

Article

Domain-Adaptive Graph Attention Semi-Supervised Network for Temperature-Resilient SHM of Composite Plates

Nima Rezazadeh ^{1,2,*} , Alessandro De Luca ² , Donato Perfetto ² , Giuseppe Lamanna ² , Fawaz Annaz ¹ and Mario De Oliveira ¹ 

¹ School of Architecture, Built Environment, Computing and Engineering, Birmingham City University, Birmingham B4 7XG, UK; fawaz.annaz@bcu.ac.uk (F.A.); mario.deoliveira@bcu.ac.uk (M.D.O.)

² Department of Engineering, University of Campania “Luigi Vanvitelli”, 81031 Aversa, Italy; alessandro.deluca@unicampania.it (A.D.L.); donato.perfetto@unicampania.it (D.P.); giuseppe.lamanna@unicampania.it (G.L.)

* Correspondence: nima.rezazadeh@mail.bcu.ac.uk

Abstract

This study introduces GAT-CAMDA, a novel framework for the structural health monitoring (SHM) of composite materials under temperature-induced variability, leveraging the powerful feature extraction capabilities of Graph Attention Networks (GATs) and advanced domain adaptation (DA) techniques. By combining Maximum Mean Discrepancy (MMD) and Correlation Alignment (CORAL) losses with a domain-discriminative adversarial model, the framework achieves scalable alignment of feature distributions across temperature domains, ensuring robust damage detection. A simple yet at the same time efficient data augmentation process extrapolates damage behaviour across unmeasured temperature conditions, addressing the scarcity of damaged-state observations. Hyperparameter optimisation via Optuna not only identifies the optimal settings to enhance model performance, achieving a classification accuracy of 95.83% on a benchmark dataset, but also illustrates hyperparameter significance for explainability. Additionally, the GAT architecture’s attention demonstrates the importance of various sensors, enhancing transparency and reliability in damage detection. The dual use of Optuna serves to refine model accuracy and elucidate parameter impacts, while GAT-CAMDA represents a significant advancement in SHM, enabling precise, interpretable, and scalable diagnostics across complex operational environments.

Keywords: structural health monitoring (SHM); composite materials; graph attention networks (GATs); domain adaptation (DA); temperature variability; explainability



Academic Editor: Jiawei Xiang

Received: 9 October 2025

Revised: 6 November 2025

Accepted: 7 November 2025

Published: 9 November 2025

Citation: Rezazadeh, N.; De Luca, A.; Perfetto, D.; Lamanna, G.; Annaz, F.; De Oliveira, M. Domain-Adaptive Graph Attention Semi-Supervised Network for Temperature-Resilient SHM of Composite Plates. *Sensors* **2025**, *25*, 6847. <https://doi.org/10.3390/s25226847>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Ultrasonic-guided waves (UGWs), particularly Lamb waves, have become fundamental in SHM due to their ability to propagate over long distances with minimal attenuation, making them highly effective for inspecting large structures such as wind turbine blades and aerospace components. Their high sensitivity to both surface and internal defects facilitates robust damage detection when interacting with structural anomalies that alter propagation characteristics [1–3]. Piezoelectric sensors, pivotal in generating and detecting Lamb waves, offer a scalable and cost-effective SHM solution, providing widespread monitoring capabilities [4,5]. However, the multimodal and dispersive nature of Lamb waves introduces significant challenges in interpreting signals and accurately localising

damage. These difficulties are compounded by environmental and operational variabilities (EOVs) such as temperature, humidity, and mechanical stress, which distort wave propagation and impair system reliability [2,6]. In composite materials such as Carbon Fibre-Reinforced Polymer (CFRP), the heterogeneity and anisotropy of the material further necessitate specialised techniques to manage failure modes, including delamination and matrix cracking [7–10].

1.1. Temperature-Induced Variability and Conventional Compensation Methods

Among the various EOVs, temperature changes can influence guided wave propagation by affecting wave speed, phase, and amplitude, potentially masking or mimicking damage signatures [11–13]. Numerous compensation strategies have been proposed to mitigate these effects. Scheerer and Lager [14] evaluated three temperature compensation techniques (Best Baseline, Signal Interpolation, and Frequency Shift), concluding that Frequency Shift was effective under moderate temperature variations (≤ 5 °C). In addition, statistical modelling approaches have been applied to tackle temperature-induced variability. For example, Silva et al. [15] used auto-regressive models with cubic spline extrapolation to track damage progression under temperature fluctuations, although the approach required extensive baseline data and thus faced scalability limitations. Ren et al. [16] adopted Gaussian Process Regression with Monte Carlo sampling to quantify uncertainty in damage detection, demanding large datasets that may not be practical for real-time SHM. Complementing these statistical methods, semi-analytical and finite element (FE) modelling also feature prominently. Ren et al. [17] explored temperature compensation in composite structures through a semi-analytical FE model, validating their findings experimentally but relying on well-defined material and geometric properties. Similarly, Perfetto et al. [18] studied guided wave propagation under thermal loads using FE modelling, deriving time compensation factors that nonetheless require extensive experimental corroboration.

1.2. Data-Driven Methods for Temperature Compensation

Given the limitations of purely model-based and statistical strategies, data-driven methods are increasingly used to compensate for temperature effects [19]. Ferreira et al. [20] proposed a Bayesian framework that integrates FE model updating with neural networks to produce synthetic datasets for temperature compensation, reducing reliance on physical experiments but introducing potential inaccuracies when dealing with poorly characterised conditions. Giannakeas et al. [21] presented an up-scaling methodology using Bayesian regression to extend temperature compensation factors across structural scales, thus improving damage detection reliability, though at the expense of comprehensive validation efforts. Meanwhile, Cheng et al. [22] combined fibre optic sensors with piezoelectric transducers in a hybrid system for guided wave monitoring, although the sensitivity adjustments required for fibre optic sensors add complexity to the system.

GATs have garnered particular interest for capturing intricate spatial-temporal dependencies in graph-structured sensor data [23,24]. Niu et al. [25], for instance, leveraged attention mechanisms within GATs to handle incomplete SHM datasets by accounting for spatial-temporal relationships among sensors. Zhao et al. [26] further demonstrated a virtual sensor for bearing load prediction using heterogeneous temporal graph neural networks that explicitly modelled sensor signal dependencies and outperformed standard convolutional neural networks in load estimation tasks.

1.3. Transfer Learning for Mitigating EOVs in SHM

While data-driven models can capture complex wave behaviours, they often depend on large, high-quality labelled datasets that may not be available across different environmental regimes. Transfer learning (TL) therefore offers a promising solution, focusing on

translating knowledge obtained in one domain to enhance performance in another [27]. Among TL methods, adversarial approaches such as Domain-Adversarial Neural Networks (DANNs) operate by minimising discrepancies between source and target distributions, thereby generating domain-invariant features. Ozdagli and Koutsoukos [28] demonstrated this technique in SHM, achieving improved cross-domain damage detection. Other feature-based alignment methods have also emerged: Zhuojun et al. [29] used MMD to reconcile simulated and experimental Lamb wave data for precise damage localisation, and Wang et al. [30] combined MMD with techniques including Variational Mode Decomposition and Transfer Component Analysis (TCA) to boost damage detection performance. Although many of these studies investigated TL under EOVs such as variations in loading or operating conditions, they rarely addressed temperature explicitly, which remains a critical and underexplored factor influencing guided wave propagation and model generalisation.

CORAL offers another feature-alignment strategy, albeit less commonly applied to UGW-based SHM. Poole et al. [31] employed CORAL in a partial DA setting, illustrating the benefit of aligning source and target covariance matrices for population-based SHM. By normalising datasets to account for anticipated EOVs, CORAL mitigates negative transfer and class imbalance, thereby improving damage localisation even with limited target-domain data. These advances collectively illustrate how TL can address data scarcity and ensure robust performance under varying environmental or operational conditions, making it an attractive approach for temperature compensation.

1.4. The Need for Explainability in ML-Based SHM

Despite these advances in data-driven and TL approaches, the opaque nature of many ML algorithms can hinder adoption in critical SHM tasks, where trust, interpretability, and actionable insights are essential. Explainable AI methods, including SHapley Additive exPlanations (SHAP) [32] and Local Interpretable Model-agnostic Explanations (LIME) [33], have been proposed to tackle the “black box” problem. As highlighted by Schnur et al. [34], neural networks are widely employed in SHM applications; however, their lack of interpretability limits their suitability for safety-critical contexts. Consequently, machine learning outcomes must be physically interpretable to facilitate their deployment in such domains.

To this end, Salih et al. [35] integrated SHAP and LIME into SHM workflows, with SHAP quantifying feature contributions and LIME providing localised signal interpretations. Tiwari et al. [36] harnessed SHAP in an ensemble approach for estimating shear strength in reinforced concrete, highlighting its utility in identifying sensor contributions. Azad and Kim [37] developed an explainable Vision Transformer for polymer composite diagnostics, employing attention mechanisms to highlight crucial decision regions, although at a substantial computational cost.

