normal condition alignment is,

1. Standardise the data from the source domain ($P(X_s)$) using,

$$\boldsymbol{z}_s = \frac{\boldsymbol{x}_s - \mu_{s,n}}{\sigma_{s,n}} \tag{1}$$

   where $\mu_{s,n}$ and $\sigma_{s,n}$ are the mean and standard deviation of the source in normal condition, respectively.
2. Align the data from each repair state (target structures $P(X_t)$) individually to the source structure using,

$$\boldsymbol{z}_t = \left(\frac{\boldsymbol{x}_t - \mu_{t,n}}{\sigma_{t,n}}\right)\sigma_{\boldsymbol{z}_s} + \mu_{\boldsymbol{z}_s} \tag{2}$$

   where $\mu_{t,n}$ and $\sigma_{t,n}$ are the means and standard deviations of the target structures in normal condition, respectively.

The above method is most suitable when the data under normal operating conditions are assumed be Gaussian, and higher order statistics are already similar across domains. By assuming Gaussianity, further alignment is not possible at this stage. It is possible to learn a nonlinear mapping to align the normal condition data instead, though there is a risk that the mapping may not generalise to future unseen states, as this would be comparable to extrapolation in nonlinear regression, which is known to be challenging.

This statistical alignment method – developed to address domain shifts in PBSHM – may be extremely helpful when dealing with post-repair data from operational structures undergoing long-term monitoring campaigns; they address the core issue of selecting a subset of the data that corresponds and do not learn nonlinear mappings that may not generalise to unseen damage-states. Statistical alignment methods are also more computationally efficient than inverting large matrices or learning deep neural networks for online solutions.

In the next section, a case study of an operational mast structure equipped with a long-term data collection system is presented. The naturally excited structure contains a number of repair and damage states that are challenging to classify using typical SHM methods, as a result of domain shifts.

3 Case study: Long-term monitoring of the LUMO structure

Located in Hannover, Germany, the **L**eibniz **U**niversity test structure for **mo**nitoring (LUMO) provides a benchmark for SHM research [25, 26]. A long-term monitoring campaign is continuously generating a dataset from the structure that is exposed to natural sources of excitation, and environmental and operational variation. The structure contains 18 reversible damage mechanisms at various locations that can be introduced and repaired. For comprehensive documentation and access to the open source measurement data, readers are referred to `https://data.uni-hannover.de/dataset/lumo`.

To date, the LUMO dataset contains data pertaining to six controlled damage states and subsequent repairs, across multiple locations along the structure. Upon repairs, it was found that the underlying distribution of the structure was affected, leading to a dataset with domain shifts in the normal condition. It is the aim of this case study to develop an online, long-term monitoring procedure for damage detection, by using the dynamic response of the structure from natural excitations as features. The developed system should be unimpeded by domain shifts post-repair. Additionally, it is the hope that the method will be sensitive to other adverse changes to the structure's dynamics behaviour as a consequence of environmental and operational variations (EOV), that can also present as domain shifts.

As the mast structure was designed to be representative of real civil structures in operation, the developed SHM system intended for damage detection should also be robust, practical and cost-effective to implement. Consequently, this study also has the following objectives:

- To develop a method that not only identifies damage, but also environmental and operational variation that affect civil structures in reality.
- The method should be sensitive to different types of damages in a variety of locations within the structure.

- Furthermore, the robustness of the techniques developed here to the measurement system is assessed, in order to identify the best combination of sensors for the task. By doing so, it may be possible to reduce the number of inputs/features and reduce the costs/efforts associated with installation, equipment, management, storage, data-processing, etc.
- Given the costs associated with data collection, labelling, storing, processing, and computational efficiency, the influence of the dataset size on the method's performance is scrutinized.

In this section, an introduction to the test structure and its reversible damage capabilities is provided, including explanations on the pre-processing techniques and feature selection procedures undertaken for this work.

3.1 The structure

LUMO is a steel lattice mast structure that stands at a height of 9 m and weighs around 90 kg. Three identical sections constitute the mast, each containing three tubular legs, seven bracing levels and short connections at the ends. At six of the bracing levels across the entire structure, there exist mechanisms that allow reversible damages to be introduced. Figures 1a to 1c presents the LUMO structure and its concrete foundation.

3.2 Reversible damage states

The ability to introduce damages and repair them in a controlled manner is a key advantage of the LUMO test structure. These reversible local damages are aimed to alter the stiffness (and possibly mass) at specific points along the mast by removing a section of the bracing supports. Damages can be introduced at six different levels of the structure (DAM1 - DAM6 in Figure 1b), where at each level, all three bracing support sections or struts

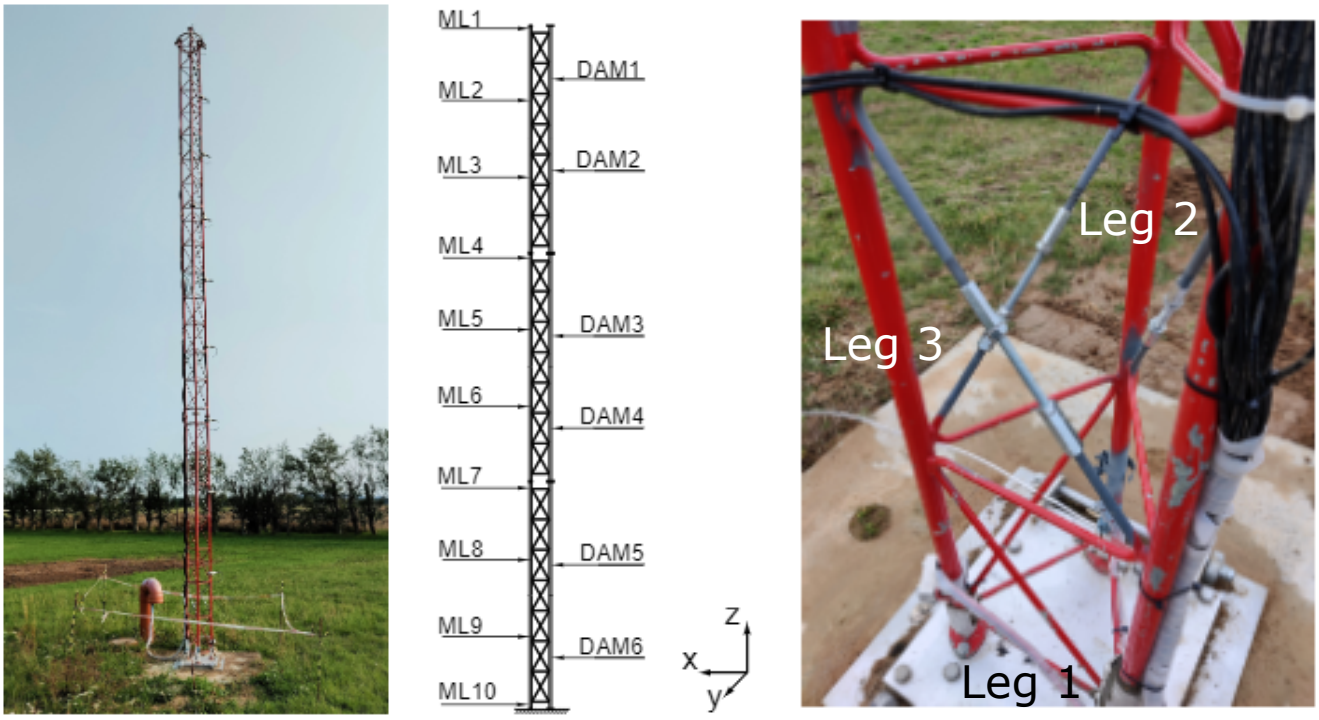

(a) (b) (c)

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 [26].

Table 1: The health states of the structure and the corresponding labels.

Condition of the structure	Label
Healthy	H
DAM6 - all struts removed	D1
Repair state after D1	R1
DAM4 - all struts removed	D2
Repair state after D2	R2
DAM3 - all struts removed	D3
Repair state after D3	R3
DAM6 - strut between legs 2 and 3 removed	D4
Repair state after D4	R4
DAM4 - strut between legs 2 and 3 removed	D5
Repair state after D5	R5
DAM3 - strut between legs 2 and 3 removed	D6
Repair state after D6	R6

(weighting 55 g each) can be removed. Figure 1c presents the bracing level containing the mechanisms that can introduce damage at level 6 (DAM6).

The data used in this paper was obtained during a 10-month monitoring campaign of LUMO, during which, six different damages were introduced. Table 1 details the location and the severity of each damage state and the repair that followed alongside an assigned label. D1, D2 and D3 introduce severe damages to the structure as the damage mechanisms in all three bracing legs are removed, leading to a significant reduction in stiffness at that location. Only one strut is removed during D4 - D6 and therefore represent comparatively smaller damages. Here, the strut between legs two and three are removed (highlighted in Figure 1c).

3.3 The collected data

A network of sensors measure the response of the structure under natural sources of excitation. At measurement levels (ML) 1 - 9 (Figure 1b), two uni-axial accelerometers measure the orthogonal deflection directions (x and y), and captures the motion in the horizontal plane. This results in a total of 18 accelerometers across the structure. ML10 is equipped with three strain gauges on each leg of the structure close to the foundations. The temperature of the structure is also measured using a thermocouple at ML10. All measurements acquired are sampled at a rate of 1651.61 Hz and saved at each 10-minute interval.

Alongside the structural measurements, a meteorological mast positioned 20 m away from the LUMO structure (by the Institute of Meteorology and Climatology (IMUK) of Leibniz University Hannover) measures the air temperature, relative humidity, sum of precipitation, global radiation, average wind speed, maximum wind speed and wind direction in one-minute mean values.

In this paper, the measurements collected via the 18 uni-axial accelerometers at ML1 - ML9 are used for analysis as they capture the dynamic response of the structure. From the time series data, the essential information about the state of the structure are extracted. For this purpose, the eigenfrequencies and mode shapes of the ambiently excitated structure are identified by assuming a linearly time-invariant (LTI) system. These features are useful for SHM because they are widely known to be sensitive to the presence of damage. When considering the LUMO structure, Wernitz et al. [26] demonstrated that the LTI approximation is valid across each 10-minute interval. Therefore, at each 10-minute interval, the eigenfrequencies and mode shapes are calculated from the time-series data collected from the LUMO structure. Next, the methods used for identifying eigenfrequencies and mode shapes through Bayesian operational modal analysis are explained.

3.4 Operational modal analysis for feature extraction

The eigenfrequencies and mode shapes used in this work were identified with frequency domain method Bayesian operational modal analysis (BAYOMA) [27]. The basis of BAYOMA is the discrete Fourier transform (DFT) of

a Gaussian distributed signal, which is statistically independent, and also Gaussian distributed for each frequency base point (given suitably long measurement times and high sampling rates) [28]. When the prior $P(\boldsymbol{\Theta})$ of the modal parameters of m dominating modes in the considered frequency range is equally distributed, the likelihood $P(D|\boldsymbol{\Theta})$ is proportional to the posterior $P(\boldsymbol{\Theta}|D)$, where D is the measured data. The modal parameters $\boldsymbol{\Theta}$ are eigenfrequencies $\mathbf{f}$, modal dampings $\boldsymbol{\zeta}$, mode shapes $\boldsymbol{\phi}_1 ... \boldsymbol{\phi}_m$, modal forces $\mathbf{S}$ and the model error S_e. The likelihood $P(D|\boldsymbol{\Theta})$ is a multivariate Gaussian distribution of the scaled DFT. The related covariance matrix is the theoretical power spectral density matrix of the modal parameters $\boldsymbol{\Theta}$,

$$\mathbf{E}_k(\Theta) = \boldsymbol{\Phi}\mathbf{H}_k(\mathbf{f}, \zeta, \mathbf{S})\boldsymbol{\Phi}^T + S_e\mathbf{I}_n, \tag{3}$$

where $\boldsymbol{\Phi}$ is the matrix of m mode shapes, $\mathbf{H}_k$ is a diagonal matrix containing the m theoretical auto power spectral densities of equivalent one-mass oscillators and $\mathbf{I}_n$ is the identity matrix.

