On the Influence of Structural Attributes for Transferring Knowledge in Population-Based Structural Health Monitoring

On the influence of structural attributes for transferring knowledge in population-based structural health monitoring

Giulia Delo[1], Daniel S. Brennan[2], Cecilia Surace[3] and Keith Worden[2]

[1]Department of Mechanical and Aerospace Engineering
Politecnico di Torino, Turin, Italy. Email: giulia.delo@polito.it.

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

[3]Department of Structural, Geotechnical and Building Engineering
Politecnico di Torino, Turin, Italy.

ABSTRACT
The recently proposed theory of Population-Based Structural Health Monitoring (PBSHM) aims at improving diagnostic inferences, by sharing damage-state knowledge across a population of structures via transfer-learning algorithms - specifically domain adaptation. Before applying these algorithms, the similarity between structures, or substructures, should be evaluated. This assessment helps prevent negative transfer, ensuring better performance and higher robustness of data-based SHM.

When structures are sufficiently similar, different transfer-learning strategies can be applied, according to the original features and the specific case study. In this framework, structural attributes play a crucial role, especially for heterogeneous populations in which the main differences can be caused by material properties, geometry or dimensions. Therefore, investigating how to consider the influence of these properties in distance metrics became necessary, and new similarity metrics have been adopted to focus on geometric features and dimensions. However, to gain a comprehensive understanding of attribute relevance, and to address it at the similarity-evaluation phase, it is necessary to evaluate the performance of transfer-learning algorithms as these structural features vary.

The present work extends this research by examining the effect of material and dimension attributes on the performance of a domain-adaptation method - the Transfer Component Analysis (TCA). This analysis is applied to an experimental population of laboratory-scale aircraft, comprising structures with different materials and dimensions, and similar topology. A confusion matrix is employed to compare the findings and show how these properties can influence the transfer-learning performance, especially for localised damage, thus highlighting the importance of their evaluation in the context of PBSHM.

Keywords: Population-Based Structural Health Monitoring (PBSHM), Aerospace Engineering, transfer-learning, Transfer Component Analysis, Similarity Assessment

INTRODUCTION

Population-based Structural Health Monitoring (PBSHM) has been introduced as a novel approach for improving the performance of Structural Health Monitoring (SHM) strategies [1–3]. It can overcome the issues related to the lack of experimental data, or their incompleteness, which often limit many data-based SHM approaches, by sharing damage-state knowledge across a population of structures. The features of interest mainly regard modal properties or dynamic transfer functions, such as Frequency Response Functions (FRFs). These features can be used in transfer-learning algorithms, specifically domain adaptation, to share knowledge between a source and a target structure: A source structure having larger and more complete experimental datasets is exploited in the algorithm training phases, and then the algorithms are specified for a target structure, which may have a dataset with fewer or incomplete samples. In the PBSHM framework, populations are distinguished as *homogeneous* if all structures in the population are nominally identical, and *heterogeneous* if they can also be different. Specifically, heterogeneous structures could differ in terms of structural topology, materials, geometrical properties, or dimensions. Therefore, before applying knowledge-transfer algorithms, the similarity between structures or substructures should be evaluated. This assessment helps prevent negative transfer, providing robust data-based SHM diagnostic inferences, even when heterogeneous populations are examined.

Different methods have been proposed to measure similarity, according to the case study and its features. First, it is required to build an abstract representation of the structures, adopting Irreducible Element (IE) models, and Attributed Graphs (AGs) [2]. These models describe the structure's topology and its attributes, i.e., all the mechanical and geometrical characteristics. Subsequently, the AGs are employed to measure similarity via distance metrics. The maximum common subgraph (MCS) between two structures has been proposed in [2] as a method to compute a similarity index based on the relative dimension of the largest subset of nodes and edges common to both structures. Subsequently, a similarity assessment method based on Graph Matching Networks (GMNs) has been proposed as an extension to GNNs [4]. This method has been introduced in PBSHM to improve accuracy and assess the impact of mechanical and geometrical properties, as shown in [5–7]. Nevertheless, to evaluate the influence of such properties, or attributes, it is necessary to choose how to embed and manage them within the similarity assessment network, and to investigate the actual impact on knowledge-sharing performance.

