13th International Conference on
Structural Health Monitoring of Intelligent Infrastructure
DOI: 10.3217/978-3-99161-057-1-092

Unsupervised Anomaly Detection for Structural Health Monitoring: A Vibration-Based Approach Using Isolation Forest

Mr Emad Soltani[1], Dr Florimond Gueniat[2], Dr Mohammad Reza Salami[3]
[1,2] Faculty of Computing, Engineering and the Built Environment , Birmingham City University, 15 Bartholomew Row, Birmingham B5 5JU , UK
[3] Department of Civil Engineering, School of Engineering, University of Birmingham, Edgbaston, Birmingham, B15 2TT,UK
email: emad.soltani@bcu.ac.uk, florimond.gueniat@bcu.ac.uk, m.r.salami@bham.ac.uk

ABSTRACT: Structural health monitoring (SHM) plays a crucial role in ensuring the safety and longevity of critical infrastructure, such as bridges. SHM refers to continuous, sensor-based, and automated monitoring that complements traditional inspection methods by providing real-time data on structural performance. This paper proposes an unsupervised machine learning approach to SHM using vibration data, aiming to address the challenges of data scarcity and the difficulty of collecting labelled damage examples. The methodology combines statistical and spectral feature extraction with an Isolation Forest anomaly detection model, trained solely on healthy data to identify potential damage. The feature extraction process includes key metrics such as root mean square, entropy, and spectral centroid, which capture both time-domain and frequency-domain characteristics of the vibration signals. The Isolation Forest model is trained on these features to distinguish between normal and anomalous patterns, making it well-suited for applications where labelled damage data is unavailable. Results from FE simulation show high accuracy (95.5%), precision (91.75%), and recall (100%), demonstrating the effectiveness of the method in distinguishing damage from healthy states. The proposed approach provides a scalable and data-efficient solution for real-time damage detection in civil infrastructure, with significant potential for deployment in large-scale monitoring systems. Future work will focus on experimental validation and improving the model's robustness in real-world conditions.

KEY WORDS: Structural Health Monitoring; Unsupervised Learning; Anomaly Detection; Vibration Data; Isolation Forest; Damage Detection; Bridges; Machine Learning; Feature Extraction; Real-Time Monitoring

1 INTRODUCTION

Structural health monitoring (SHM) is an essential aspect of maintaining the safety and integrity of civil infrastructure, particularly for critical structures like bridges. With the ageing of infrastructure, globally and increasing demands on transportation networks, the need for efficient, real-time monitoring systems has never been more pressing. SHM refers to continuous, sensor-based, automated monitoring that complements traditional inspection methods by providing real-time data on structural performance. This allows engineers to prioritise inspections and, in some cases, reduce their frequency, especially for hard-to-access structures such as long-span bridges or offshore platforms.

Machine learning-based approaches represent a more contemporary method for processing SHM data and developing models that enhance damage detection. These methods can improve both the speed and accuracy of detection and complement established techniques like system identification that are commonly used for real-time monitoring. Structural responses, such as acceleration, displacement and strain, provide quantitative measures of how a structure reacts to applied loads. Among machine learning approaches, unsupervised anomaly detection has gained interest for its ability to function without requiring labelled damage data. This is especially useful in civil infrastructure applications, where collecting labelled examples of damage is costly, time-consuming, and often impractical.

Unsupervised learning approaches are particularly advantageous in SHM applications where only healthy baseline data are available, as they do not require labelled damage examples for training. While unsupervised methods broadly include clustering and dimensionality reduction techniques, this study specifically adopts an anomaly detection approach trained solely on healthy data. This is especially relevant in large-scale infrastructure where controlled damage scenarios are infeasible [1]. Accelerometers, widely used for collecting vibration-based data in SHM, are valued for their simplicity, cost-effectiveness, and ability to capture overall structural response. Despite limitations such as temperature sensitivity and noise, they remain a preferred choice for large-scale deployment. While other sensors, such as Fibre Bragg Grating (FBG), can detect localised damage with higher precision, they require complex installation and costly equipment [2].