The most probable value (MPV) of the modal parameters for a specified frequency range, and previously defined number of dominant modes, is identified by minimising the negative log likelihood function (NLLF). The associated covariance matrix of the Gaussian approximation of the posterior distribution can be determined using inverse Hessian matrix of the NLLF at the MPV.

Table 2: Intervals of the hard criteria for the mode tracking used to cluster the eigenfrequencies of the LUMO structure throughout time.

Criteria	Values
eigenfrequencies	0.8 f - 1.1 f
min MAC	0.4

In the LUMO dataset, the largest uncertainties of this method occur when identifying the mode shapes of closely spaced modes [29]. In symmetric structures such as the LUMO mast, closely spaced modes are common. When automatically tracking these modes over a long period of time, major changes such as the introduction of damage can also pose a challenge to cluster the modes correctly, i.e., the presence of damage can shift the modes, making it difficult to assign the correct mode to a given eigenfrequency. In addition, due to the sensor setups, the torsional

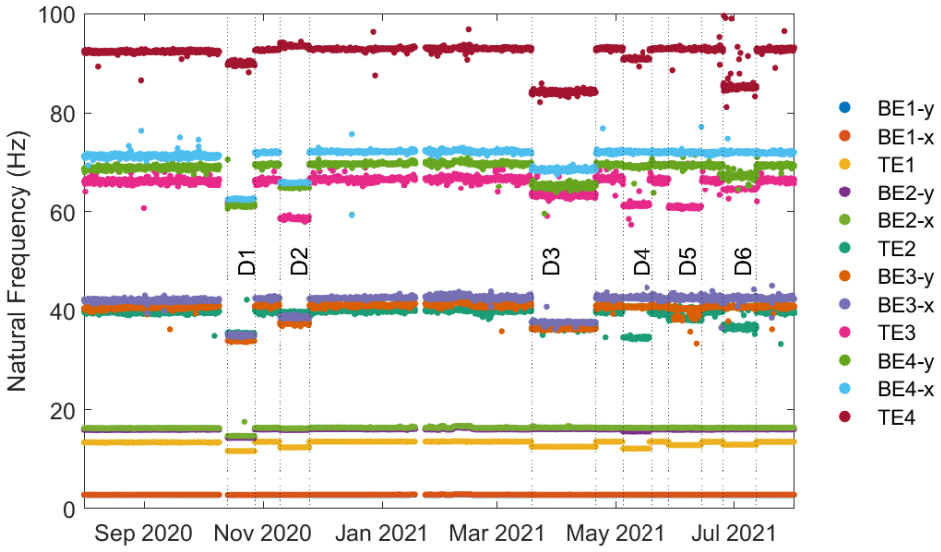

Figure 2: The 12 natural frequencies identified via BAYOMA. The points at which damage occurs is highlighted for clarity. Each datapoint represents 10 minutes of data.

modes are very similar to the bending modes in the y-direction, which further complicates the differentiation between modes. Therefore, the small intervals of the hard criterias for tracking modes across time, as described in [26], cannot be used here. Instead, large intervals (listed in Table 2) are used to cluster modes throughout time. The final mode assignment is allocated according to the best possible match of the modal assurance criterion and the eigenfrequencies. This approach can lead to errors in the mode tracking and must be taken into account in further evaluation. In future works, the mode tracking approach should be improved. However, this is not the scope of this work.

Figure 2 shows all identified natural frequencies in the frequency range between 0 to 93 Hz of the LUMO structure, across the time period used in this work. The corresponding mode shapes can be found in [26]. In Figure 2, the eigenfrequencies associated with bending and torsional modes are represented by BE and TE, respectively, where the dominating directions are also provided. For example, BE1-y represents the eigenfrequency of the first bending mode in the y-direction. Each datapoint represents the eigenfrequencies calculated from data across a 10-minute interval. The introduction of the damages described in Table 1 are also highlighted. It is clear that the natural frequencies of the LUMO structure is sensitive to the introduction of damage. Next, the selection of features from the dataset used in the long-term damage detection system is described.

3.5 Feature selection procedure

The objective of this work is to identify damage before and after repair whilst remaining sensitive to environmental and operational variations that significantly impact the normal dynamic behaviour of the structure. As a result, the eigenfrequencies used in this work should be sensitive to damage, repair and EOVs across time.

The dynamic behaviour of the LUMO structure is prone to closely spaced modes; some eigenfrequencies are found to be very close to one another, as seen in the top panel of Figure 3, where BE3-x, BE3-y and TE2 occur in a relatively small frequency range. These closely-spaced modes can be challenging to distinguish from one another, especially during damage as the eigenfrequencies of the structure shifts. As a result, the clustering technique used to assign modes to the eigenfrequencies can find it challenging to track these modes throughout time. An example of this can be seen in Figure 3 (top panel) during D5 where BE3-y and TE2 have been misclassified by the automatic mode tracking algorithm used in [29]. Large amounts of data scatter can also occur as a result of misclassifications by the mode tracking method, as observed for TE3 in Figure 3 (bottom panel).

The challenges and inaccuracies of mode tracking affects the feature selection process that follows; abundant data scatter or incomplete data can negatively impact monitoring algorithms, which expect the difference between healthy and damage data to be greater than abnormalities within the data in each health state. Otherwise, spurious effects such as data scatter – as a result of pre-processing – could be misidentified as damage, for example.

The features presented in Table 3 are chosen as the input features in this paper, as they are not significantly impacted by the drawbacks of the mode tracking method. Here, the first, second, and fourth natural frequencies (also referred to as bending eigenfrequencies) in both the x- and y-directions (labelled in this paper as BE1-x,

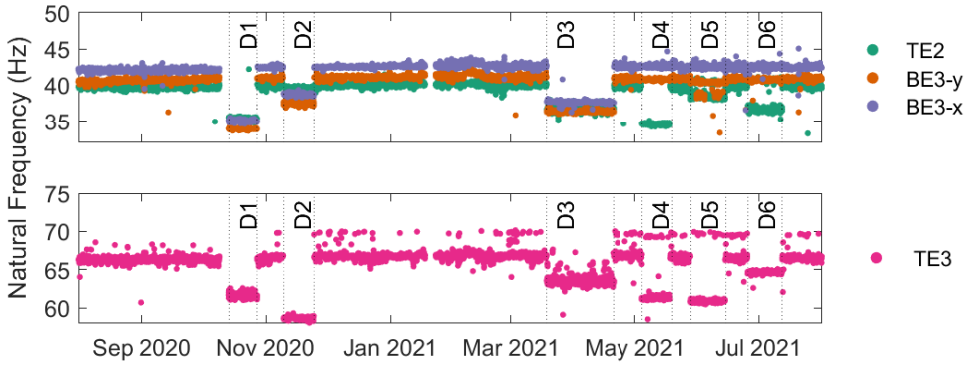

Figure 3: The eigenfrequencies that may have been misclassified by the mode tracking method.

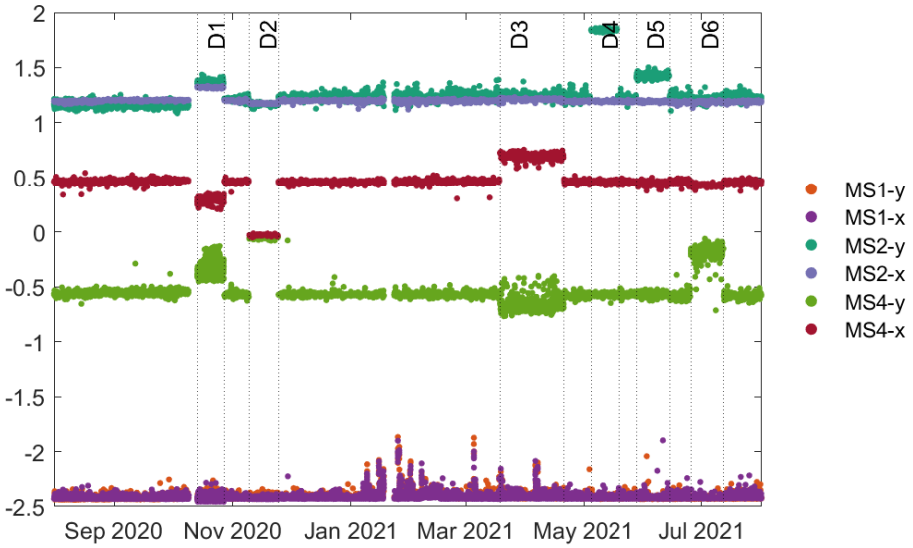

Figure 4: Sum of the mode shapes values at each 10-minutes interval.

BE1-y, BE2-x, BE2-y BE4-x, BE4-y) as well as their corresponding mode shapes (MS1-x, MS1-y, MS2-x, MS2-y MS4-x, MS4-y) are chosen. In order to represent the mode shapes by a single value, the sum of the mode shape is calculated at each 10-minute interval using,

$$\mathrm{MS}_m = \sum_{j=1}^{n} [x_j]_m \tag{4}$$

where MS_m is the value of the mode shape feature at a given 10-minute interval, m is the bending mode, x is the normalised modal amplitude evaluated at n measurement levels. In this paper, $m = \{1\text{-x},1\text{-y},2\text{-x},2\text{-y},4\text{-x},4\text{-y}\}$ and $n = 9$. Figure 4 presents the mode shapes values MS1-x, MS1-y, MS2-x, MS2-y MS4-x, MS4-y chosen as input features in this paper.

Consequently, 12 input features are chosen in total for further analysis in this paper. It is advantageous to utilise the lower modes in analysis as they can be obtained by using fewer sensors in comparison to higher modes; with the current sensor set-up – specifically, the spatial resolution of the sensors across the measurement levels – the properties of higher modes cannot be estimated with a high fidelity. The chosen modes are mostly sensitive to the damage states described in Table 1. The third bending mode is excluded from the input features as BE3-x, BE3-y and TE3 are negatively impacted by the pre-processing method (Figure 3).

All torsional modes (eigenfrequencies and mode shapes) are also excluded during the feature selection process in order to develop an SHM system that is applicable to practical scenarios; although the torsional modes are highly sensitive to the damages D1 - D6 in the LUMO structure, often in operational civil structures, such a bridges and wind turbines, torsional modes do not present a high sensitivity to damage. Additionally, some torsional modes are also negatively affected by the pre-processing method mentioned previously.

3.6 Environmental and operational variations

As previously mentioned, the objective of this paper is to identify damage prior to and post repair. However, a number of environmental and operational variations (EOVs) influence the behaviour of the structure that must also be taken into account. These EOVs are explained here and methods of addressing them are explored throughout the paper.

As the LUMO structure is excited by natural aerodynamic wind excitation, it is exposed to environmental variations such as temperature cycles that occur daily and seasonally. One of the largest influences of environmental

Table 3: The bending modes of the structure and the selected features for analysis. The average natural frequencies of the structure are presented here, which were calculated during normal operating conditions in the healthy state, H [26]. The corresponding labels for the eigenfrequencies and the mode shapes at each of the vibration modes are also stated.

Vibration mode (dominating direction)	Average natural frequency (Hz)	Eigenfrequency label	Mode shape label
1st bending mode (y)	2.76	BE1-y	MS1-y
1st bending mode (x)	2.81	BE1-x	MS1-x
2nd bending mode (y)	15.94	BE2-y	MS2-y
2nd bending mode (x)	16.28	BE2-x	MS2-x
4th bending mode (y)	69.15	BE4-y	MS4-y
4th bending mode (x)	71.51	BE4-x	MS4-x

changes to the dynamics of the structure stems from what is assumed to be stiffness changes during freezing temperatures. It is worth noting that these 'stiffening effects' can also occur in operational engineering structures such as wind turbines. As the stiffness of the structure is affected by the freezing temperatures, its dynamic properties change. It is important to distinguish between these 'stiffening effects' and damage on the structure as the type of necessary intervention can vary depending on the cause; 'stiffening effects' on a wind turbine, for example, may cause it to be temporarily shut-down, whereas damage on a wind turbine will require repairs. Consequently, an SHM method that is sensitive to the start and end of such period can be extremely helpful when managing a structure's operation. In this paper, the data collected during the assumed freezing period is manually assigned the health-state label 'SE' (stiffening effect) by identifying freezing conditions from studying the air temperature and the temperature of the structure. It should be noted that the structure was not examined visually during this time, and therefore the SE labels are assumed only.

