Article:

Giglioni, V., Poole, J. orcid.org/0000-0002-7642-9108, Venanzi, I. et al. (2 more authors) (2023) On the use of domain adaptation techniques for bridge damage detection in a changing environment. ce/papers, 6 (5). pp. 975-980. ISSN 2509-7075

https://doi.org/10.1002/cepa.2143

Proceedings
in civil engineering

Ernst & Sohn
A Wiley Brand

EUROSTRUCT

ORIGINAL

On the use of domain adaptation techniques for bridge damage detection in a changing environment

Valentina Giglioni[1] | Jack Poole[2] | Ilaria Venanzi[1] | Filippo Ubertini[1] | Keith Worden[2]

state labels to make inferences on an unlabeled monitored structure. The idea is to exploit source data to train a Machine Learning algorithm and achieve improved early-stage damage detection capabilities across a population of bridges. To account for differences in the underlying distributions of each structure, Transfer Learning is seen as a strategy enabling population-level bridge SHM. In this paper, the natural frequencies obtained from multiple vibration measurements are extracted to characterise different domains during pristine and abnormal conditions. Such damage-sensitive features are aligned via Domain Adaptation and used to train a standard classifier within a shared feature space. The methodology is validated on the heterogeneous population composed of the Z24 and S101 bridges. The results prove the effectiveness to successfully exchange damage labels, thus increasing available information for health-state inference for SHM applications with sparce datasets.

Keywords

Bridge damage detection, Transfer learning, Domain Adaptation, Population-based Structural Health Monitoring

1 Introduction

Given the broad consensus to invest in both asset management and infrastructure maintenance, Structural Health Monitoring (SHM) has emerged as a powerful tool to gain insights into structural behaviour before, during and after critical events [1,2]. Beyond visual inspections, which are often time-consuming and hazardous, vibration-based SHM systems stand out for their non-destructive character and their capability to accomplish real-time damage assessment by analysing raw vibration measurements during operational conditions [3,4]. Recently, these kinds of automated monitoring strategies are deemed particularly suitable to be adopted in conjunction with Machine Learning (ML) algorithms to address damage identification in several engineering fields [5-7]. Within the context of unsupervised learning, aiming at finding hidden patterns throughout the distribution of the unlabelled data, several works in the literature propose to consider a specific index as a damage-sensitive feature to evaluate the difference between the original and the reconstructed acceleration signal outgoing from an autoencoder network [8,9]. On the other hand, the use of supervised algorithms is not so common in practical applications since the training requires labelled data to learn a mapping function between the input and the output. In this framework, one of the main challenges of the ML techniques is the need to process large amounts of training data to build a robust algorithm. Especially in real monitoring scenarios, data collection might be discontinuous or economically unfeasible and labels are potentially unavailable or incomplete because of physical constraints or the scarcity of equipment. Population-based Structural Health Monitoring (PBSHM) attempts to expand the set of labelled data by investigating a population of structures [10-12] which may differ, however, for variations in manufacturing, materials, geometry, the monitoring setup and the surrounding environmental conditions. This translates into a shift in both data distribution and feature space between the two domains, which is the reason why a typical ML algorithm cannot be directly trained on a structure and be tested on another member of the population. Such an issue can be addressed by Transfer Learning (TL), whose goal is to improve diagnostic inferences on a target domain, for which data are missing or scarce, by transferring knowledge from a different domain [13-15]. TL methods can be broadly divided into two main categories, covering fine-tuning and domain adaptation theories. On the one hand, fine-tuning techniques perform a "surgery" on a pre-

https://doi.org/10.1002/cepa.2143

trained neural network model, where specific layers are frozen, removed or replaced with other layers for new classification tasks. Fine-tuning-based frameworks are often adopted in the literature to improve image classification for damage assessment using deep convolutional neural networks [16]. On the other hand, the focus of Domain Adaptation (DA) is to learn a mapping from source to target domains by minimising a defined statistical metric between the two distributions. Typical DA approaches, relying on kernel-based nonparametric density estimation, implement a non-linear transformation to represent source and target features in a shared latent space, using the Joint Domain Adaptation (JDA) and Domain-Adversarial Neural Networks (DANN) techniques [17].