Soltani et al, provided a comprehensive review of machine learning techniques for SHM, highlighting the increasing use of unsupervised methods such as Principal Component Analysis (PCA), Isolation Forest, and autoencoders [3]. Their study emphasised the importance of real-time, data-driven monitoring frameworks in situations where model-based or supervised methods are limited by the lack of labelled damage data. Fernandez-Navamuel and Magalhães proposed an ensemble method that combines PCA and autoencoders for feature extraction and damage detection in bridge vibration data [4]. Their hybrid approach improved sensitivity to structural changes while maintaining robustness in noisy environments, making it suitable for long-term monitoring

DOI: 10.3217/978-3-99161-057-1-092

applications. Boccagna, Bottini, Petracca, and Amelio also supported hybrid techniques by combining convolutional autoencoders with PCA and Isolation Forest to detect anomalies in railway bridge vibration data [5]. Their results demonstrated the potential of integrating deep learning and traditional unsupervised methods to detect complex structural changes in simulated datasets. This work further reinforces the value of hybrid models, particularly where conventional techniques may struggle to capture subtle anomalies. Recent work by Toufigh and Ranjbar explored a deep autoencoder–Isolation Forest framework for detecting damage in concrete structures using ultrasonic vibration responses [6]. Their method integrated automatic feature extraction with unsupervised anomaly detection, offering an efficient and scalable solution without the need for labelled damage data. Similarly, Bayane, Leander, and Karoumi developed an unsupervised SHM pipeline using vibration-based features and anomaly detection techniques to monitor bridges [7]. Their results illustrated the practicality of using data-driven methods for real-time monitoring, especially in preventing costly maintenance or failure events.

In contrast to these studies, the model developed in this paper focuses on vibration data from bridge structures, particularly using statistical and spectral feature extraction methods (e.g., RMS, entropy), combined with the Isolation Forest algorithm trained exclusively on healthy data. The novelty of this approach lies in its use of a sliding window technique to segment the signal, enabling localised anomaly detection over time. This is especially valuable for identifying slowly progressing damage, such as cracking or fatigue, which may not be visible in global features. Additionally, the implementation of a consecutive anomaly rule ensures that damage is only flagged when anomalies persist across multiple windows, which increases robustness against transient noise and false positives. Together, these design choices make the system well-suited for large-scale, real-time SHM applications where computational efficiency and scalability are essential.

In this study, a finite element (FE) model of a simply supported beam was developed to simulate the dynamic response under both healthy and damaged conditions. A moving load was applied to reflect real-world traffic scenarios, and acceleration data were collected at mid-span. The time-series signals were segmented using a sliding window, and statistical and spectral features were extracted. These were used to train an unsupervised Isolation Forest model, aiming to detect structural damage based solely on deviations from the healthy baseline vibration signature.

2 METHODOLOGY

2.1 Simulation Setup

The finite element model represents a simply supported beam subjected to a moving load, as illustrated in Figure 1. This configuration is used to emulate bridge structures, where the pinned–roller boundary condition provides a simplified yet effective representation of real-world support systems [8]. To ensure the accuracy of the model and to validate the results, the initial simulation setup was based on the approach outlined in [8]. Specifically, the model was first validated by replicating their results, ensuring that the acceleration of the mid-span of the beam matched their findings before proceeding to select the range of velocities and forces for further simulations.

The beam is discretised into ten elements, with damage introduced in a single element for selected cases[8]. In the validation setup, the damaged element is positioned at 35% of the beam's total span from the left support.