Another significant operational variation that affects the structural dynamics of LUMO are the repairs that are undertaken after each damage state. Although the repairs are designed to be reversible, in reality, small permanent changes to the system are to be expected. These changes are considered to be domain shifts in this paper. Next, the way in which the LUMO dataset is separated into source and target domains are explained, in order to undertake domain adaptation to reduce the domain shifts.

3.7 Viewing the LUMO dataset as a population

For the LUMO dataset, there are a number of domains (data generated by different marginal distributions) for each repair (and the original pre-repair); these can be treated as homogeneous structures in a PBSHM framework. Table 4 presents the new source and target structure designations in the population. Here, the assumption is that at the start of each repair state, the health status of the structure is known. This is a reasonable assumption, as a repair would require intervention to the structure and the hope is that the structure is operating under normal conditions immediately after repair; this is a common assumption used in novelty detection.

Table 4: The separation of the LUMO data into source and target domains. Refer to Table 1 for label definitions.

Source	Target 1	Target 2	Target 3	Target 4	Target 5	Target 6
H, D1	R1, D2	R2, D3	R3, D4	R4, D5	R5, D6	R6

In the next section, the repair problem is discussed in detail and methods to address it are explored in Section 5

4 The domain shift in the LUMO structure post-repair

The LUMO structure was designed to allow perfectly reversible damage states in order to introduce and reintroduce damage in a controlled manner, whilst operating in representative environmental conditions. However, slight varia-

tions to the structure are likely to be introduced during repair, leading to changes in the underlying distribution of the collected data post-repair. This effect is also known as a domain shift. Domain shifts as a result of repair can take place in real-world applications for reasons such as, the repair being conducted with different materials to the original structure, the repair process changing the mechanical properties of the structure [30–32]. Even perfectly reversible repair scenarios in controlled, experimental datasets can present domain shifts in the data [3].

Changes in the domain can reduce the effectiveness of standard supervised SHM techniques because domain shifts can lead to significant disparities between training and testing data, for example. Typical practices such as novelty detection trained on healthy data may provide misleading results during testing with repair data, as a result of domain shifts, i.e., repair data could be misdiagnosed as novel. When developing long-term monitoring campaigns, it is important to ensure that false positive flags are avoided where possible (in order to reduce unnecessary inspections), and that damage is detected and not misclassified as healthy.

In this paper, to identify the existence of domain shifts within the LUMO dataset, a distance metric named the maximum mean discrepancy (MMD) is used. The MMD is also used here to demonstrate the reduced effectiveness of standard SHM techniques such as novelty detection that transpire as a result of domain shifts. The MMD is a suitable choice here as it provides a measure of difference between data distributions. It is an intuitive metric that negates the need for density estimation in the presence of incomplete data. The MMD does not assume Gaussianity within the data, making it a flexible metric when the distribution of the testing data are unknown.

The MMD is a distance metric [1] within the integral probability metrics (IPM) family. A brief introduction to the MMD is given in this section. For a more comprehensive understanding, the reader is referred to [34, 35].

Using a kernel (otherwise known as a covariance function) to project data into a higher dimensional space (in this case, the reproducing kernel Hilbert space), the MMD provides a measure of the maximum distance between the mean embedding of two probability distributions. Consider two random variables with observations $X := \{x_1, ..., x_m\}$ and $Y := \{y_1, ..., y_n\}$ (where m and n are the number of data points in X and Y respectively), that have probability densities p and q, respectively. p and q are defined on a metric space and there exists a class of functions that map the metric space to the coordinate space over the real numbers. In order to find out if these distributions are the same, i.e, $p = q$, the MMD looks for a function in this class of functions that maximises the difference in means,

$$\mathrm{MMD}[\mathscr{F}, p, q] := \sup_{f \in \mathscr{F}} \left(\mathbb{E}_{x \sim p}[f(x)] - \mathbb{E}_{y \sim q}[f(y)] \right) \tag{5}$$

where $\mathbb{E}_{x \sim p}[f(x)]$ is the expectation with respect to p. Here, the class of continuous, smooth functions $\mathscr{F}$ is a unit ball in a reproducing kernel Hilbert space (RKHS); there exists a restriction that the norm of the function should be less than or equal to 1 (the term smooth indicates that the function is infinitely differentiable everywhere). The biased estimation of MMD^2 – to ensure that the calculated values are always positive – is given by,

$$\mathrm{MMD}_{\mathrm{b}}^2[\mathscr{F}, X, Y] = \frac{1}{m^2} \sum_{i,j=1}^{m} k(\boldsymbol{x}_i, \boldsymbol{x}_j) + \frac{1}{n^2} \sum_{i,j=1}^{n} k(\boldsymbol{y}_i, \boldsymbol{y}_j) - \frac{2}{mn} \sum_{i,j=1}^{m,n} k(\boldsymbol{x}_i, \boldsymbol{y}_j) \tag{6}$$

In equation (6), the term $\frac{1}{m^2} \sum_{i,j=1}^{m} k(\boldsymbol{x}_i, \boldsymbol{x}_j)$ can be considered to be the average similarity of every pair of samples drawn from distribution p. Term $\frac{1}{n^2} \sum_{i,j=1}^{n} k(\boldsymbol{y}_i, \boldsymbol{y}_j)$ corresponds to the average similarity of every pair of samples drawn from distribution q. The $\frac{2}{mn} \sum_{i,j=1}^{m,n} k(\boldsymbol{x}_i, \boldsymbol{y}_j)$ term is the average similarity of points drawn from p and q. In the case where p and q are dissimilar, i.e. points from p are more similar to each other than they are to points from q, then the first and second terms would produce a large value compared to the third, giving a large value for MMD. On the other hand, where p and q are similar, the first and second terms would produce a value similar to the third, giving a value close to zero for MMD.

The kernel $k(\cdot, \cdot)$ used here is a Gaussian kernel,

$$k(\boldsymbol{a}, \boldsymbol{b}) = \exp\left(- \frac{\| \boldsymbol{a} - \boldsymbol{b} \|^2}{2\sigma^2} \right) \tag{7}$$

[1] A distance metric is defined by Dudley [33] (when $x, y, z \in \mathbb{R} : x, y, z \geq 0$) as a function d that follows the set of rules: 1) The distance between a point and itself is zero, $d(x, x) = 0$, 2) The distance is symmetric: $d(x, y) = d(y, x)$, 3) The triangular inequality holds: $d(x, z) \leq d(x, y) + d(y, z)$, 4) If $d(x, y) = 0$ then this implies that $x = y$.

as it is universal, and continuous in the RKHS [34]. The width or length scale of the kernel is a hyperparameter, denoted here as σ. There is no universal method of choosing the length scale for the Gaussian kernel. Gretton *et. al.* uses the median heuristic, as explained in [34], and defined in [36]. The same method is used in this paper.

The MMD has been implemented in the past for attribute and graph matching [34], in the field of verification and validation [37], for domain adaptation [7, 38], in computer sciences to distinguish malicious users and honest users [39], for training generative adversarial networks [40] and finding similar wear mechanisms in machining tools [41], to name a few. In this paper, the MMD provides a measure of the domain shift observed as a result of damage, repair, and EOVs (stiffening effects) of the LUMO structure in Figure 5. In Figure 5 the MMD values shown are calculated between the healthy (H) state (from the source domain) and the following damage (D), repair (R) (from Table 1) and SE states. This method is equivalent to training a novelty detector on healthy training data and testing it on all data (damage, SE, and repair) that follows.

In order to evaluate whether domain shifts exist within the training and testing data, a threshold is calculated to compare the MMD values. Any data that crosses the threshold can also be considered as novel, i.e., data that crosses the threshold is sufficiently different to the healthy data used for training, suggesting that the structure is no longer similar to the original healthy state. The Monte Carlo (MC) threshold approach for the MMD introduced in this paper closely follows the method used for MC threshold calculation for the MSD described in [42]. Here, the threshold (corresponding to 95% confidence interval determined by MC approach with 1000 samples) is calculated empirically using the witness function f^* of the MMD,

$$f^* \propto \frac{1}{m}\sum_{i=1}^{m} k(\boldsymbol{x}_i, \boldsymbol{t}) - \frac{1}{n}\sum_{i=1}^{n} k(\boldsymbol{y}_i, \boldsymbol{t}) \tag{8}$$

across a variable $\boldsymbol{t}$, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are matrices with m and n observations, respectively. The witness function allows the behaviour of the MMD to be visualised. For the MC approach, in step one, a matrix $\boldsymbol{t}$ comprising the number of observations and the number of dimensions of the dataset is randomly generated from a normal distribution with zero mean and unit variance, i.e., $\mathcal{N}(0, 1)$. Then in step two, two other matrices of the same size ($\boldsymbol{x}$ and $\boldsymbol{y}$) are generated. $\boldsymbol{x}$ is generated by sampling from a normal distribution, $\mathcal{N}(0, 1)$; this distribution represents the normal condition data (healthy and repair in this case). $\boldsymbol{y}$ is sampled from a Student's t distribution with a degree of freedom (ν) of 1 and represents the damage states; the Student's t distribution is a sensible choice here as it contains heavier tails to encompass the spread of damage data. In step three, f^* is calculated using the $\boldsymbol{t}$, $\boldsymbol{x}$ and $\boldsymbol{y}$ matrices following equation (8), where the largest value of f^* is stored. Steps 2 and 3 are then repeated 1000 times and the stored largest values of f^* are sorted in order of magnitude and the value at 95% is used as the threshold, giving a 95% confidence limit. The threshold calculated by MC simulation is included as a horizontal red line in

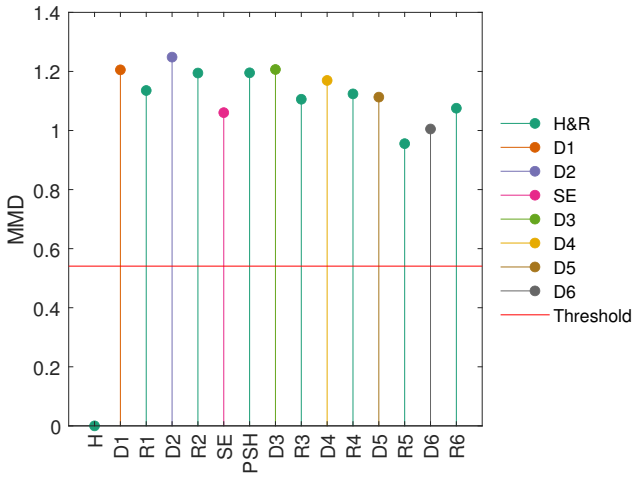

Figure 5: The maximum mean discrepancy between the first heath state and the following states of the LUMO structure. Here, PSH state represent the post stiffness effected healthy state.

Figure 5.

The possible ineffectiveness of standard SHM models (such as novelty detectors) trained on pre-repair data to only detect damage post-repair is evident in Figure 5 – although the damage and SE states have been identified as novel, so have data from the repair states; all post-repair data crosses the pre-determined threshold and will be considered novel. This is a misleading finding because the structure is behaving under normal operating conditions in the repair states, and therefore should not be considered as novel in comparison to the healthy data. Misclassifications such as these can cause unnecessary downtimes which can be extremely costly. In Appendix A, the effect of domain shift on outlier analysis (trained on 30 days of data from healthy state H) is presented to provide an example of novelty detection of the LUMO structure throughout time.

To improve these results, it is possible to retrain this procedure at each new repair state. The approach of retraining models are, however, time consuming, computationally expensive, and discards any damage state labels collected pre-repair (owing to domain shift). Furthermore, the model prior to retraining now becomes irrelevant, leading to waste of information and resources.

Recent techniques developed for transfer learning has proven to be successful in addressing the domain shift post-repair. In the next section, statistically aligning data to improve performance of SHM procedures on the LUMO dataset is explored.

5 Addressing the problem of repair via statistical alignment

It is clear from the previous section that standard novelty detection methods are insufficient post-repair as a result of domain shifts in the data; the classifier used is unable to generalise across the healthy data and the repair data. In this section, a novel approach to addressing the domain shifts caused by structural repairs is proposed by using the statistical alignment methods explained in Section 2.