However, it is worth pointing out that all the aforementioned approaches require large quantities of training labelled data to generalize well, thus increasing the risk of overfitting. Conversely, the methodology proposed in this paper exploits a DA technique based on Statistic Alignment (SA), that can be easily applied to poor and limited datasets, thereby becoming particularly suitable for SHM problems [18]. SA has also the advantage to align the lower-order statistics of source and target domains in the original features space, rather than considering a latent space that may not guarantee a human-interpretable representation.

Specifically, this paper successfully applies (i) a SA technique, called Normal Condition Alignment (NCA), to align the normal conditions of the two domains and (ii) the K-Nearest Neighbours (KNN) algorithm, which is previously trained on the source domain and afterwards adopted to perform supervised damage detection on the unknown target domain. The effectiveness of DA is demonstrated by evaluating the algorithm's accuracy before and after feature alignment.

The procedure is validated by considering the Z24 benchmark bridge as the source domain and the S101 bridge as the target domain, for which data are significantly poorer and just referred to the winter period. Experimental data allow a gain of information about the natural frequencies and mode shapes during the whole monitoring period. Natural frequencies are selected as damage-sensitive features and used to firstly align the normal conditions of the two bridges via the NCA technique and to afterwards feed the KNN algorithm for super-vised damage detection. For the current work, damage data describe the same scenarios imposed in both structures, namely the cutting and lowering of one pier and the rupture of tendons.

Numerical results show that the proposed approach enables accurate detection of structural damages on the target domain starting from the knowledge of source domain observations, being particularly useful in most real-world applications given the capability to successfully transfer damage labels within a population of heterogeneous bridges.

2 Bridge health assessment via Transfer Learning: the proposed methodology

The methodology for performing damage detection and knowledge transfer between two bridges is described in Figure 1. Raw accelerations, denoting normal or damage conditions, are acquired by the installed SHM system and afterwards employed to carry out system identification and frequency tracking. N frequencies are then selected as damage-sensitive features. These quantities are aligned via DA, processed by the ML algorithm and employed for damage detection.

2.1 K-Nearest Neighbour Machine Learning algorithm

ML theory provides a quite remarkable set of strategies for data analysis and pattern recognition [19]. Within the supervised techniques, the K-Nearest Neighbours (KNN) algorithm is adopted in this paper to predict feature labels, thereby addressing damage detection and classification. The model's performance is evaluated using the *accuracy*, a metric that quantifies the percentage of correct predictions given the number of true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN):

$$\text{Accuracy} = \frac{\text{TP+TN}}{\text{TP+TN+FP+FN}} \quad (1)$$

This index, however, may produce misleading results when working with unbalanced datasets, where each class does not include equal number of samples. In such cases, it is recommended to adopt different metrics that are not biased to favour the majority class and not sensitive to class skews. Possible solutions are represented by the balanced accuracy or the geometric mean of precision and recall.

KNN is a non-parametric method for classification and regression tasks, based on the assumption that similar points in the feature space can be found near one another. As a first step, the algorithm identifies the K nearest neighbours to a predicted datum in terms of specific distance metrics. Then, the class to which a general numerical item belongs is predicted on the basis of a majority vote, i.e. the label that is most frequently represented in the K neighbours is as-signed to a given data point.

This technique is very simple to apply, it can handle multi-class classification with few parameters to be tuned. However, the algorithm gets significantly slower as the number of features increases. It should be remarked the influence of the K value on the global performance. The lower it is, the more sensitive is the model to outliers. Conversely, a larger value may produce a neighbourhood including points of other classes. A detailed description is provided by Kramer [20].

2.2 Domain Adaptation

Domain Adaptation is a sub-discipline of TL attempting to transfer knowledge between source and target domains by reducing the distance in data distributions. To introduce the method, it is important to provide some key definitions. A domain $\mathcal{D} = \{\mathcal{X}, p(X)\}$ is described by a feature space $\mathcal{X}$ and a marginal probability distribution $p(X)$ with $X = \{x_i\}_{i=1}^{N}$ being a finite sample set from $\mathcal{X}$. A given domain is associated with a task, indicated as $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$, where $\mathcal{Y}$ is a label space and $f(\cdot)$ the predictive function learnt from the training data set.

