2nd Latin-American Workshop on Structural Health Monitoring
LATAM-SHM 2026
January, 7-9 2026

FEATURE-BASED NEURAL NETWORKS FOR IMPACT LOCALIZATION IN COMPOSITE STRUCTURES

Daniel del-Río-Velilla*, **Fernando Sánchez Iglesias***, **and Antonio Fernandez-Lopez***

* Universidad Politécnica de Madrid
Plaza del Cardenal Cisneros N3, 28040, Madrid, Spain
e-mail: daniel.delrio.velilla@upm.es, web page: https://www.upm.es/

Key words: SHM, Neural Networks, Feature Engineering, Composite, Signal Processing.

Abstract. Accurate impact localization is critical for the Structural Health Monitoring of composite materials, particularly in aerospace applications where early damage detection ensures safety and reduces maintenance costs. This study explores a neural network-based approach to localizing impacts on composite structures using engineered features extracted from piezoelectric sensor signals. Multiple MLP-based models have been trained with feature sets extracted from this signals using different supervised and unsupervised techniques. This systematic comparison highlights how the choice of feature selection strategy influences localization accuracy and model generalization. Notably, certain selection techniques yield models with stronger robustness when localizing impacts outside the original training grid. These findings suggest that careful feature selection, combined with domain-informed engineering, can enhance model performance and spatial generalization in Structural Health Monitoring applications involving composite structures.

1 INTRODUCTION

Composite materials are increasingly used in aerospace, automotive, and civil structures because of their high specific strength, stiffness, and fatigue resistance. However, their brittle and anisotropic nature makes them highly susceptible to barely visible impact damage, such as delamination and cracking of the matrix. Consequently, reliable and efficient impact monitoring and localization methods are essential within the Structural Health Monitoring (SHM) frameworks [1].

When an impact occurs, various stress waves propagate through the structure. Among them, guided acoustic waves (Lamb waves) are particularly valuable for SHM because of their sensitivity to small damages and low attenuation over long distances. Using arrays of piezoelectric transducers (PZTs), these waveforms can be recorded and analyzed to infer the location and severity of the impact [2]. Accurate localization, however, depends on extracting signal features that effectively capture the physical phenomena encoding spatial information.

Traditional impact localization methods are based on time-of-flight (ToF) models based on wave propagation velocities [3]. Although effective in simple configurations, they are sensitive to modeling errors, boundary conditions, and material variability. Machine learning (ML) methods have emerged as a promising alternative, capable of learning the nonlinear relationship between sensor data and impact position directly from experiments. In particular, artificial neural networks (ANNs) have shown strong performance in dealing with complex, high-dimensional data typical of SHM systems [4].

A key challenge in ML-based SHM is the selection and reduction of the dimensionality of the signal features. Raw sensor data often contains redundant or irrelevant information that hinders training efficiency and generalization. Unsupervised techniques such as Principal Component Analysis (PCA) and autoencoders can reduce feature dimensionality without labeled data [5], but may overlook task-relevant features. In contrast, supervised methods exploit the information in the label (e.g. impact coordinates) to extract features optimized for localization accuracy.

This work investigates how the choice between supervised and unsupervised feature extraction affects the performance of neural network–based impact localization in composite plates. An extensive experimental campaign was conducted on an AS4-8852 carbon/epoxy laminate, where the impacts of varying energies were recorded using a set of PZT sensors. The features in the time and frequency-domain were extracted and reduced using different algorithms before training densely connected neural networks under a ten-fold cross-validation scheme.

The results show that supervised feature extraction significantly improves localization accuracy and generalization, especially with a limited number of features, while unsupervised methods only achieve comparable performance when using larger feature sets. These findings highlight the importance of task-oriented dimensionality reduction in data-driven SHM and provide guidance for optimizing feature processing pipelines in impact localization systems.

2 MATERIALS AND METHODS

2.1 Specimen descriprion

The test specimen consists of a flat, square laminate plate fabricated from AS4-8852 unidirectional carbon/epoxy prepreg. The layup corresponds to a quasi-symmetric 11-ply configuration, with a nominal thickness of 2.024 mm and in-plane dimensions of 720×720 mm^2. For guided-wave sensing and impact detection, the plate is instrumented with eight circular piezoelectric transducers (PZTs), each with a diameter of 20 mm. The sensors are arranged in a regular grid of two columns and four rows. The sensor columns are positioned 188 mm from the plate edges along the Y-axis, while the rows are aligned along the X-axis with a uniform spacing of 150 mm, including the offset from the plate boundaries.