1.5. Main Contributions

This research addresses key limitations in the SHM of composite plates under temperature-induced EOVs through a novel framework, GAT-CAMDA. The main contributions are as follows:

- Combines multiple temperature domains into a single target domain, improving generalisation and reflecting real-world variability for greater robustness and practicality.
- Integrates CORAL and MMD losses to align feature distributions across temperatures, explores CORAL’s effectiveness in SHM, and employs GATs to capture complex spatial-temporal dependencies in UGW data for accurate damage detection.

- Uses GAT attention weights to visualise and quantify sensor contributions, enhancing model transparency and providing valuable insights into sensor importance for both theoretical and practical SHM applications.
- Employs Optuna [37], an automated hyperparameter optimisation framework, to systematically fine-tune model parameters within the semi-supervised process, thereby enhancing robustness and generalisability across temperature domains without requiring labelled target data for the DA stage [37].

The novelty of this work lies in its ability to generalise across diverse temperature conditions by merging multiple temperature domains into a unified framework. In contrast to previous studies restricted to narrow or isolated settings, this approach more closely represents real-world environmental variabilities. Furthermore, the thorough investigation of CORAL loss, an underexplored method in SHM applications, combined with the advanced feature extraction capabilities of GATs, marks a substantial step forward in aligning feature distributions and extracting discriminative patterns. By incorporating explainability through attention-weight visualisation and leveraging semi-supervised optimisation with Optuna, the proposed framework not only enhances damage detection accuracy but also delivers a transparent and scalable solution, setting a higher benchmark in SHM under complex environmental conditions.

The remainder of this paper is structured as follows: Section 2 presents the proposed framework, detailing its methodological advancements along with the implemented data synthesising technique. Section 3 outlines the case study, including the dataset and experimental setup. Section 4 discusses the results and evaluates the framework's performance, and Section 5 concludes with key findings and future directions.

2. Materials and Methods

2.1. Overview of the Proposed Framework

GAT-CAMDA is a semi-supervised deep learning framework designed to address the challenge of SHM across different environmental conditions subjected to data scarcity. This pipeline comprises multiple stages, beginning with feature extraction using GATs, followed by DA through DANN enhanced with MMD and CORAL losses, and concluding with damage detection in the target domain system. To address data scarcity, damaged responses at target temperatures were estimated by leveraging temperature-dependent trends in healthy signals. The objective function in the optimisation procedure (Optuna) was to maximise the classification accuracy on the validation set of the target domain. To improve the explainability of the framework, the contributions of various hyperparameters and the significance of individual sensors were calculated and visualised.

In this study, the source domain refers to the data collected from a structural system under baseline environmental conditions (e.g., at a reference temperature), where labelled observations of both healthy and damaged states are available. The target domain comprises the data from the same structural system but under different environmental conditions (e.g., at various temperatures), where labelled damaged observations are scarce or unavailable. The fundamental challenge is to transfer knowledge learned from the source domain to accurately detect damage in the target domain, despite the distribution shift caused by EOVs.

Figure 1 displays a schematic of the feature extraction and DA stages of GAT-CAMDA on the assumption that there are four GATs layers.




Figure 1. A schematic of GAT-CAMDA.

In the following parts each of these blocks are elaborated.

2.2. Graph Attention Networks

GATs were implemented as the feature extractor to capture both local and global relational structures within the input data. While conventional feature extraction methods excel at identifying localised patterns, they often struggle to comprehensively represent the complex interdependencies present in certain datasets more specifically, on the assumption that there are multi-sensor readings. GATs address this challenge by employing attention mechanisms that dynamically weight the importance of connections based on the significance of neighbouring features [38], thereby enabling a richer and more contextually informed representation. The process of a GAT can be summarised as follows [23]:

1. Feature transformation: Each node's input feature vectors x_i and x_j undergo a shared linear transformation:

$$x'_i = Wx_i, x'_j = Wx_j \quad (1)$$

where W is a learnable weight matrix; the transformed feature vectors x'_i and x'_j are then used in subsequent steps.

2. Computation of importance scores: A self-attention mechanism computes unnormalised importance scores e_{ij} , quantifying the relevance of the j -th node's features to the i -th node:

$$e_{ij} = a(x'_i, x'_j) \quad (2)$$

where a is a learned vector that parameterises the attention score.

3. Normalisation of attention coefficients: The scores e_{ij} are normalised using the softmax function to produce attention coefficients α_{ij} :

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k \in \mathcal{N}(i)} \exp(e_{ik})} \quad (3)$$

This normalisation ensures that the attention coefficients are comparable across the neighbourhood of node i . Assuming a single-layer feedforward neural network as the

attention mechanism a , and employing the LeakyReLU (Leaky Rectified Linear Unit) activation function, which introduces a small non-zero response for negative inputs to maintain gradient flow and prevent inactive neurons, the attention coefficients can be calculated as follows:

$$\alpha_{ij} = \frac{\exp\left(\text{LeakyReLU}\left(\vec{a}^\top [x'_i \| x'_j]\right)\right)}{\sum_{k \in \mathcal{N}(i)} \exp\left(\text{LeakyReLU}\left(\vec{a}^\top [x'_i \| x'_k]\right)\right)} \quad (4)$$

in which $\vec{a}^\top$ denotes the transpose of a (the attention parameter vector), and $\|$ indicates concatenation.

4. Feature aggregation: Each node's output feature is computed as a weighted sum of its neighbours' transformed features:

$$x''_i = \sigma\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij} x'_j\right) \quad (5)$$

In this equation, σ denotes an activation function (e.g., ReLU) and x''_i represents the refined feature for node i .

5. Multi-head attention: To improve stability and expressiveness, multiple attention mechanisms (heads) are employed. Each head independently computes its own set of attention coefficients and aggregated features:

$$x''_i = \parallel_{k=1}^K \sigma\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(k)} x'_j\right) \quad (6)$$

here, K denotes the number of attention heads; for the final layer, concatenation is replaced by averaging as follows:

$$x''_i = \sigma\left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(k)} x'_j\right) \quad (7)$$

Figures 2a and 2b present the attention mechanism described above, e.g., e_{ij} , and the multi-head attention when there are four heads, i.e., $K = 4$, by node 1 and its neighbour nodes, respectively.




Figure 2. A schematic of the (a) self-attention mechanism and (b) multi-head attention mechanism in the GAT layer [39].

In Figure 2b, each line style and colour represent independent (here four) attention calculations.

2.3. Domain Adaptation

The DA stage of GAT-CAMDA combined DANN with MMD and CORAL loss functions to handle more severe distribution gaps between the source and target domains.

2.3.1. Domain-Adversarial Neural Network

The DANN unit of GAT-CAMDA consists of a feature extraction unit that is responsible for projecting the source and target domain data into a shared space. Since in GAT-CAMDA an extra feature extraction stage is applied (i.e., GATs), the feature extractor unit of the DANN projects the distilled features from GATs to the shared space.

There are two heads connected to the DANN, i.e., a damage classifier and a domain discriminator. The damage classification unit predicts fault classes, while the domain discriminator determines whether input data come from the source or target domain.

Looking at Figure 1, it can be seen that a Gradient Reversal Layer (GRL) is used between the feature extractor and domain discriminator [40]; during the forward pass, the GRL acts as an identity layer, but during backpropagation, it reverses the gradients by -1 .

The damage classifier consists of three blocks, i.e., input, hidden, and addition blocks, each of which is composed of various layers such as linear layers, batch normalisation, and dropout, among others. Figure 3 presents a schematic of this neural network.

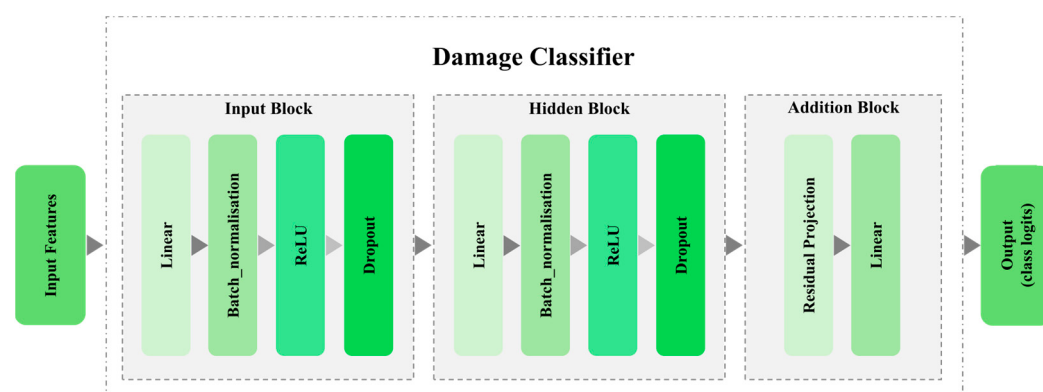


Figure 3. Schematic of damage classifier in GAT-CAMDA.