In this study, damage was introduced by locally reducing the stiffness of a single finite element in the beam by 20 per cent. The damage location remained static across all damaged simulations and was fixed at 35 per cent of the beam span from the left support, following the setup used in Mousavi and Holloway [8]. This fixed location was chosen to allow consistent comparison of model predictions across simulations and to simplify the initial sensitivity analysis. The reduction in stiffness was chosen to represent a moderate-to-severe degradation, such as advanced cracking or corrosion. While 20 per cent is a relatively large value, it was selected to ensure the damage signal was sufficiently distinct to validate the effectiveness of the detection method. Future work will explore smaller reductions and varying damage locations to assess the model's robustness in detecting more subtle or distributed damage scenarios.

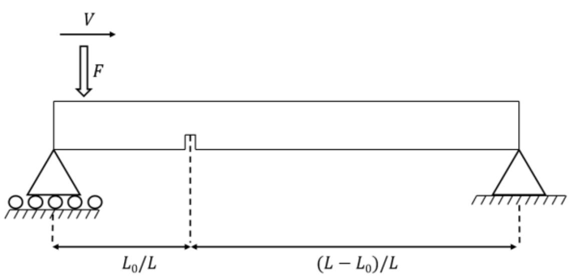

Figure 1. The schematic of the simply supported beam with a moving load

A total of 200 dynamic simulations were conducted to evaluate the beam's acceleration under varying loading conditions—100 with undamaged beams and 100 with damage applied. For each simulation, the magnitude of the vertical force and the moving velocity were randomly selected within specified ranges, using uniform random sampling. The vertical force was chosen randomly between 5000 N and 15000 N, and the velocity was selected between 13 m/s and 25 m/s. This approach introduces variability in the system, ensuring that a range of realistic vehicle loading conditions is tested.

Material properties for the beam were based on structural steel, assuming linear elastic behaviour. To simulate structural degradation, damage was introduced by locally reducing the stiffness of a single element. This modelling approach is used

DOI: 10.3217/978-3-99161-057-1-092

in the literature to represent localised damage within a linear elastic framework [9]–[11].

Huner, Irsel, Bekar, and Szala demonstrate that explicit solvers provide advantages over implicit methods in both computational efficiency and accuracy when dealing with transient impact loads [12].An explicit solver (LS-DYNA) was used to perform the simulations, selected for its effectiveness in modelling highly dynamic and nonlinear systems.

All simulations were sampled at a frequency of 1000 Hz, ensuring high-resolution capture of the system's dynamic response. The main modelling parameters are listed in Table 1.

Table 1. System Parameters for the SHM model.

Quantity	**Value**
Modulus of elasticity	200 GPa
Density	7800 kg/m³
Poisson's ratio	0.3
Beam length (L)	20 m
Cross-section width (w)	0.2 m
Cross-section depth (h)	0.2 m
Sampling frequency	1000 Hz
Load range (F)	5000 – 15000 N
Velocity range (V)	13 – 25 m/s
Number of simulations	200 (100 damaged, 100 undamaged)

2.2 *Anomaly Detection with Isolation Forest*

An unsupervised anomaly detection framework was developed to identify structural damage from acceleration time-series data obtained from simulations of both healthy and damaged beam conditions. The overall procedure is summarised in Figure 2, which presents a flowchart of the anomaly detection pipeline. Although each simulation lasts only 1.54 seconds at the lowest velocity (13 m/s), the signals were concatenated across 100 healthy and 100 damaged runs to emulate continuous traffic loading. The resulting dataset was divided into overlapping windows of two seconds, with 50% overlap. This approach enables the system to track localised signal variations while preserving frequency content and computational efficiency.

