Chapter 4
Dynamic Recognition of Damage States in Wind Turbines Using Cepstral Coefficients and ANOVA Analysis

Costanza Speciale, Sina Shid-Moosavi, Eleonora Maria Tronci, and Silvia Milana

Abstract Traditional damage detection methods for dynamic systems often use regression or classification algorithms to distinguish between damaged and intact states, but they typically require case-specific parameter tuning. This need for manual adjustment limits their generalizability across different structures and varying operational conditions, reducing diagnostic accuracy. For complex systems like wind turbines, varying conditions, such as wind speed, temperature, and load demands, can complicate the choice of optimal parameters for reliable damage-sensitive features, leading to false positives and inconsistent diagnostics. This work leverages a semi-supervised damage detection algorithm based on a strategy combining Linear Discriminant Analysis (LDA) and Probabilistic Linear Discriminant Analysis (PLDA) with cepstral coefficients for damage identification and classification in wind turbines. While LDA and PLDA are traditionally supervised, this approach integrates them into an unsupervised framework that dynamically recognizes and categorizes emerging damage states as they appear, rather than relying on predefined training phases. ANOVA-based sensitivity analysis is incorporated to evaluate how user-defined parameters impact the performance of the classification outcomes, guiding an adaptive tuning strategy that ensures consistent detection across varied operational conditions. The proposed damage detection algorithm is validated using experimental data from the Aventa AV-7 wind turbine, assessing its ability to recognize and classify multiple damage states under varied operational conditions. ANOVA-based sensitivity analysis identifies key parameters for optimal performance, demonstrating the algorithm's adaptability and robustness. This validation confirms the framework's effectiveness in achieving consistent and reliable diagnostics across real-world environmental variability.

Keywords Semi-supervised damage detection · Cepstral coefficients · ANOVA · Sensitivity analysis · Wind turbine diagnostics

Introduction

Conventional approaches to damage detection in dynamic systems often employ regression or classification techniques to distinguish between healthy and damaged states. While effective in controlled cases, these methods usually rely on manually tuned parameters optimized for specific scenarios, which limits their ability to generalize across different structures and varying operational conditions. In the context of wind turbines, where external factors such as wind speed, temperature, and load demand fluctuate continuously, parameter sensitivity poses a significant challenge. In practice, this dependence can lead to unreliable extraction of damage-sensitive features, resulting in false alarms and inconsistent diagnostic performance.

A key challenge in Structural Health Monitoring (SHM) for wind turbines is the scarcity of labeled data representing damaged states. Realistic failure cases are rare, difficult to reproduce, and often poorly documented, making supervised

Costanza Speciale · Silvia Milana
Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Rome, Italy
e-mail: costanza.speciale@uniroma1.it; silvia.milana@uniroma1.it

Sina Shid-Moosavi
Department of Civil and Environmental Engineering, Northeastern University, Boston, MA, USA
e-mail: shidmoosavi.s@northeastern.edu

Eleonora Maria Tronci
Department of Civil and Urban Engineering, Center for Urban Science and Progress, New York University, New York, NY, USA
e-mail: emt377@nyu.edu

Kristin B. Zimmerman, PhD (eds.), *Special Topics & Experimental Techniques in Structural Dynamics, 2026, Vol. 5*, Conference Proceedings of the Society for Experimental Mechanics Series, https://doi.org/10.1007/978-8-743-81444-3_4

learning approaches impractical in many scenarios [1, 2]. As a result, unsupervised methods such as novelty detection have been widely applied, where a baseline model of the healthy state is used to flag significant deviations as potential indicators of damage [3, 4, 5, 6]. While effective for distinguishing healthy from damaged conditions, these approaches are generally limited to binary classification and often lack mechanisms to adapt when new data becomes available. For dynamic systems such as wind turbines, this creates a critical need for more flexible methods capable of evolving with changing operational and environmental conditions.

As described in the previous chapter, features such as cepstral coefficients can be extracted to characterize the dynamic response of wind turbines. Building on this, the present chapter focuses on a semi-supervised damage detection framework that integrates cepstral features with a hybrid classification strategy combining Linear Discriminant Analysis (LDA) for feature compression and Probabilistic Linear Discriminant Analysis (PLDA) for adaptive classification, enabling the system to learn from healthy operational data and to detect emerging damage scenarios without requiring labeled fault data. Although LDA and PLDA are traditionally applied in supervised learning contexts, they are embedded here within an unsupervised structure that enables adaptive identification and categorization of emerging damage states without reliance on fixed training datasets.