2.2 Experimental setup and data acquisition

Impact testing and signal acquisition were conducted using a custom-built automatic impactor based on a CNC platform (AICNC), as illustrated in Figure 1. The AICNC employs a kinematic system that enables precise positioning of a drop tube along the X–Y plane. Within the drop tube, an electromagnet controls the vertical motion of the impactor object (IO), enabling automated release from variable heights.

The IO, shown in Figure 2, consists of a 20 mm-diameter hemispherical tip mounted on a stepped cylindrical body. The cylindrical housing accommodates interchangeable disks, which allows a systematic adjustment of the IO mass. To enable magnetic handling by the AICNC system, the IO is fitted with a steel cap that can be lifted and released by the electromagnet. During each test, the IO is raised to a predefined release height and then dropped under gravity to strike the plate.

Impact events are recorded using a National Instruments NI-6626 data acquisition (DAQ) card, configured to sample for one second per impact. The acquired waveforms are subsequently compressed and stored for further processing. The plate is rigidly clamped along its two edges parallel

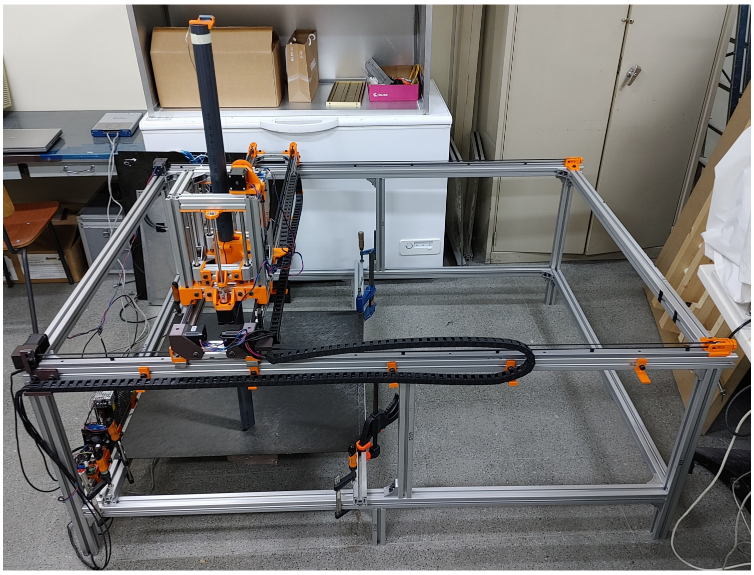

Figure 1: Automatic Impactor CNC

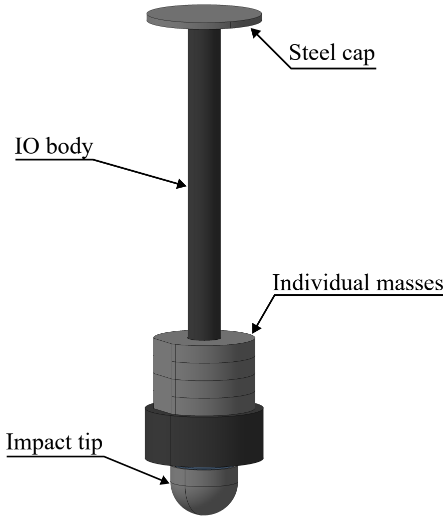

Figure 2: Impactor object

to the Y-axis, while the remaining edges remain free, replicating partially constrained boundary conditions.

For systematic impact localization experiments, the surface of the plate is discretized into a 21×21 grid of potential impact points. The tests were carried out at 32 distinct energy levels, defined by the combination of five masses (75, 100, 125, 150, and 175 g) and nine release heights ranging from 50 to 250 mm in increments of 25 mm.

2.3 Wave Phenomena Induced by Impact

An impact event excites multiple wave modes within a composite structure, which can be broadly classified as structural vibration and acoustic wave modes. Structural vibrations, associated with the global bending and flexural resonances of the plate, dominate the low-frequency spectrum and are strongly influenced by the boundary conditions of the specimen. In contrast, acoustic waves, or Lamb waves, propagate in the ultrasonic frequency range and exhibit dispersive characteristics that depend on frequency and laminate thickness. These waves are largely decoupled from structural modes and provide high-frequency information suitable for damage detection and impact localization.

