Chapter 9
Temperature-driven Unsupervised Damage Identification using Cepstral Features

Lorenzo Sclafani, Lorenzo Stagi, Silvia Milana, and Eleonora M. Tronci

Abstract In the field of Structural Health Monitoring (SHM), fast and reliable damage identification is essential to reduce maintenance costs and to ensure safety. Supervised learning techniques require extensive prior knowledge of damaged and undamaged structural conditions. However, in real applications, this information is difficult or even impossible to obtain or predict in advance. Unsupervised approaches, on the other hand, offer a promising alternative for real-time health monitoring, as they do not require pre-labeled data. Nevertheless, most existing unsupervised methods focus solely on binary classification between healthy and damaged states, lacking the ability to distinguish between different types or severities of damage and to adapt to previously unseen conditions such as environmental variability.

To address these limitations, this study presents a novel unsupervised damage detection approach based on cepstral features, which serve as damage-sensitive indicators. Originally developed for acoustic applications, these features are extracted directly from the time histories of the system's response through a straightforward process, minimizing the required user expertise and reducing computational cost.

The proposed method is designed to distinguish between undamaged and damaged conditions and classify different damage classes, while remaining robust to environmental variability. To mitigate the effect of temperature, the explicit use of temperature measurement in combination with Principal Components Analysis is investigated.

In particular, a component selection strategy based on the correlation between temperature variations and principal components is proposed to identify a subspace minimally affected by environmental changes.To highlight the advantages of exploiting temperature information, a second approach is introduced in parallel, where the subspace is instead defined purely on variance considerations, without relying on temperature data.

Once the reduced-dimensional subspace is defined, the Squared Mahalanobis Distance is employed to perform real-time novelty detection and classify new observations, by quantifying the distance of each data point from the centers of the damaged and undamaged classes.

Validation has been carried out using experimental data from the Z24 bridge, demonstrating effectiveness in detecting damage onset and recognizing different incoming unlabeled damage scenarios while showing robustness to environmental variability.

Keywords Structural health monitoring · Damage identification · Cepstral coefficients · Principal component analysis · Z24 bridge

---

Lorenzo Sclafani
Department of Physics and Astronomy "G. Galilei", University of Padua, Padua, Italy
e-mail: lorenzo.sclafani@phd.unipd.it

Lorenzo Sclafani · Lorenzo Stagi · Silvia Milana
Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Rome, Italy
e-mail: lorenzo.sclafani@uniroma1.it; lorenzo.stagi@uniroma1.it; silvia.milana@uniroma1.it

Eleonora M. Tronci
Department of Civil and Urban Engineering, New York University, Brooklyn, NY, USA
e-mail: emt377@nyu.edu

Eleonora M. Tronci
Center for Urban Science and Progress, New York University, Brooklyn, NY, USA
e-mail: emt377@nyu.edu

Kristin B. Zimmerman, PhD (eds.), *Data Science in Engineering, 2026, Vol. 9*, Conference Proceedings of the Society for Experimental Mechanics Series, https://doi.org/10.1007/978-8-743-81448-1_9

Introduction

Structural Health Monitoring (SHM) aims to assess the structural integrity of a system through continuous or periodic monitoring, enabling prompt intervention in the event of faults. Among the core tasks of SHM are: (i) damage detection, as in identifying whether the system is damaged, and (ii) damage classification, which aims to assess the type and severity of the damage. These tasks are crucial not only for ensuring safety but also for enabling optimized maintenance strategies, which carry significant economic benefits. In vibration-based SHM in particular, the problem is addressed by extracting Damage-Sensitive Features (DSFs) directly from measured dynamic responses (e.g., accelerations, velocities, and displacements) and then interrogating their behavior across operating conditions. Because damage is any change in material properties, geometry, or boundary conditions that shifts a structure away from its healthy state [1]; it is possible to compare DSFs extracted for an unknown condition against a healthy baseline to detect and interpret deviations. However, A central challenge is ensuring that such deviations are attributable to persistent structural changes rather than to exogenous environmental or operational variability.

For civil infrastructure, Environmental and Operational Variability (EOV), including seasonal and diurnal temperature cycles, humidity, traffic loading, wind, and shifts in boundary conditions, induces systematic, sometimes substantial, changes in a structure's dynamic characteristics. As a consequence, the extracted DSFs may be biased, leading to potential errors in the damage assessment. In particular, *false positives* can occur when the EOV-induced alteration is mistaken for damage, resulting in unnecessary maintenance and economic losses. On the other hand, *false negatives* appear when EOV effects mask the presence of actual damage and are especially critical, as they directly compromise structural safety. For this reason, the removal or compensation of environmental and operational variability in DSFs is considered a crucial problem in modern SHM.

In the literature, strategies addressing this issue are generally divided into either explicit or implicit approaches [2]. Explicit methods (e.g., subspace identification, regression, and temperature-frequency correlation models [3, 4, 5]) rely on the direct measurement of external variables, such as temperature or humidity, to model and explicitly account for their influence on the structural features. On the other hand, implicit approaches aim to directly extract features that are inherently robust to such variability, without requiring explicit measurement of environmental parameters. These include techniques such as cointegration, Principal Component Analysis (PCA), and various machine learning approaches [6, 7, 8, 9].

To address the EOV problem, the authors have recently proposed an implicit unsupervised approach that leverages two stages of PCA to detect and classify damage using Cepstral Coefficient (CCs) as DSFs [10]. These coefficients, originally developed to identify the presence of an echo in a sound signal [11], have recently emerged as promising alternatives to traditional features such as modal parameters. Their relationship with structural dynamics has been increasingly investigated, and several studies have demonstrated their effectiveness in both supervised and unsupervised machine learning frameworks for the health monitoring of structural and mechanical systems [12, 13, 14, 15, 16, 17, 18].

This study proposes a variation of the same framework, introducing an explicit mitigation of EOV effects on CCs. Temperature measurements are incorporated as additional input parameters, and a correlation analysis between them and the principal components, extracted via PCA on the CCs, enables the identification of the minor components representative of a less temperature-correlated subspace. The algorithm is evaluated in comparison to the implicit approach, in which the minor components are selected based solely on variance considerations, without exploiting environmental information. Both methodologies are validated using an experimental dataset with high temperature variability, and their performance is assessed with respect to both damage detection and damage classification tasks.