One of the main limitations of existing methods is the static nature of parameter selection, which typically requires offline sensitivity analysis. To address this issue, an exploratory Reinforcement Learning (RL)-based strategy is proposed to automate and optimize the tuning of parameters in cepstral-based diagnostic pipelines. In this proposed framework, an RL agent would interact with the classification system to adjust parameter settings dynamically, guided by a reward function designed to balance detection accuracy, false alarm rates, and robustness under variable operating conditions. Although not yet fully developed, this concept highlights the potential of dynamic optimization to enable continuous learning from new data in the field, thereby reducing reliance on predefined thresholds and improving adaptability to real-world variability.

The effectiveness of the proposed framework is evaluated using experimental data from the Aventa AV-7 wind turbine. Comparisons will be made between static parameter configurations and those adaptively tuned by the RL agent under diverse operational and damage scenarios. This chapter thereby addresses the limitations identified in conventional methods, advancing toward a more reliable and adaptive solution for structural health monitoring in wind farm applications.

Dataset Description

This research utilizes data from the Aventa AV-7 [7], a 7kW wind turbine operated by ETH Zurich (Figure 1), which has been extensively monitored to evaluate SHM methods as the turbine approaches the end of its design life. Since 2020, sensors have been deployed throughout the turbine to facilitate system identification, modal analysis, and failure detection studies, with data collected under four specific operational scenarios: normal operation, pitch system failure, aerodynamic imbalance, and rotor icing. For each condition, datasets include time series measurements recorded over specified dates between 2022 and 2023.

Fig. 1 Aventa AV-7 [7]

The Aventa AV-7 turbine operates with a variable-speed, variable-pitch control, beginning energy generation at a 2 m/s wind speed and ceasing at 14 m/s for a rated power of 7kW. It has a three-blade rotor with a 12.9 m diameter, positioned 18 m above ground on a concrete and steel tower. A belt-driven generator and a frequency converter enable variable-speed operation.

The turbine is outfitted with eleven accelerometers distributed along the tower, nacelle, main bearing, and generator, as well as two full-bridge strain gauges at the tower base for measuring axial and bending strains. All data is recorded at 200Hz, while SCADA (Supervisory Control and Data Acquisition) data, such as wind speed, rotor speed, and power output, is recorded at 7Hz, with temperature and humidity recorded at 1Hz. The accelerometers are paired and located at specific tower sections to capture the x- and y-axis accelerations aligned with wind direction, allowing comprehensive monitoring of the turbine's structural response.

The SCADA system provides insights into power, wind speed, rotor speed, and yaw angle, which are critical for analyzing turbine performance, especially under "normal operation" conditions. The power curve, derived from these SCADA readings, illustrates the relationship between wind speed and power output, with distinct operational regions (cut-in, rated, and cut-off speeds) [8].

The dataset further captures three specific fault conditions, providing valuable data for condition monitoring methods.

In February 2022, a failure occurred in the flexible coupling of the pitch drive system, which is critical for adjusting blade angles to optimize performance under changing wind conditions. SCADA data captured deviations in rotor speed, power output, and yaw position compared to normal operation. Before the complete failure, the pitch system became stuck, reducing the turbine's output to 1-2 kW while the rotor speed dropped to around 40 rpm. An emergency stop was triggered to prevent further damage, as indicated by sharp changes in SCADA measurements at approximately 270 seconds. Acceleration data from sensors across the tower and nacelle show a significant structural response to the failure, highlighting its impact on the turbine's dynamics.

In December 2022, rotor icing caused a near-complete halt in turbine operations, reducing rotor speed to almost zero and stopping power generation. This highlights the importance of early detection in cold climates where ice accumulation can significantly reduce aerodynamic efficiency and increase vibration.

Then, data collected after roughness tape was applied to one blade in December 2022 revealed imbalances that affected rotor speed and induced mechanical stress on key turbine components, underscoring the need to identify imbalances early to prevent failures.

This research focuses on two of these damage conditions: pitch drive failure and aerodynamic imbalance. The first corresponds to a sudden and critical rupture in the turbine's system, and is particularly challenging for SHM due to the need for immediate detection and intervention to prevent further damage. By concentrating on this sudden event, the study aims to improve SHM techniques' ability to quickly identify and respond to rapid, unforeseen failures, which are particularly difficult to manage.

The second, in contrast, typically develops gradually. It can be induced, for instance, by surface degradation or roughness accumulation on the blades. Although less abrupt than a pitch drive failure, it is still important to detect it promptly, as it increases loads and vibrations throughout the drivetrain and tower, potentially leading to more severe failures over time. Moreover, it also negatively impacts power production.

The dataset therefore provides a key resource for developing more effective methods to address these two conditions, selected as representative examples of the broader categories of sudden and gradual failures that typically characterize the operational life of wind turbines.

Semi-Supervised Damage Detection Framework

This section provides an overview of the framework employed for damage detection. After cepstral coefficients are extracted, they are used as damage-sensitive features within a statistical pattern recognition framework that integrates LDA and PLDA. The goal is to reduce the dependence on manually tuned parameters and enable adaptive identification of damage states under varying operational conditions.