To independently analyze structural and acoustic responses, the acquired time signals are processed using digital Butterworth filters: a low-pass filter extracts the structural response, while a high-pass filter isolates the Lamb-wave content. The high-pass filter employs a cutoff frequency of 5 kHz, selected based on the characteristics of the Mexican Hat (mexh) wavelet. The mexh wavelet was chosen for its strong sensitivity to impulsive, broadband events; its excellent time localization and derivative nature emphasize the short-duration acoustic waves generated by the impact while attenuating the low-frequency structural response. Notably, the wavelet representation also shows that, a few microseconds after the initial impact, the low-frequency structural modes exhibit amplitudes that are higher than at the very onset, reflecting the delayed excitation of global bending and flexural vibrations.

Representative examples of the processing results are presented in Figure 3, showing the raw signal along with the low- and high-pass filtered components, and in Figure 4, which displays the corresponding mexh wavelet power spectrum. Together, these figures illustrate the clear separation between the initial high-frequency Lamb-wave content and the subsequent low-frequency structural response. This decomposition is critical for the accurate determination of the impact onset time, defined by the arrival of the fundamental symmetric Lamb mode (S_0) in PZTs, and directly supports time-of-flight estimation for impact localization.

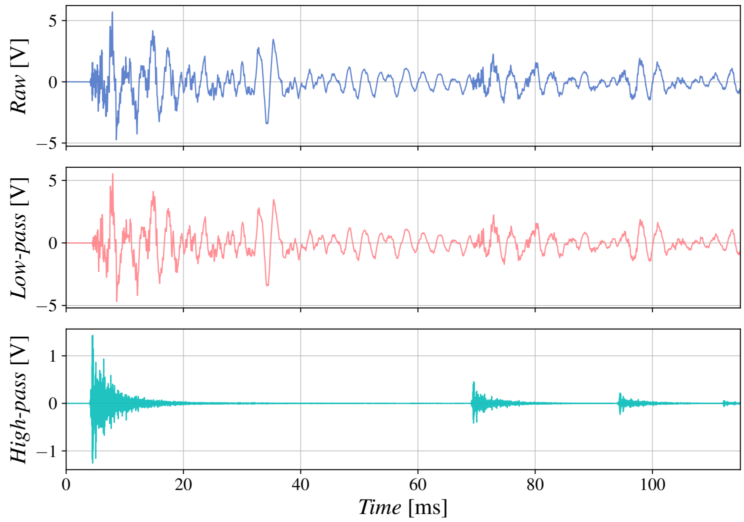

Figure 3: Representative single-channel impact signal.

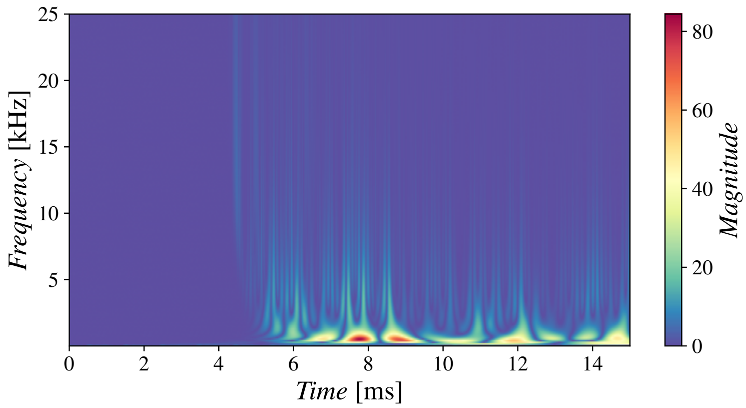

Figure 4: Continuous wavelet transform of the impact start using the Mexican Hat wavelet.

3 IMPACT LOCATION FEATURES

The feature extraction stage constitutes a critical component of the proposed impact localization framework, as it determines the quality and relevance of the information supplied to the predictive model. The primary objective of this process is to identify the most representative and discriminative set of features that can serve as the optimal input to the neural network, thereby allowing the highest possible predictive performance. Given that raw signals collected from PZT sensors are typically high-dimensional, noisy, and correlated, an effective feature extraction strategy is required to eliminate redundancy, improve signal interpretability, and improve model generalization.

3.1 Selected features

The specific features used for impact localization were selected based on established practices in SHM and impact detection [2, 3, 6, 7, 8]. These features are grouped into two main categories, time-domain and frequency-domain descriptors, which together capture both the transient and the spectral characteristics of the structural response.

