XLIII CILAMCE
21–25, November, 2022
FOZ DO IGUAÇU – PR – BRAZIL

Structural damage detection with autoencoding neural networks

Lucas V. Resende[1], Rafaelle P. Finotti[2], Flavio S. Barbosa[1], Alexandre A. Cury[1]

[1]*Graduate Program in Civil Engineering, Federal University of Juiz de Fora*
José Lourenço Kelmer St., Juiz de Fora, 36036, Minas Gerais, Brazil
resende.lucas@engenharia.ufjf.br, flavio.barbosa@engenharia.ufjf.br, alexandre.cury@engenharia.ufjf.br
[2]*Graduate Program in Computational Modeling, Federal University of Juiz de Fora*
José Lourenço Kelmer St., Juiz de Fora, 36036, Minas Gerais, Brazil
rafaelle.finotti@engenharia.ufjf.br

Abstract. Structural Health Monitoring (SHM) is a growing field in civil engineering, having relevance for the detection of changes in the state of structures, including the identification of possible damage conditions. Commonly used SHM strategies involve employing Artificial Intelligence (AI) techniques on raw dynamic data measured from structures to perform classifications or extract features from the original data. Among the AI algorithms for SHM, autoencoding neural networks, or simply autoencoders, have been identified as promising solutions, being the focus of this article. An autoencoder is designed to reproduce its inputs as closely as possible after unsupervised training. This characteristic is a useful tool to extract features that represent the original data with lower dimensionality, which facilitates classifications through statistical methods. A sparse autoencoder (SAE) is a type of autoencoder that includes a sparsity penalty at its training process. In that way, this paper presents a methodology to detect structural damage by using sparse autoencoders as parameter extractors applied to signals in the frequency domain, combined with the Mahalanobis distance to perform an unsupervised classification. Tests performed with data extracted from the Z24 bridge have shown promising results in detecting changes in the structural states, demonstrating the potential of SAE for SHM systems.

Keywords: sparse autoencoder, structural health monitoring, damage detection

1 Introduction

In recent decades, many Structural Health Monitoring (SHM) approaches have included continuous vibration monitoring in parts of the structures under analysis via accelerometers placed on strategical positions, with some methods being described by Doebling et al. [1]. The measured vibration data is often studied through modal analysis, more specifically by evaluating the evolution of natural frequencies and mode shapes, which is the case of the work of de Almeida Cardoso et al. [2]. More recently, machine learning based strategies have also been widely adopted in SHM applications (see Salehi and Burgueño [3]). Some of these strategies consist in employing artificial intelligence techniques on dynamic data obtained from structures to either perform classifications, as was done in the work of Nunes et al. [4], or extract features to allow classifications. This is the case of the Sparse Autoencoders (SAE), which have the ability to reduce the dimensionality of the analyzed data, providing interesting parameters for classification algorithms.

SAE reconstructs its inputs through an internal coding - modeled by linear and nonlinear functions - that leads them into a "new" group of variables (features). Auto-encoders are known not only for the ability to deal with large volumes of data but also for their capability to provide optimal solutions, particularly for nonlinear problems, such as structural anomaly detection.

The methodology proposed in the present work provides a possible solution for the detection of structural damage using raw dynamic data. This solution employs SAE models as feature extractors of the dynamic data collected from the Z24 bridge before and after damaging the structure. The features are then classified through an unsupervised method that uses the Mahalanobis distance to separate measurements in different groups associated with each state of structural damage. A similar strategy was proposed by Finotti et al. [5], with a few key differences described below. Finotti et al. [5] does not use dynamic signals in the frequency domain, as proposed by the present paper, but rather the raw data in the time domain. One advantage of processing the data in the frequency domain is that the transformation through the Fast Fourier Transform algorithm can already be understood as a parameter

extractor, which makes it easier for the SAE to encode the information more accurately. In addition to that, Finotti et al. [5] uses data extracted from a single accelerometer at a time, while in this work, five accelerometers are employed together, which should make the classification more precise. Finotti et al. [5] also employs the Hotelling's T-squared statistic instead of the Mahalanobis distance to perform the unsupervised classification.