The objective of this study is to facilitate knowledge sharing within a population of similar structures and explore the impact of material and geometrical properties on the transfer-learning performance, to investigate the extent to which this type of algorithm can be reliably extended to heterogeneous populations, and improve similarity assessment accordingly. The knowledge transfer process is implemented for a group of three simplified models of an aircraft. These structures have been built and experimentally analysed under undamaged and simulated damaged conditions to validate the PBSHM strategies for heterogeneous populations in the aerospace field. Transfer Component Analysis (TCA) [8], is used on each pair of structures to harmonise the original features, transforming them into a low-dimensional latent space where damage can be detected. In addition, the TCA performance is compared to a benchmark damage detection technique, as shown in [9]. The layout of the paper is as follows. Section 2 describes the TCA method, used for sharing knowledge within the population, and the conventional novelty-detection strategy based on Principal Component Analysis (PCA), used for comparison purposes. Section 3 presents the case study. Section 4 presents the results and discussion, followed by some conclusions and further developments.

METHODS

With the development of machine-learning approaches, many data-based techniques have been developed within the SHM field, exploiting dynamic features such as modal properties or transfer functions [10, 11]. However, training such models requires a large amount of data, which is not often available. In addition, the training and test data should come from the same domain, and have the same distribution [12]. Transfer-learning algorithms have been proposed to enhance diagnostic inferences by sharing knowledge from a more comprehensive source domain to a potentially different target domain. Specifically, this study focusses on TCA, which learns a nonlinear mapping between the source and target features and transforms them into a low-dimensional latent space, matching the marginal distributions [9, 12]. A brief description of the method is given below.

A domain $\mathscr{D}$ consists of a feature space χ, and a marginal probability distribution $P(X)$, where $X = \{x_1, x_2, ...x_n\} \in \chi$. A task $\mathscr{T}$ consists of a label space $\mathscr{Y}$, where $Y = \{y_1, y_2, ...y_n\} \in \mathscr{Y}$, and a predictive function $P(y_i|x_i)$, which can be learnt from the training process. For instance, these features can be portions of a transfer function such as an FRF, while the labels may be indicators of damage or of its location. Transfer-learning algorithms aim at improving the predictive function on a target domain $\mathscr{D}_T$ using the source domain $\mathscr{D}_S$ and task $\mathscr{T}_S$, assuming that $\mathscr{D}_T \neq \mathscr{D}_S$ or $\mathscr{T}_T \neq \mathscr{T}_S$. Domain adaptation is a branch of transfer-learning which assumes that $\mathscr{D}_T \neq \mathscr{D}_S$, but $P(Y_S|X_S) = P(Y_T|X_T)$. In this framework, TCA is a dimensionality-reduction method, which learns a transformation ϕ such that $P(\phi(X_S)) \approx P(\phi(X_T))$, and also assumes $P(Y_S|\phi(X_S)) \approx P(Y_T|\phi(X_T))$. The Maximum Mean Discrepancy (MMD) criterion is used in the distance minimisation, leveraging the kernel trick to avoid dealing with the nonlinear space. For further details on the kernel learning problem, the reader is referred to Pan and Yang [8]. Hence, the FRF portions related to undamaged conditions can be used as training features for the source and target domains χ_S and χ_T. Upon standardisation, these features are used to learn ϕ, by adopting a radial-basis function kernel; ϕ is used to transform the target features within the test dataset, which also includes samples of damaged conditions. Afterwards, the Mahalanobis Squared-Distance (MSD) with respect to the undamaged features is employed to detect damage [13]. Observations are labelled as damaged when the distance is higher than a threshold, which is computed via a Monte Carlo method [11]. The TCA findings are compared with those obtained via a more traditional approach, according to [9, 13]. This benchmark novelty-detection approach consists in using Principal Component Analysis (PCA) to learn a linear transformation of the features X_T, considered separately. The FRFs in X_T are transformed into a low-dimensional space. Subsequently, the same strategy based on the MSD is used to detect damage [13].

CASE STUDY

In this study, knowledge sharing is applied to a population of three laboratory-scale aircraft. These models are developed from the benchmark study conducted by the Structures and Materials Action Group (SM-AG19) of the Group for Aeronautical Research Technology in EURope (GARTEUR) [14]. They are made of beam and plate components joined together via bolted connections. A thin-section rectangular beam simulates the fuselage, and a single rectangular plate is connected to the fuselage for realising the wings. In addition, similar plate elements are used to create the vertical and the horizontal tail. Hence, variations in geometries, materials and dimensions are introduced to create a comprehensive experimental dataset, including the structures analysed in [15]. These variations result in a heterogeneous population, although with limited topological differences. The analysed population is shown in Fig.1.