The repair scenario of the long-term monitoring campaign of the LUMO structure is treated as a partial domain adaptation here, as the assumption is that the only shared class between the source and target domains is the normal condition state, and that the target domain does not contain any classes that are not in the source domain, i.e., $\mathcal{Y}_t \subset \mathcal{Y}_s$ [21]. Therefore, a subset of normal-condition data pre- and post-repair are needed to align the data to reduce the domain shifts that exist within the dataset. In this work, data from the first 30 days of operation in the healthy state H is used to obtain the normal condition statistics of the source structure. Then, the normal condition alignment detailed in Section 2 is applied to the data where the source structure (detailed in Table 4) is first normalised using Equation (1). Next, the data from the first 7 days of operation in each repair state (target structure) is used to obtain the lower order statistics in order to calculate Equation (2). As damage state labels are presumed to be unavailable, the assumption is that the structure is in normal operating conditions immediately after each repair and for the following 7 days. No assumptions regarding the state of the structure or the underlying distributions are made following day 7. As environmental variations may affect the structure adversely during this period, it is advisable to monitor EOVs and only use data collected during normal environmental conditions. This approach of using engineering expertise to select data subsets for alignment is suitable here, as naively aligning data corresponding to different health states can cause negative transfer. Also, this approach will ensure that only data from similar states are aligned.

5.1 Effect of statistical alignment on the LUMO dataset

Once the data are aligned, the healthy and repair states are expected to occupy the same space. This effect can be observed in Figure 6, which presents the relationship between normal-condition aligned eigenfrequencies in the y-direction from Table 3. In this figure, three further observations of interest can be made. Firstly, the SE state (in pink) which represents a 'stiffening affect' across the structure, moves the data in a different direction to the local damages D1-D6 that is a result of reduction in stiffness. Secondly, two sets of groupings are formed; clusters that represent severe damage states (D1 - D3) are grouped close together and clusters D4 - D6 are also grouped together. Thirdly, the patterns that have formed within each group according to damage location are also similar; for example, between BE2-y and BE4-y where damage states D1 and D4 (from location DAM6) are found below the other damage states within that group, then D2 and D5 (from location DAM4) are found slightly above them and

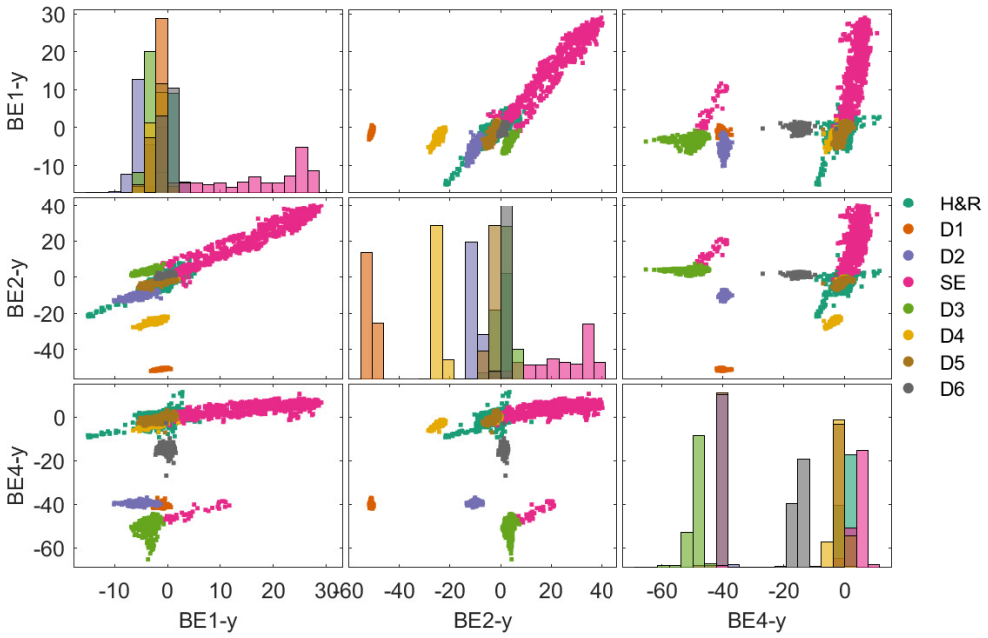

Figure 6: The normal-condition aligned natural frequencies in the y-direction from Table 3 plotted against each other.

finally D3 and D6 (from location DAM3) are found to the left. These results show that statistical alignment has provided physical interpretability of the damage states. The physical interpretability of the classes obtained from statistical alignment can be extremely helpful when validating models, as labels are unavailable in unsupervised domain adaptation. This is a major benefit over nonlinear methods where physical interpretability is not expected.

To highlight the effect of data alignment, Figure 7 presents the MMD values between the healthy state H and the following damage, repair and SE states. Here, Figure 7a presents the results when using non-aligned data (for comparison) and Figure 7b presents the statistically aligned data, following methods described in Section 2.1. These figures were obtained by using every 10^{th} datapoint from each health state to save computational time. As discussed previously, in the non-aligned case, all repair states surpass the predetermined damage threshold owing to the domain shift; a larger distance is found between the healthy and repair states compared to the repair and damage states. As the data are statistically aligned, however, the MMD (the mean embedding) of the repair states drop below the threshold while the SE and damaged states remain above the threshold. These results show that statistical alignment has significantly reduced the distance between the repair and healthy states and increased the distance between repair and damaged states.

It should be noted that given this is a partial domain adaptation problem that only contains the normal condition data as the only shared state between the source and target domains, further adaptation is not possible beyond aligning the correlations, because the normal-condition data are assumed to be Gaussian. Even if the normal conditions are not Gaussian, learning a nonlinear transformation to properly align them may lead to worse generalisation to unseen health-states such as damages (as this is essentially extrapolation of a nonlinear regression function). The non-zero distances between the repair states and the healthy states Figure 7b suggests that the normal-condition data are not entirely Gaussian here. However, it is clear that normal condition alignment has, in this instance, reduced the effect of domain shift significantly, therefore alleviating any need for further adaptation.

The MMD has proved to be a fast and easy procedure to obtain an indication of the similarity of measured data and the effect of domain adaptation methods, post-repair. However, the MMD is a summary metric that represents all datapoints within each class, and does not indicate the novelty of each individual datapoint. For long-term, online monitoring applications, a continuous method of damage detection is, therefore, required. In the next section,

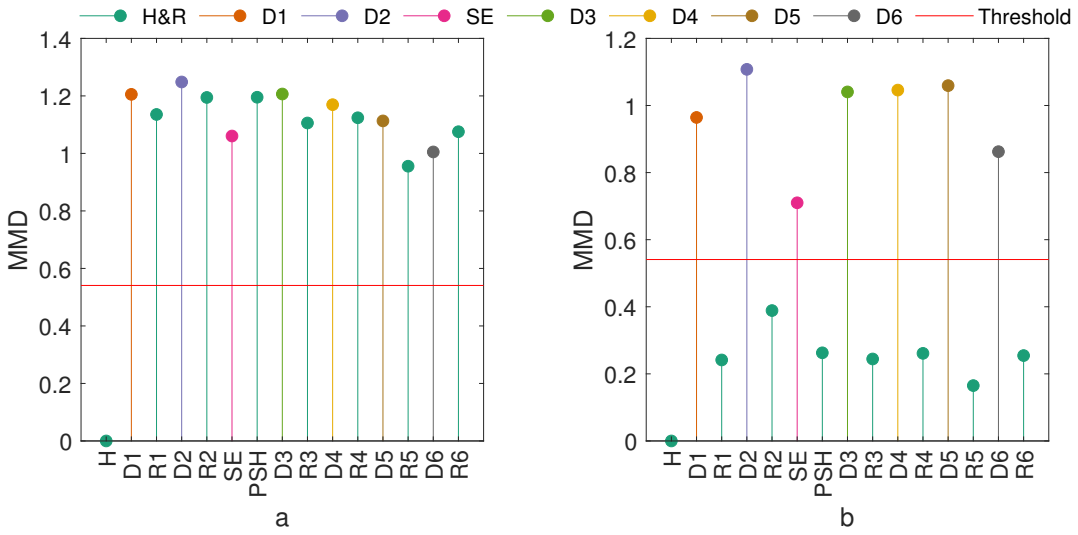

Figure 7: The maximum mean discrepancy between the first heath state and the following states of the LUMO structure when using normal condition aligned data is presented in (b). For comparison purposes, (a) presents the MMD values when using non-aligned data (Figure 5). Here, PSH state represent the post 'stiffening effect' healthy state.

the use of Dirichlet process mixture models (DPMMs) [43] is used to cluster the data in an unsupervised manner, as a method of long-term monitoring of the LUMO structure.

6 Dirichlet process mixture models for damage detection

Dirichlet process mixture models (DPMMs) is a technique that clusters unlabelled data according to similarity. Unlike the novelty detection procedure discussed previously, a major benefit here is that it does not require a training phase based on healthy normal condition data that may not be representative of future observations.

A simplified explanation of the DPMM theory is provided in this section. For a comprehensive understanding the reader is referred to [23] and [43].

The DPMM uses mixture of Gaussian distributions to cluster Gaussian and non-Gaussian data, i.e., it can use an unlimited number of independent Gaussian distributions to represent data of any shape. An explanation of the infinite Gaussian mixture model is given in [23]. The DPMM learns information about the probability of each data point belonging to each cluster. As the LUMO dataset is affected by EOVs, such as stiffening due to freezing temperatures, that influences the shape of the data (as evidenced by the elongated cluster in Figure 6), the ability of the DPMM to cluster Gaussian and non-Gaussian data is valuable.

A collapsed Gibbs sampler is used to make inferences over the DPMM online [44]. The Gibbs sampler can initiate new clusters if new data is sufficiently different to previous clusters, avoiding the need to set the number of expected clusters *a-priori*. The Gibbs sampler therefore allows for an online implementation of the DPMM which is well-suited for long-term monitoring of the LUMO structure, as new damage states are introduced.

A visual explanation of the DPMM and Gibbs sampling is provided using Figures 8a to 8c in sequential order. In these figures, the Gibbs sampler is clustering dataset $\boldsymbol{x}$. The number of possible clusters is K, where $k = 1, ..., K$.

1. Suppose in Figure 8a, that the model has already seen data that is sufficiently different, and has established two Gaussian clusters (green and orange), i.e, $k = 2$.
2. A new datapoint (black square) is then introduced. There is a prior belief that the new datapoint will belong to the prior Gaussian cluster (purple with no data). Therefore, the observations x_i (conditioned on the class label c_i) – where $i = 1, ..., N$, and N is the number of data points – have a Gaussian distribution with mean $\boldsymbol{\mu}_{c_i}$ and covariance Σ_{c_i}. Here, conjugate priors are placed over the mean (Normal distribution) and the covariance

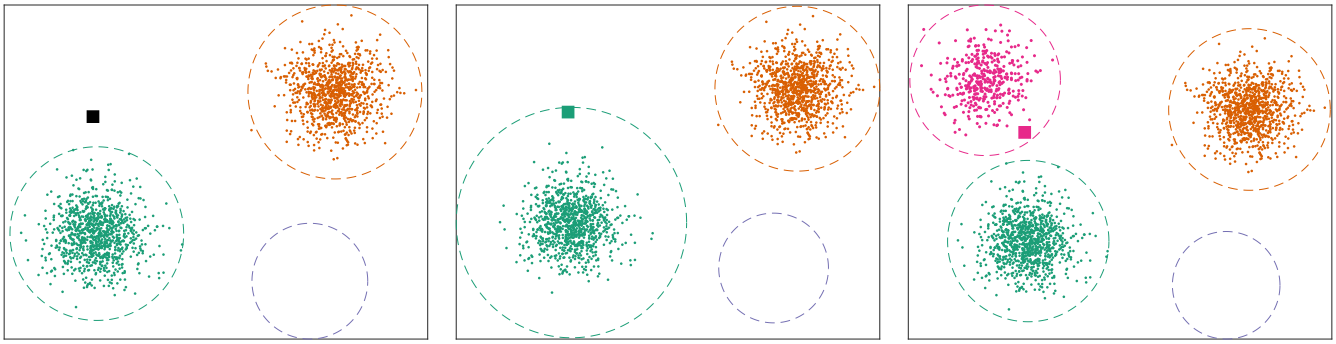

(a) (b) (c)

Figure 8: A visual explanation of the DPMM and Gibbs sampling process where in (a) a new datapoint (black square) is introduced to a model that has two existing classes (green and orange). Then in (b) the square datapoint is assigned to the green cluster based on the likelihood. In (c) new data (pink) forms a new cluster to which the square datapoint is assigned.