LDA [9] is a supervised machine learning method commonly applied for dimensionality reduction and feature extraction. Its aim is to project high-dimensional data into a lower-dimensional latent space that maximizes class separability. This is achieved by identifying projection directions that maximize the distance between class means (between-class scatter) while minimizing variance within each class (within-class scatter) [10].

Let the training dataset consist of N samples $\{\mathbf{x}_n\}_{n=1}^{N}$, where each observation $\mathbf{x}_n \in \mathbb{R}^{D\times 1}$ belongs to one of K known classes. The algorithm begins by computing the class-wise means $\mathbf{m}_k$, the global mean $\mathbf{m}$, the within-class scatter matrix $\mathbf{S}_W$, and the between-class scatter matrix $\mathbf{S}_B$. The optimal projection matrix is obtained by solving the generalized eigenvalue problem:

$$\mathbf{S}_B\mathbf{w} = \lambda\,\mathbf{S}_W\mathbf{w} \tag{1}$$

The eigenvectors $\{\mathbf{w}_i\}$ define the directions of maximum separability, and their eigenvalues $\{\lambda_i\}$ quantify the discriminative strength of each direction. The g leading eigenvectors are selected to construct the LDA projection matrix, which maps the original features into a reduced space optimized for classification.

A summary of the full LDA procedure is provided in Algorithm 1.

Algorithm 1 Linear Discriminant Analysis

Require: Training data $\{\mathbf{x}_i, y_i\}_{i=1}^N$ with K classes
1: Compute class means: $\mathbf{m}_k = \frac{1}{n_k}\sum_{\mathbf{x}_i \in C_k} \mathbf{x}_i$
2: Compute global mean: $\mathbf{m} = \frac{1}{N}\sum_{i=1}^N \mathbf{x}_i$
3: Compute within-class scatter: $\mathbf{S}_W = \sum_{k=1}^K \sum_{\mathbf{x}_i \in C_k} (\mathbf{x}_i - \mathbf{m}_k)(\mathbf{x}_i - \mathbf{m}_k)^T$
4: Compute between-class scatter: $\mathbf{S}_B = \sum_{k=1}^K n_k(\mathbf{m}_k - \mathbf{m})(\mathbf{m}_k - \mathbf{m})^T$
5: Solve generalized eigenvalue problem: $\mathbf{S}_B\mathbf{w} = \lambda\mathbf{S}_W\mathbf{w}$
6: Select top g eigenvectors: $\mathbf{W} = [\mathbf{w}_1, ..., \mathbf{w}_g]$
7: **return** LDA projection matrix $\mathbf{W}$

The first step in the framework is training the classification model using data representing the various operational classes of the undamaged state. In this study, 70% of the undamaged data were randomly selected for training, while the remaining 30% were reserved for testing. Figure 2 illustrates the overall process.

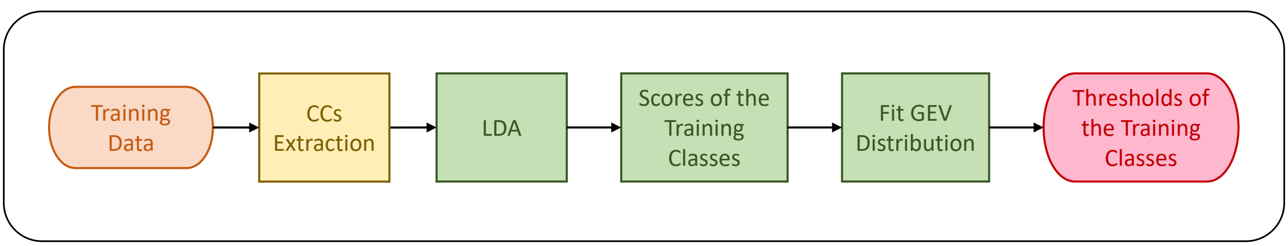

Fig. 2 Training phase flowchart

Classification is performed using a PLDA model [11], a probabilistic extension of LDA that provides robustness in complex classification scenarios, including the recognition of previously unseen classes. Unlike conventional LDA, which classifies based on distance to centroids in the latent space, PLDA models both samples and class centers as Gaussian realizations. This allows likelihood-based evaluation of class membership, enabling the model to identify new patterns beyond the training set. Such flexibility is especially important for SHM, where operational variability and emerging damage states are expected.