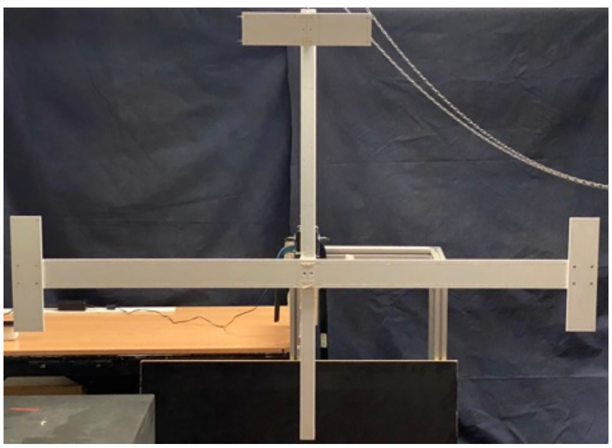

(a) Large steel model.

(b) Small steel model.

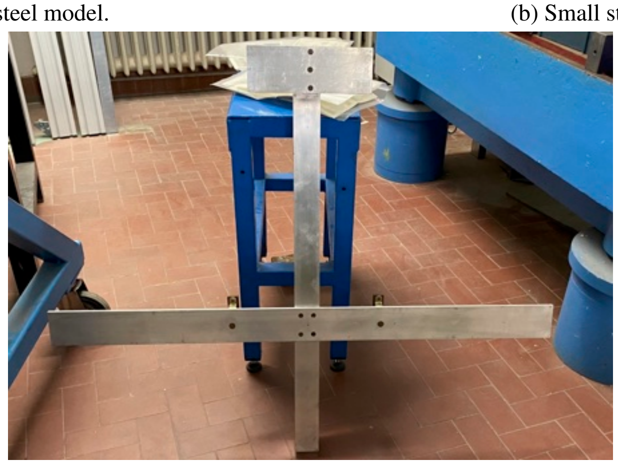

(c) Small aluminium model.

Figure 1: Population of three laboratory-scale aircraft.

The first structure of the population is made of steel and has a wingspan of 2.0 m (Fig.1a). It can be tested in the variant with or without winglets. In the present study, the experimental data are acquired without winglets. The second is of the same material, but has a wingspan of 1.0 m, and all dimensions are overall reduced by 50% (Fig.1b). The third structure is made of aluminium and has the same dimensions as the previous one, with a wingspan of 1.0 m (Fig.1c). Each structure has been tested by means of an Experimental Modal Analysis (EMA), performed using a Scanning Laser Doppler Vibrometer (SLDV) and applying a periodic chirp input via a shaker attached to the fuselage. The FRFs have been acquired, using 2048 spectral lines, in undamaged conditions and in six damaged conditions. Damage has been simulated by applying concentrated masses (around 2% of the aircraft mass), in multiple positions over the structure: one at the wing tip, two along the wing, one on the horizontal tail and two on the fuselage, at the tip and at the end. The experimental tests have been carried out in the “LAQ-AERMEC Aeromechanical Structural Systems” laboratory of the Department of Mechanical and Aerospace Engineering, Politecnico di Torino. The measured transfer functions are shown in Fig.2, and a dataset is computed by polluting the normalised FRFs with $\sigma = 0.2$ Gaussian Noise. The extracted features consist of FRF segments, being equivalent in the three models, evaluated around the peaks of the main vibration modes identified in the measured frequency range.

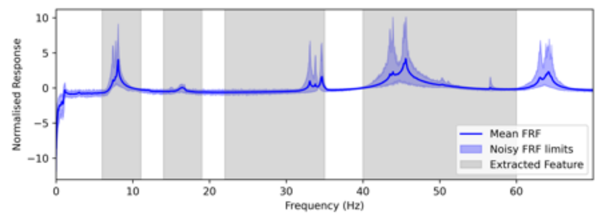

(a) Large steel model.