At first, a brief overview of the cepstral coefficients and their extraction process is provided, followed by the description of the proposed PCA-based unsupervised algorithm. Subsequently, the case study is introduced, and the procedures adopted for training and testing are discussed. Finally, the results of the application are analyzed before drawing the main conclusions.

Damage sensitive features: Cepstral Coefficients

Cepstral coefficients are a compact numerical representation of the spectral content of a signal. Despite their seemingly abstract nature, their definition and extraction process are simple and straightforward. Given a discrete signal $a[n]$ (e.g., the acceleration time history of the system of interest), a set of D cepstrals is extracted by doing the Inverse Discrete Fourier Transform (IDFT) of the logarithm of the squared magnitude of the Discrete Fourier Transform (DFT) of the signal. Therefore,

$$CC[q] = \mathcal{F}^{-1}\{log(|\mathcal{F}\{a[n]\}|^2)\}, \quad q = 0, \ldots, D-1 \tag{1}$$

where $\mathcal{F}$ is DFT and $\mathcal{F}^{-1}$ represents the IDFT.

The compactness of the representation allows for high computational efficiency, showing good compatibility with large-scale applications. Meanwhile, the straightforward extraction process does not require highly-experienced users, as the decision making is limited to the number D of coefficients to extract, for which some criteria are available in the literature [17]. Furthermore, the CCs require only a small section of the signal to be extracted, allowing for a description of how the system evolves over time. Indeed, the actual feature extraction is usually preceded by two processing phases: *framing* of the measured signal into equal-length segments, and *windowing* of these frames by multiplying them with a hamming window to reduce the onset and offset ripple's effect. Of the D coefficients extracted from each frame, it is common practice in SHM to neglect the first, as its strong sensitivity to the input excitation may hide the damage information.

Overall, their advantages place the cepstral features as a fairly competitive alternative to more commonly used DSFs (e.g., natural frequencies). Because CCs are sensitive to shifts in resonant frequencies and mode shapes, any structural damage alters the coefficients. However they do have a shortcoming: when dealing with complex systems, the link between the coefficients changes and the variations in the structural properties becomes harder to interpret, and with a more complex analytical formulation [17].

A PCA-based unsupervised algorithm

In this work, a real-time unsupervised approach is proposed to classify new incoming data. The algorithm leverages two stages of principal component analysis to reduce the dataset to its most relevant information.

PCA is a widely known unsupervised ML algorithm that reduces the dimensionality of large datasets by projecting them into a lower-dimensional subspace defined by the uncorrelated directions of maximum variance.

After a standardization of the data, the covariance matrix is calculated and the relative eigenproblem is solved to find the N eigenvectors and associated eigenvalues. The eigenvectors represent the orthogonal directions associated with the largest variance, while the eigenvalues quantify the amount of variance carried by each direction. At this point, the Principal Components (PCs) can be defined as the first $p < N$ directions contributing the most to the variance, and the transformation matrix T is obtained simply using the selected eigenvectors, ordered with decreasing eigenvalue, as columns.

The remaining $m = N - p$ components are commonly referred to as Minor Components (MCs), and they describe the directions associated with the lowest variability. When dealing with CCs, projecting the data onto these components mitigates the highly variable effects of external inputs (e.g., environmental influences), which are predominant in the PCs. Meanwhile, the contribution of damage remains appreciable, as fundamental changes to the structural properties are less susceptible to such variability. Although the distinction between principal and minor components is conceptually clear, there is no proper threshold that separates the two, as multiple perspectives can be considered. In this work, as will be further discussed later on, two approaches to achieve this separation are compared: one in which the threshold is defined based on a cumulative percentage of the total variance, and another in which it is determined according to the correlation of the CCs with an external input variable (i.e., temperature).

Given that the detailed description of the algorithm exceeds the purposes of this study, only a brief overview of its main steps is provided here. For an in-depth explanation and step-by-step account of the classification procedure, the reader is referred to the authors' previous work [10]. The overall process can be divided into two stages: a training phase and a subsequent testing phase.

Following the multivariate novelty detection approach, the algorithm is first trained using a dataset representative of the undamaged system to create a reference model of the healthy state. This is done by feeding the training dataset to the PCA algorithm and identifying the minor components and the related transformation matrix.

Once trained, the algorithm transitions to the testing phase. Here, unlabeled observations are considered one at a time to simulate the real-time monitoring of the system. Each incoming datapoint (i.e., a set of cepstral coefficients extracted from a single observation) is first projected into the minor components' subspace defined during training. At this point, a second PCA, this time focusing on the principal components, is employed on the minors to enhance the differentiation between the new incoming datapoint and the previously classified data.

The resulting components, denoted as PC_m, are then used in an outlier analysis based on the Squared Mahalanobis Distance (SMD) [19] to assess whether the new datapoint belongs to any of the known classes. Specifically, the SMD of the unlabeled data is compared against dynamically defined thresholds, which are continuously recalibrated as classes evolve. If the datapoint falls within the threshold of a class, it is assigned accordingly; otherwise, it is marked as an outlier. A new class is created when a predefined number of nearly consecutive datapoints are identified as outliers.

Overall, the main parameters that influence the classification are two:

- N_{PC2}, which refers to the number of PC_m selected during the second PCA;
- $N_{NewClass}$, which is the number of almost consecutive points recognized as outliers needed to create a new class.

Figure 1 shows a schematic overview of both training and testing phases.