Time-Domain Features

- *Time of Arrival (ToA):* The time of arrival in each sensor, defined as the first instant when the signal amplitude exceeds a threshold of 6σ relative to the baseline noise level. ToA is extracted from high-pass filtered (HPF) signals to emphasize high-frequency Lamb-wave content.
- *Impact Duration:* The temporal interval between the first ToA and the time when the last sensor signal decays below the 6σ threshold (HPF).
- *Maximum Rate of Change:* The maximum derivative of the signal amplitude with respect to time, along with its corresponding timestamp (HPF).
- *Statistical Descriptors:* Conventional parameters—maximum, minimum, mean, standard deviation, root-mean-square value (RMSV), kurtosis, and skewness—computed for each channel between the first ToA and the end of the impact. These are evaluated on low-pass-filtered (LPF) signals, where the structural response dominates.
- *Energy-Related Parameters:* The signal energy momentum (EM) and the average signal level (ASL), both derived from the LPF signals.

Frequency-Domain Features

- *Spectral Response:* The amplitude and frequency of the dominant peak obtained from the Fourier spectrum of each LPF signal.
- *Power Spectral Density Parameters (PSD):* a set of descriptors (A–E) that characterize the distribution of spectral energy across the frequency band of interest (LPF).
- *Derived Spectral Metrics:* Additional indicators such as the positive zero-crossing frequency $E(0)$ and the expected number of spectral peaks $E(p)$, which describe the distributional characteristics of the frequency content.

These features collectively describe both the temporal evolution and the spectral distribution of the impact response, forming a high-dimensional feature space that captures the dynamics of wave propagation and structural interaction.

3.2 Feature Engineering and Dimensionality Reduction

The feature engineering process aims to identify the most representative and discriminative subset of features that can serve as an optimal input to the neural network, thus maximizing predictive performance. Because the raw feature matrix derived from subsection 3.1 is highly dimensional and potentially redundant, an appropriate dimensionality reduction strategy is required to enhance the generalization of the model, mitigate noise effects, and reduce computational complexity.

The implemented pipeline evaluates a comprehensive range of unsupervised, supervised, and neural-network-based feature extraction techniques. All approaches begin with feature normalization using a Min–Max scaler to ensure numerical consistency and comparability among variables.

Unsupervised methods, such as Principal Component Analysis (PCA), Truncated Singular Value Decomposition (SVD), Independent Component Analysis (ICA), Factor Analysis (FA) and Kernel PCA (KPCA) are employed to uncover latent structures in the data without using label information. These techniques identify low-dimensional subspaces that retain most of the intrinsic variance and physical information in the original dataset. In addition, nonlinear models such as Autoencoders (AE) and Variational Autoencoders (VAE) are incorporated to capture complex, nonlinear relationships among features, providing more expressive latent representations.

Complementary to these, several supervised feature extraction approaches explicitly exploit the target information (impact coordinates) to guide the reduction of dimensions. These include Mutual Information (MI) and Random Forest (RF) feature importance ranking, Correlation-based selection, and model-driven methods such as Recursive Feature Elimination (RFE) and Partial Least Squares (PLS). A Supervised Autoencoder (SAE) is also implemented to jointly optimize reconstruction and prediction losses, enforcing that the learned latent space remains compact and discriminative with respect to the localization targets. This combination of statistical and deep learning approaches allows for a balanced exploration of feature relevance from both data-driven and physics-informed perspectives.

To ensure statistical robustness and assess the reliability of each method, the complete feature extraction process is performed under a 10-fold cross-validation scheme ($K = 10$). This procedure enables the evaluation of the performance of the true statistical feature-model, reducing the dependence on specific data splits and delivering a more reliable assessment of the effectiveness of each method. The dimensionality of the reduced feature spaces is systematically varied (typically between 5 and 100 components) to identify the configuration that provides the optimal trade-off between compactness and predictive capability for the neural network.

4 RESULTS

4.1 Neural Network Architecture for Impact Localization

The objective of this study is to evaluate the effectiveness of various signal features for the localization of impact events in composite structures. A multilayer perceptron (MLP) was selected as the learning model for this task. This type of architecture has demonstrated strong performance in multiple SHM applications, including impact localization, damage detection, and guided-wave signal interpretation. Representative studies include works on damage location with Lamb wave features based on deep learning [9] or neural networks for the location of the impact source [10, 4]. These and other contributions highlight the suitability of fully connected networks for modeling nonlinear relationships between measured sensor features and spatial coordinates of impacts.

The main hyperparameters defining the MLP architecture are the number of layers, the number of neurons per layer, and activation functions. To ensure architectural consistency between models trained with different feature sets, the number and size of the hidden and output layers were kept constant. Only the size of the input layer was modified to match the dimensionality of the corresponding feature vector.