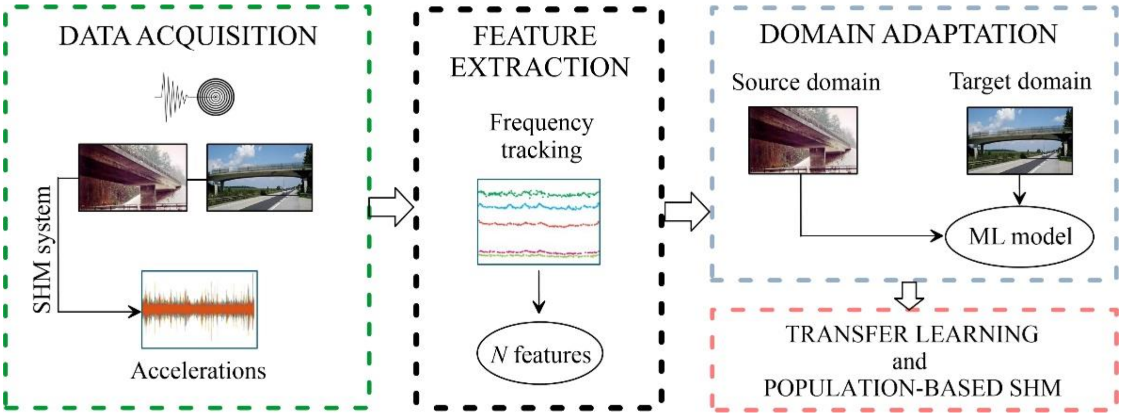

Figure 1 A general overview of the proposed framework.

With these premises, given a source domain $\mathcal{D}_s$ with task $\mathcal{T}_s$ and a target domain $\mathcal{D}_t$ with task $\mathcal{T}_t$,DA is defined as the process of improving the target predictive function $f_t(\cdot)$ in $\mathcal{T}_t$ using the knowledge extracted from $\mathcal{D}_s$.

Classic DA methods match data distributions by using non-parametric distance metrics, which require enough data to perform accurate density estimation. Moreover, they project data into a latent space, decreasing the interpretability of the results. In contrast, this paper uses a SA technique, named as Normal Condition Alignment (NCA), working in the original feature space to align just a subset of data corresponding to normal conditions, which are usually easier to obtain. Precisely, these approaches align the lower-order statistics of the features, quantities that should be able to be estimated with a limited data set. The source domain is firstly standardised as:

$$z_s^{(i)} = \frac{x_s^{(i)} - \mu_s}{\sigma_s} \qquad (2)$$

where μ_s and σ_s are the mean and standard deviation of the source. Normal conditions of the target domain are then aligned with those of the source using Equation 3,

$$z_t^{(i)} = \left(\frac{x_t^{(i)} - \mu_{t,n}}{\sigma_{t,n}}\right)\sigma_{s,n} + \mu_{s,n} \qquad (3)$$

where $\mu_{s,n}$, $\mu_{t,n}$ and $\sigma_{s,n}$, $\sigma_{t,n}$ indicate the means and standard deviations extracted from normal condition data of the source and target domain, respectively.

The idea presented in this work is to match the normal conditions of two bridges in a single healthy class. The KNN algorithm is therefore implemented to evaluate damage detection capabilities before and after DA.

3 Case study: Transfer Learning between two real bridges

The applicability of DA for bridge SHM is demonstrated via a case study involving knowledge transfer between the Z24 and the S101 benchmark bridges.

3.1 General description of the Z24 bridge and S101 bridge

The Z24 bridge, built in 1963 and demolished at the end of 1998, was a post-tensioned RC bridge linking the villages of Koppigen and Utzenstorf in Switzerland [21]. It was characterised by a main span of 30 m and two side spans of 14 m, for a global length of 60 m. Two rectangular concrete piers were located at the limits of the main span and clamped into the deck's girder, whose cross-section was made of two box cells, for a global width of 8.6 m.

The bridge was continuously monitored from November 1997 to September 1998 by measuring accelerations and various environmental parameters. Specifically, vibration data were recorded by eight sensors every hour, for about 10 minutes, at a sampling frequency of 100 Hz.

In order to investigate the dynamic response towards multiple damage scenarios, the bridge was subjected to progressive damage scenarios, carried out in August 1998, shortly before its demolition. A detailed description of the damage test, the experimental data and the monitoring setup can be found in [22].