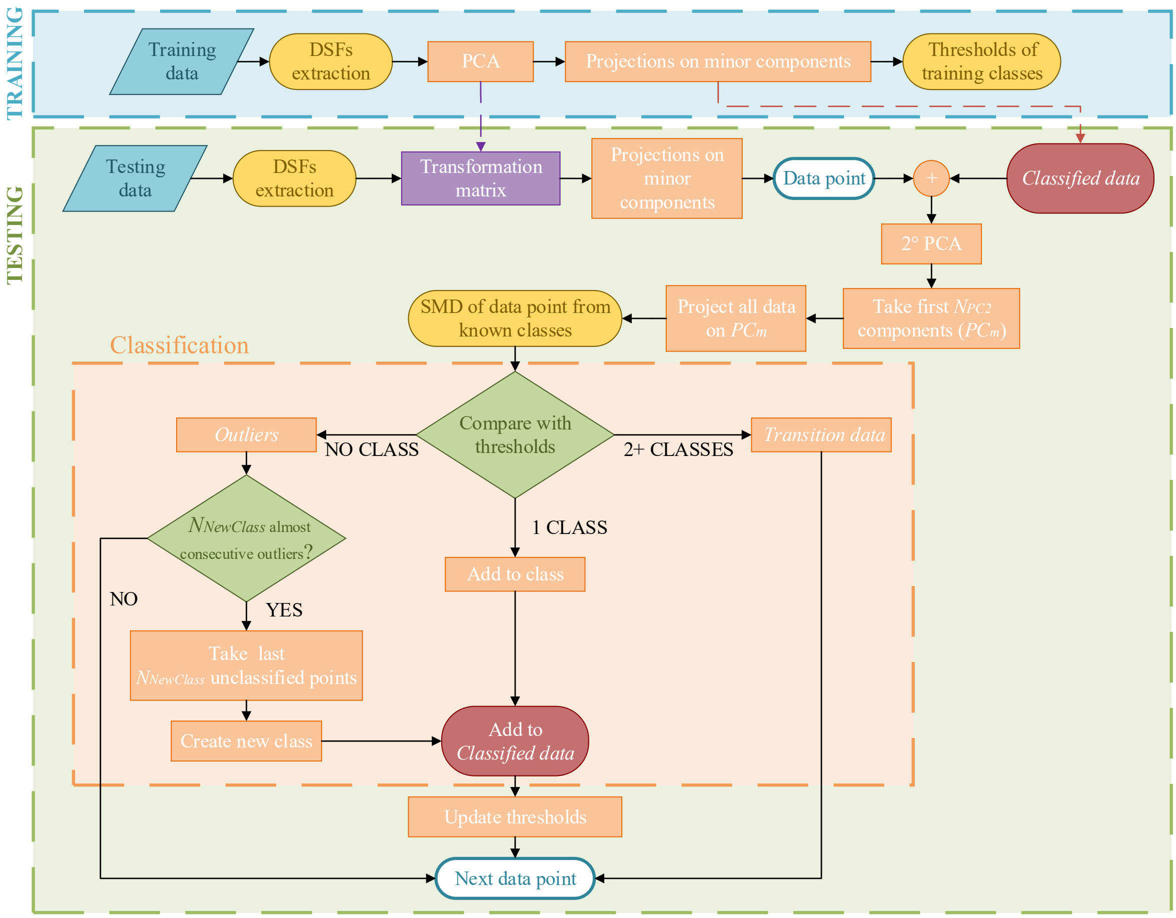

Fig. 1 Flowchart of the proposed algorithm

Finally, the algorithm's performance is evaluated on two tasks: damage detection, i.e., its ability to distinguish between damaged and undamaged states, and damage classification, i.e., its ability to identify and differentiate among different damage scenarios. Accordingly, two performance indices are defined. The *Binary Performance Index* PI_{Binary} corresponds to the accuracy of the algorithm in detecting the presence of damage, expressed as:

$$PI_{Binary} = \frac{\text{Correctly Classified D/U Data}}{\text{Total Predictions}} \tag{2}$$

The multiclass performance index $PI_{MultiClass}$ instead represents the algorithm's ability to classify previously unseen classes. It is defined as the classification accuracy penalized by a factor that depends on the number of extra or missed classes:

$$PI_{MultiClass} = \frac{\text{Correctly Classified Data}}{\text{Total Predictions}} \cdot s \tag{3}$$

where $s = \frac{1}{1+\text{Missed or Extra Classes}}$. Thus, if the algorithm generates an excessive number of classes or fails to recognize existing ones, the performance index decreases accordingly, reflecting the degradation in classification performance.

Case study

An experimental dataset was employed to assess the performance of the proposed methodology. The selected dataset originates from a major European civil engineering project: the Brite EuRam research project BE-3157, *System Identification to*

Monitor Civil Engineering Structures (SIMCES). This project involved the long-term monitoring of a representative structure, the Z24 Bridge in Switzerland, carried out between November 1997 and September 1998, followed by progressive damage tests during the final month. Thanks to its extensive and distinctive nature, the Z24 dataset has become a benchmark in vibration-based structural health monitoring studies [20]. A comprehensive description of the SIMCES project can be found in De Roeck [21].

The Z24 Bridge was a two-cell post-tensioned box-girder concrete bridge spanning the national highway A1 between Utzenstorf and Koppigen. It featured a main span of 30 m and two side spans of 14 m each, with a slightly skewed alignment. The superstructure was supported by triple concrete columns connected to the girder through concrete hinges at the abutments, while the intermediate supports consisted of concrete piers clamped into the girder.

The first stage of the project consisted of long-term continuous monitoring of the bridge in its undamaged condition. The primary objective was to investigate the influence of environmental variability (e.g., temperature, humidity, and wind) on the dynamic properties of the structure. Structural responses were recorded through 16 accelerometers strategically positioned to capture the lowest mode shapes with sufficient resolution. However, due to malfunctions, eight accelerometers failed before the end of the campaign, and one of the remaining eight provided intermittent data. Consequently, only the reliable signals were retained for this analysis. Acceleration measurements were collected every hour in 10-minute time windows, while environmental parameters were sampled in 10 scans immediately before and after each acquisition.

In the final month, just prior to the scheduled demolition of the bridge, a series of 15 progressively induced damage scenarios were implemented to study the behavior of the structure under different degraded conditions. A complete overview of these scenarios is reported in Table 1. The induced damage conditions can be categorized into two groups: scenarios $D01$–$D05$, classified as reversible since the pier settlements were later restored to their original position, and scenarios $D07$ onward, which were irreversible and cumulative [22].