2 The Z24 bridge

The Z24 bridge was a highway overpass located at Canton Bern near Solothurn, in Switzerland, built in the early 1960s to connect Koppigen and Utzenstorf. This post-tensioned structure was 58m long composed of three continuous spans with 14m, 30m and 14m, supported on four piers. Due to the construction project of a new railway line underneath the highway, the bridge had to be demolished. Before the structure was completely knoked down, it was instrumented and gradually damaged (De Roeck et al. [6]). The experimental setup is summarized in Fig. 1.

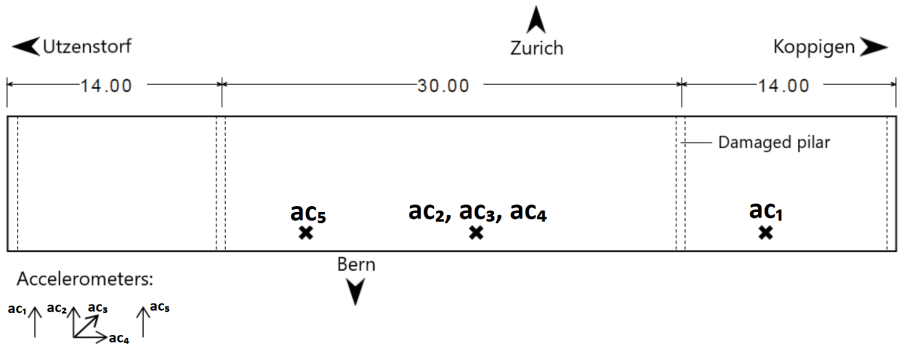

Figure 1. Top view and experimental setup of the Z24 Bridge.

In this work, the proposed SHM strategy is assessed using dynamic signals from forced vibration tests conducted in three structural states of the Z24, which will be called as "nd" (no damage), "d1" (first damage scenario), and "d2" (second damage scenario). Vibration responses from each of the five accelerometers have 65535 data points, acquired at a frequency of 100 Hz. Nine dynamic tests were performed at each damage state, with measurements being made by five accelerometers; ac_1 to ac_5, positioned in the locations and directions shown in Fig. 1. Figure 2 shows a typical response from a dynamic test in the undamaged condition, measured by accelerometer ac_1.

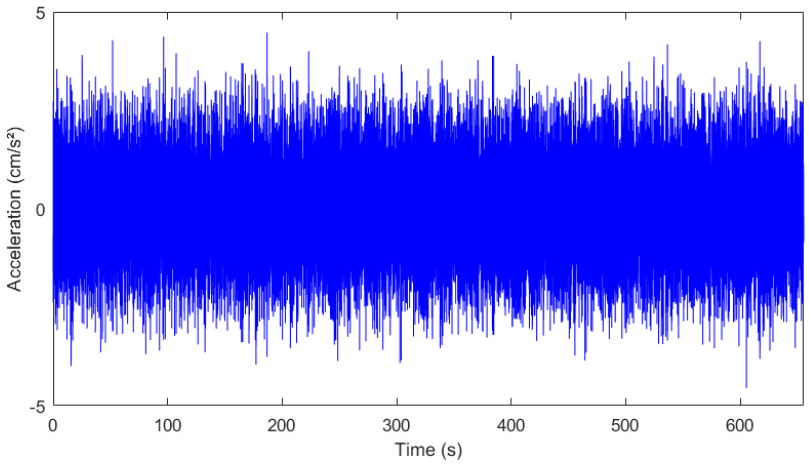

Figure 2. Typical response of dynamic test in undamaged condition - Accelerometer ac_1.

3 Tools used in the proposed methodology

3.1 Autoencoding neural networks

An autoencoding neural network, or just autoencoder, is a type of artificial neural network designed to reproduce its inputs as closely as possible after unsupervised training. An autoencoder is made up of two parts: an encoder, which receives the input data $\mathbf{X}$ and transforms it into a lower dimensional code $\mathbf{h}$; and a decoder, responsible for decoding $\mathbf{h}$ into an output $\mathbf{X}^{'}$ that closely matches $\mathbf{X}$. In other words, autoencoders attempt to reconstruct their inputs as outputs, while generating in the process a vector $\mathbf{h}$ that holds much of the information contained within the training data in fewer dimensions. Therefore, autoencoders are useful tools for extracting features that represent the original data with lower dimensionality, which facilitates a classification through statistical methods. Figure 3 shows a schematic of the basic structure of an autoencoder that encodes a 5-dimension input vector into a 3-dimension code.