Structuring an overly complex damage classifier in DA tasks can lead to overfitting to the source domain [41]. Consequently, it is crucial to establish a trade-off between the model's complexity and its capacity to generalise effectively to target domains. This balance ensures that the classifier maintains robust performance across diverse domains without being unduly influenced by the intricacies of the source data. As a result, a simple structure was chosen as the classifier in GAT-CAMDA.

In contrast, the domain discriminator comprises five blocks. Each block contains a linear layer followed by a LeakyReLU activation to introduce non-linearity while avoiding the dying ReLU problem. Dropout layers are included for regularisation, and the final block maps the learned features to domain-specific logits, enabling discrimination between source and target domains. A deeper discriminator architecture was selected to ensure sufficient representational power to capture subtle domain shifts and to generate a meaningful adversarial signal for the feature extractor during training. This design choice establishes an equilibrium between the discriminator and the classifier: the classifier remains generalisable, while the discriminator is strong enough to challenge it, thereby promoting robust and domain-invariant feature learning.

The loss functions for damage classifier (L_D) and domain discriminator (L_T) are well discussed in the literature [40].

2.3.2. Maximum Mean Discrepancy

The MMD quantifies the distance between the average feature representations. This metric is pivotal in evaluating how effectively the feature distributions of the two domains align, which is essential for successful DA. The MMD (L_M) between two sets of observations belonging to domain s and t can be computed as follows [42]:

$$L_M = \left\| \frac{1}{n} \sum_{i=1}^n \phi(x''_{s_i}) - \frac{1}{m} \sum_{j=1}^m \phi(x''_{t_j}) \right\|^2 \quad (8)$$

in which x''_{s_i} and x''_{t_j} represent the sets of features from the source and target domains, respectively. The function ϕ maps these features into a higher-dimensional feature space, and n and m are the number of samples in each domain.

By mapping the features into a higher-dimensional space, MMD loss effectively captures more intricate patterns and discrepancies between the domains; a lower MMD loss indicates a greater similarity between the domains, which is indicative of improved DA performance.

When applying MMD, it is essential to select an appropriate kernel type, such as linear or Radial Basis Function (RBF). Specifically, for RBF kernels, adjusting the Gamma parameter is crucial, as it governs the kernel's bandwidth and subsequently influences the sensitivity of the MMD measure to discrepancies in feature distributions.

2.3.3. Correlation Alignment

The CORAL loss functions as a metric for evaluating the alignment between two domains by quantifying the discrepancies in their feature covariance matrices. Within the scope of DA, it is employed to align the data distributions of the source and target domains, thereby reducing statistical variations to enhance generalisation. By minimising this loss, the features that are learned become more consistent across both domains, thereby increasing the prediction accuracy for the target domain. The central concept of CORAL loss involves addressing domain shift by aligning second-order statistics, specifically through the minimisation of the Frobenius norm of the difference between the covariance matrices of the source and target domains. The CORAL loss (L_C) function can be expressed as follows [43]:

$$L_C = \frac{1}{4d^2} \|\text{Cov } x''_s - \text{Cov } x''_t\|_F^2 \quad (9)$$

where $\text{Cov } x''_s$ and $\text{Cov } x''_t$ denote the covariance matrices of the source and target features, respectively, d represents the dimensionality of the features, and $\|\cdot\|_F$ refers to the Frobenius norm. CORAL loss can be computed on a per-class basis or as an overall measure. Typically, the aggregate CORAL loss is more indicative, as it captures the general alignment across all classes, thereby ensuring that the adaptation generalises well across the entire domain rather than merely fitting specific classes.

2.4. Training Process of GAT-CAMDA

The total training loss comprises weighted losses from damage classification, domain discrimination, MMD, and CORAL, expressed as follows:

$$L_{total} = \lambda_1 \times L_D + \lambda_2 \times L_T + \lambda_3 \times L_M + \lambda_4 \times L_C \quad (10)$$

where λ_1 , λ_2 , λ_3 , and λ_4 are the weights assigned to the classification, discriminator, MMD, and CORAL losses, respectively.

In this hybrid DA framework, unlike the damage classification loss (L_D) and the domain discriminator loss (L_T), which each have explicit backpropagation paths due to their association with trainable network components, MMD (L_M) and CORAL (L_C) losses act as regularisation terms applied directly to shared feature representations. Consequently, their gradients are propagated only through the feature extractor to encourage domain-invariant feature learning, rather than through separate trainable branches.

Each training iteration comprises forward and backward propagation within a standard optimisation loop. During the forward pass, data from both domains are processed through the feature extractor to generate embeddings, which are subsequently fed into the damage classifier and the domain discriminator. The GRL ensures that these embeddings do not reveal their domain origin to the discriminator, thereby promoting domain-invariant feature learning. The MMD and CORAL losses are also computed once per iteration to minimise discrepancies between the feature distributions of the two domains.

During the backward pass, the combined losses update the model parameters as defined in Equation (10). This iterative process is fully automated and follows conventional neural network training procedures, without any manual adjustment or heuristic matching between the source and target domains. The parameter γ regulates the strength of the GRL by scaling the gradients associated with the domain discriminator loss L_T . Initially, γ is small to maintain training stability and is progressively increased according to a sigmoidal schedule defined as follows:

$$\gamma = \gamma_{max} \left(\frac{2}{1 + \exp(-10p)} - 1 \right) \quad (11)$$

in which $p = \frac{\text{Epoch}}{\text{Epoch}_{max}}$, γ_{max} is the maximum value of γ , and Epoch_{max} is the total number of training epochs. This adaptive scheduling strategy enables gradual and effective TL while preserving the discriminative capacity of the extracted features for damage classification.

Designed as a semi-supervised framework, GAT-CAMDA does not use the labels from the target domain observations during the DA process. Additionally, given that there are multiple target domains, i.e., signals recorded at various temperatures, these are consolidated into a single target domain to further assess the generalisability of the framework. Algorithm 1 denotes a pseudocode for GAT-CAMDA.

2.5. Data Synthesising

As will be shown in the following section, the employed dataset contains observations for all health states at a baseline temperature (T_0 , e.g., 30 °C), whilst for the remaining temperatures only healthy measurements are available. To address the lack of data at other temperatures and for the damaged plates, a data augmentation stage was applied. The methodology extrapolates the behaviour of damaged signals at new target temperatures (T^*) by leveraging the temperature dependence observed in healthy signals.

First, the data for the healthy and damaged plates were separated based on their labels, ensuring a clear delineation of reference (healthy) and target (damaged) conditions. The average healthy response $\bar{x}(T_i)$ was computed for each available temperature T_i ; a spline interpolation function was then fitted to these mean signals, providing a continuous approximation of $\bar{x}(T)$ across the temperature range. Using this interpolation, the mean healthy signal at the target temperature $\bar{x}(T^*)$ was estimated as follows:

$$\bar{x}(T) = \text{Interp}(\{\bar{x}(T_i)\}, \{T_i\}) \quad (12)$$

Algorithm 1. GAT-CAMDA framework for SHM.

Input: Source domain data and labels, target domain data, target domain labels (held out for the validation and testing).

Output: Trained GAT-CAMDA model, feature-space alignment and sensor-importance visualisations, final classification performance on source test and target test.

1. Configuration and Setup

1.1 Set the device to GPU if available.

1.2 Initialise random seeds for reproducibility.

1.3 Define global configuration (e.g., hidden dimensions, batch size).

2. Data Preprocessing

2.1 Load source and target data.

2.2 Split source data into training, validation, and testing sets.

2.3 Split target data into training, validation, and testing sets, supporting stratification.

2.4 Convert labels to tensors and create graph-based representations of the data.

3. Model Initialisation

3.1 Define the GNN-based feature extractor using GAT.

3.2 Define the discriminator for domain classification.

3.3 Define the classifier for damage classification.

3.4 Initialise the DANN.

4. Training the Model

4.1 For each epoch:

4.1.1 Compute the adaptive weight for domain-adversarial loss.

4.1.2 For each batch of source and target data:

a. Forward pass through the feature extractor, classifier, and discriminator.

b. Compute classification loss, domain loss, MMD loss, and CORAL loss.

c. Backpropagate the combined loss and update model parameters.