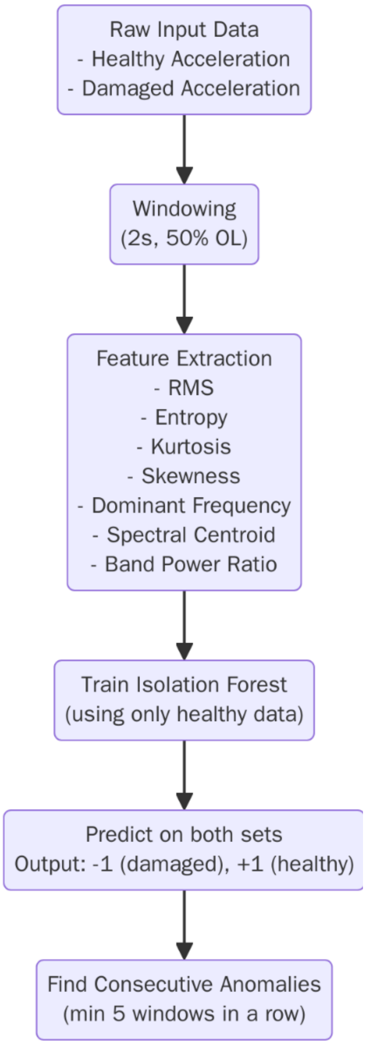

Figure 2. Flowchart of the anomaly detection model.

This configuration was selected to balance temporal and frequency resolution and improve the model's sensitivity to dynamic structural changes. While smaller window sizes were considered, they reduced the effectiveness of frequency-domain features and increased susceptibility to noise. The adopted strategy aligns with recent studies in structural health monitoring that use similar windowing to support reliable anomaly detection. An example of the sliding window approach used to segment the time-series data is illustrated in Figure 3.

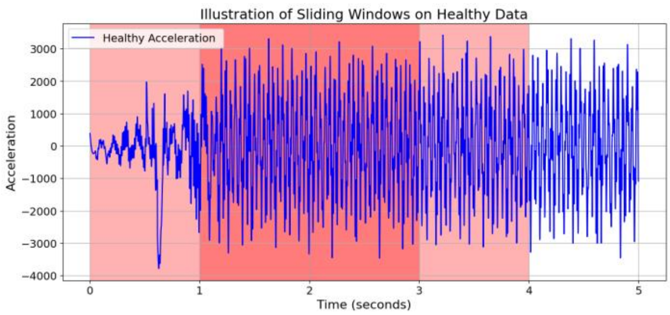

Figure 3. Sliding windows illustration

DOI: 10.3217/978-3-99161-057-1-092

For each window, a set of statistical and spectral features was extracted to characterise the dynamic response. The features used in this study include root mean square (RMS), entropy, kurtosis, skewness, dominant frequency, spectral centroid, and band power ratio. These features were selected for their proven utility in vibration-based structural health monitoring and are supported by recent benchmark studies.

An Isolation Forest model was trained exclusively on the healthy feature set to establish a baseline for normal structural behaviour. This algorithm builds an ensemble of binary trees (isolation trees), each of which recursively partitions the data space using randomly selected features and split values. The key idea is that anomalies differ sufficiently from the bulk of the data to be isolated more quickly. For any input vector x, the anomaly score is computed as:

$$s(x,n) = 2^{-\frac{E(h(x))}{c(n)}}$$

where:

- $h(x)$ is the path length, that is, the number of splits required to isolate the data point within a tree.
- $E(h(x))$ represents the average path length of x across all trees in the forest.
- $c(n)$ is the expected average path length in a binary search tree built from n samples. It is approximated by:

$$c(n) = 2H(n-1) - \frac{2(n-1)}{n}$$

Here, $H(i)$ is the ith harmonic number, approximated by ln(i) $+\gamma$, with $\gamma \approx 0.5772$ being the Euler–Mascheroni constant. Data points with short path lengths (i.e. those isolated quickly) receive anomaly scores close to 1, indicating a higher likelihood of being anomalous. Normal points, which are harder to isolate, tend to have longer path lengths and receive scores closer to 0.

In this study, the IsolationForest implementation from the scikit-learn library (v1.6) was used. The number of trees (n_estimators) was set to 100 to ensure a stable estimation of anomaly scores. The random_state was fixed at 42 to ensure reproducibility. The contamination parameter, which estimates the expected proportion of anomalies in the dataset, was tuned to 0.085 based on preliminary experiments. This value provided a good balance between capturing true positives and minimising false detections.