The PLDA framework assumes that each observation $\mathbf{x}$ and its class center $\mathbf{z}$ are generated from latent variables $\mathbf{u}$ and $\mathbf{v}$ through linear mappings $\mathbf{x} = \mathbf{m} + \mathbf{A}\mathbf{u}$ and $\mathbf{z} = \mathbf{m} + \mathbf{A}\mathbf{v}$, with $\mathbf{u} \sim \mathcal{N}(\mathbf{v}, \mathbf{I})$ and $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Psi})$. Model parameters $\{\mathbf{m}, \mathbf{A}, \mathbf{\Psi}\}$ are estimated from the latent representations derived via LDA, as summarized in Algorithm 2.

Algorithm 2 Probabilistic Linear Discriminant Analysis Training

Require: Projected training data $\{\mathbf{u}_i\}$ from LDA
1: Estimate global mean $\mathbf{m}$ and projection matrix $\mathbf{A}$
2: Assume latent class means $\mathbf{v} \sim \mathcal{N}(0, \mathbf{\Psi})$
3: **for** each training class k **do**
4: Compute PLDA scores: $\log P(\mathbf{u}_p|\mathbf{u}_1^k, ..., \mathbf{u}_{n_k}^k)$
5: Fit a GEV distribution to scores
6: Set threshold at 95th percentile
7: **end for**
8: **return** Parameters $\{\mathbf{m}, \mathbf{A}, \mathbf{\Psi}\}$ and score thresholds

Once trained, the model computes a PLDA score for each new observation, corresponding to the log-likelihood of class membership. Thresholds are established by fitting score distributions with a Generalized Extreme Value (GEV) model [12] and setting cutoffs at GEV_{Th}, ensuring robustness against variability and facilitating detection of previously unseen conditions.

The testing phase evaluates both undamaged (remaining 30%) and damaged data. Each data point, corresponding to a set of cepstral coefficients from one signal frame, is processed sequentially, making the method suitable for real-time SHM.

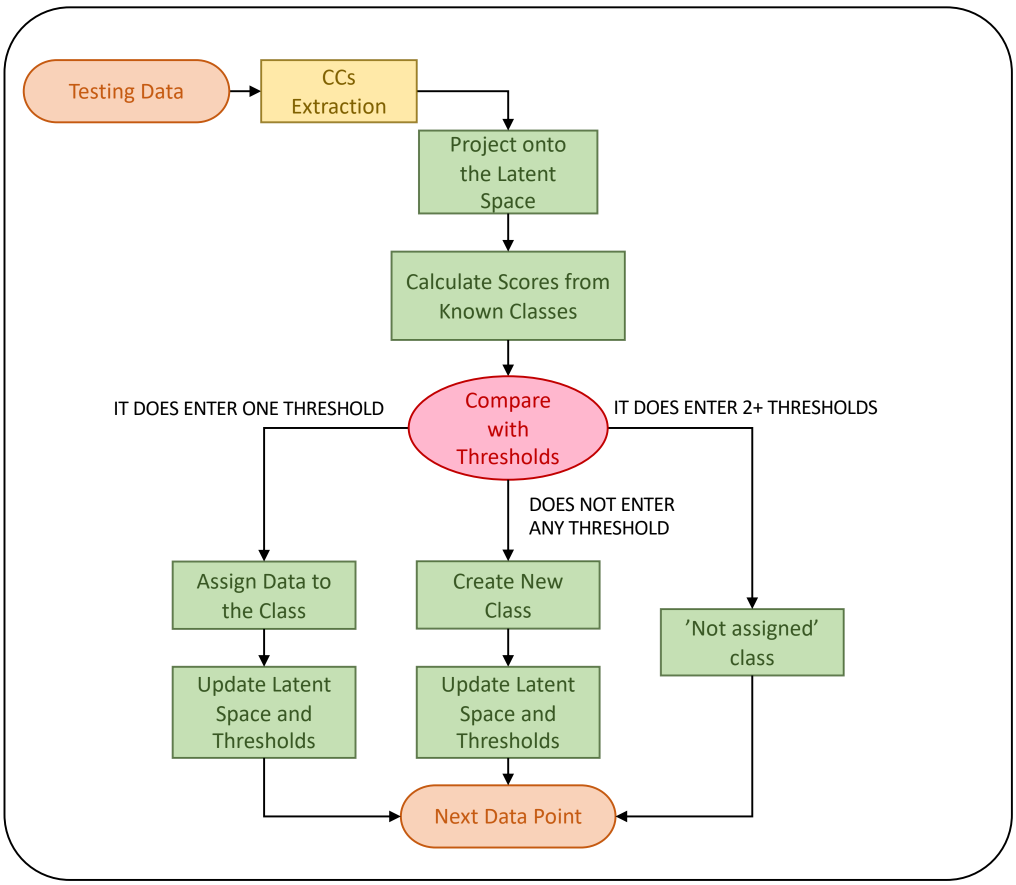

Fig. 3 Testing phase flowchart

For each frame, cepstral coefficients are projected into the latent space defined by LDA, ensuring dimensionality reduction while preserving discriminative features. The projected vector is evaluated using the trained PLDA model, which computes likelihood scores. Classification proceeds as follows: (i) if a point matches one class above its threshold, it is assigned to that class and thresholds are updated once a counter N_{UP} is reached; (ii) if a point fails to meet any threshold, it is used to define a new class once counter N_{NC} is reached; (iii) if a point exceeds thresholds for multiple classes, it is left unassigned.

