Enhancing Structural Health Monitoring with Machine Learning and Data Surrogates: A TCA-Based Approach for Damage Detection and Localisation

R. S. BATTU, K. AGATHOS and E. PAPATHEOU

ABSTRACT

Structural health monitoring (SHM) involves constantly monitoring the condition of structures to detect any damage or deterioration that might develop over time. Machine learning methods have been successfully used in SHM, however, their effectiveness is often limited by the availability of data for various damage cases. Such data can be especially hard to obtain from high-value structures. In this paper, transfer component analysis (TCA) with domain adaptation is utilised in conjunction with high-fidelity numerical models to generate surrogates for damage identification without the requirement for high volumes of data from various damaged states of the structure. The approach is demonstrated on a laboratory structure, a nonlinear Brake-Reuß beam, where damage scenarios correspond to different torque settings on a lap joint. It is shown that, in a three-class scenario, machine learning algorithms can be trained using numerical data and tested successfully on experimental data.

INTRODUCTION

Efficient Structural Health Monitoring (SHM) systems can improve the safety, reliability and service life of structures. Additionally, they aid in lowering the cost and time required for maintenance procedures. There are diverse advancements in recent research in SHM, particularly employing machine learning (ML) algorithms. Typically, ML algorithms which are used in SHM can be categorised into supervised, unsupervised and semi-supervised. Supervised ML algorithms require data from all possible damaged states for successful damage identification in SHM. While it is extremely simple to damage and subsequently obtain data from inexpensive structures, damage to high-value structures, is nearly unattainable.

It is a challenging task to solve this problem of data insufficiency. Experimental approaches to use 'damage proxies' in the form of added masses have been shown in [1, 2], but they are generally limited to accessible areas of a structure. Another possible solution is to create surrogates with the assistance of high-fidelity physical models. However,

Raja Sekhar Battu, PhD Student, Email: rb756@exeter.ac.uk.
Dr Konstantinos Agathos, Dr Evangelos Papatheou. Department of Engineering, University of Exeter, North Park Road, EX4 4QF, Exeter, UK

damage modelling for the purposes of generating surrogate models is also challenging. Here, a novel technique is employed with the integration of high-fidelity physical models and Transfer Component Analysis (TCA) with domain adaptation. TCA with domain adaptation was initially introduced by Pan et al. [3], who provided a unique feature extraction approach as well as an extensive understanding of the theoretical foundations and mathematical equations of TCA with domain adaptation. Chakraborty et al. [4] demonstrated a transfer learning strategy in SHM employing time-frequency features on a classification problem. In one of the applications of domain adaptation in SHM, Gardner et al. [5] provided a simple implementation with feature sets of damped natural frequencies and damping ratios in the application of population-based SHM. The results outperformed conventional supervised learning algorithms in terms of classification rates. In another investigation, Ozdagli and Koutsoukos [6] used a domain-adversarial neural network to achieve domain adaptation, which provided evidence for improved prediction accuracy. Poole et al. [7] conducted domain adaptation with statistical alignment as the initial step and proved that the normal correlation alignment is robust to solve the problem of class imbalance.

In this work, the core objective is to acquire data surrogates for diverse structural damage classes, which are then used to train efficient multiple classifiers. The validation of the classifiers is done with the use of actual damage data from laboratory experiments. The structure of this article begins with an introductory section, followed by a description of the research strategy and the techniques adopted herein. The third section describes the test setup for the experimental work as well as the development of the numerical model. The fourth section presents the three-class problem examined here and the approach used to attain efficiency. In the last section, results of trained classifiers are reported alongside concluding observations.