A compressor-style architecture was adopted, in which the number of neurons initially gradually expands to 512 in the third hidden layer and then is progressively reduced through successive layers until reaching the two-neuron output layer representing the impact coordinates (X,Y). Among various available activation functions [11], the Rectified Linear Unit (ReLU) was selected for its computational simplicity and efficient convergence properties [12]. The latter is particularly advantageous in this study, given the need to train approximately $2\,400$ individual models, which imposes substantial computational demands. The complete training campaign required a cumulative runtime of 72 hours with a NVIDIA 4070 GPU and a total memory usage of 3.38 GB.

All models were trained using the Adam optimization algorithm with a learning rate of 10^{-3}, a total of 50 training epochs, and a batch size of 32 samples. These training parameters were selected to balance convergence stability and computational efficiency in all feature configurations evaluated. The loss function used was the Mean Squared Error (MSELoss), which is particularly suitable for this task because it directly quantifies the squared Euclidean distance between the predicted and true impact coordinates. Consequently, minimizing this loss function corresponds to minimizing the physical localization error in millimeters, thereby aligning the training objective with the ultimate performance metric of interest.

To ensure robust convergence, the learning rate was dynamically adjusted using PyTorch's learning rate scheduler, which decreases the learning rate when the validation loss plateaus. This adaptive scheduling mechanism helps the optimizer escape shallow minima and improve the generalization of the final model. The final model retained for each training instance corresponds to the epoch that yields the lowest validation error throughout the training process, rather than necessarily the final epoch. This selection strategy ensures that the performance reported reflects the best generalization capability of the model rather than its last recorded state.

The final DNN configuration used in this work is summarized in Table 1.

Table 1: Nerual network architecture.

Layer	**Num. neurons**	**Act. funct.**
In	number of features	ReLU
1	128	ReLU
2	256	ReLU
3	512	ReLU
4	256	ReLU
5	128	ReLU
6	64	ReLU
7	32	ReLU
Out	2	ReLU

4.2 Evaluation Metrics and Statistical Analysis

The performance of each trained model was evaluated using three statistical indicators: the mean prediction error, the standard deviation, and the 95th percentile error. The prediction error was quantified as the Euclidean Error (EE), defined as the Euclidean distance between the actual and predicted impact coordinates, expressed in millimeters. Given the large number of trained networks, the reported values correspond to the average across all ten cross-validation folds for each specific combination of feature set and dimensionality-reduction method.

Figures 5, 6, and 7 summarize the performance metrics obtained from the test datasets of each fold. It is important to note that these test impacts were never used during training, ensuring that all reported results represent the generalization of the models. Each heatmap presents the corresponding statistical metric while differentiating between unsupervised and supervised feature extraction techniques. This distinction allows for a direct comparison of their relative effectiveness as the number of input features increases.

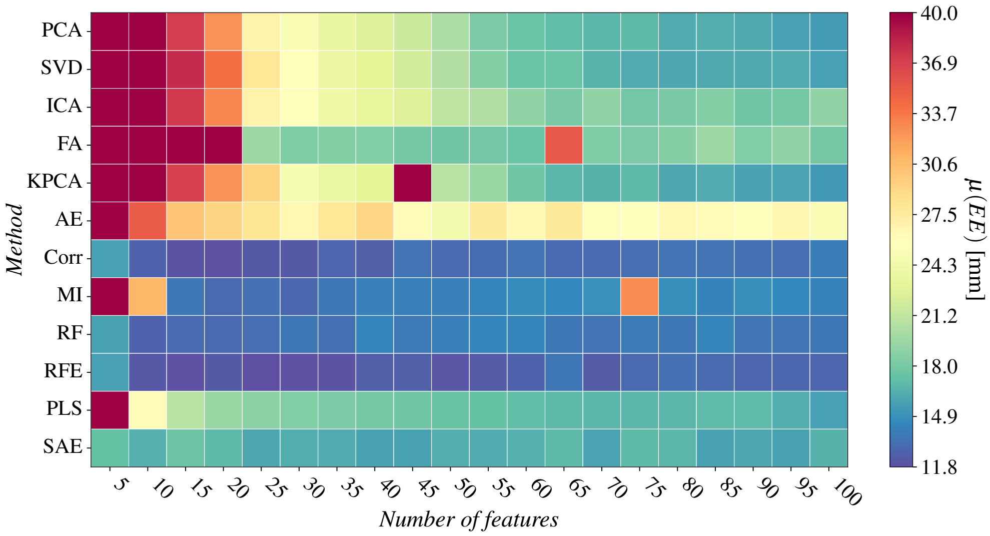

Figure 5: Mean prediction error of models as a function of the number of features for supervised and unsupervised feature extraction methods.