4.2 Perform early stopping based on validation loss.

5. Hyperparameter Optimisation

5.1 Use Optuna for hyperparameter tuning with a defined search space.

5.2 Optimise learning rate, weight decay, loss weights, and model architecture based on validation set of target domain.

5.3 Train and evaluate the final model with the best hyperparameters.

6. Model Evaluation

6.1 Evaluate the model on the source test set for classification accuracy.

6.2 Evaluate the model on the target validation set for domain adaptation performance.

6.3 Generate confusion matrices and classification reports.

7. Feature-Space Visualisation

7.1 Extract feature embeddings using the trained model.

7.2 Apply t-SNE for dimensionality reduction.

7.3 Visualise embeddings before and after domain alignment.

8. Sensor Importance Analysis

8.1 Compute attention scores for each sensor from the GAT layers.

8.2 Normalise and visualise sensor importance scores.

9. Final Outputs

9.1 Trained GAT-CAMDA model.

9.2 Visualisations of feature alignment and sensor importance.

To derive the temperature-specific scaling factor for the damaged data, the ratio between the mean healthy signals at the target and baseline temperatures was employed:

$$S = \frac{\bar{x}(T^*)}{\bar{x}(T_0)} \quad (13)$$

This scaling factor S was then applied to the damaged signals at the baseline temperature, $x_{\text{damaged}}(T_0)$, to generate the corresponding signals at the target temperature. A small stochastic term ϵ , drawn from a Gaussian distribution with an incredibly low variance, was included to introduce realistic variability. Thus, the synthesised damaged data were obtained as follows:

$$x_{\text{damaged}}(T^*) = x_{\text{damaged}}(T_0) \cdot S \cdot (1 + \epsilon) \quad (14)$$

Following this procedure, new labels were assigned to the synthesised datasets to maintain consistency with the classification scheme.

2.6. Hyperparameter Optimisation

The optimisation process applies Optuna [44] to maximise the classification accuracy on the validation set of the target domain; this technique employs the Tree-structured Parzen Estimator (TPE) to model the probability distribution of $J(\theta)$ being maximised. This approach allows the optimiser to focus on hyperparameter regions that are more likely to enhance the classification accuracy. The TPE method effectively constructs a model of J and updates this model iteratively as new data (hyperparameter values and their corresponding A_t) are observed.

The objective function J is designed to maximise the classification accuracy of the validation subset of the target domain (A_t); this function can be formulated as follows:

$$J(\theta) = A_t(\theta) \quad (15)$$

where θ and $A_t(\theta)$ present the vector of model hyperparameters and the classification accuracy on the target test set, respectively.

The goal of the optimisation is to find the optimal set of hyperparameters θ^* that maximises J :

$$\theta^* = \underset{\theta}{\operatorname{argmax}} J(\theta) \quad (16)$$

This optimisation problem involves dynamically adjusting θ to achieve the highest possible A_t , which indicates the model's performance in effectively transferring learned knowledge from the source domain to the target.

2.7. Computing Sensor Importance

The computation of sensor importance in GAT-CAMDA directly relies on the attention coefficients (α_{ij}) defined in Section 2.2 and illustrated in Figure 2. These coefficients represent the learned weights that quantify the relative contribution of neighbouring nodes (sensors) to each node's feature aggregation. For clarity, the symbol $\zeta(\zeta_e^l)$ is subsequently used to denote the same attention weights when aggregated at the edge level in layer l , where each edge $e \in E$ connects two sensors i and j .

To determine which sensor exerts the greatest influence on the final damage-detection outcome, the attention weights from the GAT are interpreted as indicators of each sensor's contribution to the overall task. For each GAT layer, the attention weights corresponding to the edges connecting sensors are extracted and evenly distributed between the

two connected nodes. These per-sensor contributions are then aggregated to obtain a batch-wise score, normalised within each batch, and subsequently accumulated and renormalised across all batches to yield a global sensor-importance metric.

$$\begin{aligned}
 S_i^{batch} &+ = \frac{\zeta_e^i}{2} \\
 S_j^{batch} &+ = \frac{\zeta_e^j}{2} \\
 S^{batch} &\leftarrow \frac{S^{batch}}{\sum_i S_i^{batch}} \\
 S &\leftarrow \frac{S}{\sum_i S_i}
 \end{aligned} \tag{17}$$

In the following section, the calculated weights are visualised as a gradient bar chart for the three receiving sensors.

3. Case Study

A publicly available dataset, CONCEPT: CarbON-epoxy ComposiTE PlaTe, was employed to evaluate the effectiveness of the SHM framework designed in this study. The dataset includes Lamb wave measurements captured on a CFRP plate made of unidirectional plies, in both healthy and damaged states. The experiments focused on the effects of temperature variations and damage progression on the structural integrity of the plate.

Four Lead Zirconate Titanate transducers from Acellent Technologies were bonded to the plate, with one (PZT1) serving as an actuator and the others (PZT2, PZT3, and PZT4) as receiving sensors. To minimise wave propagation constraints, the laminate was mounted under free–free boundary conditions. Figure 4 shows the setup along with the test rig and instrumentation; Table 1 presents the test rig and the instruments used.

The experiments were conducted in a Thermotron thermal chamber to precisely control temperature. The chamber uses an integrated cascade refrigeration system, with the 0 °C condition achieved through mechanical cooling rather than ambient freezing. A sinusoidal tone burst served as the excitation signal, with response signals sampled accordingly. Signal generation and measurement were managed by designated data acquisition systems, controlled via LabVIEW software. For the intact plate, 100 measurements were conducted at each of seven temperatures, ranging from 0 °C to 60 °C in 10 °C increments. For the damaged scenarios, 100 observations were recorded only at 30 °C, which was considered the baseline, with no measurements for damaged cases at other temperatures.

Table 1. Summary of the test rig and instruments used in the CONCEPT experiments.

Category	Parameter	Value
Laminate plate	Dimensions (L × W × T)	500 mm × 500 mm × 2 mm
	Number of Plies	10
Transducers	Type	PZT
	Diameter	6.35 mm
	Configuration	One actuator, three sensors
Mounting conditions	Boundary condition	Free–Free
Temperature control	Range	0 °C to 60 °C
	Increment	10 °C
Excitation signal	Type	Five-cycle sinusoidal tone burst
	Frequency	250 kHz
Data sampling	Sampling rate	5 MHz
	Duration per measurement	100 ms
Data acquisition	Generation system	NI USB 6353
	Measurement system	Keysight DSO7034B
	Control software	LabVIEW




Figure 4. The test rig setup with CFRP plate and PZT transducers and temperature chamber from the CONCEPT experiments [45].

Table 2. Damage scenarios and severities for simulated defects in the CONCEPT experiment.

Damage Scenario	Severity (Area Covered)	Class Label	Description	Temperature (Degree Celsius)	Temperature Label
Healthy	0%	C0	No damage	0	0
				10	1
				20	2
				30	3
				40	4
				50	5
Damaged D1	0.196%	C1	Industrial putty	30	6
Damaged D2	0.282%	C2	Increased coverage of putty		3
Damaged D3	0.384%	C3	Further increase in coverage		
Damaged D4	0.502%	C4	Progressive increase		
Damaged D5	0.785%	C5	Larger area covered		
Damaged D6	1.13%	C6	Substantial coverage		
Damaged D7	1.53%	C7	Continued increase		
Damaged D8	1.95%	C8	Different progression pattern		
Damaged D9	2.01%	C9	Extensive coverage		
Damaged D10	2.27%	C10	High severity		
Damaged D11	2.54%	C11	Maximum simulated severity		

The damage scenarios were simulated by applying industrial adhesive putty to the composite plate's surface, representing delamination-like defects. The damage was progressively increased in a localised region between PZT1 and PZT2 to examine the attenuation and propagation changes caused by the defects. Various health scenarios along with their severities and brief descriptions are summarised in Table 2, which also contains the labels allocated to the different temperatures. This dataset was employed in multiple research studies that have represented its capability in evaluating SHM frameworks in mitigating EOVs [15,46,47].

4. Result and Discussion

This section aims to show that the proposed SHM system, i.e., GAT-CAMDA, can accurately identify different health states of the structure, even with temperature changes that occur during operation. As previously mentioned, the source domain dataset comprises measurements from all health scenarios at a baseline temperature of 30 °C, with one PZT functioning as the actuator positioned at the centre of the plate (PZT1) and the remaining three PZTs (PZT2, PZT3, and PZT4) capturing the received Lamb waves. Initially, a data synthesis approach was employed as detailed in the preceding sections. Subsequently, the performance of the end-to-end hybrid DA and classification model was evaluated under the assumption that the testing data encompassed a range of background temperatures. Finally, a comparative study was performed to assess the superiority of GAT-CAMDA. Additionally, the analysis identified which receiver sensors were more significant and which hyperparameters had a substantial impact on the results.

4.1. Dataset Complementation

To produce data for target temperatures, e.g., 0 °C, 10 °C, 20 °C, 40 °C, 50 °C, and 60 °C, the interpolation and extrapolation processes were executed for each temperature and sensor using MATLAB (R2023a)[®]. The spline interpolation technique was employed to estimate signal values at unmeasured temperatures. In synthesising the damaged scenario data for new temperatures, a scaling factor, which was empirically derived from the healthy signals across different temperatures to reflect realistic temperature-dependent attenuation trends in guided wave propagation, was applied to adjust the signals. To add realism and account for natural variations in the data, a controlled variability of $\pm 1\%$ was introduced to the scaled signals. For each of the damaged health scenarios at the target temperatures, 100 observations were generated.