(Inverse-Wishart distribution) – with a number of hyperparameters – to achieve a closed form solution,

$$\boldsymbol{x}_i \,|\, c_i \;\sim\; \mathcal{N}(\boldsymbol{x}_i \,|\, \boldsymbol{\mu}_{c_i},\, \Sigma_{c_i}) \tag{9a}$$

$$\boldsymbol{\mu}_{c_i} \,|\, \Sigma_{c_i},\, c_i \;\sim\; \mathcal{N}\left(\boldsymbol{\mu}_{c_i} \,|\, \boldsymbol{\mu}_0,\, \frac{\Sigma_{c_i}}{\kappa_0}\right) \tag{9b}$$

$$\Sigma_{c_i} \,|\, c_i \;\sim\; \mathcal{IW}(\Sigma_{c_i} \,|\, \Sigma_0\,,\, \nu_0) \tag{9c}$$

In this work, the prior distribution has a zero mean and unit variance Gaussian.

3. The likelihood of the current datapoint (in black) belonging to each existing cluster is calculated by assessing their data and shape.
4. A sample label is then assigned to the datapoint based on the calculated likelihoods. In this example, the datapoint is assigned to the green cluster in Figure 8b. The cluster labels c_i are sampled from a multinomial distribution with the mixing proportion, $\boldsymbol{\pi}$, the probability of data belonging to each cluster. In order to calculate these probabilities, the Dirichlet distribution is used as it is the conjugate prior to the multinomial distribution,

$$c_i \,|\, \boldsymbol{\pi} \;\sim\; \text{Mult}(\boldsymbol{\pi}) \tag{10a}$$

$$\boldsymbol{\pi} \;\sim\; \text{Dir}(\alpha) \tag{10b}$$

where $\boldsymbol{\pi}$ is controlled by the strength parameter α, the Dirichlet process prior. The number of clusters identified by the DPMM increases with the value of α. However, if the clusters are well-separated, the number of clusters can become insensitive to α.
5. The Gibbs sampler now re-evaluates all the data in the model by following the process described in Figure 9.
6. Over time, new data arrives (pink) in Figure 8c, which are assigned to a new cluster (so that $k = K + 1$) by the Gibbs sampler. The number of clusters are now updated, $k = 3$.
7. The square datapoint is re-evaluated again by the Gibbs sampler and is assigned a new cluster label (pink) according to the calculated likelihoods.
8. The parameters of all clusters are updated.
9. The entire process is continued until all data are sampled.

The aim of this process is to find the posterior distribution over the cluster labels and parameters, to chose the most likely label. To this end, the collapsed Gibbs sampler sequentially samples new sets of cluster parameters based on samples of labels, and new sets of labels based on parameters. It is a valid Markov chain Monte Carlo method that finds the distribution over the cluster labels and the distribution over the cluster parameters. The

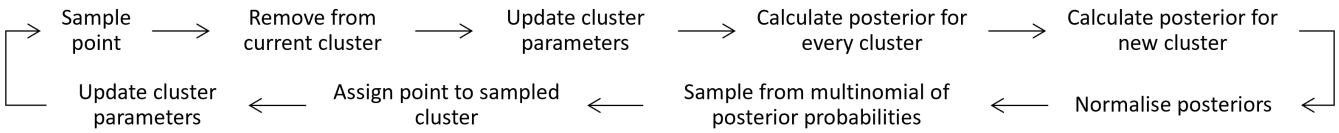

Figure 9: The process of Gibbs sampling. Image reconstructed from [23].

collapsed Gibbs sampler is performed over a window of data in this work, the length of which is specified according to available computation power.

The three main benefits of using this model on an ever-evolving dataset from a test structure such as LUMO are:

- Knowledge of the number of clusters is not required *a-priori* and, therefore, any number of damage states can be included in the model.
- Threshold tuning, and calibration is not required.
- DPMMs permits the covariance function to change with the input data [43], resulting in a model that may be able to handle dissimilarity in the healthy and repair states.

The results presented in Figure 10 show the DPMM clustering of the natural frequencies and their corresponding mode shapes from Table 3. Only the results for the first four input features (BE1-y, BE1-x, MS1-y and MS1-x) are presented for visual clarity. The data is downsampled here to increase computational efficiency; 150 datapoints from each condition described in Table 1 is used in the formulation. Here an $\alpha = 1$ is used (reader is referred to Section 6.2 for parameter choice explanation) to limit the number of clusters formed. For the Gibbs sampler, a window length of 200 is used to save on computational burden as observed for a dataset of similar size [23].

Out of the six clusters identified in Figure 10, five correspond to damage states and one corresponds to the healthy and repair state data. This is a promising start as the DPMM does not find the difference between the repair states and the healthy states to be large in comparison to the damaged states. The DPMM is, however, unable to separate the damage state five and the shift in the features as a result of freezing conditions (Figure 11), because their underlying distributions are similar to the healthy and repair states, i.e., these clusters are not well

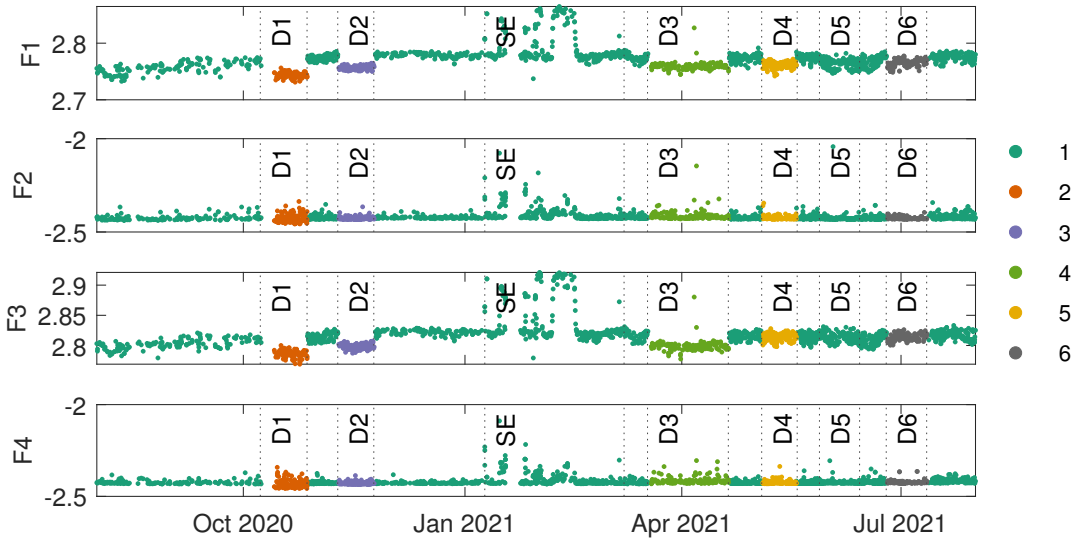

Figure 10: DPMM clustering of features in Table 3. Results for bending mode 1 is displayed here for clarity (F1 = BE1-y, F2 = MS1-y, F3 = BE1-x and F4 = MS1-x) . Each colour represents a cluster. Six clusters are identified by the DPMM.

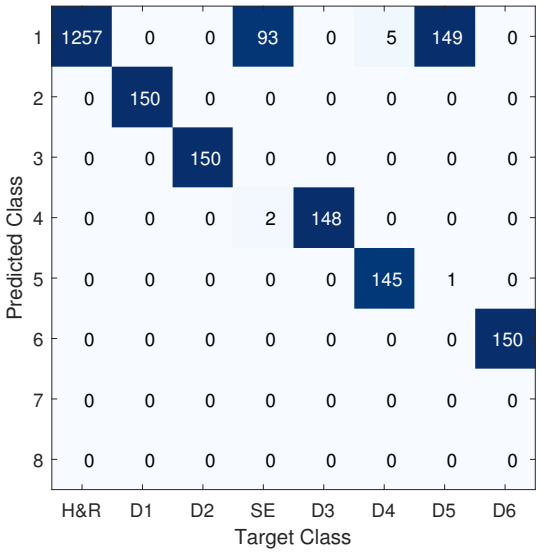

Figure 11: The confusion matrix showing the target class and the resulting predicted class of the DPMM when using 150 datapoints in each class (except for the SE class which contains 95 datapoints). Here, the SE and D5 classes are misclassified to class 1.

separated. Consequently, the true positive rate (TPR) of the DPMM process is 0.75 when using non-aligned data. The TPR is calculated by,

$$\text{TPR} = \frac{\text{TP}}{(\text{TP} + \text{FN})} \tag{11}$$

where TP is the number of damage data that are classified as such, and FN are the number of incorrectly classified damage state data. The DPMM has, however, achieved a FPR of 0. The false positive rate (FPR) is calculated here by,

$$\text{FPR} = \frac{\text{FP}}{(\text{FP} + \text{TN})} \tag{12}$$

where FP is the number of normal condition data that are misclassified, and TN are the number of correctly classified normal condition data.

These results show that when using non-aligned data, even with a sophisticated model such as the DPMM, damage can be misclassified as healthy, which could lead to unsafe operation. Next, the normal condition aligned data are used as inputs to the DPMM in an attempt to reduce the distance between the healthy and repair states, and hopefully further separate the damage and SE states, and in turn, increase the TPR of the model.

6.1 Statistical alignment prior to DPMM clustering

Figures 12 and 13 presents the DPMM results and the confusion matrix (respectively) when using normal condition aligned data as input features. Eight clusters are formed in this case, where all damage states, as well as the SE state, are identified as separate clusters, and the repair and healthy states have been clustered together in a single cluster. The TPR has increased from 75% (non-aligned) to 99% (aligned). The FPR rate has, however, slightly increased from 0% to 1.6% as a result of normal condition alignment.

Normal condition alignment has proved to be a vital step in the unsupervised DPMM clustering procedure for damage detection in the presence of domain shifts. This is especially advantageous to long-term monitoring campaigns, as expensive-to-collect damage state labels are not required to achieve these high classification accuracies. Arguably, more information has been provided to the DPMM by conducting statistical alignment prior to clustering, as some knowledge of the repair states have been included in the data. This alignment approach can be extremely helpful for industrial applications where valuable engineering knowledge should be included in the

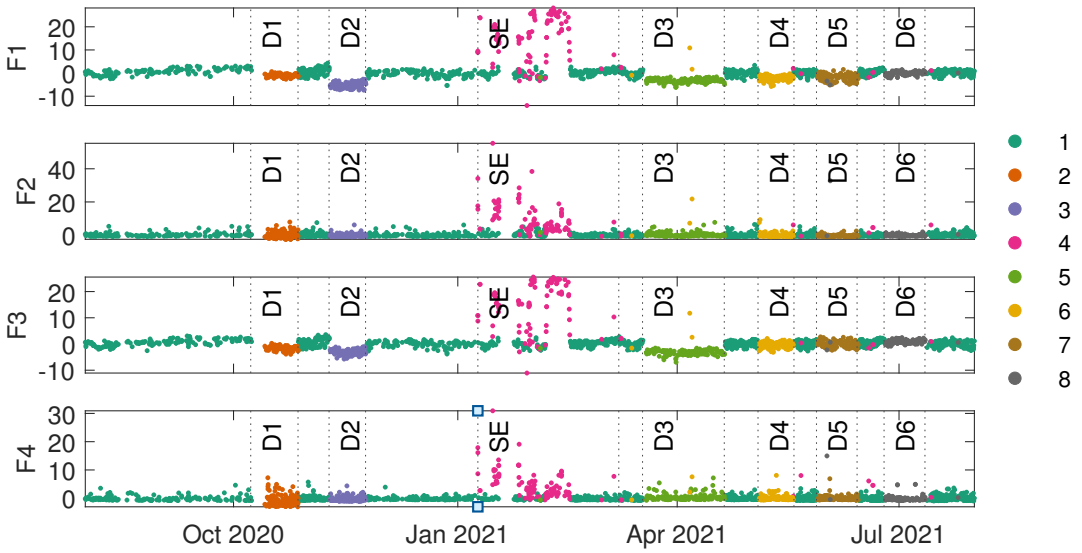

Figure 12: DPMM clustering of normal condition aligned features from Table 3. Results for bending moment 1 is displayed here for clarity (F1 = BE1-y, F2 = MS1-y, F3 = BE1-x and F4 = MS1-x) . Each colour represents a cluster. Eight clusters are identified by the DPMM.