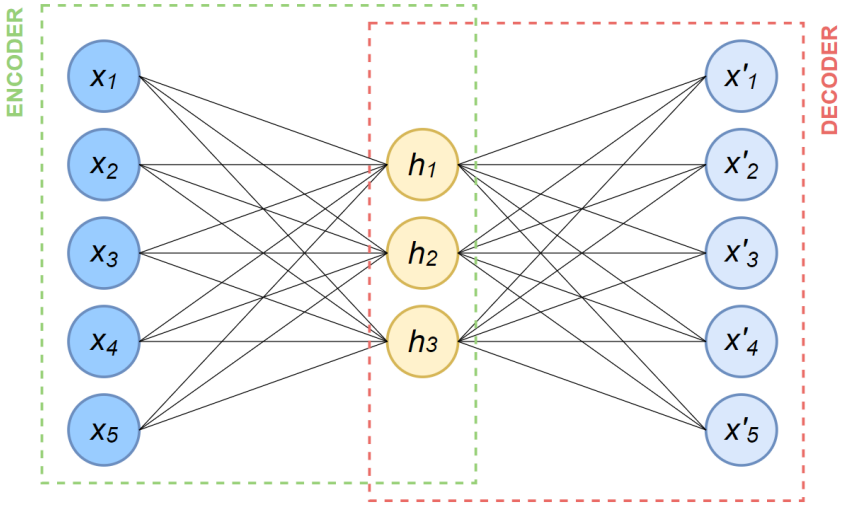

Figure 3. Basic structure of an autoencoding neural network.

A Sparse Autoencoder (SAE) is a type of autoencoding neural network that includes a sparsity penalty term at its training function, being particularly useful to represent large datasets with a small number of $\mathbf{h}$ components. The function that calculates the error between the input and the output is minimized in the training process (see Goodfellow et al. [7]). In the present work, the chosen cost function for the SAE was the mean squared error with sparsity regularizers and the training algorithm was the scaled conjugate gradient backpropagation. The adopted autoencoder had only three layers: an input layer $\mathbf{X}$, a hidden layer $\mathbf{h}$ containing 10 neurons, and an output layer $\mathbf{X}'$.

There are a few hyperparameters that must be optimized for SAE models to be able to represent the training data while still being able to generalize. In this work, those hyperparameters are the Sparsity Proportion (ρ), the L2 Weight Regularization (λ), and the Sparsity Regularization (β). The optimization was performed empirically through a grid search process, as described in section 4.1.

3.2 Control Chart

The Control Chart is a graphical statistical tool used to monitor the variability of a problem's parameters over time. The charts usually depict several data points, which are formed by a specific statistical characteristic and horizontal lines (control limits) responsible for indicating the extreme values of such characteristic when the problem is in-control state. Any point that is beyond these predetermined limits, on the other hand, reveals unusual sources of variability, suggesting an out-of-control situation (Montgomery [8]).

Due to their relatively simple implementation, intuitive interpretation, and effective results, control charts are considered suitable tools for structural on-line monitoring and anomaly detection. In the present work, the characteristic plotted in this chart is the Mahalanobis distance, which is a unitless measure of the distance between an observation and a reference distribution. This metric can be calculated as shown in eq. (1).

$$T = \sqrt{(\mathbf{x} - \boldsymbol{\mu})^{\mathrm{T}} \, \mathbf{S}^{-1} \, (\mathbf{x} - \boldsymbol{\mu})}. \tag{1}$$

In eq. (1), T is the Mahalanobis distance between an observation $\mathbf{x} = (x_1, x_2, x_3, \ldots, x_N)^{\mathrm{T}}$ on $\mathbb{R}^N$ and a distribution Q with mean $\boldsymbol{\mu} = (\mu_1, \mu_2, \mu_3, \ldots, \mu_N)^{\mathrm{T}}$ and positive-definite covariance matrix $\mathbf{S}$ (De Maesschalck et al. [9]). It is worth noting that in the calculations presented in this paper, $\mathbf{S}$ was often found to be singular and thus could not be inverted. For that reason, $\mathbf{S}^{-1}$ was approximated by the Moore-Penrose pseudoinverse of matrix $\mathbf{S}$ (see Ben-Israel and Greville [10]). In this work, the Upper Control Limit (UCL) is defined as the 95$^{\text{th}}$ percentile of the T values of the reference distribution, meaning values greater than the UCL should be observed by chance only 5% of the time. The Lower Control Limit (LCL) is zero.