To evaluate the effectiveness of the applied data synthesis, the source temperature, e.g., 30 °C, was employed. To this end, among the 100 observations belonging to the healthy situation, 50 observations were selected randomly and employed to produce synthesised data, while the remaining 50 samples were used for the validation phase. Figure 5 displays an experimental observation of the intact plate, i.e., C0 (recorded at the baseline temperature) alongside its synthesised counterpart for the three channels.

Plots in Figure 5 show that the experimental and the synthesised data are close to each other for all three sensors in the time domain. To further analyse whether the synthesised data mimic the behaviour of the experimental signals, two metrics, namely Dynamic Time Warping (DTW) [48] and Cross-Correlation (CC) [49], were employed. Data was generated for the base temperature of 30 °C, which includes observations for all health scenarios. From each health state, 50 observations were used for data synthesising, while the remaining 50 observations per class were reserved for comparison with the synthesised data; this approach ensured a balanced and consistent evaluation process.




Figure 5. Time-domain signal of the experimental and synthesised data for the healthy plate for (a) PZT2, (b) PZT3, and (c) PZT4.

DTW measured the optimal alignment between temporal sequences, with lower DTW values indicating greater similarity between the real and synthesised signals. CC assessed the linear relationship between the signals, where values close to -1 or $+1$ demonstrated higher similarity in amplitude and phase. The results of this validation process are presented in the bar graphs of Figure 6a,b, providing a visual representation of the synthesised data's fidelity to the original observations by measuring the average values of DTW and CC metrics for the three sensors.




Figure 6. Comparison of synthesised and real data across sensors (PZT2, PZT3, PZT4) at 30 °C using (a) mean DTW values and (b) mean CC values.

The acquired results demonstrate the effectiveness of the data synthesising process, as indicated by the low DTW values and CC values near 1 across all sensors. For PZT2, the DTW value is slightly higher (0.122) compared to PZT3 (0.034) and PZT4 (0.0321), which can be attributed to the larger amplitude of signals recorded by this sensor. However, the CC values for all sensors, e.g., PZT2 (0.99994), PZT3 (0.99977), and PZT4 (0.99979) are consistently close to 1, highlighting a strong similarity between the synthesised and real signals.

in terms of their temporal structure and amplitude relationships. These results validate the ability of the data synthesising process to generate realistic signals that align closely with the real observations, even in the presence of sensor-specific amplitude differences.

4.2. Damage Detection

To assess the performance of the proposed SHM framework, i.e., GAT-CAMDA, for detecting damage in composite plates subjected to EOVs, the data acquired at 30 °C was designated as the source domain. This source dataset was split into training (63%), validation (27%), and testing (10%). The remaining temperatures (0, 10, 20, 40, 50, and 60 °C) were combined to form a single dataset serving as the target domain, of which 70% was allocated to DA (with equal representation from each target temperature ensured by the ‘stratify’ option) and 15% was assigned to the validation set. The remaining 15% of the target dataset was reserved for testing; as a result, each of the 12 damage classes (C0–C11) in the target test set comprises 90 samples. Among these 90 observations, 15 belonged to 0 °C, 15 to 10 °C, and the others belonged to the remaining target temperatures. Table 3 provides a summary of the source and target datasets.

Table 3. Splitting the dataset for source and target domains.

Domain	Number of Observations per Class		
	Training	Validation	Testing
Source	63	27	10
Target	420	90	90

GAT-CAMDA was executed for 100 trials using an Optuna optimisation schedule to identify the framework’s optimal configuration, guided by 13 hyperparameters outlined in Table 4. The number of epochs was fixed at 50, and the objective function was defined as the maximum classification accuracy achieved on the target validation set.

Table 4. Hyperparameter values for executing GAT-CAMDA.

Hyperparameter	Value	Hyperparameter	Value
Learning rate	$(1 \times 10^{-4}, 1 \times 10^{-2})$	Number of GAT heads	1, 2, 4
Weight decay	$(1 \times 10^{-5}, 1 \times 10^{-3})$	Pooling option	Max, Mean, Sum
Adversarial weight	(0, 0.3)	Normalisation	True, False
MMD weight	(0, 0.3)	MMD kernel	Linear, RBF
CORAL weight	(0, 0.3)	Gamma parameter	(0.1, 10)
Hidden dimension	128, 256, 512	Number of GAT layers	4, 8
Dropout rate	(0.1, 0.5)	Batch size	32

During the DA phase, a random seed of 42 was then fixed to ensure reproducibility; the entire process of running GAT-CAMDA was performed through Jupyter Notebook in Python 3.11 (64-bit).

Executing GAT-CAMDA through the mentioned optimisation scheme, the highest damage detection accuracy was found at trial 56, as can be observed in the optimisation history plot (e.g., Figure 7).

Results from Figure 7 illustrate the extent to which choosing suitable hyperparameters can affect the outcome. For instance, at trial 20, the accuracy reached nearly 20%, while in the most optimal scenario (i.e., trial 56), it developed to 95.83%.




Figure 7. Hyperparameter optimisation history.

The slice plot in Figure 8 shows the values and options (for the categorical hyperparameters such as pooling function, normalisation approach, and the type of MMD kernel) throughout the optimisation process and, more specifically, the trial in which the most optimal objective function was discovered; the hyperparameter values for this trial are surrounded by the red rectangle.




Figure 8. Slice plot of hyperparameters during optimisation.

The prominent results presented in Figure 8 and more specifically red box indicate that implementing a normalisation process within GAT-CAMDA led to inferior outcomes compared to using the original signals during the signal processing stage. Additionally, employing lower learning rate values resulted in improved performance. Furthermore, the mean pooling function was found to be unsuitable for GAT-CAMDA, as the highest accuracy achieved in this scenario was approximately 75%.

Before representation of the damage detection results utilising the full version of GAT-CAMDA and to demonstrate the effectiveness of the feature extraction phase and the capability of the classifier, the trained GAT-CAMDA framework was evaluated using the test subset of the source domain. Moreover, the DA component of GAT-CAMDA was removed, and the trained network was subsequently tested on the target test subset. The confusion matrices presented in Figure 9a,b display the damage detection outcomes for both of these scenarios, respectively.




Figure 9. Classification outcomes through GAT-CAMDA without DA section for (a) source domain test set and (b) target test set.

The confusion matrix presented in Figure 9a demonstrates that the implemented feature extraction method can effectively differentiate between the various damage scenarios. Furthermore, these results confirm that the classifier is trained impeccably and has a high capability in the classification phase. Figure 9b, on the contrary, reveals that this well-trained network fails to accurately identify the different damage classes (approximately 36% accuracy) when the feature distribution space of the test set is altered due to EOVs. Consequently, these findings underscore the need for a DA phase to ensure that the framework trained on the source domain functions properly when applied to data from the target domain.

The outcomes of the most optimal framework in classifying the target test set are depicted in Figure 10a,b; the latter shows the classification metrics, i.e., Precision, Recall, and F1-Score [50].

The classification metrics indicate that the proposed GAT-CAMDA framework achieves a robust differentiation among the 12 structural health scenarios, maintaining high precision and recall values across most classes (with an average accuracy of 95.83%). In particular, classes exhibiting a precision of 1 reveal that the model consistently classifies those specific damage conditions without false positives, demonstrating its capacity to accurately identify even subtle differences in wave propagation patterns. Moreover, the consistently elevated recall across all classes signifies a low rate of missed damage instances, a critical requirement in damage detection tasks where undetected flaws may compromise structural integrity. The marginally reduced precision for C10 nonetheless retains a high recall, suggesting that, although the model occasionally misassigns instances into this category, it rarely overlooks genuine damage of that type.




Figure 10. Classification outcomes produced by the optimum GAT-CAMDA framework on the target test set, showing (a) the confusion matrix and (b) the classification performance metrics.

To visualise the impact of GAT-CAMDA on feature distribution, t-distributed Stochastic Neighbour Embedding (t-SNE) [51] was employed. This dimensionality reduction technique projected the extracted features into a two-dimensional space, both before and after the hybrid DA stage within GAT-CAMDA. The resulting data distributions are displayed in Figures 11a and 11b, respectively, with each scatter point representing an observation belonging to a specific temperature and damage scenario.




Figure 11. Two-dimensional t-SNE projections illustrating the distribution of extracted features (a) before DA and (b) after DA using the GAT-CAMDA framework.