RESEARCH FRAMEWORK

Broadly speaking, common ML techniques employed in SHM depend crucially on a feature selection/extraction stage. The higher the sensitivity of the feature, the more precise any classifier is in categorising damage [8]. In general, one could perform a manual or automatic feature selection (e.g. see [9] for the latter). Here, the research

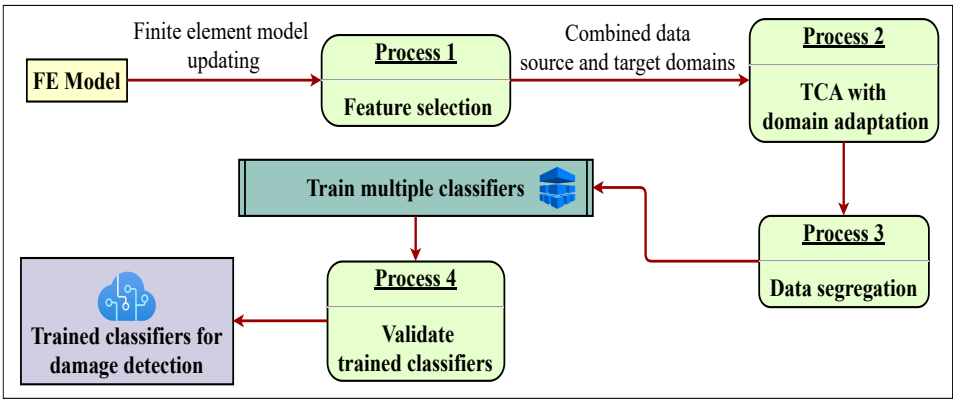

Figure 1. The proposed framework

framework (shown in Figure 1) begins with a manual feature selection on a highly accurate FE model. Features are primarily extracted through Frequency Response Functions generated from the FE model, and their selection follows a similar approach to one used in [2], while at least one feature per damage-class is extracted. The next step is to incorporate TCA with domain adaptation to transfer knowledge from one domain to another. Following data segregation, the next step of the framework involves deploying multiple classification algorithms for training on FE data and validating on experimental data-features.

The main objective of TCA with domain adaptation is to align the marginal distributions of the source and target domains. A domain can be defined as a feature space and a marginal probability distribution over the feature data [10]. In this article, the source and target domains are the FE and the experimental-data, respectively - aiming at transfer learning between FE and experiment. TCA translates the feature sets of the source and target domains to a single latent feature space, given the input data for the domains and their labels. This transformation yields two new feature sets for the source and target domains. The projection of the input feature sets is accomplished by minimising the distance between the empirical means of the two distributions (source and target), as assessed by the distance metric called maximum mean discrepancy (MMD). This distance metric quantifies the distance in a reproducing kernel Hilbert space (RKHS), where the RKHS depends primarily on the kernel function [3].

Let ϕ be the kernel-induced feature map, and $L = \{l_i\}$, $M = \{m_i\}$ the samples from the two distributions. The empirical estimation of the distance between two samples is

$$MMD(L, M) = ||1/n_1 \sum_{i=1}^{n_1} \phi(l_i) - 1/n_2 \sum_{i=1}^{n_2} \phi(m_i)||^2_{\mathcal{H}} \tag{1}$$

where $||.||_{\mathcal{H}}$ is the RKHS norm. Thus, from equation (1) it is evident that kernel mapping plays a vital role in reducing the distance between $P(\phi(L))$ and $P(\phi(M))$. The new feature space is identified using the MMD and a kernel trick [3]. The chosen kernel function is utilised to assess a kernel matrix, which is then used to conduct the parametric kernel mapping. This mapping, which is an optimisation problem, results in a new feature dataset and the transformation matrix. The hyperparameters required for all of these operations are the kernel function, kernel sigma, and regularisation parameter. The selection of hyperparameters for the application of TCA in SHM can be influenced by the type of features, the complexity of the structure, and the dimensionality of the input dataset. Besides, kernel functions that are often utilised are linear, polynomial, and radial basis function (RBF). A linear function is widely utilised due to its efficiency and ease of computing.