Overall, the results indicate that supervised feature extraction techniques consistently outperform unsupervised ones when using a small number of features. This shows that supervised methods are better suited to identify the most informative variables directly related to the impact localization task. However, as the number of features increases, the performance gap between the two approaches progressively narrows, and both methods exhibit comparable prediction accuracy at higher feature dimensionalities. This convergence suggests that with sufficient input information, even unsupervised techniques can approximate the relevant relationships captured by supervised approaches, although less efficiently.

It should be noted that in Figures 5–7, the maximum values of the colorbars are not the absolute global maxima of the datasets. Some models exhibited higher error values, but the colorbar limits were intentionally capped to preserve visual resolution and contrast in the lower-error regions, where most results are concentrated. This approach enhances interpretability while maintaining a consistent visual comparison between the different configurations.

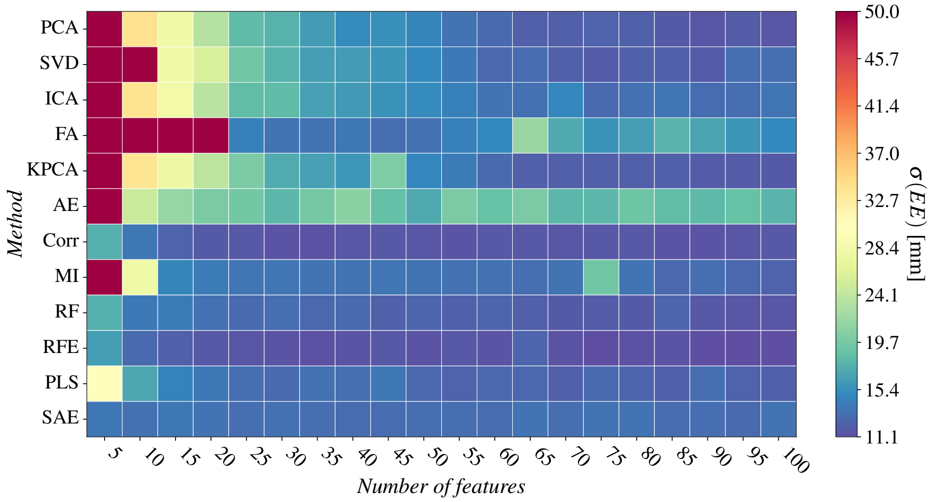

Figure 6: Standard deviation of prediction errors as a function of the number of features for supervised and unsupervised feature extraction methods.

4.3 Overfitting performance

Overfitting arises when a learning model continues to improve its performance in training data, while its accuracy in unseen data—represented by the validation or test sets—deteriorates. This behavior occurs when the model memorizes training-specific patterns rather than learning generalizable relationships between the input features and the target variables. Consequently, the training error decreases while the validation error begins to increase, signaling a loss of generalization capability.

To examine this phenomenon, two representative configurations were selected: an unsupervised feature extraction model using Principal Component Analysis (PCA) and a supervised model based on correlation-derived features (Corr). For both cases, the average prediction errors were calculated across all datasets (training, validation and test), cross-validation folds, and number of features. The resulting trends are depicted in Figure 8.

As shown in Figure 8, the PCA-based model exhibits consistently higher prediction errors in all datasets, confirming that unsupervised feature extraction provides less discriminative information for the localization task. However, as the number of features increases, the gap between training and validation errors progressively decreases, indicating a partial improvement in the model's capacity to generalize when provided with a richer set of characteristics.

In contrast, the Corr-based model achieves substantially lower prediction errors overall, demonstrating the advantage of supervised feature extraction in aligning the reduced feature space with the localization objective. Nevertheless, when the number of features exceeds approximately 15, a gradual divergence appears between the training and validation curves: the training error continues to decrease, while the validation error begins to rise. This behavior is a classical signature of overfitting, suggesting that beyond this feature dimensionality, the model starts to capture noise or redundant information that does not contribute to predictive generalization.

These observations confirm that, while supervised feature extraction significantly improves localization accuracy, careful tuning of the feature dimensionality remains essential to avoid overfitting and preserve model robustness.