When comparing the t-SNE scatter plots before and after the GAT-CAMDA-based DA, a marked improvement in feature-space alignment is evident. Specifically, observations corresponding to identical temperature-damage scenarios become more cohesively

clustered post-adaptation, illustrating that GAT-CAMDA effectively compensates for the distributional shifts introduced by varying temperature. Consequently, data points representing identical classes converge into denser, well-separated regions, underscoring the framework's ability to disentangle and preserve class-discriminative characteristics. These findings underscore GAT-CAMDA's capacity to capture subtle signal features and thereby distinguish healthy from damaged states with high fidelity, which is an essential requirement for a robust and generalised SHM system. By mitigating the confounding effects of temperature variation, GAT-CAMDA enhances inter-class separability while maintaining intra-class consistency, providing a scalable and reliable platform for advanced SHM applications.

4.3. Comparative Study

The performance of the proposed GAT-CAMDA framework was compared with multiple established feature-based DA methods. All models used an equal number of observations, 70% of the combined target domain, for the feature alignment phase. The hyperparameters of these baseline models were tuned by trial and error to determine their best performance. In total, 10 DA methods, i.e., TCA, Feature Selection with MMD (fMMD), CORAL, Deep CORAL, SrcOnly Prediction (PRED), Subspace Alignment (SA), Domain-Adversarial Training of Neural Networks (DANN), DUA (Dynamic Unsupervised Adaptation), and BAR (Balanced Adaptation Regularisation-based Transfer Learning) were evaluated using the augmented version of the CONCEPT dataset. Among these, two methods are supervised, while the remaining eight are UDA approaches.

Optuna-based optimisation was not employed for the baseline models to maintain consistency with their originally reported configurations and to avoid introducing optimisation biases unrelated to their native designs. Similarly, GAT-based feature extraction was not applied to these methods, as doing so would alter their underlying architectures and compromise the validity of the comparison. Table 5 presents the classification metrics for all approaches alongside the proposed GAT-CAMDA framework.

Table 5. Comparison of GAT-CAMDA with other feature-based methods.

Model	Accuracy (%)	Precision	F1-Score
TCA [52]	52	0.63	0.48
fMMD [53]	58	0.67	0.54
CORAL [43]	59	0.66	0.56
Deep CORAL [54]	63	0.63	0.62
PRED [55]	79	0.80	0.79
SA [56]	79	0.81	0.79
DUA [57]	82	0.81	0.79
DANN [40]	82	0.83	0.82
BAR [58]	91.02	0.92	0.92
GAT-CAMDA	95.83	0.96	0.96

The results in Table 5 show that, although it is a challenging case study because of the integration of disjoint target domains, GAT-CAMDA sharply outperformed almost all of the comparative approaches. It achieved an accuracy of 95.83%, with precision and F1-scores each at 0.96. This is further improved compared to BAR, which, despite outperforming the other methods assessed, with an accuracy of 91.02%, precision and F1-scores of 0.92 and 0.92, respectively, was not as effective as GAT-CAMDA. The better performance underlines the robustness of the procedure proposed within this work. Confusion matrices in Figure 12a to Figure 12d display the classification results of Deep CORAL, PRED, SA,

and DANN, respectively, on the target test set, providing further visual insight into the class-wise performance.




Figure 12. Damage detection result utilising (a) Deep CORAL, (b) PRED, (c) SA, and (d) DANN.

From Figure 12, it is evident that GAT-CAMDA exhibits exceptional precision in classifying the intact plate scenario (C0) and distinct damage states (C5 and C11) compared to other methods such as SA and DANN, which frequently misclassify closely related classes (e.g., C2 and C3). However, it is important to acknowledge that C6 presented considerable classification challenges for all evaluated methods, although GAT-CAMDA managed to resolve this challenge more effectively.

4.4. Hyperparameter Importance

To illustrate the relative influence of the 13 hyperparameters involved in the optimisation process on the outcomes of GAT-CAMDA, their importances were calculated using Optuna and are depicted in Figure 13. This visual representation facilitates informed decision-making by highlighting the hyperparameters with the most significant impact, allowing future researchers to prioritise these for fine-tuning to enhance the model's accuracy and efficiency in future applications.




Figure 13. The relative influence of hyperparameters on model performance in damage detection.

From Figure 13, it can be observed that the learning rate is the most critical hyperparameter, with a significant importance value of 0.59, underlining its pivotal influence on model performance. The weight of the MMD loss and the number of heads in the GAT also hold considerable importance, with values of 0.10 and 0.08, respectively, highlighting their substantial roles in shaping the model's effectiveness. Additionally, the dimension of the hidden layer, though less influential at 0.04, is crucial for fine-tuning the model's operational capabilities in damage detection. This analysis underscores the necessity of prioritising these hyperparameters during the optimisation process to maximise the model's accuracy and efficiency.

It is important to note that the hyperparameter importance scores in Figure 13 represent global aggregated effects estimated by Optuna's TPE across all 100 optimisation trials, rather than performance trends observed in individual configurations such as those shown in Figure 8. Consequently, parameters exhibiting consistent influence across multiple trials, such as the learning rate, receive higher importance values, even if certain categorical settings (e.g., normalisation = true or specific MMD kernels) occasionally produce pronounced local variations in accuracy. This explains why the learning rate dominates the global importance ranking, whereas other factors such as the MMD kernel, pooling type, and normalisation show significant but less consistent effects. The two figures are therefore complementary: Figure 8 visualises trial-level sensitivity, while Figure 13 summarises overall influence across the entire optimisation landscape.

4.5. Sensor Importance

In the GAT-CAMDA framework, the attention scores for each sensor were normalised to effectively evaluate their significance in the damage detection process. These scores were determined by the extent to which each sensor contributed to identifying structural anomalies, with all calculations having been performed under the optimal configuration settings of the framework. The normalised scores ranged from 0 to 1, providing a consistent and comparative metric of sensor importance. Figure 14 displays these scores, visualised as a gradient bar chart, which illustrates the relative contribution of each sensor to the overall damage detection capability.




Figure 14. Sensor contribution levels determined by attention mechanism for the normalised signals.

As previously noted, the most optimal framework identified was for scenarios in which the signals were not normalised. Figure 5 demonstrates that the amplitude of the signals recorded by PZT2 is approximately eight times greater than those from the other two receivers, PZT3 and PZT4. Furthermore, the signal amplitude captured by PZT4 is marginally higher than that of PZT3. These differences in signal amplitude play a pivotal role in determining the final outcome, as evidenced by the calculated sensor importance scores, whereby PZT2, PZT4, and PZT3 exhibit progressively higher importance in the final output, respectively.

The same optimal framework was assessed under the assumption that normalisation was not applied, while maintaining the other hyperparameter values unchanged; the computed sensor importances are displayed in Figure 15.




Figure 15. Sensor contribution levels determined by attention mechanism for unnormalised signals.

The pronounced significance of PZT3 within the designed SHM framework on the assumption that the unnormalised versions of the data were fed into the signal processing and learning network is consistent with the findings in [47]. It was confirmed that PZT3

exhibits a wider confidence interval in its signal data compared to PZT2, indicating greater variability. This increased variability is influenced by temperature changes, which affect the material's mechanical properties. Specifically, as temperature rises, the stiffness of the CFRP laminate decreases, with the shear modulus G_{12} reducing by 7%. This decline is much higher than the 1.3% decrease observed in the elastic modulus E_1 . This difference is crucial because PZT3 is positioned at a 45-degree angle from the excitation point, making it particularly sensitive to changes in the shear modulus due to its orientation relative to the main direction of wave propagation. The heightened signal dispersion at PZT3 results from its unique placement and the viscoelastic characteristics of the epoxy resin, which are not present in the carbon fibres [47]. Consequently, despite a lower signal-to-noise ratio, PZT3 shows an enhanced ability to detect subtle anomalies in wave propagation that may indicate structural damage. Figure 15 reflects the intricate relationship between sensor placement, material properties, and wave behaviour in composite structures.

5. Conclusions and Future Work

The GAT-CAMDA framework was designed to address the challenge of variability in SHM caused by temperature fluctuations, with its effectiveness demonstrated on a CFRP plate. Leveraging GATs for advanced feature extraction and incorporating MMD, CORAL, and adversarial losses for discriminative DA, the framework successfully aligned feature distributions across diverse temperature domains. This alignment significantly improved the accuracy of damage detection within the system. Comparative studies with established methods further validated the framework's superior performance, highlighting its ability to handle the complexities of temperature-integrated target domain datasets.

To address the challenge of data scarcity, particularly for unmeasured damage conditions at different temperatures, a synthetic data generation algorithm was employed. This algorithm enriched the dataset by creating synthetic samples, helping the validation of the framework's generalisability and robustness. As a result, the framework was extensively evaluated and refined for conditions that were not physically measured, ensuring consistent performance across a theoretical range of operational scenarios.

Explainability techniques played a crucial role in clarifying the decision-making processes of GAT-CAMDA. These techniques visualised the contributions of various sensors, identifying the most influential one for damage detection. Furthermore, the study displayed the impact of key hyperparameters on damage detection accuracy, providing valuable insights into optimal model configurations.