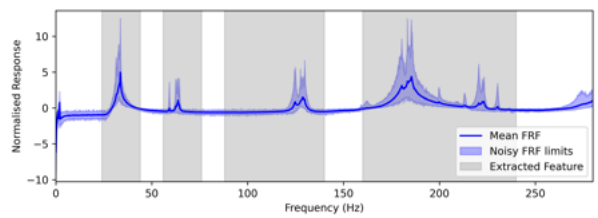

(b) Small steel model.

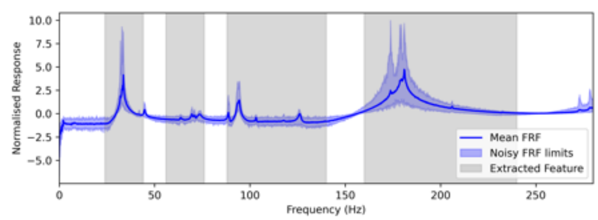

(c) Small aluminium model.

Figure 2: Normalised noisy FRFs, computed by polluting the measurements with Gaussian noise, and extracted features in shaded grey.

RESULTS AND DISCUSSION

The first investigation concerns the use of the PCA and the MSD on the individual FRFs of each model, to provide a baseline comparison, and understand the performance that could be achieved in the absence of knowledge transfer. For each target domain, a dataset is produced by splitting the target features into a training set, a validation set, and a test set. Assuming that only a few experimental data can be measured, the training set only includes five samples from the undamaged condition, and the validation set includes five samples from the undamaged condition and one sample for each damaged condition. The test dataset includes ten samples for the undamaged and each damaged condition. The PCA produces a transformation of the features into a low-dimensional space, and even if two-dimensional features could be used to aid visualisation, the number of components to keep can be optimised with respect to the validation set. The results of this benchmark analysis on the large steel model are shown in Fig.3. In this case, four principal components are considered. It can be noted that only some of the samples from damaged conditions in the test set are correctly labelled. Similar results can be found for the other two models. A summary of the performance is shown in Fig.4. Although the undamaged condition samples are correctly identified, according to the true negative rates (TNR), the true positive rates (TPR), i.e., the correctly labelled samples from the damaged conditions, are significantly lower.

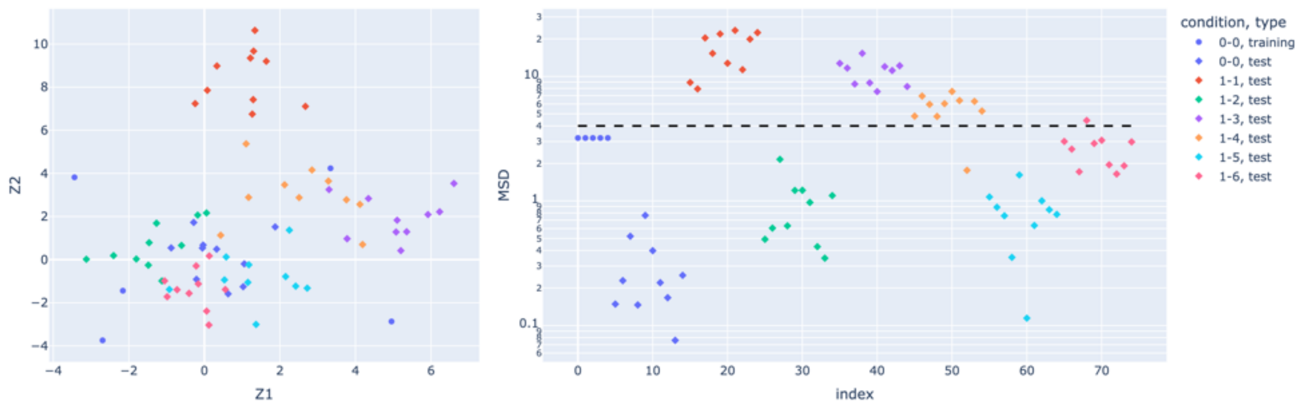

(a) First two principal components from the PCA

(b) Mahalanobis Square Distance

Figure 3: Results of the benchmark PCA and MSD analysis applied on the large steel model. The label “0-0” refers to the undamaged condition, the labels “1-n” refer to the damaged condition, where n defines the damage location.