After training on the healthy data, the model was applied to both healthy and damaged datasets. The .predict() method classified each window as either normal (+1) or anomalous (−1), depending on whether its anomaly score exceeded the threshold determined by the contamination setting. To enhance reliability, only groups of five or more consecutive anomalous windows were treated as an indication of structural damage. This post-processing step helped reduce the risk of false positives caused by transient fluctuations or signal noise, ensuring that only sustained deviations from the healthy baseline were flagged as damage.

Once trained, the model was used to classify both healthy and damaged windows as either inliers (labelled +1) or anomalies (labelled −1). To enhance the robustness of the classification and reduce false positives, a post-processing step was applied whereby only groups of five or more consecutive anomalous windows were considered indicative of actual damage. This thresholding logic aligns with practices in unsupervised SHM where transient anomalies or noise could otherwise trigger misleading alerts [17].

This decision was motivated by the observation that isolated anomalous predictions frequently arose due to short-lived signal fluctuations or imperfect feature generalisation. By requiring a minimum streak of five consecutive anomalies, the model avoids false alarms while still being sensitive to sustained deviations caused by damage. Additionally, the simulation data were not shuffled during concatenation. Instead, the 100 healthy simulations were placed first, followed by 100 damaged ones, ensuring a continuous transition from undamaged to damaged conditions in the time series. This ordering reflects a realistic monitoring scenario in which damage develops after a prolonged healthy period and also allows for visual and algorithmic evaluation of detection accuracy at the transition point.

3 RESULTS

The system's performance was evaluated using accuracy, precision, recall and F1-score metrics, with the confusion matrix providing a summary of classification outcomes.
Results
The performance of the anomaly detection framework was evaluated using both visual comparison of signals and quantitative metrics derived from classification results.

Figure 4 shows a side-by-side comparison of raw acceleration data from the damaged and undamaged simulations. It is evident that the damaged signal (red) exhibits a higher density of peaks and more abrupt variations in magnitude than the undamaged signal (blue). This variation highlights the physical impact of stiffness reduction on dynamic response.

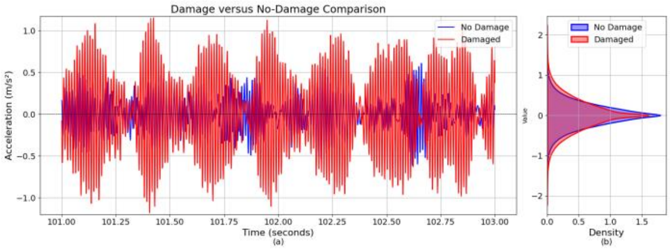

Figure 4. (a) Time-series comparison of acceleration signals between damaged and undamaged beams at mid-span from 101 to 103 seconds under a 10,000 N moving load at 15 m/s. (b) Kernel density estimate of full signal distributions, showing increased spread in the damaged case.

 DOI: 10.3217/978-3-99161-057-1-092

Figure 5 illustrates the point of transition between healthy and damaged data in the concatenated signal, where the black dashed line marks the onset of damage. This figure does not represent a change in damage location, but rather demonstrates the model's ability to correctly identify the onset of anomalous behaviour. The Isolation Forest model successfully distinguishes between the two conditions, with no false

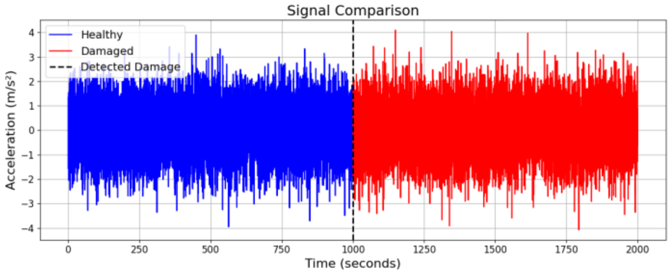

Figure 5. Mid-span acceleration response under moving load for healthy and 20% stiffness-reduced beam. Dashed line marks detected damage.