Table 1 Chronological overview of the applied scenarios. Classes used to test the algorithm are highlighted in bold

#	Start date	Description	#	Start date	Description
U00	Nov, 11 1997	Undamaged condition	D07	Aug, 26 1998	Spalling of concrete at soffit, 12 m^2
U01	Aug, 4 1998	Undamaged condition	D08	Aug, 26 1998	Spalling of concrete at soffit, 24 m^2
U02	Aug, 9 1998	Installation of pier settlement system	D09	Aug, 27 1998	Landslide of 1 m at abutment
D01	Aug, 10 1998	Lowering of pier, 20 mm	D10	Aug, 31 1998	Failure of concrete hinge
D02	Aug, 12 1998	Lowering of pier, 40 mm	D11	Sep, 2 1998	Failure of 2 anchor heads
D03	**Aug, 17 1998**	**Lowering of pier, 80 mm**	D12	Sep, 3 1998	Failure of 4 anchor heads
D04	**Aug, 18 1998**	**Lowering of pier, 95 mm**	**D13**	**Sep, 7 1998**	**Rupture of 2 out of 16 tendons**
D05	Aug, 19 1998	Lifting of pier, tilt of foundation	**D14**	**Sep, 8 1998**	**Rupture of 4 out of 16 tendons**
D06	Aug, 20 1998	New reference condition	**D15**	**Sep, 9 1998**	**Rupture of 6 out of 16 tendons**

Training

The undamaged periods used for the training phase were selected based on temperature variations. To properly capture the correlation between temperature and the principal components, both time windows characterized by different absolute temperature levels and those exhibiting significant temperature changes were considered. Figure 2 illustrates the temperature evolution during the monitoring campaign, with the selected training windows highlighted by colored rectangles. The yellow, blue, and red rectangles correspond to periods of nearly stable temperatures at approximately 4, –4, and $20\,°C$, respectively, whereas the green rectangle denotes a single interval with pronounced temperature variations.

Once the training data were defined, two different strategies were adopted for the subsequent analysis. The first strategy consisted of selecting, from the first PCA, only those components contributing less than 0.5% of the total variance. The second strategy, which explicitly exploits the temperature information, involved discarding all components whose correlation coefficient with temperature exceeded 0.2, measured through the non-linear Spearman correlation coefficient. The rationale behind these two strategies, both aimed at mitigating the effect of temperature on the DSFs, is illustrated in Figure 3. In particular, Figure 3a reports the bar plot of the temperature-correlation coefficients across the components, highlighting the discarded ones (i.e., those above the threshold) and showing the effects of the correlation-driven selection. The insets compare one of the rejected components with one of the retained ones, clearly revealing the non-linear dependence of the former on environmental variability, which is no longer appreciable in the latter. The idea behind this selection is to discard,

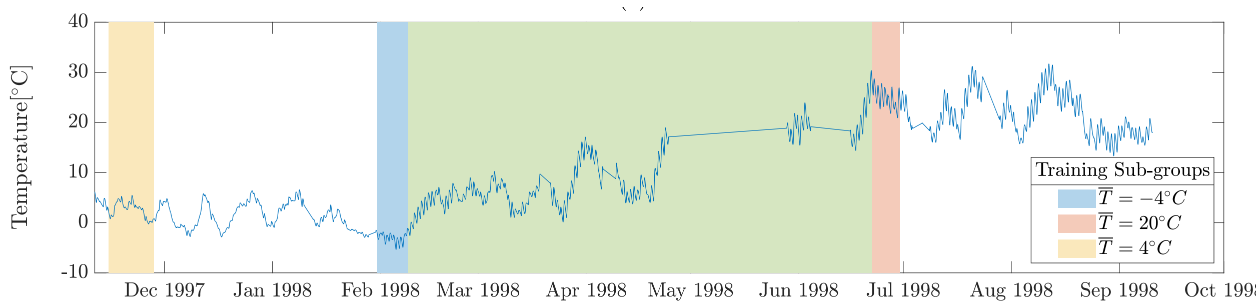

Fig. 2 The temperature during the monitoring period. The colored rectangles show the periods considered in the training phase

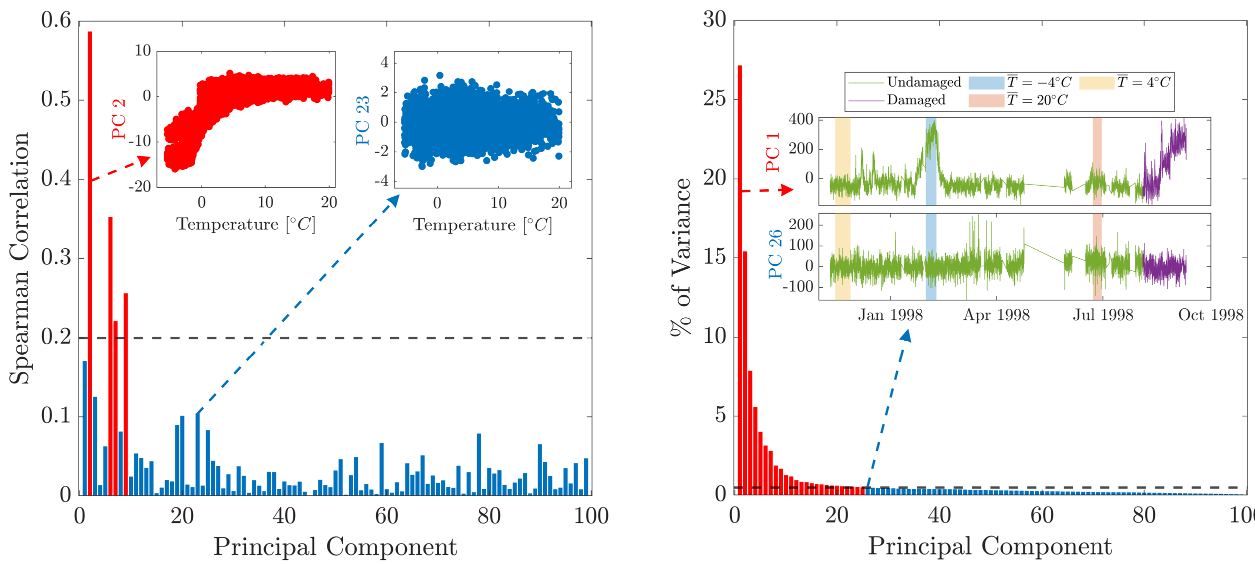

(a) (b)

Fig. 3 (a) Spearman correlation between each principal component and temperature. Insets show the projections of the cepstral coefficients onto two selected components: PC 2, which exhibits strong nonlinear dependence on temperature, and PC 23, which is nearly uncorrelated. (b) Percentage of variance explained by each principal component. Insets illustrate the temporal evolution of two projections: PC 1, which mainly reflects environmental variability and shows marked fluctuations at low temperatures (blue window), and PC 26, which is largely insensitive to temperature changes

in a selective manner, the components most affected by temperature, thereby mitigating environmental sensitivity while preserving the overall information.