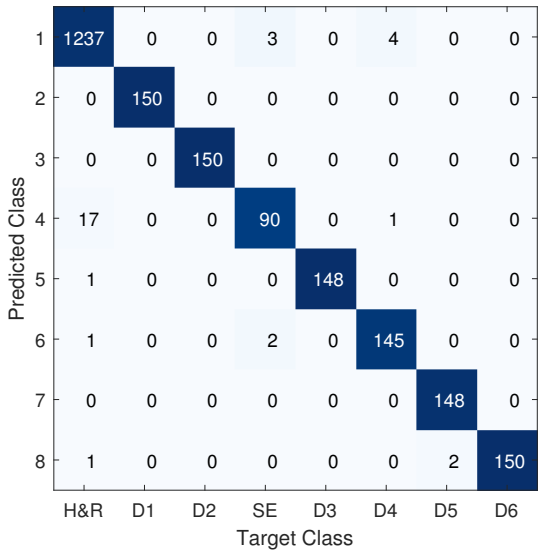

Figure 13: The confusion matrix showing the target class and the resulting predicted class of the DPMM when using statistically aligned data. Here, all damage and SE classes have been identified.

process. Nevertheless, a technique such as active learning could also be used alongside the DPMM to facilitate the inclusion of damage labels, bringing physical meaning to the clusters [45]. Importantly, as the repair and healthy state data are represented by a single cluster, the number of potential false alarms can be reduced and historic data can be retained in the model. It is clear that normal condition alignment has effectively addressed the problem of domain shift post repair.

Figure 12 demonstrates that even relatively small damages can be identified by a combination of DPMMs and normal condition alignment, when using the features described in Table 3. This is a significant advantage because these features represent the bending modes that are widely used in industry, suggesting that this method may be

applicable to other civil structure that collect features which are sensitive to damage.

The method suggested in this paper for long-term online monitoring – using damage sensitive features collected from the structure – is,

1. Initiate the DPMM using suitable hyperparameters. The first cluster will be formed automatically
2. After a suitable period of time, standardise the data in the normal condition (first healthy state in the source domain)
3. Standardise each new datapoint using the parameters of the healthy state
4. If a new cluster is formed, investigate the structure for damages
5. If damage has occurred, repair the damage
6. Take a subset of data immediately after each repair
7. Statistically align the subset post-repair to the healthy state
8. Align each new datapoint to the healthy state
9. Repeat steps 5 to 9

Following brief discussions on the effect of hyperparameters used in the DPMM in the next section, the sensitivity of data to normal condition alignment is then explored.

6.2 The effect of altering parameters of the DPMM

Figures 14a and 14b present the true positive rate (TPR) and true negative rate (TNR) of the DPMM respectively, with varying hyperparameter α and the percentage of datapoints included in the formulation, when using statistically aligned data. The box plots here present the median values (the red line), the 25^{th} and 75^{th} percentiles (the top and bottom of each blue box, respectively), as well as the maximum and minimum values (whiskers) and any outliers (red +) after 100 repeats of each combination. The TNR is given by,

$$\text{TNR} = \frac{\text{TN}}{(\text{TN} + \text{FP})} \tag{13}$$

where FP is the number of normal condition data that are misclassified, and TN are the number of correctly classified normal condition data.

It is clear that both the TPR and TNR do not change significantly (the rate averages do not drop below 96%) as the variables are altered, suggesting that the current formulation of the DPMM is well suited to the repair problem explored in this paper. The figures show that the number of datapoints in the model has a larger influence on the results in compassion to the value of α, which does not have a significant affect. The insensitivity to α is a sign that the healthy and repair states are well separated from the damage states when using normal-condition aligned data. Increasing the number of datapoints in the model increases the average TPR whilst reducing the variance. The opposite effect is observed when considering the average TNR. This result suggests that as the number of datapoints is increased, the model gets better at detecting damage but simultaneously misclassifies more healthy data. A closer inspection finds that the healthy and repair data is grouped into the cluster containing the SE state data more readily, as the number of datapoints in the model are increased. Given there is some overlap between these classes as seen in Figure 6, this result is not surprising.

The parameters tested here do, however, affect the number of clusters initiated by the DPMM as seen in Figure 15. Consequently, the number of health states that exist does not always directly correspond to the number of clusters initiated by the algorithm. However, from Figures 14a and 14b, it is clear that the model is able to distinguish the difference between healthy/repair states and damage. Given the false positive rates are low, this suggests that some damage states initiate more than one cluster or that some clusters contain data from more than one damage state. For the purpose of damage detection when considering the repair problem, the number of cluster initiations is, therefore, not as important as the rate of true positives and negatives. On the other hand, if the cluster initiations are used as a detection threshold as suggested in [23], using small values of alpha in practice can limit the number of clusters initiated.

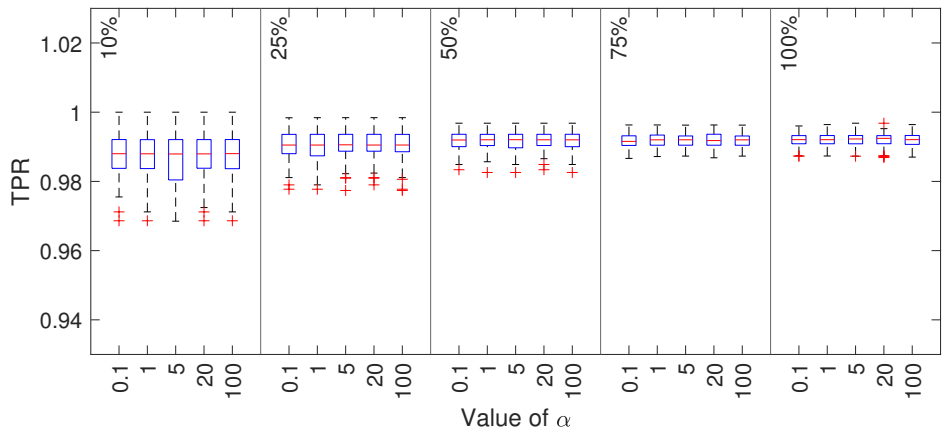

(a) TPR

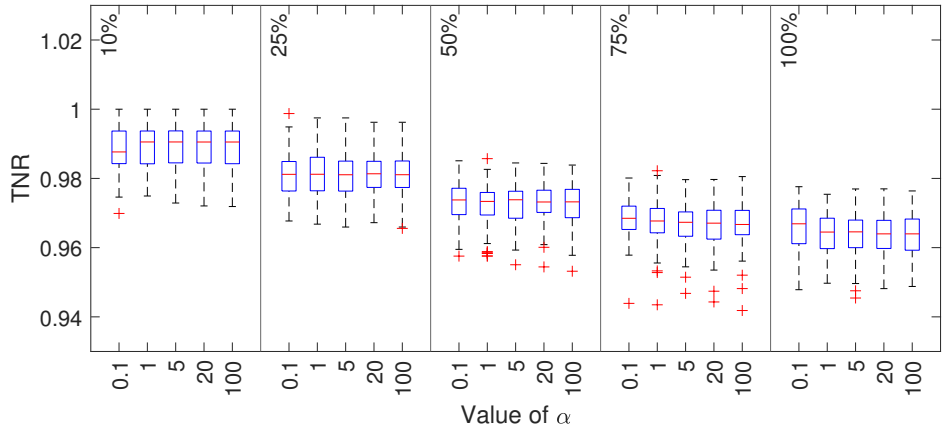

(b) TNR

Figure 14: The effect altering the value of α and the percentage of datapoints included in the model on the (a) true positive rate and the (b) true negative rate of the DPMM when using normal-condition aligned data. Here, the results represent 100 repeats conducted in each state.

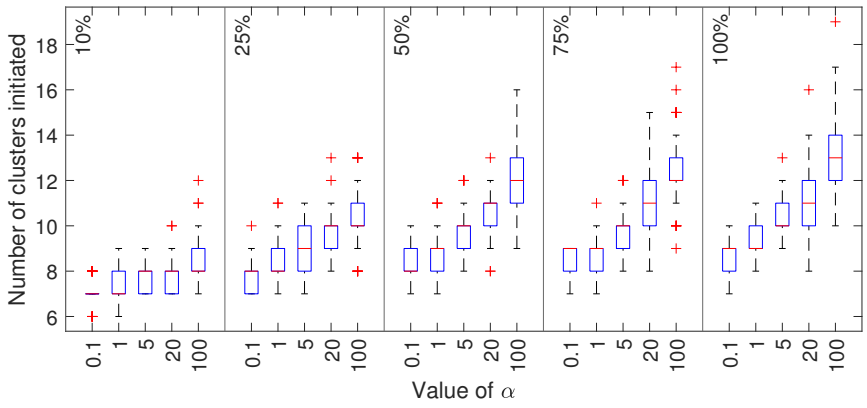

Figure 15: The effect altering the value of α and the percentage of datapoints included in the DPMM model on the number of initiated clusters.

6.3 Sensitivity of normal condition alignment to number of data points

Normal condition alignment has allowed the TPR of Dirchlet process clustering results to improve significantly. In this section, the sensitivity of the normal condition alignment to the number of data points required will be assessed. By doing so, it is possible to evaluate how many data points are required to obtain a sufficient classification accuracy for the LUMO dataset. Consequently, a time window within each repair state can be established for normal condition alignment.

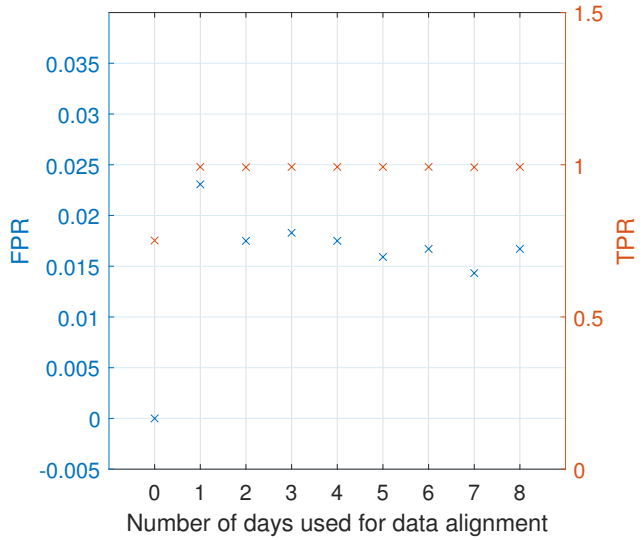

Figure 16: The performance of the DPMM clustering method with respect to the number datapoints used in the normal-condition alignment.

Figure 16 presents the FPR and TPR results of the Dirichlet process mixture model. Here, the number of days used for aligning the normal condition is increased from one to eight. The largest window was set to eight days as that is the duration of the smallest repair state. For comparison, day zero presents the FPR and TPR when using non-aligned data. In Figure 16, the FPR and TPR are affected as the data are aligned using the normal condition. As the number of days are increased from one to two, the lower order statistical alignment allows the damage state data to move further away from the repair states, compared to the distance between the repair states to the healthy state. Increasing the number of days for alignment does not significantly impact the FPR as the lower order statistics are relatively consistent throughout.

These results not only show that the combination of statistical alignment and DPMM clustering are well suited for the long-term monitoring of the mast structure, but also that the amount of data required for alignment is very small. However, as discussed previously, care must be taken when choosing the data for alignment as they should represent only the normal condition without anomalous behaviours.

In the next section, the effect of further reducing the number of input features is investigated in order to minimise the costs associated with sensor installation, data collection and storage.