The major benefit of the TCA is that when knowledge is transferred from one domain to another, the physical structure of the datasets is preserved [10]. Domain adaptation is performed after the selection of the necessary hyperparameters (kernel, kernel sigma, latent subspace dimension, regularisation parameter). This study uniquely utilises the output transformation matrix from the healthy class to transform the feature sets of damaged classes into the new feature sets of the corresponding classes. For the purpose of performing this transformation, a combined dataset (damage) including merged datasets for each feature from the source (FE) and target (Exp) domains is passed through a kernel function. The projected dataset in the new latent feature space is obtained by performing

matrix multiplication between the resulting kernel dataset and the transformation matrix (healthy). This approach is carried out for every feature in each class on a consistent basis to deliver datasets for all of the specified damage classes. Once the feature sets for healthy and damaged classes are computed, the Mahalanobis distance between normal (healthy) and abnormal classes can be calculated [11] and act as a measure of 'abnormality' in the dataset.

Initially, a group of 800 multi-layer perceptrons (MLPs) with 20 initialisations and hidden units ranging from 5 to 200 serves as multiple classifiers. Furthermore, multiple K-nearest neighbour (KNN) classifiers were deployed, each with its own set of neighbour values and distance metrics. The Mahalanobis distance FE features from the previous stage serve as input to train the deployed classifiers. Finally, the performance of these trained classifiers on experimental data was tested.

EXPERIMENTAL SETUP AND NUMERICAL MODEL

The experimental structure chosen in this study was a Brake-Reuß beam [12] with a lap joint secured with bolts. Two stainless steel 304 beam components were fastened together using MS M8 bolts to produce a prismatic beam having a length of 108 cm and a square cross-section with a side of 2.54 cm. The bolts were secured with a pretension of 20 Nm, which is considered a baseline or healthy (normal) state throughout the article. Loosening these bolts can thus facilitate a series of tests for various damage classes. The experimental test apparatus is depicted in Figure 2. To approximate free-free boundary conditions, the beam was suspended at two locations by means of fishing lines. Uniaxial piezoelectric accelerometers of the PCB type were used to record accelerations at seven locations on the beam (see again 2).

The data acquisition was performed with an NI DAQ and Labview (signal express) application running on a Dell PC. A random noise generator was adopted to create the excitation force for 20 seconds. Throughout the experiments, a sample rate of 25.6 kHz was used. Figure 2 displays the sensor positions, and the data obtained from these locations include force and acceleration measurements throughout this article. Frequency

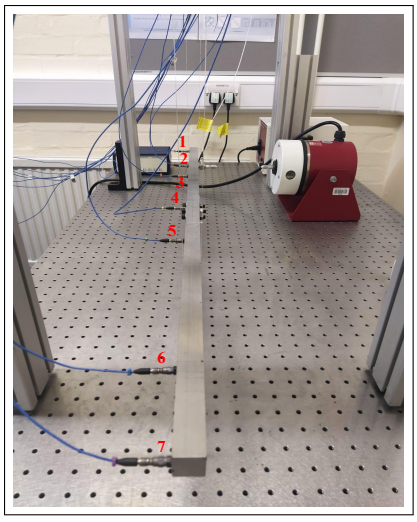

Figure 2. Experimental setup with sensor locations of the Brake-Reuß beam

TABLE I. ORDER OF EXPERIMENTAL TESTS AND BOLT CONFIGURATIONS

Test	Bolt configuration
Normal	All bolts Torque = 20Nm
Damage class 1	All bolts Torque = 10Nm
Damage class 2	All bolts Torque = 5Nm

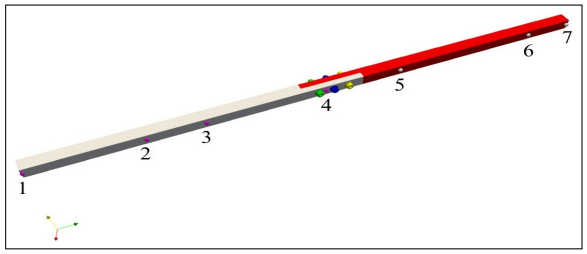

Figure 3. FE model Brake-Reuß beam

response functions (FRFs) were estimated using the H1 frequency response estimator with a frequency resolution of 0.3906 Hz and a bandwidth of 0-12800 Hz. Furthermore, experimental modal analysis was carried out on the SDTools Matlab toolbox to identify modal parameters such as natural frequencies and mode shapes of the first 9 modes. A series of tests were conducted via loosening bolts for two damage classes. The designated bolt configurations for these tests are illustrated in Table I.