The confusion matrix in Figure 6 provides a quantitative view of the model's predictive accuracy. Out of 945 damaged windows, all were correctly identified, yielding a recall of 1.000. Meanwhile, 860 out of 945 healthy windows were correctly classified, producing a precision of 0.917. The overall accuracy of the model was 0.955, and the F1 score was 0.957, indicating a strong balance between precision and recall.

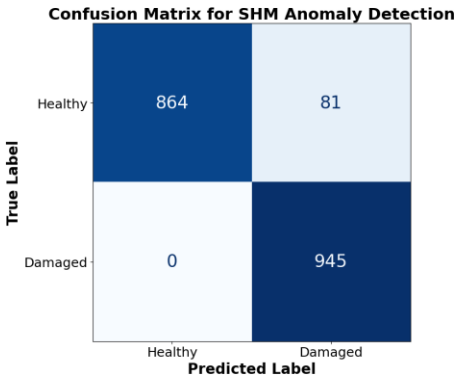

Figure 6. Confusion matrix for SHM anomaly detection.

Table 2. Anomaly detection performance metrics.

Metric	Value
Accuracy	0.955
Precision	0.917
Recall	1.000
F1 Score	0.957
True Positives (TP)	945
True Negatives (TN)	860
False Positives (FP)	85
False Negatives (FN)	0

These results, summarised in Table 2, demonstrate that the proposed combination of statistical and spectral features with an unsupervised Isolation Forest algorithm is effective in identifying structural damage using only healthy training data. The perfect recall of 1.000 confirms the model's ability to detect all instances of damage under the simulated conditions. While this result is promising, it may also reflect the relatively distinct nature of the simulated damage (a 20% stiffness reduction), which provides a clear contrast to the healthy baseline. In real-world scenarios, where damage may be more subtle or masked by noise, recall performance may vary. Nonetheless, achieving full sensitivity in this setup is an important step toward validating the model's potential for practical SHM applications.

The absence of false negatives indicates that the selected features, including root mean square, kurtosis and spectral centroid, are sensitive enough to detect changes in the mid-span acceleration response associated with damage. Permutation feature importance analysis in Figure 7 supports this, showing that RMS had the greatest influence on the model's predictions, while entropy, kurtosis and skewness contributed less. This highlights the importance of signal energy in distinguishing damaged from undamaged states and suggests that the model relied primarily on RMS to detect anomalies.

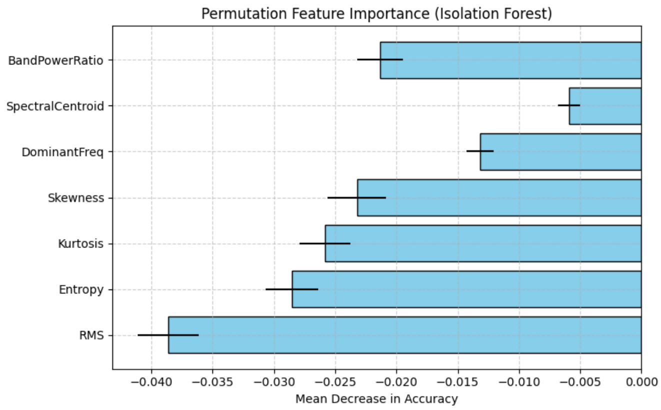

Figure 7. Permutation feature importance showing the relative contribution of all features to the Isolation Forest model's anomaly detection performance.

However, the presence of 85 false positives, reflected in a precision of 0.917, suggests that some healthy windows exhibited irregular but non-damaging patterns. These anomalies may be due to transient structural behaviours, elevated loads, or complex vibration modes that were not well represented in the training data.