4 Proposed methodology

4.1 Data division

Initially, all acceleration signals were normalized between -1 and 1 through the division of all points by the maximum absolute acceleration value measured by each accelerometer. For the SAE training, the normalized dynamic signals were divided in smaller subsets of 501 data points (5 second signals). Given that 9 vibration tests were performed in each structural state of the Z24, a total of 1170 vectors of 501 dimensions were generated per accelerometer per structural state. These vectors were then converted to the frequency domain through a Fast Fourier Transform algorithm. By excluding the negative frequency values, the resulting vectors had 251 positions. Figure 4 shows a typical 251 dimension signal in the frequency domain.

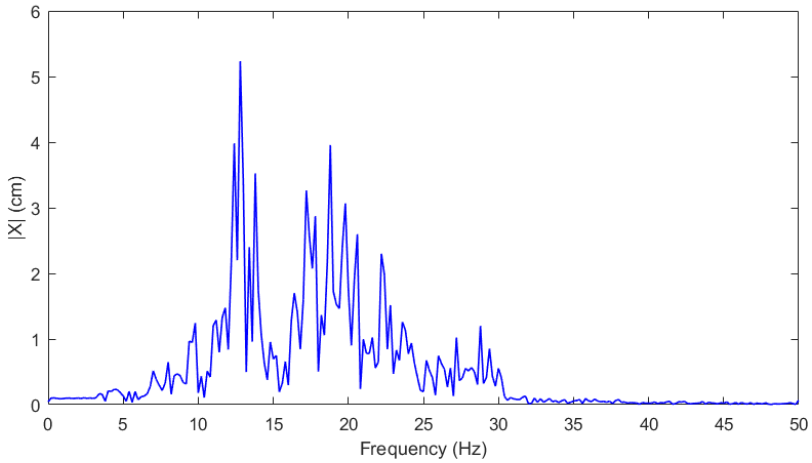

Figure 4. Typical frequency domain signal from accelerometer ac_1.

The signals in the frequency domain from each accelerometer were organized as in eq. (2), where $\mathbf{X}_n$ is a 251×3510 matrix containing the 3510 vectors of samples in the frequency domain for accelerometer ac_n. $X^{\text{nd}}_{i,j}$ refers to the no damage structural condition, whereas $X^{\text{d1}}_{i,j}$ and $X^{\text{d2}}_{i,j}$ are the two defined damage conditions, respectively. The data contained in $\mathbf{X}_1$, $\mathbf{X}_2$, $\mathbf{X}_3$, $\mathbf{X}_4$, and $\mathbf{X}_5$ will be used for SAE training.

$$\mathbf{X}_n = \begin{bmatrix} X^{\text{nd}}_{1,1} & X^{\text{nd}}_{1,2} & \cdots & X^{\text{nd}}_{1,1170} & X^{\text{d1}}_{1,1} & \cdots & X^{\text{d1}}_{1,1170} & X^{\text{d2}}_{1,1} & \cdots & X^{\text{d2}}_{1,1170} \\ X^{\text{nd}}_{2,1} & X^{\text{nd}}_{2,2} & \cdots & X^{\text{nd}}_{2,1170} & X^{\text{d1}}_{2,1} & \cdots & X^{\text{d1}}_{2,1170} & X^{\text{d2}}_{2,1} & \cdots & X^{\text{d2}}_{2,1170} \\ X^{\text{nd}}_{3,1} & X^{\text{nd}}_{3,2} & \cdots & X^{\text{nd}}_{3,1170} & X^{\text{d1}}_{3,1} & \cdots & X^{\text{d1}}_{3,1170} & X^{\text{d2}}_{3,1} & \cdots & X^{\text{d2}}_{3,1170} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ X^{\text{nd}}_{251,1} & X^{\text{nd}}_{251,2} & \cdots & X^{\text{nd}}_{251,1170} & X^{\text{d1}}_{251,1} & \cdots & X^{\text{d1}}_{251,1170} & X^{\text{d2}}_{251,1} & \cdots & X^{\text{d2}}_{251,1170} \end{bmatrix} \tag{2}$$