Acquiring a high-quality baseline numerical model was critical for creating damage surrogates for experimentally tested damage scenarios. A CAD model is assembled with five components, including two beam components and three bolts. The current study incorporated solid tetrahedron elements in the FE mesh. Figure 3 depicts the model assembly of the beam. Modelling of the joint was done by representing bolt pretensions as directional springs, as in [12]. In this case, five spring interfaces were created by splitting lap joint faces into five divisions. This model was then updated to improve accuracy based on experimental modal parameters. Specifically, the spring interfaces of the modelled joint, were updated using the sensitivity model updating technique [13] by comparing the first five bending modes. The full approach for the joint modelling and FE model updating of the beam model can be found in [14].

THREE CLASS PROBLEM

In a three-class classification problem, the damage classes listed in Table I were taken into account. As a result, the damage classes of bolt pretensions 10 Nm and 5 Nm were assumed to be 50% and 25% of the baseline model, respectively. Damage can be modelled in the FE model by varying the 8 parameters, which include 5 interface stiffnesses and 3 material stiffnesses of the bolts. Two features are chosen for each damage class through feature ranking (see [1]), which yields one peak and three troughs from sensors 2,3,4, and 7. Moreover, a feature set with 500 copies added with white Gaussian noise was implemented to achieve feature variability (see also [1]). This process was repeated consistently for each damage class. Subsequently, as part of the proposed framework (see Figure 1), TCA with domain adaptation was applied. The TCA parameters are de-

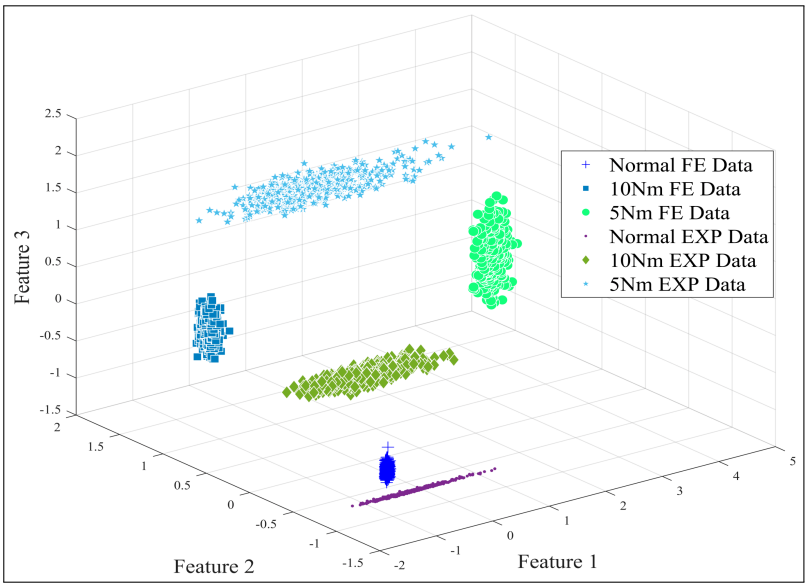

Figure 4. Comparison of FE and experimental features of the three-class problem (Bolt pretension: 20Nm, 10Nm, 5Nm) before TCA application.

termined by multiple factors, including the source and target domains, type of problem, and structure. The ideal parameters for this study were determined iteratively as a linear kernel with a sigma of length 15 and a latent subspace dimension of 25. The Mahalanobis distance was computed using the output features of all damage classes from the previous stage [2, 11]. The comparison of the Mahalanobis distance features before and after TCA is illustrated in Figures 4 and 5.