The choice of N_{UP} and N_{NC} is critical, as they govern the adaptability of the framework. To guide parameter selection, two quality indices are defined. These indices provide a quantitative measure of the framework's performance:

$$I_{BIN} = \frac{\text{number of correctly classified data}}{\text{total predictions}} \tag{2}$$

$$I_{MULT} = \frac{\text{number of correctly classified data}}{\text{total predictions}} \cdot \frac{1}{1 + \text{number of missed or extra classes}} \tag{3}$$

The binary index I_{BIN} evaluates the accuracy of distinguishing between damaged and undamaged states, while the multi-class index I_{MULT} extends this evaluation to multiple damage scenarios. The latter penalizes cases in which additional spurious classes are created or true classes are not identified, thereby encouraging both accuracy and stability in the classification results.

Sensitivity Analysis

After presenting the framework, it was applied to the experimental data from the Aventa AV-7 turbine to evaluate wind turbine damage detection. As anticipated, for the training phase only undamaged data from normal turbine operation were used. These data were partitioned into four classes according to wind speed and rotor speed values, each representing a distinct operating region. For the testing phase, instead, damaged data were also included. Specifically, four additional classes were defined: three derived from measurements collected on the day of the pitch failure, corresponding to different damaged states (incipient failure, stuck pitch drive, and complete failure), and one representing the aerodynamic imbalance condition. Within this framework, several parameters were identified as influential: N_{CC}, the number of cepstral coefficients used as input features; N_{LDA}, the number of LDA components retained; GEV_{Th}, the percentile threshold derived from the GEV distribution; and the counters N_{UP} and N_{NC}, which regulate class updating and new-class creation, respectively, and can be further tuned during RL optimization.

In previous work [13], it was observed that the two input parameters N_{NC} and N_{UP} significantly affect the algorithm's performance. An analysis was carried out by performing multiple simulations with different combinations of these parameters. As shown in Figure 4, the mean and standard deviation of the two performance indices, I_{BIN} and I_{MULT}, vary substantially across the different cases.

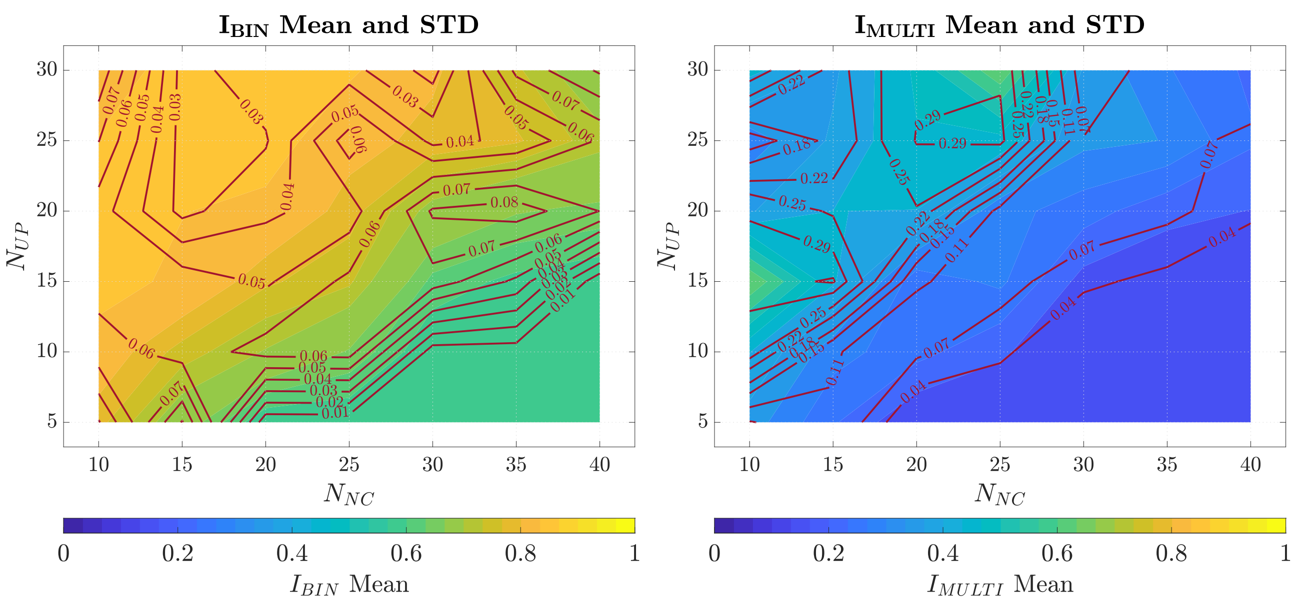

Fig. 4 Mean and standard deviation (red lines) of the binary and multi-class performance indices [13]