The second benchmark case study is represented by the S101 bridge, built in the 1960s and located across the A1 Westautobahn in Austria [23]. It was a post-tensioned three-span bridge, composed of a main span of 32 m and two 12 m long side spans. The cross-section was 7.2 m wide and was designed as a double-webbed t-beam, with the height varying from 0.9 m in the mid-span to 1.7 m over the piers. Before the demolition of the bridge, a monitoring campaign was carried out from 10th to 13th December 2008 by using a permanent system characterised by forty-five channels, whose task was to acquire acceleration measurements at a frequency of 500 Hz. Since there were hardly any temperature changes during the measurement period, being freezing conditions dominant, it is conceivable to neglect temperature effects on modal responses. In fact, unlike the Z24, the S101 was deliberately damaged in the winter season for only three days, after collecting one day of healthy measurements. Despite the different monitoring period, both bridges experienced similar damage scenarios, which are divided into two macro-categories, the former describing the lowering of one pier and the latter involving the cutting of tendons along the deck.

3.2 Feature selection

After signal processing, needed to re-sample and filter the original monitoring data, a complete dynamic characterisation of the two bridges is carried out by exploiting the Subspace Identification (SSI) technique within MOSS, a software recently implemented by the SHM group of the Department of Civil and Environmental Engineering at the University of Perugia [24]. System identification results are listed in Table 1, containing the natural frequencies identified for the Z24 and the S101 bridges through an Ambient Vibration Test (AVT).

Table 1 Natural frequencies [Hz] of the Z24 and the S101 bridges

Mode	Z24 bridge	S101 bridge
1	3.85	4.04
2	4.91	6.28
3	9.77	9.72
4	-	13.27
5	12.46	15.81

Although the two structures present some differences in design, belonging thus to a heterogeneous population, there are similarities in the modal responses. Despite different absolute values in natural frequencies, the Z24 and the S101 bridge have in common three bending modes (modes 1, 3 and 5), and a torsional mode (mode 2). In this paper, the first two natural frequencies (describing a bending and a torsional mode, respectively) are considered as damage-sensitive features and therefore employed as inputs for domain adaptation and TL-based damage detection.

After frequency tracking, the datasets are split into training and testing data for the ML classifier. Training data include normal and damage conditions of the Z24: healthy data refer to the first eight months of monitoring and include several temperature ranges (from November 1997 to June 1998), while the investigated damage period covers two weeks during August 1998. On the contrary, testing data are fully represented by the S101 dataset, where normal conditions are exclusively collected in the first day of monitoring.

3.3 Domain Adaptation results

When different domains, characterised by the distribution represented in Figure 2a, are adopted for training and testing, the classifier shows an insufficient performance, less than random guessing, with an accuracy of 49%. This outcome clearly motivates the need for performing feature alignment. NCA is therefore applied to align the healthy data of the Z24 and the S101 bridge, denoted as "0 Z24" and "0 S101", respectively, thereby creating a single healthy-class cluster and removing the offset between the source and target distribution (Figure 2b). Once the mapping has been inferred, any future instances can be projected onto the shared feature space.

Note that the widespread distribution of the Z24 healthy instances is because of the huge amount of available data for that specific case study, covering a broad range of weather conditions. Moreover, looking at Figure 2b, it is interesting to point out that damage instances of both bridges, labelled with "1" (lowering of one pier) and "2" (rupture of tendons), have a similar distribution and belong to the same clusters. This result demonstrates the power of DA, enabling one to classify specific scenarios based on damage experienced by another bridge.

To provide a damage detector and prove the capability of DA to share knowledge between the two bridges, the KNN algorithm is afterwards trained on the Z24 (labelled source domain) and tested on the S101 (labelled target domain) using statistically-aligned feature data and setting the K parameter equal to 1.