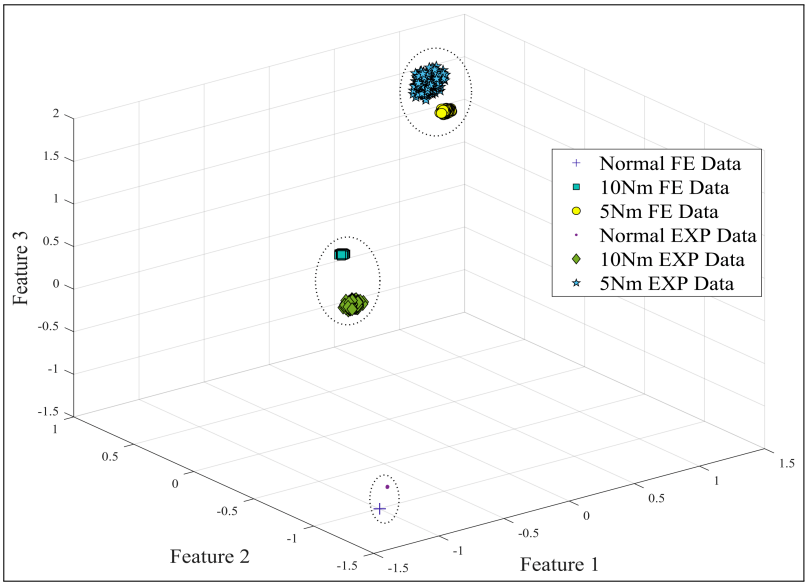

Figure 5. Comparison of FE and experimental features of the three-class problem (Bolt pretension: 20Nm, 10Nm, 5Nm) after TCA application.

TABLE II. TRAINING AND TESTING RESULTS OF KNN CLASSIFIERS BEFORE APPLICATION OF TCA

Name of KNN	K	Distance metric	Training accuracy FE Data	**Testing accuracy Experimental data**
Fine	1	Euclidian	100%	**50.2%**
Medium	10	Euclidian	100%	**48.3%**
Coarse	100	Euclidian	100%	**47.5%**
Cosine	10	Cosine	100%	**47.8%**
Cubic	10	Minkowski (cubic)	100%	**46.9%**
Weighted	10	Euclidian	100%	**48.5%**
Optimisable	2	Mahalanobis	100%	**22.2%**

RESULTS AND CONCLUSION

The FE feature dataset from the previous stage is split into 70% training, 15% validation, and 15% testing sets. These datasets are delivered to the deployed 800 MLPs as input for training. After training, features from the experimental dataset are supplied into the networks for experimental validation. Prior to the application of TCA, only **2** of the 800 neural networks scored adequately with the experimental data. Following the application, **569** out of 800 networks on experimental data worked well (with at least 90% classification accuracy). From this result, the efficiency of the fraction of networks that function effectively on experimental data increases from **0.25%** to **71%**.

Furthermore, when dealing with the three-class problem, KNN classifiers perform similarly to neural networks. The performance of the classifiers on experimental data increased significantly from 50% (approx) to 100% with the application of the proposed framework. Tables II and III exhibit the performance of various KNN classifiers before and after TCA application.