To determine which parameters are most critical and to investigate their potential interactions, a factorial Analysis of Variance (ANOVA) was performed. The analysis was conducted using both the binary index I_{BIN} and the multi-class index I_{MULT} as performance metrics. A full factorial design was employed to span the parameter space, testing all possible combinations of factor levels. The factor levels are summarized in Table 1. The total number of experimental runs is equal to the product of the level counts, yielding 5,040 unique parameter combinations.

Table 1 Factor levels and number of levels in the full factorial design

Factor (Parameter)	Levels	No. of Levels
N_{CC} (Number of cepstral coefficients)	6, 10, 15, 20, 30, 50, 75, 100	8
N_{LDA} (Number of LDA components)	1, 2, 3, 4, 5	5
GEV_{Th} (Percentile threshold)	0.85, 0.90, 0.95	3
N_{UP} (Update counter)	5, 10, 15, 20, 25, 30	6
N_{NC} (New-class counter)	10, 15, 20, 25, 30, 35, 40	7

The ANOVA was conducted using the `pg.anova` function from the `Pingouin` Python package [14, 15]. Effect sizes were quantified as partial eta-squared (η_p^2), which represents the proportion of variance explained by a factor after accounting for other sources of variation. For a one-way ANOVA, the decomposition of sums of squares is defined as:

$$\begin{gathered} SS_{\text{total}} = SS_{\text{effect}} + SS_{\text{error}}, \\ SS_{\text{total}} = \sum_i \sum_j (Y_{ij} - \overline{Y})^2, \quad SS_{\text{effect}} = \sum_i n_i(\overline{Y_i} - \overline{Y})^2, \\ SS_{\text{error}} = \sum_i \sum_j (Y_{ij} - \overline{Y}_i)^2 \end{gathered} \tag{4}$$

with the F-statistic computed as:

$$F^* = \frac{MS_{\text{effect}}}{MS_{\text{error}}} = \frac{SS_{\text{effect}}/(r-1)}{SS_{\text{error}}/(n_t - r)}, \tag{5}$$

where r is the number of groups and n_t is the total sample size. Partial eta-squared is then given by:

$$\eta_p^2 = \frac{SS_{\text{effect}}}{SS_{\text{effect}} + SS_{\text{error}}}. \tag{6}$$

The η_p^2 values serve as sensitivity indices, providing a quantitative ranking of parameter contributions to variability in the output. This methodology follows best practices for reporting effect sizes in experimental design [16, 17]. Results for both I_{BIN} and I_{MULT} are presented in Figure 5. The analysis shows that N_{LDA} is the most influential parameter, followed by GEV_{Th} and N_{CC}. Importantly, part of the sensitivity of all parameters arises from interaction effects with other parameters, underscoring the need to account for parameter correlation rather than evaluating each factor in isolation. For instance, the effect of N_{CC} is strongly tied to the number of LDA components retained, as the discriminatory power of additional cepstral coefficients can only be realized if sufficient LDA dimensions are available. Similarly, the influence of GEV_{Th} depends on the interplay with both N_{UP} and N_{NC}, since the threshold for defining new classes and updating existing ones directly shapes how probabilistic scores are interpreted in practice.