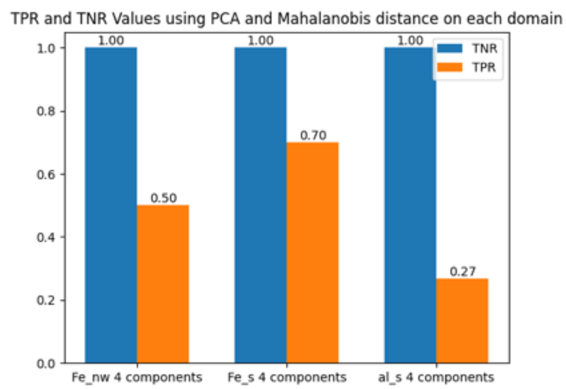

Figure 4: Summary of the results of the benchmark PCA and MSD analysis applied to the three models, with the optimised number of components. The TNR is shown in blue, and the TPR is shown in orange.

The damage detection performance can be improved by transferring knowledge from a similar structure, used as a source. Specifically, this study aims to investigate to what extent structures must be similar in order to boost damage detection, and what is the relevance of different attributes on the final performance. Thus, each combination of the three structures is analysed. The first structure is considered as a source, whose dataset is assumed to be twice as large as that of the target. The second structure is considered a target, using the previously-defined dataset. In addition, TCA also consists of a dimensionality reduction of the features, and the choice of the reduced dimension can be considered as a hyperparameter of the method. Therefore, in this study, the number of components to be used in the TCA is defined as equal to the number of components evaluated in the PCA for the source structure. Fig.5 shows the TCA findings computed using the large steel aircraft model as the source and the small steel aircraft model as the target domain. The first two components of the transformed features in the latent space are shown in Fig.5a, while Fig.5b shows the MSD values. It can be seen that samples from damaged conditions are more distant than those from normal conditions, making it easier to identify them by MSD.

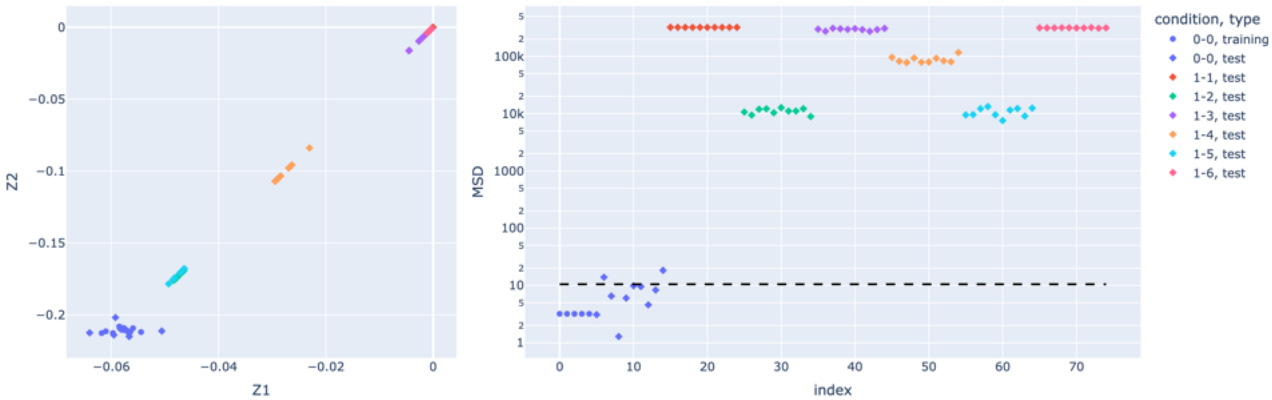

(a) First two principal components from the TCA

(b) Mahalanobis Square Distance

Figure 5: Results of the TCA and MSD analysis applied using the large steel model as the source domain and the small steel model as the target domain. The label “0-0” refers to the undamaged condition, the labels “1-n” refer to the damaged condition, where n defines the damage location.

Similar results can be found for the other pairs of source and target structures. A summary of the performance is shown in Fig.6 in terms of TNR and TPR, and in Fig.7 in terms of the F1 score, a metric used to assess the accuracy and recall of a classification model in a single value. It can be noted that the TNR (Fig.6b) is almost always equal to 1, i.e., all undamaged condition samples are correctly classified. The only cases where there is a lower value is when combining the datasets of the two steel models, which could be related to the sample biases because of the small dimension of the dataset. Furthermore, the TPR (Fig.6a) matrix shows a significant increase in the number of samples correctly identified as belonging to the damaged condition, for each combination of structures. However, it is interesting to note that these results, for each target, are variable depending on the type of source used, and thus, according to the material and size attributes, which are the main difference within the population. This effect can also be seen from the F1 score matrix (Fig.7). For instance, when using the large steel model as the target, the F1 score is slightly higher when the source dataset is extracted from the small model of the same material, and decreases when the small aluminium model is used as the source. On the other hand, any starting structure gives good results on the intermediate model, the small steel aircraft, although slightly better when using the model with equal dimensions. These results may be subject to sampling bias and the TCA may be sensitive to the choice of hyperparameters, so future developments will focus on optimising these factors.