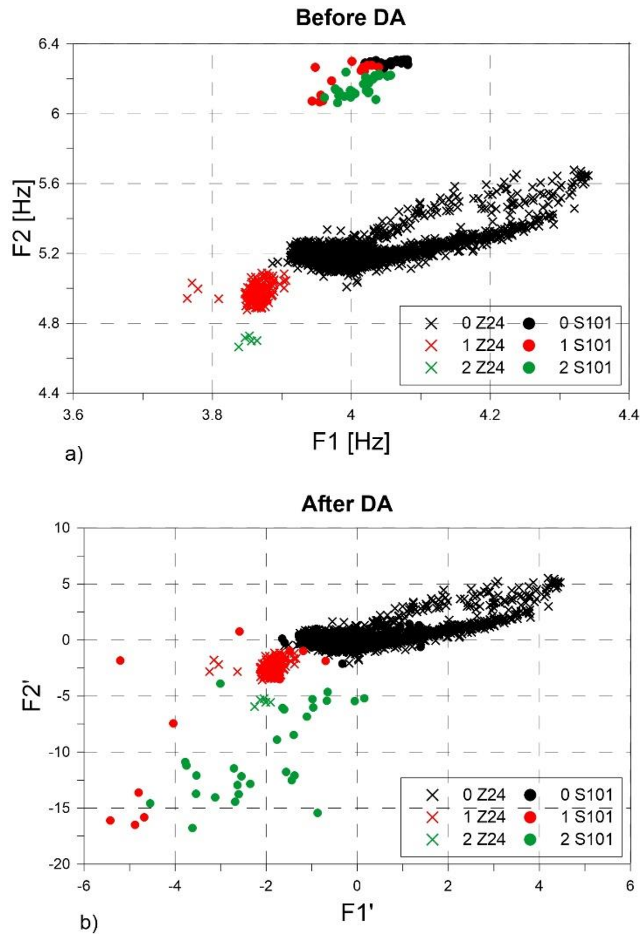

Figure 2 The first (F1) and second (F2) natural frequencies of the Z24 and the S101 bridge are plotted in the same feature space before DA (a) and after aligning the normal conditions with NCA (b). Training and testing data points are denoted with "×" and "o", respectively.

Such a classifier is used in a supervised way to: (i) detect damage conditions and (ii) discriminate between different types of damage, providing 96% and 88% accuracy, respectively. It means that damage labels can be successfully transferred within the population. Beyond this, the model's performance on the target domain alone is clearly improved with DA, yielding 3.3% and 17% improvements in accuracy for tasks (i) and (ii), respectively. It follows that major advantages are visible when the intention is to identify specific damage scenarios on the target domain (S101 bridge), using the supervised KNN algorithm trained on the source domain (Z24 bridge). These results are summarised in Figure 3 in terms of confusion matrices. It is possible to notice the ability of the KNN to well recognize the rupture of tendons in the S101 bridge, given that the features of both domains are contained in a well-defined cluster. In addition, particularly important is the absence of false positives or false alarms, which can cause management issues within SHM activities.

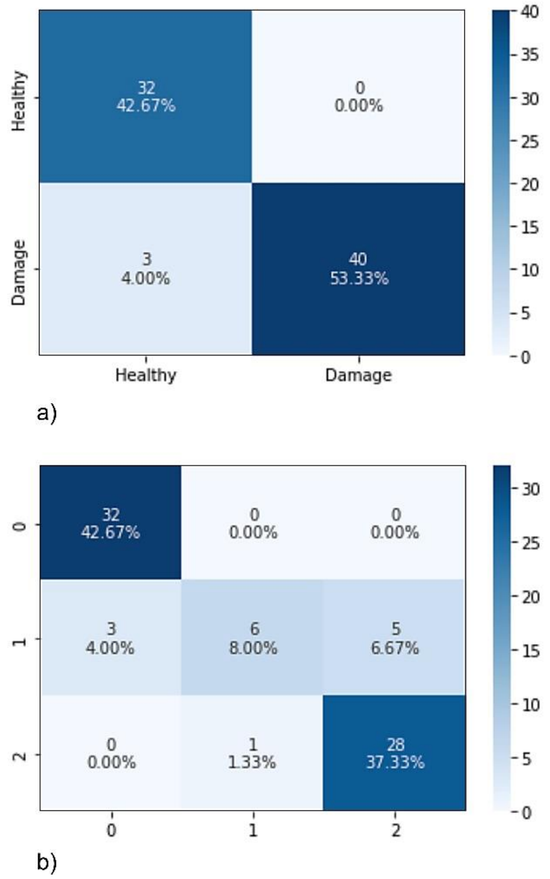

Figure 3 Model performance after DA when classifying between healthy and damage conditions (a) and between healthy and different damage scenarios, indicated with labels "1" and "2" (b).