6.4 The effect of reducing the features of the DPMM

In SHM campaigns, the cost, labour, and storage capacity requirements associated with measurement sensors can be substantial to owners. Unfortunately, for fault analysis (and other techniques such as modal analysis), a network of sensors is often necessary. For instance, the analysis conducted in the previous sections uses data from all 18 uni-axial accelerometers in the form of mode shapes and natural frequencies. In this section, the performance of the DPMM is assessed against the number of input features to the model, as the number of input features is correlated to the number of sensors required for analysis. For example, removing the mode shapes from the input features reduces the spacial resolution required from the sensor network (to determine the mode shapes), hence reducing the

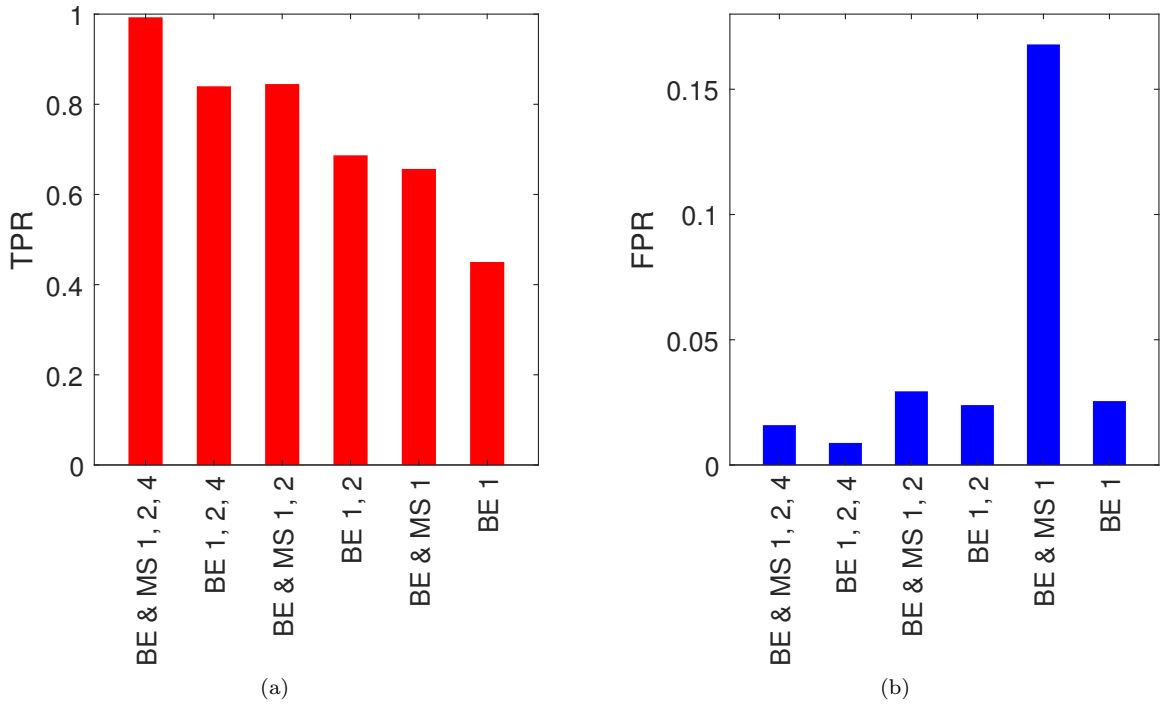

(a) (b)

Figure 17: The DPMM performance as the number of features are reduced from analysis. (a) The true positive rate showing the amount of correctly labelled damaged and SE data and (b) The false positive rate showing the amount of mislabelled healthy and repair data. The input features that are included in the model are described in Table 5.

number of necessary sensors required for analysis. By studying the true positive rate and the false positive rate of the DPMM, this section explored the optimal feature requirements necessary for a classifier to detect damage (and EOVs such as SE) whist remaining insensitive to healthy and repair data.

The performance of the DPMM as the number of input features (described in Table 5) are reduced is studied in Figure 17. Here Figure 17a shows the behaviour of the true positive rate of the DPMM as the input features are reduced. This plot is, therefore, an indication of how well the model is able to detect data related to damage or EOVs such as SE. It is clear that as the input features are removed from analysis – specifically, as the higher eigenfrequencies and mode shapes are removed – the model's ability to detect damage and SE data reduces. This is not a surprising result because smaller damages are usually more sensitive to higher modes. It should be noted,

Table 5: The input features included in the DPMM and the corresponding performance of the model. The labels of each input feature combination used in Figure 17 is also presented. Here, each bending eigenfrequency and mode shape contains components in the x and y directions.

Number of input features	Bending eigenfrequencies (BE)	Mode shapes (MS) included	Label	TPR	FPR
12	1, 2 and 4	Yes	BE & MS 1, 2, 4	0.99	0.016
6	1, 2 and 4	No	BE 1, 2, 4	0.84	0.009
8	1 and 2	Yes	BE & MS 1, 2	0.84	0.029
4	1 and 2	No	BE 1, 2	0.69	0.024
4	1	Yes	BE & MS 1	0.66	0.168
2	1	No	BE 1	0.45	0.026

however, that the first two bending eigenfrequencies and mode shapes (BE & MS 1, 2) capture all damages and the SE state except for D6, leading to a TPR of 0.84 (and a small FPR of 0.029). As the first two bending modes occurs at low frequencies, this result shows that the sampling frequencies and the number of accelerometers can be reduced from the LUMO structure, whilst achieving a relatively well separated healthy and damage states, when using statistically aligned data. This result is a helpful argument for lowering installation, storage and computational costs associated with sensors.

Interestingly, removing the mode shapes within a given combination has a similar effect to removing a higher bending eigenfrequency altogether (for example, the TPR of combination BE 1, 2, 4 is the same as the combination BE & MS 1, 2). Given, technically, the eigenfrequencies of a structure can be calculated with fewer sensors than the mode shapes (calculating mode shapes require a network of well placed sensors, whereas a single sensor can be used to calculate the eigenfrequencies with a high resolution), this may be a helpful finding, as costs of sensor installation could be reduced by focusing solely on bending eigenfrequencies for input features. In the LUMO structure, however, the sensor resolution is not sufficient to determine the higher modes with adequate fidelity, and as a result, it is not possible to assess whether the input feature combination BE 1, 2, 4, 5 would give a similar TPR as using features BE & MS 1, 2, 4.

Figure 17b presents the false positive rate of the DPMM as the input features are reduced. The number of misclassified normal condition healthy or repair data does not seem to vary with the input features in any meaningful way, i.e., there is no indication of a clear correlation. In general, however, reducing the number of features does not negatively impact the FPR, which remains relatively small.

7 Conclusions and Future work

During their lifetimes, the behaviour of civil structures can be affected by environmental and operational variations (such as freezing conditions), defects and damages, and subsequent structural repairs, that makes long-term data-driven monitoring challenging. These variables often change the underlying distributions of the data collected from the structure, thus introducing domain shifts between the healthy data the model was trained on, and the testing data where the structure is in operation. Disparities between training and testing data can significantly reduce the model performance or even invalidate the model entirely.

This paper proposed a method for long-term damage detection of a mast structure under natural excitation, that can adapt to domain shifts caused by structural repairs, or environmental variations such as stiffening effects. A novel method was proposed here that combines domain adaptation techniques and fully-online clustering methods (DPMM) that do not require a training phase or damage labels. First, the domain shifts within the normal condition data were reduced using normal condition alignment by considering a partial domain adaptation approach – where the target domain is a subspace of the source domain, and the only shared class between the two domains are the normal condition. Then, Dirichlet process mixture models were used to cluster the data online, because they can automatically increase the number of clusters as new data affected by structural changes – due to damage and EOVs such as stiffening effects – are introduced. By using this novel combination of methods, this paper demonstrated that it is possible to achieve a true positive rate of 0.99 and false positive rates of 0.016 when detecting damage on a dataset from a mast structure.

As the statistical alignment has provided good physical interpretability of the damage information, the next steps of this work is to apply these findings in a population-based framework via domain adaptation in order to achieve tasks in level two of the Rytter's hierarchy: damage location.

Acknowledgements

We gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - SFB-1463 - 434502799. CW 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.

References

[1] C. R. Farrar and K. Worden, *Structural health monitoring: a machine learning perspective*. John Wiley & Sons, 2012.

[2] K. Worden and J. M. Dulieu-Barton, "An overview of intelligent fault detection in systems and structures," *Structural Health Monitoring*, vol. 3, pp. 85–98, 2004.

[3] P. Gardner, L. Bull, N. Dervilis, and K. Worden, "Overcoming the problem of repair in structural health monitoring: Metric-informed transfer learning," *Journal of Sound and Vibration*, vol. 510, p. 116245, 2021.

[4] P. Gardner, L. A. Bull, N. Dervilis, and K. Worden, "Domain-adapted gaussian mixture models for population-based structural health monitoring," *Journal of Civil Structural Health Monitoring*, pp. 1–11, 2022.

[5] J. Poole, P. Gardner, N. Dervilis, L. Bull, and K. Worden, "On statistic alignment for domain adaptation in structural health monitoring," *Structural Health Monitoring*, pp. 1–20, 2022.

[6] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, pp. 1345–1359, 2009.

[7] P. Gardner, X. Liu, and K. Worden, "On the application of domain adaptation in structural health monitoring," *Mechanical Systems and Signal Processing*, vol. 138, p. 106550, 2020.

[8] P. Gardner, L. Bull, J. Gosliga, N. Dervilis, and K. Worden, "Foundations of population-based SHM, Part III: Heterogeneous populations–mapping and transfer," *Mechanical Systems and Signal Processing*, vol. 149, p. 107142, 2021.

[9] G. Csurka, "A comprehensive survey on domain adaptation for visual applications," *Domain adaptation in computer vision applications*, pp. 1–35, 2017.

[10] S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira, "Analysis of representations for domain adaptation," *Advances in neural information processing systems*, vol. 19, 2006.

[11] X. Wang and F. Liu, "Triplet loss guided adversarial domain adaptation for bearing fault diagnosis," *Sensors*, vol. 20, p. 320, 2020.

[12] X. Li, W. Zhang, Q. Ding, and J.-Q. Sun, "Multi-layer domain adaptation method for rolling bearing fault diagnosis," *Signal processing*, vol. 157, pp. 180–197, 2019.

[13] Y. Li, Y. Song, L. Jia, S. Gao, Q. Li, and M. Qiu, "Intelligent fault diagnosis by fusing domain adversarial training and maximum mean discrepancy via ensemble learning," *IEEE Transactions on Industrial Informatics*, vol. 17, pp. 2833–2841, 2020.

[14] J. Jiao, M. Zhao, J. Lin, and K. Liang, "Residual joint adaptation adversarial network for intelligent transfer fault diagnosis," *Mechanical Systems and Signal Processing*, vol. 145, p. 106962, 2020.

[15] Y. Ding, M. Jia, and Y. Cao, "Remaining useful life estimation under multiple operating conditions via deep subdomain adaptation," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–11, 2021.

[16] L. Bull, P. Gardner, N. Dervilis, E. Papatheou, M. Haywood-Alexander, R. Mills, and K. Worden, "On the transfer of damage detectors between structures: an experimental case study," *Journal of Sound and Vibration*, vol. 501, p. 116072, 2021.

[17] P. Gardner, L. Bull, J. Gosliga, J. Poole, N. Dervilis, and K. Worden, "A population-based shm methodology for heterogeneous structures: Transferring damage localisation knowledge between different aircraft wings," *Mechanical Systems and Signal Processing*, vol. 172, p. 108918, 2022.

[18] G. Michau and O. Fink, "Domain adaptation for one-class classification: monitoring the health of critical systems under limited information," *In press in the International journal of prognostics and health management, deep learning and emerging analytics special issue*, 2019.

[19] S. Xu and H. Y. Noh, "Phymdan: Physics-informed knowledge transfer between buildings for seismic damage diagnosis through adversarial learning," *Mechanical Systems and Signal Processing*, vol. 151, p. 107374, 2021.

[20] M. Long, J. Wang, G. Ding, J. Sun, and P. S. Yu, "Transfer feature learning with joint distribution adaptation," in *Proceedings of the IEEE international conference on computer vision*, 2013, pp. 2200–2207.

[21] Z. Cao, M. Long, J. Wang, and M. I. Jordan, "Partial transfer learning with selective adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2724–2732.

[22] C. T. Wickramarachchi, J. Poole, E. J. Cross, and K. Worden, "On aspects of geometry in shm and population-based shm," in *Data Science in Engineering, Volume 9*, 2022, pp. 67–77.