Although the GAT-CAMDA framework effectively addresses variability in SHM caused by temperature changes, further improvements are needed to enhance its computational efficiency for real-time deployment. This includes optimising graph processing algorithms and integrating hardware accelerations, such as Tensor Processing Units, to reduce processing times. Additionally, increasing the realism and diversity of synthetic data through advanced generative models could further strengthen the framework's practical utility. As the current experiments were conducted using synthetic data, future work should include comprehensive validation on real-world experimental datasets to confirm the framework's robustness and generalisability under practical operating conditions. Extending the framework's adaptability to include a wider range of environmental conditions beyond temperature, such as humidity or load variations, could broaden its applicability across different industrial settings.

Author Contributions: Conceptualisation, N.R. and F.A.; methodology, N.R., D.P., A.D.L., and M.D.O.; investigation, N.R. and A.D.L.; writing—original draft, preparation, review, and editing, N.R., M.D.O., D.P., F.A., and G.L.; supervision, A.D.L. and M.D.O. All authors have read and agreed to the published version of the manuscript.

Funding: This research study was carried out in the framework of the project “TU-LEARN—sTrUctural Life Extension enhAnced by aRtificial iNtelligence”, funded by Unione Europea—Next Generation EU, as part of the Prin 2022 PNRR—D.D. n. 1409 del 14-09-2022 programme.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The employed dataset is available upon request via https://github.com/shm-unesp/DATASET_PLATEUN01?tab=readme-ov-file (accessed on 20 December 2024).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BAR	Balanced Adaptation Regularisation-based Transfer Learning
CC	Cross-Correlation
CFRP	Carbon Fibre-Reinforced Polymer
CORAL	Correlation Alignment
DANN	Domain-Adversarial Training of Neural Networks
DA	Domain Adaptation
DTW	Dynamic Time Warping
DUA	Dynamic Unsupervised Adaptation
EOVs	Environmental and Operational Variabilities
FE	Finite Element
fMMD	Feature Selection with MMD
GATs	Graph Attention Networks
GRL	Gradient Reversal Layer
ML	Machine Learning
MMD	Maximum Mean Discrepancy
PRED	SrcOnly Prediction
PZT	Lead Zirconate Titanate
RBF	Radial Basis Function
ReLU	Rectified Linear Unit
SA	Subspace Alignment
SHM	Structural Health Monitoring
TCA	Transfer Component Analysis
TL	Transfer learning
t-SNE	t-distributed Stochastic Neighbour Embedding
TPE	Tree-structured Parzen Estimator
UDA	Unsupervised Domain Adaptation

References

1. Cawley, P. Guided Waves in Long Range Nondestructive Testing and Structural Health Monitoring: Principles, History of Applications and Prospects. *NDT E Int.* **2024**, *142*, 103026. [CrossRef]
2. Gorgin, R.; Luo, Y.; Wu, Z. Environmental and Operational Conditions Effects on Lamb Wave Based Structural Health Monitoring Systems: A Review. *Ultrasonics* **2020**, *105*, 106114. [CrossRef] [PubMed]
3. De Luca, A.; Greco, A.; Rezazadeh, N.; Perfetto, D.; Aversano, A. Guided Waves Propagation in Additively Manufactured GF30-PA6 Panel. *Macromol. Symp.* **2024**, *413*, 2400042. [CrossRef]
4. Giurgiutiu, V. Structural Health Monitoring with Piezoelectric Wafer Active Sensors—Predictive Modeling and Simulation. *INCAS Bull.* **2010**, *2*, 31–44. [CrossRef]
5. De Luca, A.; Perfetto, D.; Polverino, A.; Aversano, A.; Caputo, F. Finite Element Modeling Approaches, Experimentally Assessed, for the Simulation of Guided Wave Propagation in Composites. *Sustainability* **2022**, *14*, 6924. [CrossRef]
6. Zheng, S.; Luo, Y.; Xu, C.; Xu, G. A Review of Laser Ultrasonic Lamb Wave Damage Detection Methods for Thin-Walled Structures. *Sensors* **2023**, *23*, 3183. [CrossRef]

7. Yin, Z.; Tie, Y.; Duan, Y.; Li, C. Optimization of Nonlinear Lamb Wave Detection System Parameters in Cfrp Laminates. *Materials* **2021**, *14*, 3186. [[CrossRef](#)]
8. Rezazadeh, N.; Perfetto, D.; Polverino, A.; De Luca, A.; Lamanna, G. Guided Wave-Driven Machine Learning for Damage Classification with Limited Dataset in Aluminum Panel. *Struct. Health Monit.* **2024**, *24*, 3138–3161. [[CrossRef](#)]
9. Zeng, X.; Liu, X.; Yan, J.; Yu, Y.; Zhao, B.; Qing, X. Lamb Wave-Based Damage Localization and Quantification Algorithms for CFRP Composite Structures. *Compos. Struct.* **2022**, *295*, 115849. [[CrossRef](#)]
10. Zou, L.; Gong, Y.; Tian, D.; Zhao, L.; Zhang, J.; Hu, N. Effect of Interface Angle on Mode I Delamination Damage Behavior of Multidirectional Fully Isotropic Laminates with the Same Global Stiffness. *Thin-Walled Struct.* **2023**, *182*, 110211. [[CrossRef](#)]
11. Konstantinidis, G.; Wilcox, P.D.; Drinkwater, B.W. An Investigation into the Temperature Stability of a Guided Wave Structural Health Monitoring System Using Permanently Attached Sensors. *IEEE Sens. J.* **2007**, *7*, 905–912. [[CrossRef](#)]
12. Moll, J.; Kexel, C.; Pötzsch, S.; Rennoch, M.; Herrmann, A.S. Temperature Affected Guided Wave Propagation in a Composite Plate Complementing the Open Guided Waves Platform. *Sci. Data* **2019**, *6*, 191. [[CrossRef](#)]
13. Rezazadeh, N.; De Luca, A.; Perfetto, D.; Salami, M.R.; Lamanna, G. Systematic Critical Review of Structural Health Monitoring under Environmental and Operational Variability: Approaches for Baseline Compensation, Adaptation, and Reference-Free Techniques. *Smart Mater. Struct.* **2025**, *34*, 073001. [[CrossRef](#)]
14. Scheerer, M.; Lager, D. Validation of Temperature Compensation Techniques for Impact Damage Detection and Localization Using Ultrasonic Sparse Arrays. In Proceedings of the 7th European Workshop on Structural Health Monitoring, EWSHM 2014—2nd European Conference of the Prognostics and Health Management (PHM) Society, Nantes, France, 8–11 July 2014.
15. da Silva, S.; Paixão, J.; Rébillat, M.; Mechbal, N. Extrapolation of AR Models Using Cubic Splines for Damage Progression Evaluation in Composite Structures. *J. Intell. Mater. Syst. Struct.* **2021**, *32*, 284–295. [[CrossRef](#)]
16. Ren, F.; Giannakeas, I.N.; Sharif Khodaei, Z.; Aliabadi, M.H.F. Sensitivity Analysis of Temperature Effects on Guided Wave-Based Damage Detection. *Mech. Syst. Signal Process* **2023**, *196*, 110322. [[CrossRef](#)]
17. Ren, F.; Giannakeas, I.N.; Sharif Khodaei, Z.; Aliabadi, M.H.F. Theoretical and Experimental Investigation of Guided Wave Temperature Compensation for Composite Structures with Different Thicknesses. *Mech. Syst. Signal Process* **2023**, *200*, 110594. [[CrossRef](#)]
18. Perfetto, D.; Sharif Khodaei, Z.; De Luca, A.; Aliabadi, M.H.; Caputo, F. Experiments and Modelling of Ultrasonic Waves in Composite Plates under Varying Temperature. *Ultrasonics* **2022**, *126*, 106820. [[CrossRef](#)]
19. Rezazadeh, N.; Annaz, F.; Jabbar, W.A.; Vieira Filho, J.; De Oliveira, M. A Transfer Learning Approach for Mitigating Temperature Effects on Wind Turbine Blades Damage Diagnosis. *Struct. Health Monit.* **2025**. [[CrossRef](#)]
20. de Ferreira, L.P.S.; de Teloli, R.O.; da Silva, S.; Figueiredo, E.; Maia, N.; Cimini, C.A. Bayesian Data-Driven Framework for Structural Health Monitoring of Composite Structures under Limited Experimental Data. *Struct. Health Monit.* **2024**, *24*, 738–760. [[CrossRef](#)]
21. Giannakeas, I.N.; Sharif Khodaei, Z.; Aliabadi, M.H. An Up-Scaling Temperature Compensation Framework for Guided Wave-Based Structural Health Monitoring in Large Composite Structures. *Struct. Health Monit.* **2023**, *22*, 777–798. [[CrossRef](#)]
22. Cheng, X.; Sharif Khodaei, Z. Temperature Calibration for Guided Wave Hybrid System with FBG-Based Fabry-Perot Sensor. *J. Intell. Mater. Syst. Struct.* **2023**, *34*, 766–784. [[CrossRef](#)]
23. Abdel-Basset, M.; Moustafa, N.; Hawash, H.; Tari, Z. *Responsible Graph Neural Networks*; Chapman and Hall/CRC: New York, NY, USA, 2023.
24. Vrahatis, A.G.; Lazaros, K.; Kotsiantis, S. Graph Attention Networks: A Comprehensive Review of Methods and Applications. *Future Internet* **2024**, *16*, 318. [[CrossRef](#)]
25. Niu, J.; Li, S.; Li, Z. Restoration of Missing Structural Health Monitoring Data Using Spatiotemporal Graph Attention Networks. *Struct. Health Monit.* **2022**, *21*, 2408–2419. [[CrossRef](#)]
26. Zhao, M.; Taal, C.; Baggerohr, S.; Fink, O. Virtual Sensor for Real-Time Bearing Load Prediction Using Heterogeneous Temporal Graph Neural Networks. *PHM Soc. Eur. Conf.* **2024**, *8*, 8. [[CrossRef](#)]
27. Rezazadeh, N.; Perfetto, D.; de Oliveira, M.; De Luca, A.; Lamanna, G. A Fine-Tuning Deep Learning Framework to Palliate Data Distribution Shift Effects in Rotary Machine Fault Detection. *Struct. Health Monit.* **2024**. [[CrossRef](#)]
28. Ozdagli, A.; Koutsoukos, X. Domain Adaptation for Structural Health Monitoring. *Annu. Conf. PHM Soc.* **2020**, *12*, 9. [[CrossRef](#)]
29. Zhuojun, X.; Hao, L.; Jianbo, Y.; Jingwen, Y. A Transfer Learning Approach for Data-Driven Localization of Damage Areas in Plate-like Structures of CFRP Materials. *Eng. Struct.* **2024**, *314*, 118352. [[CrossRef](#)]
30. Wang, Y.; Cui, X.; Liu, Q.; Zhao, B.; Liao, Y.; Qing, X. An Adaptive Damage Monitoring Method Based on Transfer Features Mapped for Advanced Composite Structures. *Compos. Struct.* **2024**, *329*, 117742. [[CrossRef](#)]
31. Poole, J.; Gardner, P.; Dervilis, N.; Bull, L.; Worden, K. On Statistic Alignment for Domain Adaptation in Structural Health Monitoring. *Struct. Health Monit.* **2023**, *22*, 1581–1600. [[CrossRef](#)]