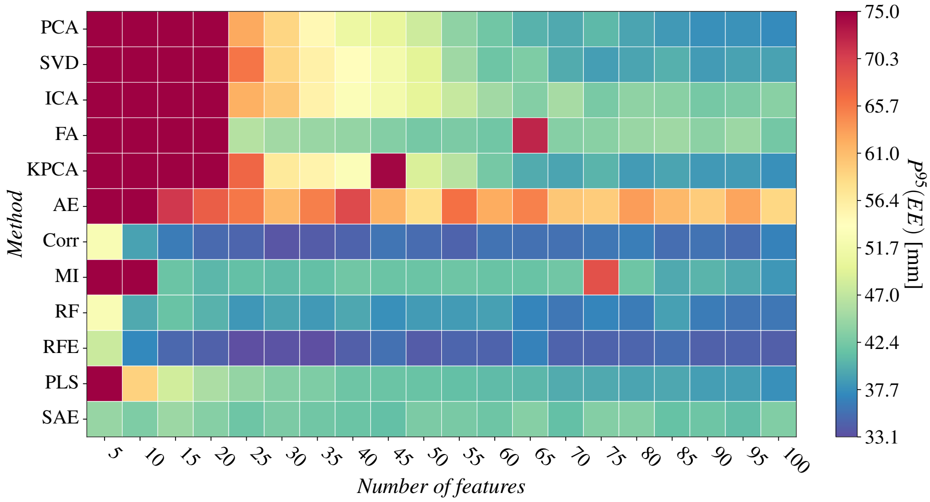

Figure 7: 95th percentile of prediction errors as a function of the number of features for supervised and unsupervised feature extraction methods.

5 DISCUSSION

The results obtained in this study provide a complete understanding of how different feature extraction strategies influence the performance of neural network–based impact localization in composite structures. In particular, comparative analysis between supervised and unsupervised dimensionality-reduction techniques reveals significant differences in both prediction accuracy and generalization capability.

When using a small number of features, the models trained with supervised feature extraction demonstrate a clear performance advantage. These techniques explicitly incorporate the relationship between the extracted features and the impact coordinates, enabling the neural network to learn a more relevant representation of the physical phenomena involved in the impact event. Consequently, the mean prediction errors, as well as the dispersion metrics (standard deviation and 95th percentile), are consistently lower compared to those obtained using unsupervised approaches.

In contrast, unsupervised feature extraction methods—such as PCA—tend to emphasize global data variance rather than the specific information necessary for localization. As a result, these models exhibit higher overall prediction errors. However, as the number of features increases, their performance progressively approaches that of supervised methods. This convergence suggests that with sufficient feature dimensionality, unsupervised techniques can indirectly capture part of the relevant information, although they remain less efficient in doing so.

The analysis of overfitting behavior further highlights the advantages of supervised feature extraction. Models trained with correlation-based (Corr) features achieved lower prediction errors and maintained better consistency across training, validation, and test datasets. Nevertheless, when the number of features exceeded approximately 15, the validation error began to increase while the training error continued to decrease, indicating the onset of overfitting. This phenomenon reflects the increasing sensitivity of the model to noise or redundant variables introduced at higher dimensions. In contrast, PCA-based models showed higher overall errors but a reduced gap between training and validation errors as the number of features grew, suggesting a limited yet more stable generalization pattern.

Collectively, these findings confirm that supervised dimensionality-reduction techniques enable

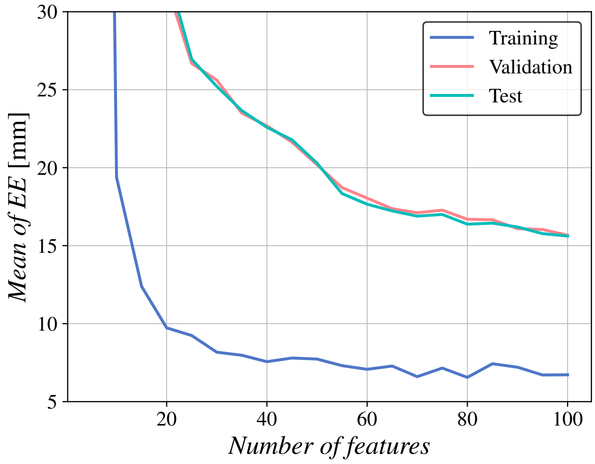

(a) Model predictions using PCA-based features.

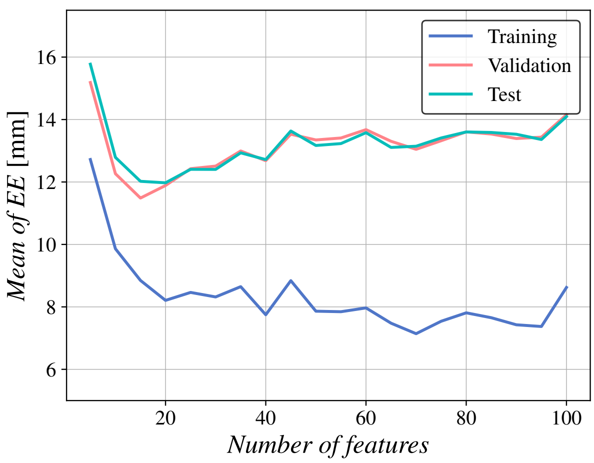

(b) Model predictions using correlation-based (Corr) features.