This result reflects a trade-off between sensitivity and specificity. In structural monitoring applications, prioritising recall is often preferred, as missing damage poses a greater risk than raising a false alarm. The high F1 score of 0.957 supports the strength of this balance. Nevertheless, repeated false positives may reduce user trust in the system and increase inspection costs. To address this, further work is needed to

DOI: 10.3217/978-3-99161-057-1-092

investigate the nature of these misclassifications and explore feature refinement or adaptive thresholding strategies that can improve specificity without compromising detection accuracy.

The simulations used in this study were designed to represent realistic bridge conditions, based on established finite element modelling techniques reported in structural health monitoring literature. A simply supported beam was subjected to a moving load, with force and velocity ranges chosen to reflect typical vehicle traffic scenarios. By randomly sampling these parameters across simulations, the model captures a level of variability comparable to that seen in practice. However, while the simulations reproduce essential structural behaviours, certain complexities that exist in real-world monitoring, such as sensor noise, thermal drift, and environmental changes, were not included. Although the method performs well in this controlled environment, additional validation using experimental or field data is needed to confirm its applicability in real conditions.

In this study, damage was introduced by reducing the stiffness of a single element in the finite element model by 20%. This level of degradation represents a significant structural change, such as might result from cracking or corrosion. The resulting increase in vibration energy was clearly detectable in the acceleration signal, particularly at the beam's mid-span. While the model achieved full sensitivity to this level of damage, its performance with smaller changes remains to be assessed. Reductions in stiffness of 5%-10%, for example, may produce more subtle variations in the signal, making them harder to distinguish from normal fluctuations. Future work will investigate how sensitive the model is to such smaller degradations and how well it performs under more realistic conditions that include noise and environmental variation.

All results presented in this study are based solely on finite element simulations. While these simulations provide a controlled and repeatable environment for evaluating the proposed method, they do not capture the full complexity of real-world monitoring scenarios. It is well known that vibration-based damage detection techniques often face challenges in practical applications due to noise, environmental variability, temperature effects, and operational conditions. These factors can introduce variability that may obscure subtle signs of damage or increase the rate of false positives.

We acknowledge this as a limitation of the current study and plan to address it in future work. Specifically, we aim to validate the approach using real-world sensor data collected from instrumented laboratory-scale bridge models or in situ field deployments. This would involve applying accelerometers to physical structures subjected to controlled damage and comparing the model's predictions against ground truth. Such validation is essential for assessing the robustness and transferability of the proposed method and would provide critical insights into how it performs under realistic operating conditions.

4 CONCLUSION

This study set out to develop and evaluate an unsupervised machine learning framework for structural health monitoring (SHM), with the specific aim of detecting damage using only vibration data from healthy structural conditions. The motivation was to address the limitations of supervised approaches that rely on labelled damage data, which is often unavailable or impractical to collect in real-world scenarios.

The proposed method combined lightweight time-domain and frequency-domain features with an Isolation Forest anomaly detection model, further enhanced by a consecutive anomaly rule to reduce false positives. The approach was tested on simulated acceleration data from both healthy and damaged beam configurations. Results showed that the model achieved high detection accuracy, with perfect recall and no false negatives, indicating strong sensitivity to damage. Precision and F1 scores also demonstrated the model’s ability to reliably distinguish between normal and abnormal structural behaviour, despite being trained exclusively on healthy data.

These findings confirm the method’s suitability for low-cost, real-time deployment on civil infrastructure, particularly in settings where computational resources and data availability are constrained. Although this study relied on simulated data, it provides a solid foundation for future experimental validation using real-world sensor inputs. Further refinement may focus on improving robustness to environmental variability and exploring hybrid models to enhance detection performance.