Similarly, Figure 3b shows the selection of components based on the variance-driven threshold. In this case, the effect is illustrated by comparing the temporal evolution of the first principal component with that of the first minor component. The discarded component strongly reflects the influence of temperature, exhibiting significant deviations during the low-temperature period (blue window). By contrast, the projection onto the first retained PC (i.e., the 26th component) appears insensitive to environmental variability. However, it is also evident that, compared to the previous method, this approach leads to a greater loss of information, and as a result, the deflections associated with damage are mitigated as well, as shown in the insets.

Therefore, the main advantage of the temperature-driven strategy lies in its ability to selectively choose the components to retain. This makes it possible to preserve some of the leading components characterized by large variance (i.e., high information content) but low temperature correlation, thereby constructing a reduced subspace that is insensitive to environmental effects yet highly sensitive to structural damage.

Testing

For the testing phase, the two approaches described during training, one exploiting temperature information and the other based solely on the minor components, were compared using two sets of testing conditions. Undamaged data from the three temperature subgroups, not used during the training phase, were also included to assess the algorithm's sensitivity to EOV effects.

The first set (Case 1) considered two specific damage classes: D04 and D15. These scenarios were chosen because they represent, respectively, the most severe reversible and irreversible damage cases, thus providing a suitable benchmark to preliminarily assess the discriminative capability of the algorithm before performing more refined analyses.

The second set (Case 2) was designed to be more challenging for the algorithm. In this case, the undamaged testing data were considered alongside grouped damage classes: D03 and D04 were merged into a single class, while D13, D14, and D15 formed a second class. The reason behind this grouping is that the structural differences among the damages within each group are minimal, making it unnecessary to distinguish them individually. Consequently, when data from D03 are tested, the algorithm is expected to create a new class and subsequently associate the D04 data with it. Similarly, when analyzing data from D13, the algorithm should define a new class and recognize subsequent damages D14 and D15 as part of the same class.

Results

Different combinations of the parameters of interest (N_{PC2} and $N_{NewClass}$) were investigated to assess the algorithm's performance. Each combination was tested over 10 repetitions to obtain a more reliable indicator. In the plots, the color scale represents the mean value of the binary or multiclass index across the 10 repetitions, while the contour lines indicate the standard deviation, thus providing insight into the variability of the results. The x-axis corresponds to the number of components retained in the second PCA, while the y-axis represents $k = N_{NewClass} - N_{PC2}$. This definition ensures that the number of observations remains larger than the number of features (i.e., N_{PC2}), which is a strict requirement for performing PCA.

Case 1.

Figure 4 reports the results obtained when only damage classes D04 and D15 were considered in the testing phase. The first row shows the performance with the binary index, whereas the second row refers to the multiclass index. For the binary case, no substantial differences are observed, confirming that both approaches successfully distinguish between damaged and undamaged data. In addition, stable performance ($< 94\%$) is achieved across all parameter combinations, with very low variance across repetitions, indicating both robustness and reliability of the results. In contrast, major differences between the two approaches arise for the multiclass index. When temperature information is used to guide the selection of the principal components, better performance is achieved: the index exceeds 0.5 for several parameter combinations, reaching a maximum value of 73%,indicating that the algorithm correctly identifies the three bridge conditions. Conversely, when the variance criterion is adopted, the index remains below 0.5 with low variability for a wide range of parameters (see the blue area in the bottom-right plot of Figure 4), with a maximum value of 68%, suggesting that the algorithm systematically fails to recognize one of the damage classes. This effect is explained by the term s in equation 3, which halves the index value in the case of a missed class. Only for a few parameter combinations the multiclass index reaches values comparable to those obtained when temperature is considered.

To gain further insight into the results, Figure 5 illustrates the evolution of the algorithm, datapoint by datapoint, during the testing phase. The plots were obtained using the temperature-driven strategy, although analogous results are observed when the variance-based strategy is adopted. In particular, cases where the performance index drops below 0.5—such as those forming the blue plateau in Figure 4—correspond to parameter combinations for which the algorithm fails to reliably distinguish class D15 from class D04 (Figure 5a). Conversely, high performance index values indicate a successful separation of the three structural conditions, namely the undamaged state, damage class D04, and damage class D15 (Figure 5b).

A key aspect emerging from the plots concerns the role of *outliers* and *true outliers*. The algorithm classifies as *outliers* all datapoints that cannot be assigned to any class, i.e., those lying outside the defined thresholds, negatively affecting the performance index. However, in experimental datasets, sensor reliability and measurement accuracy must be carefully considered. In the Z24 bridge case, several records are affected by measurement noise or acquisition errors, likely due to malfunctions of the data acquisition system. Such anomalies produce datapoints that markedly deviate from other observations of the same class. These datapoints were identified as *true outliers* when deviating more than three median absolute deviations from the class median. Treating these samples as misclassifications would distort the evaluation of the algorithm.