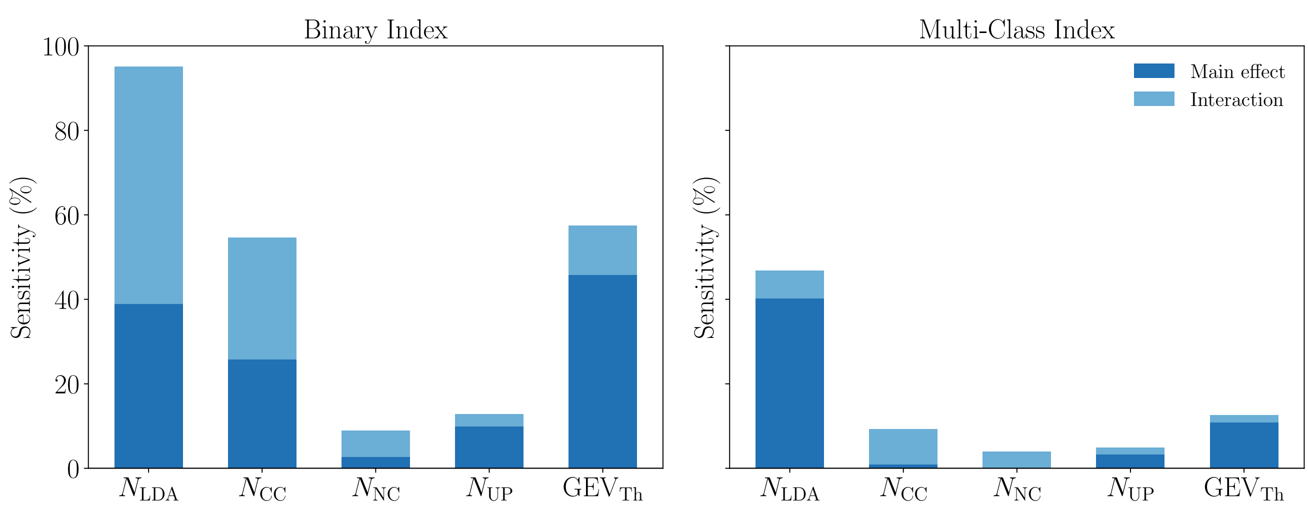

Fig. 5 Sensitivity analysis results for framework parameters based on binary and multi-class performance indices

To further explore these relationships, pairwise interaction effects were visualized using sensitivity matrices, shown in Figure 6. These matrices display η_p^2 values for both main effects and two-way interactions across all parameter pairs. The binary index highlights strong interactions between N_{LDA} and both N_{CC} and GEV_{Th}, indicating that the influence of cepstral resolution and threshold settings depends heavily on how many LDA dimensions are retained. In the multi-class case, the same pairs also show moderate interaction effects, though GEV_{Th} emerges as more influential overall than N_{CC}.

These findings reinforce that parameter interactions, not just main effects, play a decisive role in shaping performance. Therefore parameter optimization for damage detection cannot be approached as a purely one-dimensional tuning exercise. Instead, parameter sets must be tuned collectively, as the performance of the framework emerges from the balance and

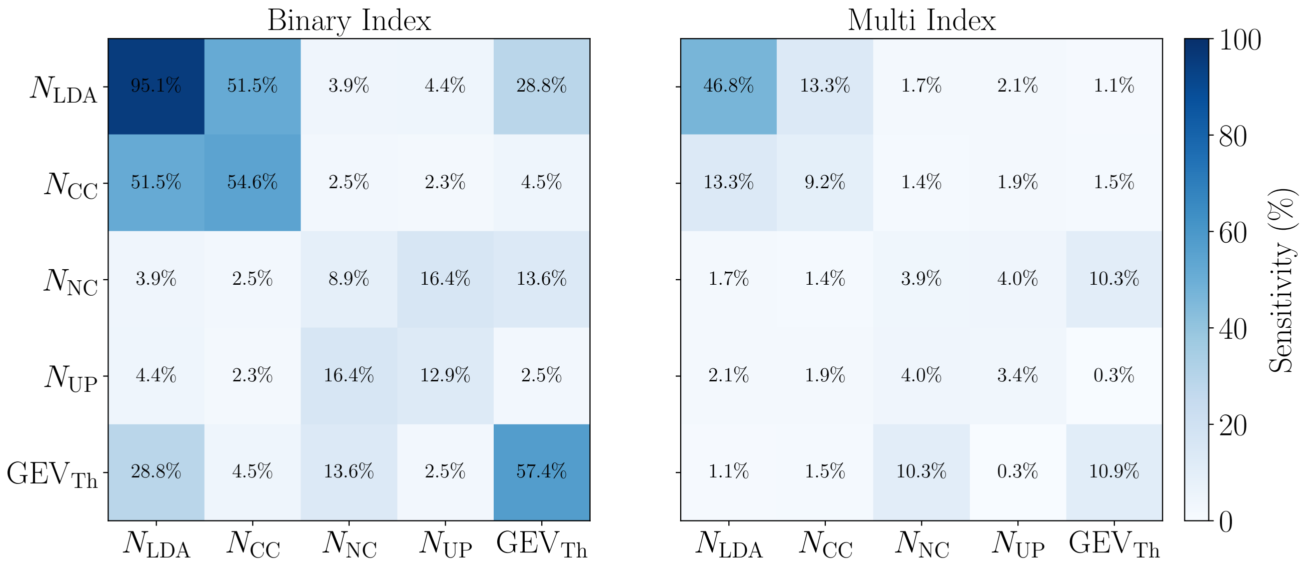

Fig. 6 Sensitivity matrices showing partial η_p^2 for main and pairwise interaction effects across all parameters on binary and multi-class performance indices

interaction among them. This interdependence further motivates the adoption of adaptive optimization strategies, such as reinforcement learning, which can dynamically adjust multiple parameters simultaneously in response to evolving operating conditions and data characteristics. Such an approach offers the potential to avoid suboptimal configurations that might result from static or independent parameter tuning, thereby enhancing the robustness and reliability of the detection system.