4.2 Grid search

As explained in section 3.1, a grid search process had to be performed in order to find the best hyperparameters ρ, λ, and β for the SAE model, as well as the best encoder activation function. A grid search for a SAE consists in training the autoencoders with different combinations of parameters until finding a combination that provides the best model, that is, until the SAE model is able to generate as outputs, artificial signals that reproduce the inputs with the smallest average error (Touati et al. [11]). In this work, the lists of hyperparameters for evaluation were: $\rho = \{0.0063,\ 0.0125,\ 0.025,\ 0.05,\ 0.1,\ 0.2,\ 0.4,\ 0.8\}$; $\lambda = \{0.0001, 0.001, 0.01, 0.1\}$; $\beta = \{0.5, 1, 2, 4, 8\}$. The encoder transfer functions tried were the log-sigmoid transfer function and the saturating linear transfer function. Therefore, the grid search was done through 320 iterations ($8\rho \times 4\lambda \times 5\beta \times 2$ activation functions).

A tenth of the signals from accelerometer ac_1 was randomly selected for training, with the remaining 90% being reserved for validation. Each SAE model was defined as having ten hidden neurons, thus producing a vector $\mathbf{h}$ with ten dimensions containing the features that represent each 251-dimensional signal. Model training of a model ended after 200 epochs. After the iterative process, it was found that the SAE model that resulted in the lowest mean squared error for the validation data had the following hyperparameters: $\rho = 0.4$, $\lambda = 0.001$, $\beta = 4$. The best encoder transfer function was found to be the log-sigmoid function, and the chosen decoder transfer function was a simple linear transfer function.

4.3 SAE training and feature extraction

SAE training was divided in two phases. In phase 1, it was necessary to train the SAE model with data from only one damage condition at a time, in order to make sure the proposed model would correctly classify those signals as belonging to a same class. 10% of the signals from one damage scenario were randomly selected for the training, totaling 117. Five SAE models with the parameters defined in the grid search were then trained, one for each of the five accelerometers. Afterwards, the trained SAE models were used to extract features from all of the 1170 signals available for a structural state. The features were organized in matrix $\mathbf{H}$, shown in eq. (3),

$$\mathbf{H} = \begin{bmatrix} \mathbf{h}_1^1 & \mathbf{h}_2^1 & \mathbf{h}_3^1 & \cdots & \mathbf{h}_{1170}^1 \\ \mathbf{h}_1^2 & \mathbf{h}_2^2 & \mathbf{h}_3^2 & \cdots & \mathbf{h}_{1170}^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{h}_1^5 & \mathbf{h}_2^5 & \mathbf{h}_3^5 & \cdots & \mathbf{h}_{1170}^5 \end{bmatrix} \tag{3}$$

where $\mathbf{h}_i^n$ is a vector containing 10 features corresponding to accelerometer ac_n extracted from a frequency domain sample of index i. This procedure was repeated 30 times per structural state.

Phase 2 consisted in utilizing data from all three damage conditions at the same time in order to analyze the model's capacity to separate the structural conditions. The methodology was basically the same as the one described in phase 1, except for the number of available samples. In that way, 10% of the 3510 signals were randomly selected for training, which was also repeated 30 times.

4.4 Mahalanobis distance classification

After extracting the features by, the Mahalanobis distance T between every sample feature vector and the feature vectors of a reference distribution was computed. The idea is that if a distance calculated for a sample is lower than the UCL, this sample probably belongs to the same structural state class as the reference distribution. On the other hand, if the sample has a higher T value than the UCL, it most likely belongs to a class different to that of the reference distribution.

The reference distribution was defined as containing 351 samples, or 10% of the total. The samples in the reference distribution must belong to the same group and be chronologically continuous. The Mahalanobis distances of the points relative to the reference distribution were computed for both phases of training, with typical results presented in section 5.