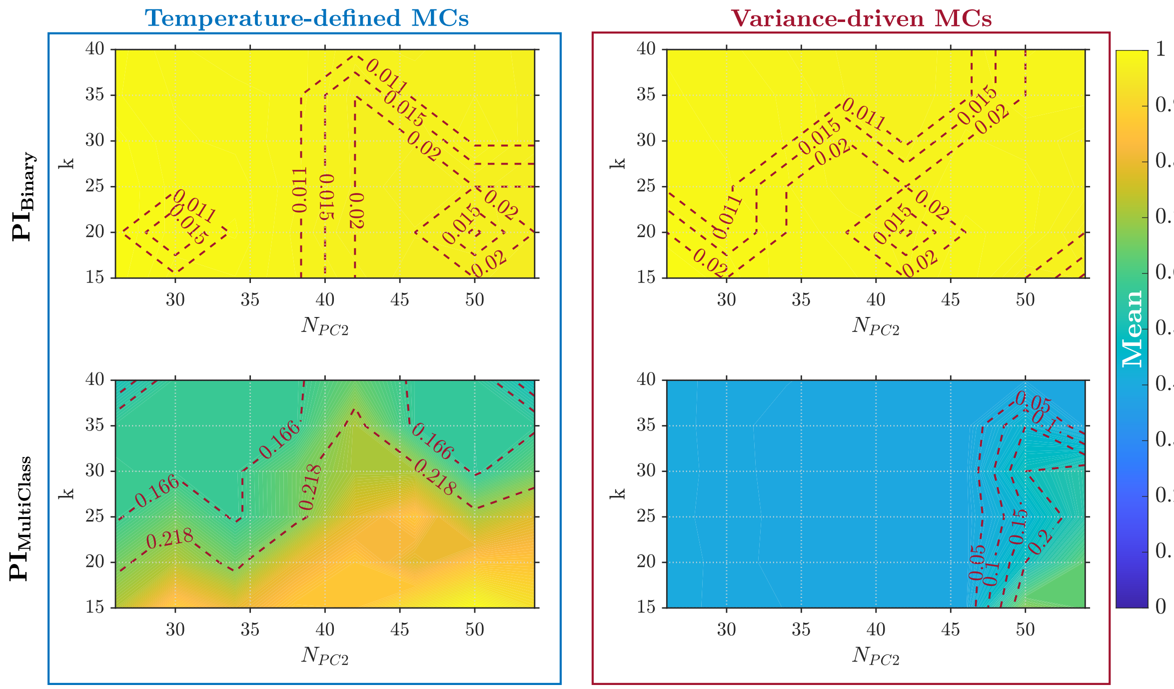

Fig. 4 Contour plots of the mean (color scale) and standard deviation (contour lines) of the performance indices obtained over 10 repetitions for different combinations of N_{PC2} and k in Case 1 (damage classes D04 and D15). The first row refers to the binary index (PI_{Binary}), while the second row refers to the multiclass index ($PI_{MultiClass}$). The left column corresponds to the case where temperature is used as input for the selection of the components, whereas the right column corresponds to the variance-based criterion

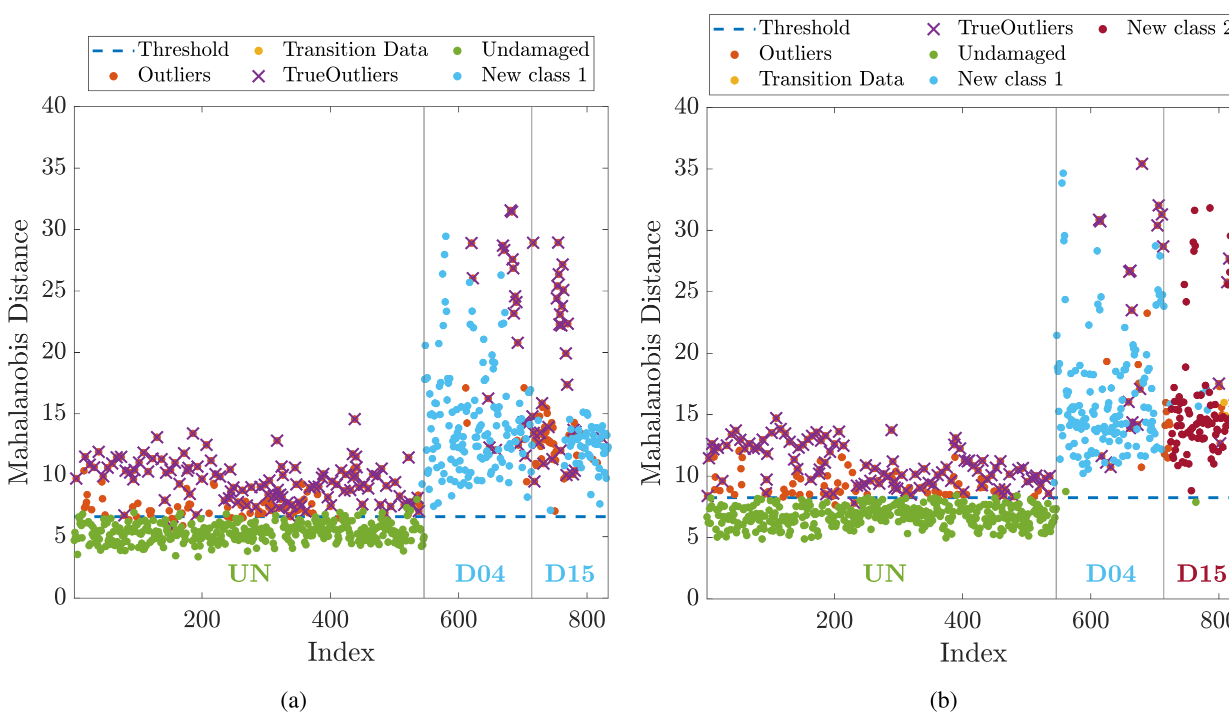

Fig. 5 Squared Mahalanobis distance of points to the undamaged class and their classification for temperature-driven strategy. Results from a single repetition for (a) worst case $N_{PC2} = 30$, $k = 35$; (b) best case $N_{PC2} = 50$, $k = 15$

Therefore, they were included among the *Correctly Classified Data*. Their explicit representation in the plots provides the reader with an indication of their influence on the overall results.

Case 2.

Figure 6 shows the performance indices obtained when classes D03, D04, D13, D14, and D15 are considered in the testing phase, in addition to the undamaged data. It is important to recall that the algorithm is not required to distinguish every individual class, but rather to separate two macro-classes: one grouping data from D03 and D04, and the other grouping the remaining damage classes.

As in Case 1, the binary performance index demonstrates consistently good results ($< 98\%$) across all parameter combinations and repetitions (low variance), both when temperature information is used and when the variance-driven strategy is adopted. In contrast, for the multiclass index, the performance gap between the temperature-based and variance-based approaches is more pronounced than in Case 1. When temperature is employed, good and homogeneous results are obtained across all parameter combinations (maximum 88%). Conversely, the variance-based approach exhibits the same trend already observed in Case 1: a large blue area in the plot indicates low index values with low variance, corresponding to poor performance. Only when a large number of components are retained in the second PCA, combined with small values of k, does the variance-based strategy achieve results comparable to those obtained using the temperature information, with a maximum value of 73%.