Conclusions

This work presented a semi-supervised framework for damage detection in wind turbines that integrates cepstral coefficients with LDA and PLDA. The proposed approach addresses key limitations of conventional supervised and unsupervised strategies by enabling dynamic recognition of damage states without relying on large labeled datasets or static parameter settings. Through application to experimental data from the Aventa AV-7 wind turbine, the framework demonstrated its ability to classify different failure modes, capturing the variability of real-world operating conditions.

A sensitivity analysis based on ANOVA quantified both single–parameter and interaction effects. In isolation, performance is dominated by the number of LDA components (N_{LDA}), followed by the GEV percentile threshold (GEV_{Th}) and the number of cepstral coefficients (N_{CC}), with the update and new–class counters (N_{UP}, N_{NC}) exerting smaller but non-negligible influence. At the same time, several interaction terms are significant—most notably $N_{LDA} \times N_{CC}$ and $N_{LDA} \times GEV_{Th}$ (and, to a lesser extent, $GEV_{Th} \times N_{UP}/N_{NC}$)—highlighting that multi-parameter tuning is necessary and providing clear directions for adaptive optimization.

Overall, the presented methodology advances structural health monitoring practices by combining interpretable feature extraction with adaptive classification and systematic parameter evaluation. Future work will focus on integrating an automated tuning strategy for online parameter adjustment, extending validation to larger datasets, and additional failure types.

Acknowledgments The authors would like to sincerely thank ETH Zurich and its collaborators for providing the "Aventa AV-7" dataset, which was invaluable for our study.

References

1. Farrar, C.R. and Worden, K. *Structural health monitoring: a machine learning perspective*. John Wiley & Sons (2012)
2. Regan, T., Beale, C., and Inalpolat, M. "Wind turbine blade damage detection using supervised machine learning algorithms". *Journal of Vibration and Acoustics*, 139(6):061010 (2017)

3. Eltouny, K., Gomaa, M., and Liang, X. "Unsupervised learning methods for data-driven vibration-based structural health monitoring: A review". *Sensors*, 23(6) (2023)
4. Castellani, F., Garibaldi, L., Daga, A.P., Astolfi, D., and Natili, F. "Diagnosis of faulty wind turbine bearings using tower vibration measurements". *Energies*, 13(6) (2020)
5. Movsessian, A., García Cava, D., and Tcherniak, D. "An artificial neural network methodology for damage detection: Demonstration on an operating wind turbine blade". *Mechanical Systems and Signal Processing*, 159:107766 (2021)
6. Feijóo, M.d.C., Zambrano, Y., Vidal, Y., and Tutivén, C. "Unsupervised damage detection for offshore jacket wind turbine foundations based on an autoencoder neural network". *Sensors*, 21(10) (2021)
7. Chatzi, E., Abdallah, I., Hofsäß, M., Bischoff, O., Barber, S., and Marykovskiy, Y. "Aventa AV-7 ETH Zurich Research Wind Turbine SCADA and high frequency Structural Health Monitoring (SHM) data" (2023)
8. Bilendo, F., Meyer, A., Badihi, H., Lu, N., Cambron, P., and Jiang, B. "Applications and modeling techniques of wind turbine power curve for wind farms—a review". *Energies*, 16(1):180 (2022)
9. Xanthopoulos, P., Pardalos, P. M., and Trafalis, T. B. *Robust Data Mining*. Springer (2013)
10. Tharwat, A., Gaber, T., Ibrahim, A., and Hassanien, A.E. "Linear discriminant analysis: A detailed tutorial". *Ai Communications*, 30:169–190, (2017)
11. Ioffe, S. "Probabilistic linear discriminant analysis". In Leonardis, A., Bischof, H., and Pinz, A., editors, *Computer Vision – ECCV 2006*, pages 531–542, Berlin, Heidelberg (2006) Springer Berlin Heidelberg.
12. Kotz, S. and Nadarajah, S. *Extreme value distributions: theory and applications*. World Scientific (2000)
13. Speciale, C., Tronci, E.M., Betti, R., and Milana, S. "Cepstral coefficients for wind turbine damage detection: Recognizing unknown classes in real-time". *Available at SSRN 5173243*.
14. Vallat, R. "Pingouin: statistics in python". *Journal of Open Source Software*, 3(31):1026 (2018)
15. Team, P.D. "`pingouin.anova` documentation". https://pingouin-stats.org/build/html/generated/pingouin.anova.html Accessed September 5, 2025.
16. Olejnik, S. and Algina, J. "Generalized eta and omega squared statistics: Measures of effect size for some common research designs". *Psychological Methods*, 8(4):434–447 (2003)
17. Lakens, D. "Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and anovas". *Frontiers in Psychology*, 4:863 (2013)