5 Results

With the T values of all points calculated for different scenarios, it was possible to plot the results in the form of control charts. Figures 5a, 5c, and 5e depict typical results for phase 1 of training, with only one damage scenario at a time. Meanwhile, figures 5b, 5d, and 5f, show typical results of the proposed strategy for phase 2, which includes all three damage scenarios at the same time. In the charts, orange points are the samples that belong to the reference distribution, blue points are the monitoring data, the red line represents the UCL (95th percentile of the samples in the reference distribution), and the green line is the moving median (taken from an equal number of samples on both sides of a central value) of the T values of 351 samples. The black vertical lines in figures 5b, 5d, and 5f are meant to clearly separate the three structural states.

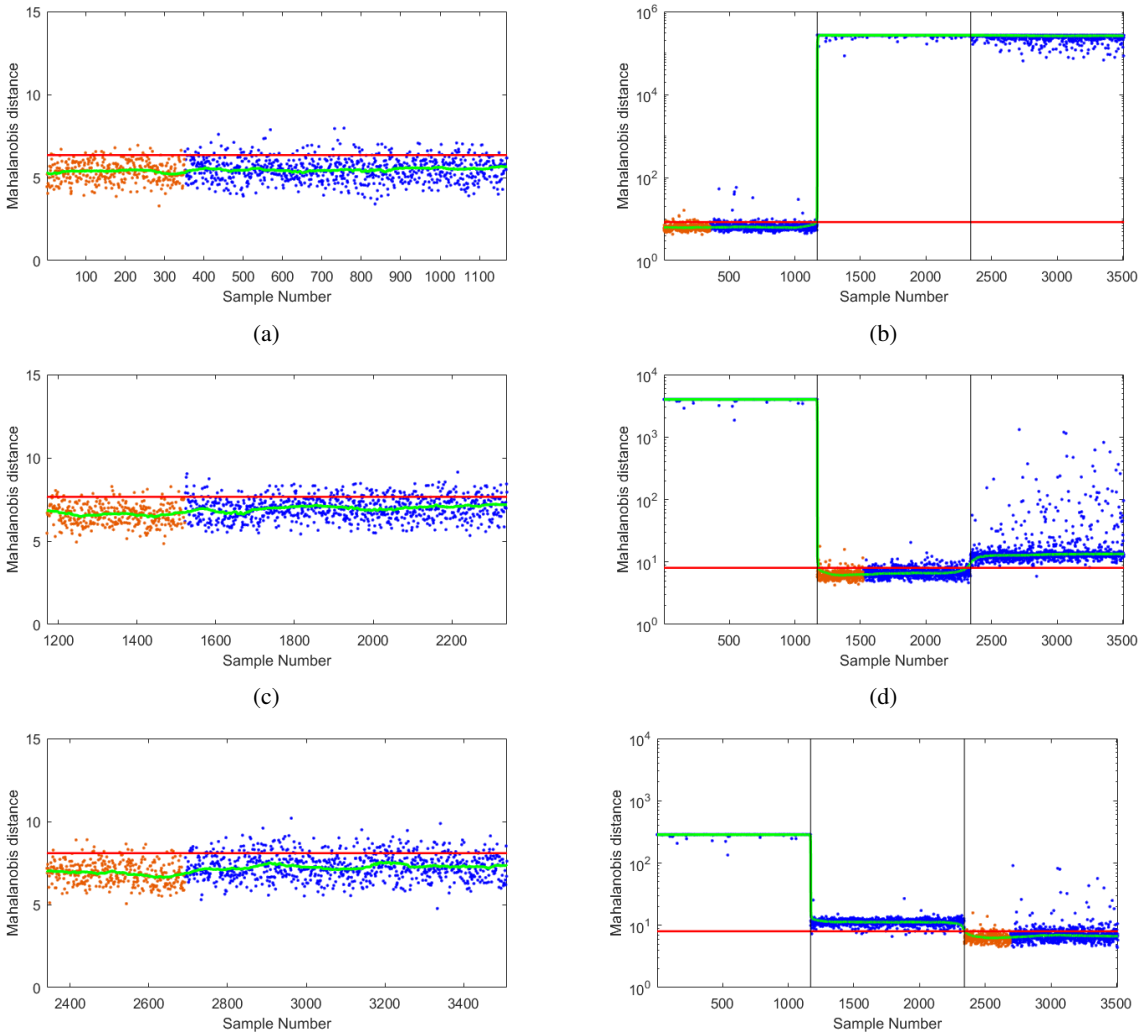

Figure 5. Typical results of the proposed methodology. (a) Training phase 1, scenario "nd", reference distribution indices: 1 to 351. (b) Training phase 2, scenario "nd", reference distribution indices: 1 to 351 (c) Training phase 1, scenario "d1", reference distribution indices: 1171 to 1521. (d) Training phase 2, scenario "d1", reference distribution indices: 1171 to 1521. (e) Training phase 1, scenario "d2", reference distribution indices: 2341 to 2691. (f) Training phase 2, scenario "d2", reference distribution indices: 2341 to 2691.

6 Conclusions

The results obtained with the strategy proposed in this work show that it was possible to separate the three damage scenarios of the Z24 Bridge, thus exemplifying the power of SAE algorithms in feature extraction for

SHM problems. However, the methodology presented in this paper does still have a few shortcomings that should be addressed in future works. For instance, the numbers of samples used for SAE training and for the reference matrix for calculating the Mahalanobis distance, as well as the size of the hidden layer $\mathbf{h}$, were mostly arbitrary. Additionally, while most trained SAE models correctly identified the different structural conditions, not all models did so convincingly, and there were significant differences in the order of magnitude of the computed Mahalanobis distances. Despite that, this work shows that machine learning-based strategies have much potential yet to be unlocked in Structural Health Monitoring.