[23] T. Rogers, K. Worden, R. Fuentes, N. Dervilis, U. Tygesen, and E. Cross, "A bayesian non-parametric clustering approach for semi-supervised structural health monitoring," *Mechanical Systems and Signal Processing*, vol. 119, pp. 100–119, 2019.

[24] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, pp. 43–76, 2020.

[25] S. Wernitz, B. Hofmeister, C. Jonscher, T. Grießmann, and R. Rolfes, "Dataset: Lumo - leibniz universtity test structure for monitoring," 2021, https://doi.org/10.25835/0027803.

[26] S. Wernitz, B. Hofmeister, C. Jonscher, T. Grießmann, and R. Rolfes, "A new open–database benchmark structure for vibration–based structural health monitoring," *Structural Control and Health Monitoring*, vol. 29, p. e2660, 2022.

[27] S.-K. Au, F.-L. Zhang, and Y.-C. Ni, "Bayesian operational modal analysis: Theory, computation, practice," *Computers & Structures*, vol. 126, pp. 3–14, 2013.

[28] K.-V. Yuen, L. S. Katafygiotis, and J. L. Beck, "Spectral density estimation of stochastic vector processes," *Probabilistic Engineering Mechanics*, vol. 17, pp. 265–272, 2002.

[29] C. Jonscher, B. Hofmeister, T. Grießmann, and R. Rolfes, "Influence of environmental conditions and damage on closely spaced modes," in *European Workshop on Structural Health Monitoring*. Springer International Publishing and Imprint Springer, 2023, vol. 270, pp. 902–911.

[30] M. Brown and F. Proschan, "Imperfect repair," *Journal of Applied Probability*, vol. 20, pp. 851–859, 1983.

[31] T. Moan, "Reliability-based management of inspection, maintenance and repair of offshore structures," *Structure and Infrastructure Engineering*, vol. 1, pp. 33–62, 2005.

[32] D. Radford, D. Van Goethem, R. Gutkowski, and M. Peterson, "Composite repair of timber structures," *Construction and Building Materials*, vol. 16, pp. 417–425, 2002.

[33] R. Dudley, "Real analysis and probability , cambridge uni," *Press, Cambridge, UK*, vol. 68, 2002.

[34] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *The Journal of Machine Learning Research*, vol. 13, pp. 723–773, 2012.

[35] A. Gretton, "Introduction to RKHS, and some simple kernel algorithms," pp. 1–33, 2016.

[36] D. Garreau, W. Jitkrittum, and M. Kanagawa, "Large sample analysis of the median heuristic," *arXiv preprint arXiv:1707.07269*, 2017.

[37] P. Gardner, C. Lord, and R. J. Barthorpe, "A unifying framework for probabilistic validation metrics," *Journal of Verification, Validation and Uncertainty Quantification*, vol. 4, 2019.

[38] Y. Chen, S. Song, S. Li, and C. Wu, "A graph embedding framework for Maximum mean discrepancy-based domain adaptation algorithms," *IEEE Transactions on Image Processing*, vol. 29, pp. 199–213, 2020.

[39] S. Yuan, L. Li, and C. Chigan, "Maximum mean discrepancy based secure fusion strategy for robust cooperative spectrum sensing," in *IEEE International Conference on Communications*, 2018.

[40] G.K. Dziugaite, D.M. Roy, and Z. Ghahramani, "Training generative neural networks via maximum mean discrepancy optimization," in *Uncertainty in Artificial Intelligence - Proceedings of the 31st Conference, UAI 2015*, 2015, pp. 258–267.

[41] C. T. Wickramarachchi, W. Leahy, K. Worden, and E. J. Cross, "On metrics assessing the information content of datasets for population-based structural health monitoring," in *European Workshop on Structural Health Monitoring*, 2020, pp. 494–504.

[42] N. Dervilis, E. Cross, R. Barthorpe, and K. Worden, "Robust methods of inclusive outlier analysis for structural health monitoring," *Journal of Sound and Vibration*, vol. 333, no. 20, pp. 5181–5195, 2014.

[43] C. Rasmussen, "The infinite gaussian mixture model," *Advances in neural information processing systems*, vol. 12, 1999.

[44] R. M. Neal, "Markov chain sampling methods for dirichlet process mixture models," *Journal of computational and graphical statistics*, vol. 9, pp. 249–265, 2000.

[45] L. A. Bull, N. Dervilis, K. Worden, E. J. Cross, and T. J. Rogers, "A sampling-based approach for information-theoretic inspection management," *Proceedings of the Royal Society A*, vol. 478, p. 20210790, 2022.

[46] K. Worden, G. Manson, and N. R. Fieller, "Damage detection using outlier analysis," *Journal of Sound and vibration*, vol. 229, pp. 647–667, 2000.

A Outlier analysis of the LUMO dataset

In order to demonstrate the ineffectiveness of conventional long-term SHM methods when considering repair, outlier analysis (first introduced for damage detection in [46]) is used in this appendix as a demonstration. Here, the eigenfrequencies of the LUMO dataset (Table 3) are used as input features. To perform multivariate outlier analysis, first the Mahalanobis squared distance (MSD) is calculated on a sample of the data (training set). Then, during the testing procedure, the MSD of each new datapoint is assessed using a threshold corresponding to 95% confidence interval determined by Monte Carlo (MC) simulation with 1000 samples [46]. The data that crosses the threshold are presumed to be novel in comparison to the data used to train the model.

The MSD is a distance metric and is defined as,

$$D_\zeta = (\boldsymbol{x}_\zeta - \bar{\boldsymbol{x}})^{\mathrm{T}} \boldsymbol{S}^{-1} (\boldsymbol{x}_\zeta - \bar{\boldsymbol{x}}) \tag{A.1}$$

where $\boldsymbol{x}_\zeta$ is the current data point (dimension $1 \times n$), $\bar{\boldsymbol{x}}$ is the mean vector of the training data (dimension $1 \times n$) and $\boldsymbol{S}$ is the covariance matrix (dimension $n \times n$), where n is the number of features.

The results of the multivariate outlier analysis performed on the features in Table 3 are presented in Figure A.1. The data used for training is displayed in blue whilst the remaining healthy and repair (H&R) data is presented in green. The model is trained using data from the first month of operation where a healthy, normal condition is assumed. This assumption is sensible as damage is unlikely to prevail during the first month. In real-world applications, the validity of the assumed normal condition can be ensured by inspection of the structure. The damage states (D1 - D6) and the SE state are also highlighted manually in this figure for clarity. The threshold calculated by MC simulation is included as a horizontal red line. The ineffectiveness of standard SHM models trained on pre-repair data to only detect damage post-repair is clearly evident in Figure A.1 – although the damage and SE states have been identified as outliers, so have data from the repair states. As a result, this model is insufficient in practice, as it is prone to false positives at a rate of 0.55.

To assess the effect of harmonising lower order statistics according to normal condition, the outlier analysis procedure is repeated. As the MSD is defined completely by the correlation (Equation (A.1)), by using statistical alignment, the correlations can be aligned directly, leading to a higher confidence in the ability of the model to generalise. The results of the outlier analysis following statistical alignment are presented in Figure A.2. Data alignment has enabled the majority of repair state data to move below the threshold whilst the SE and damage state data remain above it. The FPR has dropped from 0.55 pre-alignment to 0.22 post alignment, demonstrating that data alignment has helped improve the outlier analysis results. This method does not involve retraining the outlier analysis model at each repair state, and allows the model to retain all the information collected up to the current date. However the FPR is still high because there are drifts in the data during the testing period as a result of environmental and operational variations.

These results express the need for a more robust method of damage detection when considering repair, which has been demonstrated by the combination of Dirichlet process mixture models and statistical alignment techniques in Section 6.1.

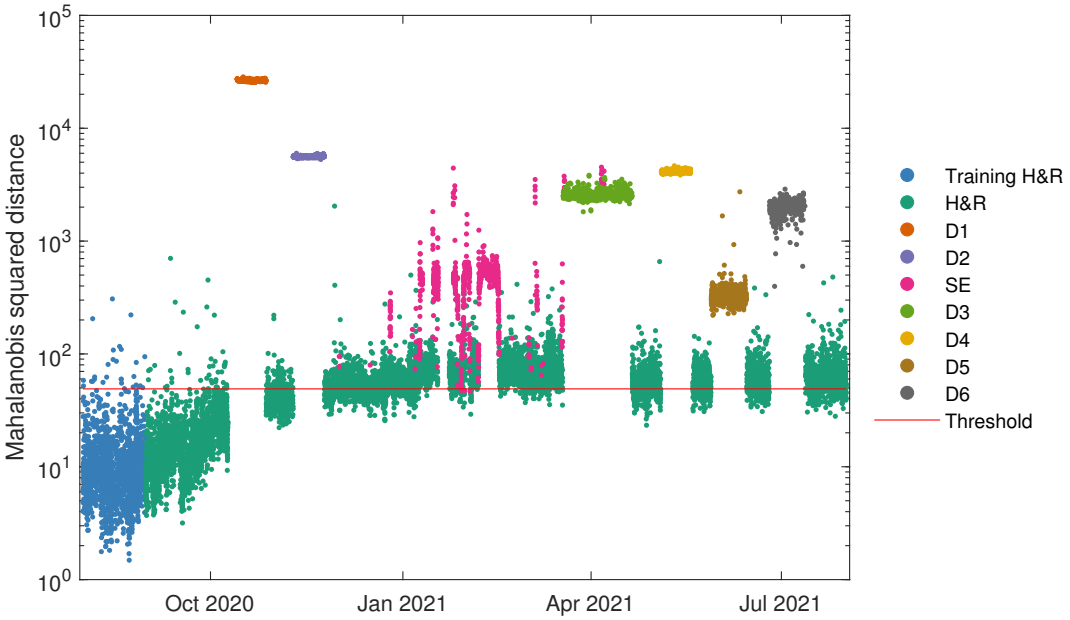

Figure A.1: Outlier analysis results when using the first month of data for training.

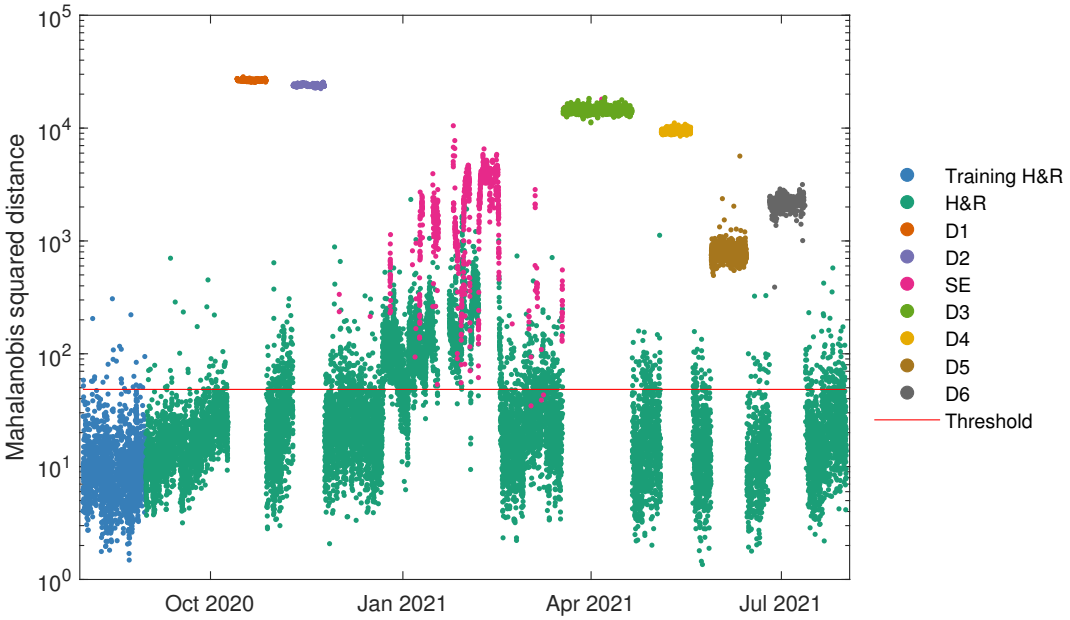

Figure A.2: Outlier analysis on the harmonised data when using the first month of data for training.