The primary goal of this research was to deal with the problem of data scarcity in supervised machine learning algorithms. Additionally, the aim was to generate surrogates without relying on a large amount of experimental data. This problem is addressed here

TABLE III. TRAINING AND TESTING RESULTS OF KNN CLASSIFIERS AFTER APPLICATION OF TCA

Name of KNN	K	Distance metric	Training accuracy FE Data	**Testing accuracy Experimental data**
Fine	1	Euclidian	100%	**100%**
Medium	10	Euclidian	100%	**100%**
Coarse	100	Euclidian	100%	**100%**
Cosine	10	Cosine	100%	**100%**
Cubic	10	Minkowski (Cubic)	100%	**100%**
Weighted	10	Euclidian	100%	**100%**
Optimisable	1	Cityblock	100%	**100%**

as a three-class problem (normal, damage scenario 1, damage scenario 2), where damage was represented by different torque settings on the lap joint of a laboratory structure. Overall, results reveal that the accuracy of the multiple classifiers greatly increased when TCA with domain adaptation was implemented, while the use of experimental data from damage cases was completely avoided in the training stage of the ML algorithms. In the case of neural networks, the proportion of neural networks that perform well on experimental data has improved from **0.25%** to **71%**. KNN classifiers exhibited similar performance by improving their performance from **50%** (approx) to **100%** on experimental data. In the future, this technique might be used for localisation problems as well as damage classification in more complicated structures.

REFERENCES

1. Papatheou, E., G. Manson, R. J. Barthorpe, and K. Worden. 2010. "The use of pseudo-faults for novelty detection in SHM," *Journal of Sound and Vibration*, 329(12):2349–2366.
2. Papatheou, E., G. Manson, R. J. Barthorpe, and K. Worden. 2014. "The use of pseudo-faults for damage location in SHM: An experimental investigation on a Piper Tomahawk aircraft wing," *Journal of Sound and Vibration*, 333(3):971–990.
3. Pan, S. J., I. W. Tsang, J. T. Kwok, and Q. Yang. 2011. "Domain adaptation via transfer component analysis," *IEEE Transactions on Neural Networks*, 22(2):199–210.
4. Chakraborty, D., N. Kovvali, B. Chakraborty, A. Papandreou-Suppappola, and A. Chattopadhyay. 2011. "Structural damage detection with insufficient data using transfer learning techniques," in *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, vol. 7981, pp. 1175–1183.
5. Gardner, P., X. Liu, and K. Worden. 2020. "On the application of domain adaptation in structural health monitoring," *Mechanical Systems and Signal Processing*, 138.
6. Ozdagli, A. I. and X. Koutsoukos. 2020. "Domain Adaptation for Structural Health Monitoring," in *Annual Conference of the PHM Society*, PHM 2020, pp. 1–9.
7. Poole, J., P. Gardner, N. Dervilis, L. Bull, and K. Worden. 2022. "On statistic alignment for domain adaptation in structural health monitoring," *Structural Health Monitoring*, 0(0):1–20.
8. Farrar, C. R. and K. Worden. 2012. *Structural Health Monitoring: A Machine Learning Perspective*, Wiley, 1st edition edn.
9. Manson, G., E. Papatheou, and K. Worden. 2008. "Genetic optimisation of a neural network damage diagnostic," *Aeronautical Journal*, 112(1131):267–274.
10. Pan, S. J. and Q. Yang. 2010. "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
11. Worden, K., G. Manson, and N. R. Fieller. 2000. "Damage detection using outlier analysis," *Journal of Sound and Vibration*, 229(3):647–667.
12. Lacayo, R., L. Pesaresi, J. Groß, D. Fochler, J. Armand, L. Salles, C. Schwingshackl, M. Allen, and M. Brake. 2019. "Nonlinear modeling of structures with bolted joints: A comparison of two approaches based on a time-domain and frequency-domain solver," *Mechanical Systems and Signal Processing*, 114:413–438.
13. Mottershead, J. E., M. Link, and M. I. Friswell. 2011. "The sensitivity method in finite element model updating: A tutorial," *Mechanical Systems and Signal Processing*, 25(7):2275–2296.
14. R. S. Battu, K. Agathos, C. Smith, and E. Papatheou. 2022. "Robust training databases for supervised learning algorithms in structural health monitoring applications," in *Conference Proceedings of ISMA2022 - USD2022*, K U Leuven, Leuven, pp. 3671–3681.