Acknowledgements. The authors would like to thank CAPES (Coordenação de Aperfeiçoamento de Pessoal de Nível Superior), CNPq (Conselho Nacional de Desenvolvimento Científico e Tecnológico, grants 308008/2021-9-PQ and 304329/2019-3-PQ), and FAPEMIG (Fundação de Amparo à Pesquisa de Estado de Minas Gerais, grants PPM00106-17 and PPM-00001-18).

Authorship statement. The authors hereby confirm that they are the sole liable persons responsible for the authorship of this work, and that all material that has been herein included as part of the present paper is either the property (and authorship) of the authors, or has the permission of the owners to be included here.

References

[1] S. Doebling, C. Farrar, and M. Prime. A summary review of vibration-based damage identification methods. *The Shock and Vibration Digest*, vol. 30, pp. 91–105, 1998.
[2] de R. Almeida Cardoso, A. Cury, and F. Barbosa. A clustering-based strategy for automated structural modal identification. *Structural health monitoring*, vol. 17, n. 2, pp. 201–217, 2018.
[3] H. Salehi and R. Burgueño. Emerging artificial intelligence methods in structural engineering. *Engineering Structures*, vol. 171, pp. 170–189, 2018.
[4] L. Nunes, R. Piazzaroli Finotti, F. Barbosa, and A. Cury. A hybrid learning strategy for structural damage detection. *Structural health monitoring*, vol. 20, n. 4, pp. 2143–2160, 2021.
[5] R. Finotti, C. Gentile, F. Barbosa, and A. Cury. Vibration-based anomaly detection using sparse auto-encoder and control charts. *XI International Conference on Structural Dynamics*, vol. , pp. 1335–1347, 2020.
[6] G. De Roeck, B. Peeters, and J. Maeck. Dynamic monitoring of civil engineering structures. *Computational Methods for Shell and Spatial Structures*, vol. , 2000.
[7] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
[8] D. C. Montgomery. *Introduction to Statistical Quality Control*. John Wiley & Sons, sixth edition, 2013.
[9] R. De Maesschalck, D. Jouan-Rimbaud, and D. Massart. The Mahalanobis distance. *Chemometrics and Intelligent Laboratory Systems*, vol. 50, n. 1, pp. 1–18, 2000.
[10] A. Ben-Israel and T. N. E. Greville. *Generalized inverses: theory and applications*. Springer, Springer New York, NY, second edition, 2003.
[11] R. Touati, M. Mignotte, and M. Dahmane. Anomaly feature learning for unsupervised change detection in heterogeneous images: A deep sparse residual model. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 588–600, 2020.