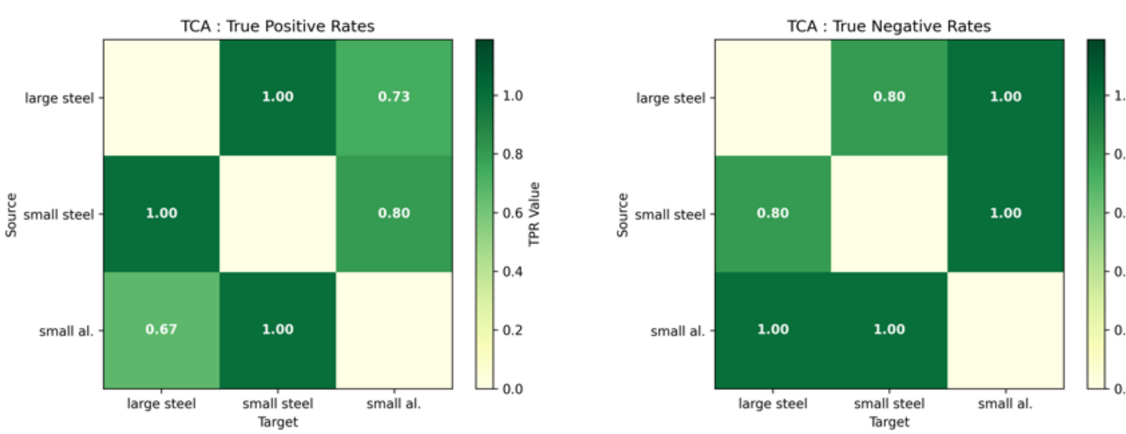

(a) True Positive Rates

(b) True Negative Rates

Figure 6: Results of the TCA and MSD analysis applied on the aircraft population, in terms of TPR and TNR.

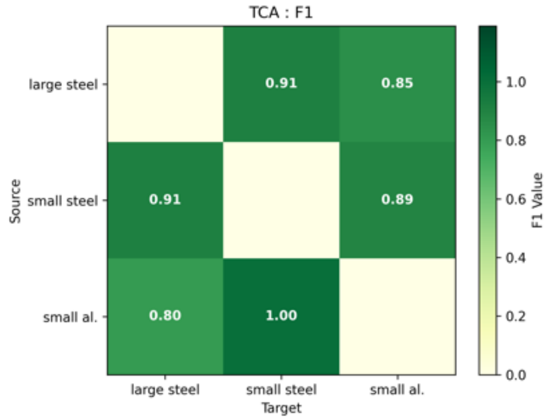

Figure 7: Results of the TCA and MSD analysis applied on the aircraft population, in terms of F1 score.

CONCLUSION

The present study investigates the challenge of sharing knowledge between “similar”” structures via transfer-learning approaches, specifically domain adaptation. Indeed, one of the objectives of PBSHM is to understand what structures are suitable for knowledge transfer, reducing the issues related to negative transfer, especially within heterogeneous populations. To do so, it is required to distinguish the structures with respect to their topology, but also consider the differences in materials, geometry and dimensions. This is already relevant in the first phase of the PBSHM approach, that is, the similarity assessment. However, it is likewise necessary to study the impact of the attribute differences on the transfer-learning performance, because it may provide insight into dealing with these attributes for assessing structure similarity, and more broadly, on the possible extent of the use of PBSHM.

Thus, a small population of laboratory-scale aircraft models is presented as an experimental case study. An EMA has been performed for each structure, achieving the FRFs under undamaged conditions and six simulated damaged conditions. These aircraft models have similar topologies but different properties, such as materials and dimensions, to create a heterogeneous population and facilitate the analysis of the attribute impact on the similarity assessment process and the transfer-learning results. The method used to share damage-sensitive knowledge is the TCA transformation, followed by a classification method based on the MSD. The analysis is performed for each pair of structures, computing the effect of every possible source on the novelty-detection performance for a specific target domain. Moreover, the performance of TCA is compared with that of a benchmark strategy, employing PCA on each domain individually to prove whether sharing knowledge could enhance novelty detection.