3.4 Advantages and limitations

DA is particularly useful when the availability of monitoring data is limited and not sufficient for implementing a reliable and robust damage-detection tool. This is the case for the S101 bridge, whose data are not suitable to train a supervised/unsupervised ML algorithm with good generalisation. Therefore, data stemming from the Z24 monitoring campaign and properly transformed with DA techniques, represent a considerable source to enlarge the original S101 dataset.

Prior to DA and damage detection, one of the main issues to face is the number of features to align and afterwards adopt for classification purposes. In this work, the first two natural frequencies, describing a bending and a torsional mode, are deemed effective to capture the similarities between the two bridges as well as able to indicate the discrepancies between healthy and damage states. However, the way by which a large number of modes (and the corresponding features) affects TL outcomes deserves future research.

Moreover, it should be stressed that a good feature alignment facilitates the classifier's implementation. In this context, it could be interesting to investigate other ML algorithms beyond the KNN, such as Artificial Neural Networks, to check their capability in handling aligned features and bringing any advantage in terms of accuracy and false-detection errors.

The most critical aspects underlying the application of TL are represented by the similarity degree between different structures and the type of information that can be successfully transferred. The current paper underlines TL performances when considering that the same damage scenarios occurred within the population of bridges. However, the effectiveness of the proposed approach should be explored by introducing different kinds of damage, to observe if the corresponding features can be sufficiently aligned in the shared feature space. The importance to understand which damage scenarios are transferable could represent a fundamental step towards the implementation of supervised ML algorithms able to assess several bridges with different characteristics.

4 Conclusions

Within the field of SHM, this paper proposes a strategy of DA to statistically align different domains and provide a population-level damage detector by transferring labels across a population of bridges; this can be particularly useful when the target domain is described by a limited quantity of vibration data. The proposed approach is tested on two heterogeneous real structures, the Z24 and the S101 bridges, thus representing an enrichment of the current literature. Given the acceleration responses collected by multiple sensors, system identification provides a set of natural frequencies, considered as damage-sensitive features. The NCA technique is then applied to align the features characterising normal conditions into a shared bi-dimensional space. Common clusters, containing healthy and different damage instances, can be easily identified, thereby yielding clear improvements on KNN performance when classifying target data (S101 bridge), based on the knowledge learnt from source domain observations (Z24 bridge). The supervised damage-identification procedure is deemed to be effective to transfer damage labels between the two bridges, bringing important advantages to practical SHM applications and speeding up the decision-making process; this means that if a bridge undergoes a damage event that has been previously learnt and classified on another domain, the recovery activities could be scheduled and well-focussed on that particular problem. For more reliable outcomes, it is worth mentioning that the concept of similarity and the definition of transferable information are fundamental and deserve particular attention and deeper investigations. Overall, such a promising tool would potentially allow higher-level diagnostic assessment across members of a population before/during/after critical events that worsen the bridges' modal responses.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) and by the University of Perugia via the funded projects "Math4Bridges-Machine Learning and Transfer Learning approaches for Structural Health Monitoring and risk assessment of bridges" in the program "Fondo Ricerca di base di Ateneo, edizione 2020" and "AIDMIX-Artificial Intelligence for Decision making:

Methods for Interpretability and eXplainability" within the program "Fondo Ricerca di Ateneo, 2021".

References

[1] García-Macías E.; Ubertini F. (2022) *Integrated SHM Systems: Damage Detection Through Unsupervised Learning and Data Fusion.* Structural Integrity 21, pp. 247-268.

[2] Brownjohn J.; De Stefano A.; Xu Y. L.; Wenzel H.; Aktan A. (2011) *Vibration-based monitoring of civil infrastructure: Challenges and successes*. Journal of Civil Structural Health Monitoring 1, pp. 79–95.

[3] Fritzen C. P. (2005) *Vibration-based structural health monitoring - concepts and applications*. Key Engineering Materials 293–294:3–18

[4] Deraemaeker A.; Reynders E.; De Roeck G.; Kullaa J. (2008) *Vibration-based structural health monitoring using output-only measurements under changing environment*. Mechanical System and Signal Processing 22, pp. 34–56.