To further investigate these outcomes, Figure 7 illustrates examples of both successful and unsuccessful classifications. As in the previous case, the plots are obtained with the temperature-driven strategy, but analogous results hold when the variance-based approach is used. In the worst case (Figure 7a), when the algorithm processes data from class D03, it creates a new class and assigns data from class D04 to it as expected. However, when data from class D13 are analyzed, instead of creating a new class, the algorithm assigns them to the previously created class. The same misclassification occurs for classes D14 and D15. The failure to create a new class halves the multiclass performance index. On the other hand, in the best case (Figure 7b), the algorithm correctly distinguishes between the two macro-classes, namely one comprising classes D03 and D04 and the other including classes D13, D14, and D15, leading to a high performance index.

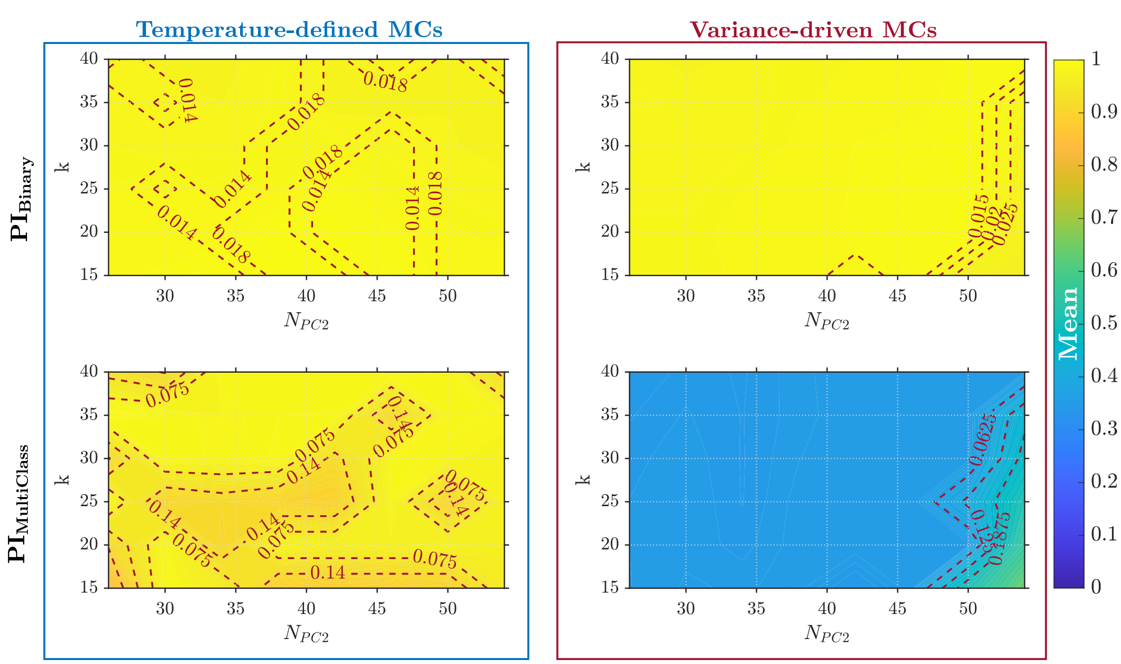

Fig. 6 Contour plots of the mean (color scale) and standard deviation (contour lines) of the performance indices obtained over 10 repetitions for different combinations of N_{PC2} and k in Case 2 (damage classes D03, D04 and D13, D14, and D15). The first row refers to the binary index (PI_{Binary}), while the second row refers to the multiclass index ($PI_{MultiClass}$). The left column corresponds to the case where temperature is used as input for the selection of the components, whereas the right column corresponds to the variance-based criterion

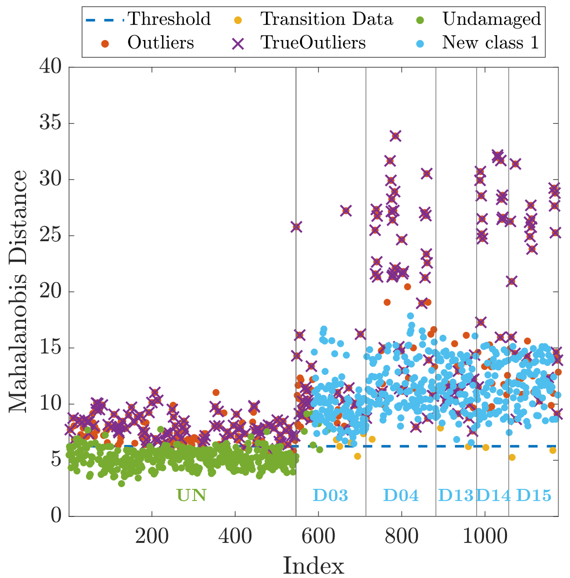

(a)

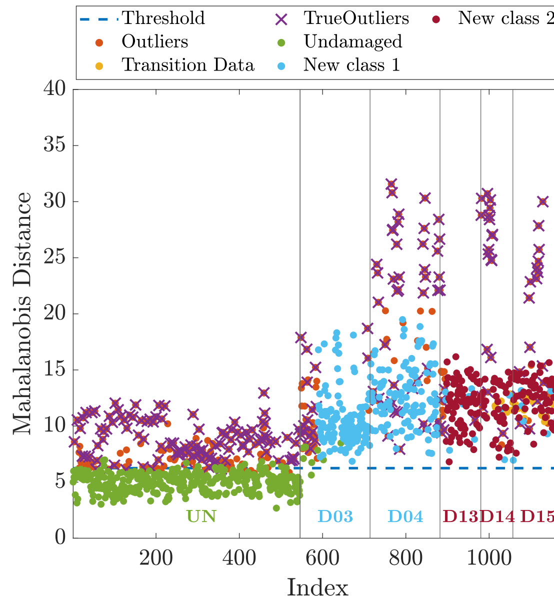

(b)

Fig. 7 Squared Mahalanobis distance of points to the undamaged class and their classification for temperature-driven strategy. Results from a single repetition for (a) worst case $N_{PC2} = 26$, $k = 15$; (b) best case $N_{PC2} = 26$, $k = 35$

Conclusions

This study presented an unsupervised methodology for quasi-real-time damage detection and classification based on cepstral features as damage-sensitive indicators. The approach was designed to mitigate environmental influences while enabling the recognition of new damage conditions. Undamaged data were first used to define a reduced-dimensional subspace via Principal Component Analysis. Two strategies for component selection were investigated: a temperature-driven approach, which discards the components most affected by nonlinear correlations with temperature (i.e., explicit approach), and a variance-driven approach, which excludes components carrying more than 0.5% of the variance without using environmental information (i.e., implicit approach). Testing data were then progressively projected into this reduced space, and a second PCA step was employed to enhance separability and support new class creation.