The findings show that using a population-based approach broadly improves novelty detection results. By exploiting knowledge from a source dataset, it is possible to improve the diagnostic capability of each target structure, even by embracing a simple MSD-based classification. Indeed, there are no cases where the use of TCA leads to negative transfer compared to the baseline analysis. This outcome could be expected, as the population members share almost equal topology. Moreover, the results are also affected by attribute differences in the aircraft models, which shows the relevance of attributes within PBSHM and suggests the necessity to embed them in the similarity assessment, especially when analysing more localised damage, or specific conditions that might cause damage to the target structure. However, the weight of each attribute over another is strongly related to the specific target case study. Moreover, it would be interesting to analyse the sensitivity of other techniques to attribute differences. Therefore, upcoming research will focus on these effects on different transfer-learning approaches. Additionally, further insights will regard subsequent monitoring steps such as damage localisation and extent.

ACKNOWLEDGEMENTS

The authors of this paper gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) via grant references EP/W005816/1 and EP/S001565/1. The authors also express their gratitude for the TCA implementation made available through the GitHub repository (https://github.com/labull/EngineeringPatternRecognition).

REFERENCES

[1] Bull, L.A., Gardner, P.A., Gosliga, J., Rogers, T.J., Dervilis, N., Cross, E.J., Papatheou, E., Maguire, A.E., Campos, C., and Worden, K. "Foundations of population-based SHM, Part I: Homogeneous populations and forms". *Mechanical Systems and Signal Processing*, 148:107141 (2021)

[2] Gosliga, J., Gardner, P.A., Bull, L.A., Dervilis, N., and Worden, K. "Foundations of population-based SHM, Part II: Heterogeneous populations–graphs, networks, and communities". *Mechanical Systems and Signal Processing*, 148:107144 (2021)

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

[4] Li, Y., Gu, C., Dullien, T., Vinyals, O., and Kohli, P. "Graph matching networks for learning the similarity of graph structured objects". In *International Conference on Machine Learning*, pages 3835–3845. PMLR (2019)

[5] Brennan, D.S., Rogers, T.J., Cross, E.J., and Worden, K. "On quantifying the similarity of structures via a graph neural network for population-based structural health monitoring". In *Proceedings of the 30th International Conference on Noise and Vibration Engineering, ISMA 2022*. KU Leuven Department of Mechanical Engineering (2022)

[6] Brennan, D.S., Rogers, T.J., Cross, E.J., and Worden, K. "Calculating structure similarity via a graph neural network in population-based structural health monitoring: Part II". In *Proceedings of the 40th IMAC, A Conference and Exposition on Structural Dynamics 2022* (2022)

[7] Delo, G., Surace, C., Worden, K., and Brennan, D.S. "On the influence of structural attributes for assessing similarity in population- based structural health monitoring". In *Proceedings of the 14th International Workshop on Structural Health Monitoring 2023: Designing SHM for Sustainability, Maintainability, and Reliability*. DEStech Publications, Inc. (2023)

[8] Pan, S.J., Tsang, I.W., Kwok, J.T., and Yang, Q. "Domain adaptation via transfer component analysis". *IEEE Transactions on Neural Networks*, 22(2):199–210 (2010)

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

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

[11] Farrar, C.R. and Worden, K. *Structural Health Monitoring: a Machine Learning Perspective*. John Wiley & Sons (2012)

[12] Pan, S.J. and Yang, Q. "A survey on transfer learning". *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359 (2009)

[13] Worden, K., Manson, G., and Fieller, N.R. "Damage detection using outlier analysis". *Journal of Sound and Vibration*, 229(3):647–667 (2000)

[14] Balmes, E. and Wright, J.R. "Garteur group on ground vibration testing: results from the test of a single structure by 12 laboratories in europe". In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 80401, page V01AT03A004. American Society of Mechanical Engineers (1997)

[15] Delo, G., Mattone, M., Surace, C., and Worden, K. "Novelty detection across a small population of real structures: A negative-selection approach". In *XII International Conference on Structural Dynamics* (2023)