32. Lundberg, S.M.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017.
33. Ribeiro, M.T.; Singh, S.; Guestrin, C. “Why Should i Trust You?” Explaining the Predictions of Any Classifier. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016.
34. Schnur, C.; Goodarzi, P.; Lugovtsova, Y.; Bulling, J.; Prager, J.; Tschöke, K.; Moll, J.; Schütze, A.; Schneider, T. Towards Interpretable Machine Learning for Automated Damage Detection Based on Ultrasonic Guided Waves. *Sensors* **2022**, *22*, 406. [\[CrossRef\]](#) [\[PubMed\]](#)
35. Salih, A.M.; Raisi-Estabragh, Z.; Galazzo, I.B.; Radeva, P.; Petersen, S.E.; Lekadir, K.; Menegaz, G. A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME. *Adv. Intell. Syst.* **2024**, *7*, 2400304. [\[CrossRef\]](#)
36. Tiwari, A.; Gupta, A.K.; Gupta, T. A Robust Approach to Shear Strength Prediction of Reinforced Concrete Deep Beams Using Ensemble Learning with SHAP Interpretability. *Soft Comput.* **2024**, *28*, 6343–6365. [\[CrossRef\]](#)
37. Azad, M.M.; Kim, H.S. An Explainable Artificial Intelligence-based Approach for Reliable Damage Detection in Polymer Composite Structures Using Deep Learning. *Polym. Compos.* **2024**, *46*, 1536–1551. [\[CrossRef\]](#)
38. Li, X.; Xu, L.; Guo, H.; Yang, L. Application of Graph Convolutional Neural Networks Combined with Single-Model Decision-Making Fusion Neural Networks in Structural Damage Recognition. *Sensors* **2023**, *23*, 9327. [\[CrossRef\]](#)
39. Veličković, P.; Casanova, A.; Liò, P.; Cucurull, G.; Romero, A.; Bengio, Y. Graph Attention Networks. In Proceedings of the 6th International Conference on Learning Representations, ICLR 2018—Conference Track Proceedings, Vancouver, BC, Canada, 30 April–3 May 2018.
40. Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; Lempitsky, V. Domain-Adversarial Training of Neural Networks. *J. Mach. Learn. Res.* **2016**, *17*, 2096–3030.
41. Huang, J.; Guan, D.; Xiao, A.; Lu, S. RDA: Robust Domain Adaptation via Fourier Adversarial Attacking. In Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 10–17 October 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 8968–8979.
42. Kim, B.; Pineau, J. Maximum Mean Discrepancy Imitation Learning. In Proceedings of the Robotics: Science and systems, Ann Arbor, MI, USA, 18–22 June 2016.
43. Sun, B.; Feng, J.; Saenko, K. Return of Frustratingly Easy Domain Adaptation. In Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI 2016, Phoenix, AZ, USA, 12–17 February 2016.
44. Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; Koyama, M. Optuna: A Next-Generation Hyperparameter Optimization Framework. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Anchorage, AK, USA, 4–8 August 2019.
45. da Silva, S. Data-Driven Model Identification of Guided Wave Propagation in Composite Structures. *J. Braz. Soc. Mech. Sci. Eng.* **2018**, *40*, 543. [\[CrossRef\]](#)
46. de Almeida, E.F.; Chavarette, F.R.; Merizio, I.F.; Gonçalves, A.C. Artificial Immune System for Fault Detection and Localization in a Composite Material Plate with Temperature Variation. *J. Braz. Soc. Mech. Sci. Eng.* **2024**, *46*, 717. [\[CrossRef\]](#)
47. de Ferreira, L.P.S.; de Teloli, R.O.; da Silva, S.; Figueiredo, E.; Moldovan, I.D.; Maia, N.; Cimini, C.A. Bayesian Calibration for Lamb Wave Propagation on a Composite Plate Using a Machine Learning Surrogate Model. *Mech. Syst. Signal Process* **2024**, *208*, 111011. [\[CrossRef\]](#)
48. Senin, P. Dynamic Time Warping Algorithm Review. *Science* **2008**, *855*, 40.
49. Habermehl, S.; Schlesinger, C.; Prill, D. Comparison and Evaluation of Pair Distribution Functions, Using a Similarity Measure Based on Cross-Correlation Functions. *J. Appl. Crystallogr.* **2021**, *54*, 612–623. [\[CrossRef\]](#)
50. Rezazadeh, N.; De Oliveira, M.; Lamanna, G.; Perfetto, D.; De Luca, A. WaveCORAL-DCCA: A Scalable Solution for Rotor Fault Diagnosis Across Operational Variabilities. *Electronics* **2025**, *14*, 3146. [\[CrossRef\]](#)
51. van der Maaten, L.; Hinton, G. Visualizing Data Using T-SNE. *J. Mach. Learn. Res.* **2008**, *9*, 2579–2605.
52. Pan, S.J.; Tsang, I.W.; Kwok, J.T.; Yang, Q. Domain Adaptation via Transfer Component Analysis. *IEEE Trans. Neural Netw.* **2011**, *22*, 199–210. [\[CrossRef\]](#)
53. Uguroglu, S.; Carbonell, J. *Feature Selection for Transfer Learning*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 430–442.
54. Sun, B.; Saenko, K. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In Proceedings of the Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Amsterdam, The Netherlands, 11–14 October 2016; Volume 9915. LNCS.
55. Daumé, H., III. Frustratingly Easy Domain Adaptation. In Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, Prague, Czech Republic, 23–30 June 2007.
56. Fernando, B.; Habrard, A.; Sebban, M.; Tuytelaars, T. Unsupervised Visual Domain Adaptation Using Subspace Alignment. In Proceedings of the Proceedings of the IEEE International Conference on Computer Vision, Sydney, Australia, 1–8 December 2013.

57. Mirza, M.J.; Micorek, J.; Possegger, H.; Bischof, H. The Norm Must Go On: Dynamic Unsupervised Domain Adaptation by Normalization. In Proceedings of the Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022.
58. Hu, Q.; Si, X.; Qin, A.; Lv, Y.; Liu, M. Balanced Adaptation Regularization Based Transfer Learning for Unsupervised Cross-Domain Fault Diagnosis. *IEEE Sens. J.* **2022**, *22*, 12139–12151. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.