The methodology was validated on experimental data from the monitoring campaign of the Z24 bridge. A multivariate analysis of parameter combinations provided insights into the influence of the number of retained components in the second PCA (N_{PC2}) and the number of data points needed to create new classes.

Results indicate that both strategies achieve near-perfect binary classification between undamaged and damaged conditions. However, when the classification of multiple damage scenarios is required, the temperature-driven strategy clearly outperforms the variance-driven one. In the case of damage classes D04 and D15, the temperature-based approach reached a maximum accuracy of 73%, with robust performance across a broad range of parameter combinations, whereas the variance-based method achieved lower accuracy (68%) and only in a narrow parameter region. The performance gap becomes even more evident when classes D03, D13, and D14 are included: the explicit approach achieved results comparable to binary classification (88%), while the variance-driven approach again provided satisfactory results (65%) only in limited parameter settings.

In summary, the proposed algorithm demonstrates excellent performance when temperature information is exploited during training to guide component selection, highlighting a promising direction toward robust and reliable strategies for structural damage detection and classification.

Acknowledgments The authors sincerely acknowledge the Katholieke Universiteit Leuven, Structural Mechanics Section, for providing the data collected within the SIMCES project.

References

1. Farrar, C.R. and Worden, K. “An introduction to structural health monitoring”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851):303–315 (2007).
2. García Cava, D., Avendaño-Valencia, L.D., Movsessian, A., Roberts, C., and Tcherniak, D. “On explicit and implicit procedures to mitigate environmental and operational variabilities in data-driven structural health monitoring”. In *Structural health monitoring based on data science techniques*, pages 309–330. Springer (2021).
3. Peeters, B. and De Roeck, G. “Reference-based stochastic subspace identification for output-only modal analysis”. *Mechanical systems and signal processing*, 13(6):855–878 (1999).
4. Zhang, Y.M., Wang, H., Bai, Y., Mao, J.X., and Xu, Y.C. “Bayesian dynamic regression for reconstructing missing data in structural health monitoring”. *Structural Health Monitoring*, 21(5):2097–2115 (2022).
5. Roberts, C., Avendaño-Valencia, L.D., and Cava, D.G. “Robust mitigation of eovs using multivariate nonlinear regression within a vibration-based shm methodology”. *Mechanical Systems and Signal Processing*, 208:111028 (2024).
6. Deraemaeker, A., Reynders, E., De Roeck, G., and Kullaa, J. “Vibration-based structural health monitoring using output-only measurements under changing environment”. *Mechanical systems and signal processing*, 22(1):34–56 (2008).
7. Deraemaeker, A. and Worden, K. “A comparison of linear approaches to filter out environmental effects in structural health monitoring”. *Mechanical systems and signal processing*, 105:1–15 (2018).
8. Cross, E.J., Worden, K., and Chen, Q. “Cointegration: a novel approach for the removal of environmental trends in structural health monitoring data”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2133):2712–2732 (2011).
9. Tronci, E.M., Betti, R., and De Angelis, M. “Damage detection in a rc-masonry tower equipped with a non-conventional tmd using temperature-independent damage sensitive features”. *Developments in the Built Environment*, 15:100170 (2023).
10. Sclafani, L., Stagi, L., Tronci, E.M., Betti, R., and Milana, S. “Real-time unsupervised structural damage detection using cepstral features and principal component analysis”. *Structural Health Monitoring*, page 14759217251355946 (2025).
11. Bogert, B.P. “The quefrency alanysis of time series for echoes: Cepstrum, pseudoautocovariance, cross-cepstrum and saphe cracking”. In *Proc. Symposium Time Series Analysis, 1963*, pages 209–243 (1963).
12. Tronci, E.M., Beigi, H., Betti, R., and Feng, M.Q. “A damage assessment methodology for structural systems using transfer learning from the audio domain”. *Mechanical Systems and Signal Processing*, 195:110286 (2023).
13. Stagi, L., Sclafani, L., Tronci, E.M., Betti, R., Milana, S., Culla, A., Roveri, N., and Carcaterra, A. “Enhancing the damage detection and classification of unknown classes with a hybrid supervised–unsupervised approach”. *Infrastructures*, 9(3):40 (2024).
14. Stagi, L., Sclafani, L., Tronci, E.M., Betti, R., Milana, S., Culla, A., Roveri, N., and Carcaterra, A. “An unsupervised damage detection strategy for recognizing unseen health conditions in monitoring bridges”. In *International Operational Modal Analysis Conference*, pages 196–207. Springer (2024).
15. Li, L., Morgantini, M., and Betti, R. “Structural damage assessment through a new generalized autoencoder with features in the quefrency domain”. *Mechanical Systems and Signal Processing*, 184:109713 (2023).
16. Balsamo, L., Betti, R., and Beigi, H. “A structural health monitoring strategy using cepstral features”. *Journal of Sound and Vibration*, 333(19):4526–4542 (2014).
17. M.Morgantini, R.Betti, and L.Balsamo. “Structural damage assessment through features in quefrency domain”. *Mechanical Systems and Signal Processing*, 147:107017 (2021).
18. Tronci, E.M., Shid-Moosavi, S., and Speciale, C. “Pattern recognition and damage detection in wind turbine monitoring using cepstral coefficients”. In *International Conference on Experimental Vibration Analysis for Civil Engineering Structures*, pages 981–990. Springer (2025).
19. McLachlan, G.J. “Mahalanobis distance”. *Resonance*, 4(6):20–26 (1999).
20. Peeters, B. and De Roeck, G. “One-year monitoring of the z24-bridge: environmental effects versus damage events”. *Earthquake engineering & structural dynamics*, 30(2):149–171 (2001).
21. Roeck, G.D. “The state-of-the-art of damage detection by vibration monitoring: the simces experience”. *Journal of Structural Control*, 10(2):127–134 (2003).
22. Brincker, R., Andersen, P., and Cantieni, R. “Identification and level i damage detection of the z24 highway bridge”. *Experimental techniques*, 25:51–57 (2001).