REFERENCES

[1] A. Calderon Hurtado, K. Kaur, M. Makki Alamdari, E. Atroshchenko, K.C. Chang, and C.W. Kim, Unsupervised learning-based framework for indirect structural health monitoring using adversarial autoencoder, Journal of Sound and Vibration, vol. 550, p. 117598, 2023.
[2] D. Heider, Fibre Bragg grating sensors for aircraft structural health monitoring, in Proceedings of SPIE 7292: Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems, paper 72920H, 2009.
[3] E. Soltani, E. Ahmadi, F. Gueniat, and M.R. Salami, A review of bridge health monitoring based on machine learning, Proceedings of the Institution of Civil Engineers - Bridge Engineering, 2022.
[4] A. Fernandez-Navamuel and F. Magalhães, Deep learning enhanced principal component analysis for structural health monitoring, Journal of Intelligent Material Systems and Structures, 2022.
[5] R. Boccagna, M. Bottini, M. Petracca, and A. Amelio, Unsupervised deep learning for structural health monitoring, Big Data and Cognitive Computing, 2023.
[6] V. Toufigh and I. Ranjbar, A deep learning framework for ultrasonic-based distributed damage detection in concrete: integration of a deep auto-encoder and Isolation Forest for anomaly detection, Journal of Intelligent Material Systems and Structures, 2024.
[7] I. Bayane, J. Leander, and R. Karoumi, An unsupervised machine learning approach for real-time damage detection in bridges, Engineering Structures, 2024.
[8] M. Mousavi and D. Holloway, Using a moving load to simultaneously detect location and severity of damage in a simply supported beam,

DOI: 10.3217/978-3-99161-057-1-092

Journal of Vibration and Control, vol. 25, no. 19–20, pp. 2737–2752, 2019.

[9] M. Najafgholipour and S. Dehghan, Finite element analysis of reinforced concrete beam with damage simulation, International Journal of Civil Engineering, vol. 18, no. 5, pp. 567–579, 2020.

[10] M.A.S. Bahri, M.M. Ratnam, and H.P.S. Abdul Khalil, Functionally graded wood filler–recycled polypropylene composite: Effect of mechanical loading on deflection of cantilever beam, Advanced Composites Letters, vol. 29, pp. 1–9, 2020.

[11] H. Luo, J. Fu, P. Wang, and W. Zhou, Design optimization of the ram structure of friction stir welding robot, Journal of Mechanical Science and Technology, vol. 33, no. 1, pp. 123–130, 2019.

[12] U. Huner, G. Irsel, U. Bekar, and M. Szala, A comprehensive case study to the implicit and explicit approach in finite element analysis, Cogent Engineering, DOI: 10.1080/23311916.2024.2399737, 2025.

[13] K.P. Anastasiadis, C.S. Sakaris, R. Schlanbusch, and J.S. Sakellariou, Vibration-based SHM in the synthetic mooring lines of the semisubmersible OO-Star Wind Floater under varying environmental and operational conditions, Sensors, vol. 24, no. 2, p. 543, 2024.

[14] K. Eltouny, M. Gomaa, and X. Liang, Unsupervised learning methods for data-driven vibration-based structural health monitoring: a review, Sensors, vol. 23, no. 6, p. 3290, 2023.

[15] K. Kritikakos, X. Konstantinou, C.F. Lee, C.S. Sakaris, J.S. Sakellariou, and R. Schlanbusch, Vibration-based structural health monitoring of the mooring lines in a floating offshore wind turbine under varying environmental conditions: NN vs STS based methods, ResearchGate, 2025.

[16] T. Buckley, B. Ghosh, and V. Pakrashi, A feature extraction & selection benchmark for structural health monitoring, Journal of Intelligent Material Systems and Structures, 2023.

[17] K. Yang, T. Liu, Z. Yang, Y. Zhou, and Z. Tian, Baseline optimized autoencoder-based unsupervised anomaly detection in uncontrolled dynamic structural health monitoring, Journal of Intelligent Material Systems and Structures, 2025.