[5] Zhang Y.; Burton H. V.; Sun H.; Shokrabadi M. (2018) *A machine learning framework for assessing post-earthquake structural safety*. Structural Safety 72, pp. 1-16.

[6] Bao Y.; Li H. (2021) *Machine learning paradigm for structural health monitoring.* Structural Health Monitoring 20, pp. 1353-1372.

[7] Worden K.; Manson G. (2007) *The application of machine learning to structural health monitoring*. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 365(1851), pp. 515-537.

[8] Giglioni V.; Venanzi I.; Poggioni V.; Milani A.; Ubertini F. (2022) *Autoencoders for unsupervised real-time bridge health assessment.* Computer-Aided Civil and Infrastructure Engineering.

[9] Giglioni V.; Venanzi I.; Baia A. E.; Poggioni V.; Milani A.; Ubertini F. (2023) *Deep Autoencoders for Unsupervised Damage Detection with Application to the Z24 Benchmark Bridge.* Lecture Notes in Civil Engineering 254 LNCE, pp: 1048-1057.

[10] Bull L.A.; Gardner P.A.; Gosliga J.; Rogers T.J.; Dervilis N.; Cross E.J.; Papatheou E.; Maguire A.E.; Campos C.; Worden K. (2021) *Foundations of population-based SHM, Part I: Homogeneous populations and forms*. Mechanical Systems and Signal Processing 148.

[11] Gosliga J.; Gardner P.A.; Bull L.A.; Dervilis N.; Worden K. (2021) *Foundations of Population-based SHM, Part II: Heterogeneous populations – Graphs, networks, and communities.* Mechanical Systems and Signal Processing 148.

[12] Gardner P.; Bull L. A.; Gosliga J.; Dervilis N.; Worden K. (2021) *Foundations of population-based SHM, Part III: Heterogeneous populations – Mapping and transfer*. Mechanical Systems and Signal Processing 149.

[13] Pan S. J.; Yang Q. (2010) *A survey on Transfer Learning.* IEEE Transactions on Knowledge and Data Engineering 22, pp. 1345-1359.

[14] Gardner P.; Bull P. L. A.; Dervilis N.; Worden K. (2022) *On the application of kernelised Bayesian transfer learning to population-based structural health monitoring*. Mechanical Systems and Signal Processing 167.

[15] Tronci E. M.; Beigi H.; Feng M. Q.; Betti R. (2022) *A transfer learning SHM strategy for bridges enriched by the use of speaker recognition x-vectors*. Journal of Civil Structural Health Monitoring.

[16] Ogunjinmi P. D.; Park S.; Kim B.; Lee D. (2022) *Rapid Post-Earthquake Structural Damage Assessment Using Convolutional Neural Networks and Transfer Learning*. Sensors 22.

[17] Gardner P.; Liu X.; Worden K. (2020) *On the application of domain adaptation in structural health monitoring.* Mechanical Systems and Signal Processing 138.

[18] Poole J.; Gardner P.; Dervilis N.; Bull L.; Worden K. (2023) *On the Application of Partial Domain Adaptation for PBSHM*. Lecture Notes in Civil Engineering 270 LNCE, pp. 408-418.

[19] Olivas E. S.; Guerrero J. D. M.; Sober M. M.; Benedito J. R. M.; Lopez A. J. S. (2009) *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*. Hershey: Information Science Reference.

[20] Kramer O. *K-Nearest Neighbors, In: Dimensionality Reduction with Unsupervised Nearest Neighbors.* Berlin: Intelligent Systems Reference Library, 51 Springer.

[21] Maeck J.; De Roeck G. (2003) Description of Z24 benchmark. Mechanical Systems and Signal Processing 17, pp. 127-131.

[22] Swartz R. A.; Lynch J. P. (2003) *Damage characterization of the Z24 bridge by transfer function pole migration*. Conference Proceedings of the Society for Experimental Mechanics Series 13.

[23] Döhler M.; Hille F.; Mevel L.; Rücker W. (2014) *Structural health monitoring with statistical methods during progressive damage test of S101 Bridge*. Engineering Structures 69, pp. 183-193.

[24] García-Macías E.; Ubertini F. (2020) *MOVA/MOSS: Two integrated software solutions for comprehensive Structural Health Monitoring of structures*. Mechanical Systems and Signal Processing 143.