Figure 8: Comparison of the mean prediction error for training, validation, and test sets in models trained with PCA and Corr features.

the neural network to better focus on the most physically meaningful aspects of sensor signals, facilitating both accurate and generalizable impact localization. However, excessive feature dimensionality can introduce overfitting even in supervised configurations, emphasizing the need for a balanced feature selection process.

6 CONCLUSIONS

Based on the experimental and computational analysis presented, the following conclusions can be drawn.

- Supervised feature extraction techniques provide substantially better localization accuracy than unsupervised methods, particularly when the number of available features is small.
- Unsupervised methods such as PCA require a larger number of features to achieve comparable accuracy, as they are unable to prioritize the task-relevant components of the signal.
- Overfitting behavior was observed in supervised models when the number of features increases, indicating that excessive dimensionality introduces redundant or noisy information that reduces generalization.
- Supervised feature extraction leads to more efficient use of available data, resulting in smaller mean errors and improved consistency across training, validation, and test sets.
- The study highlights the importance of balancing the relevance of features and dimensionality to achieve optimal performance in neural network–based impact localization frameworks.

In summary, the integration of supervised dimensionality-reduction techniques within neural network architectures significantly enhances the ability of data-driven models to locate impacts in composite structures accurately. These results underline the importance of task-oriented feature extraction in SHM applications and provide a foundation for further research into adaptive, physics-informed feature selection methods that could further improve localization robustness and scalability.

REFERENCES

[1] Cecilia L. Wilson, Kuldeep Lonkar, Surajit Roy, Fotis Kopsaftopoulos, and Fu Kuo Chang. 7.20 Structural Health Monitoring of Composites. *Comprehensive Composite Materials II*, pages 382–407, 1 2018.

[2] Michal Blahacek, M. Chlada, and Z. Prevorovský. Acoustic Emission Source Location Based on Signal Features. *Advanced Materials Research*, 13-14:77–82, 2 2006.

[3] M. Shehadeh, J. A. Steel, and R. L. Reuben. Acoustic emission source location for steel pipe and pipeline applications: The role of arrival time estimation. *Proceedings of the Institution of Mechanical Engineers, Part E: Journal of Process Mechanical Engineering*, 220(2):121–133, 2006.

[4] Daniel del Río-Velilla, Andrés Pedraza, and Antonio Fernández-López. Impact localization in composite structures with Deep Neural Networks. *Structural Health Monitoring*, 2024.

[5] Chengjia Han, Zixin Wang, Yuguang Fu, Shirley Dyke, and Adnan Shahriar. Transfer-AE: A novel autoencoder-based impact detection model for structural digital twin. *Applied Soft Computing*, 166:112174, 11 2024.

[6] Jiao Jingpin, Wu Bin, and He Cunfu. Acoustic emission source location methods using mode and frequency analysis. *Structural Control and Health Monitoring*, 15(4):642–651, 6 2008.

[7] Matthew Geoffrey Baxter, Rhys Pullin, Karen M. Holford, and Sam L. Evans. Delta T source location for acoustic emission. *Mechanical Systems and Signal Processing*, 21(3):1512–1520, 4 2007.

[8] Pietro Pedemonte, Wieslaw J. Staszewski, Francesco Aymerich, Mike S. Found, and Pierluigi Priolo. Signal processing for passive impact damage detection in composite structures. *https://doi.org/10.1117/12.436470*, 4326:169–178, 8 2001.

[9] Yumeng Gao, Lingyu Sun, Ruijie Song, Chang Peng, Xiaobo Wu, Juntao Wei, Mingshun Jiang, Qingmei Sui, and Lei Zhang. Damage localization in composite structures based on Lamb wave and modular artificial neural network. *Sensors and Actuators A: Physical*, 377:115644, 10 2024.

[10] Iuliana Tabian, Hailing Fu, and Zahra Sharif Khodaei. A Convolutional Neural Network for Impact Detection and Characterization of Complex Composite Structures. *Sensors 2019, Vol. 19, Page 4933*, 19(22):4933, 11 2019.

[11] Prajit Ramachandran, Barret Zoph, and Quoc V Le Google Brain. SEARCHING FOR ACTIVATION FUNCTIONS.

[12] Chun-Nan Chou, Chuen-Kai Shie, Fu-Chieh Chang, Jocelyn Chang, and Edward Y Chang. Representation Learning on Large and Small Data.