

Damage detection in composite structures utilising machine learning and cross- domain adaptation



By

Nima Rezazadeh

Submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Supervised by

Dr. Mario de Oliveira

Dr. Fawaz Annaz

School of Architecture, Built Environment, Computing and Engineering

Birmingham City University

2026

Keywords

Capsule networks, Composite materials, Damage detection, Data augmentation, Environmental variability, Graph attention networks, Prototype-based learning, Structural health monitoring, Unsupervised domain adaptation, Wavelet transform scattering

Abstract

Structural health monitoring of composite structures in realistic industrial settings is constrained by 3 persistent difficulties. First, labelled data describing damaged conditions are inherently scarce because deliberate damage is costly, time-consuming and often impractical. Second, sensor responses are strongly influenced by environmental and operational variability, such as changes in temperature and loading, which alter the measured signals even when the structural state is unchanged. Third, models that are trained under a particular set of conditions often fail to generalise when deployed under different conditions, so that performance degrades precisely when reliability is most needed. This thesis addresses these challenges by developing and evaluating a coherent set of machine learning frameworks that jointly tackle data scarcity, shifting operating conditions and the need for interpretable, computationally practical diagnosis of damage in composite laminates and a small-scale wind turbine blade.

The work begins by strengthening the data foundation through a suite of augmentation and dataset completion procedures tailored to structural health monitoring signals. Windowing is used to segment long records into multiple shorter samples in a manner that preserves temporal coherence. Spline-based interpolation across temperature is introduced to generate physically plausible signal variations between a limited number of measured operating points. A convolutional conditional variational autoencoder is designed to learn a latent representation of the signals that is explicitly conditioned on both damage state and temperature, and this representation is then used to synthesise new signals under unmeasured conditions whilst preserving damage-sensitive patterns in the time, frequency and time–frequency domains. In the feature space, Synthetic Minority Over-sampling Technique is employed to rebalance classes when damaged conditions are underrepresented. These components are used selectively to expand the effective training set without distorting the underlying physics of wave propagation and vibration.

The novel contribution of this thesis lies in the development of three complementary damage detection frameworks, each tailored to different regimes of domain shift and supervision availability, combined with interpretability and computational tractability as primary design constraints. The first framework is designed for mild shifts with a small number of labelled examples in the target condition. It introduces a novel integration of manifold learning based on uniform manifold approximation and projection with a capsule-based neural feature extractor and an instance reweighting strategy that downweights misleading source examples whilst upweighting informative ones, followed by an ensemble classifier. The second framework addresses multi-sensor, semi-supervised adaptation. It proposes a novel architecture that employs a graph attention-based encoder to model the spatial relationships between sensors and actuator–sensor paths, and aligns source and target feature distributions by combining adversarial training with statistical alignment of their means and covariances. The third framework is intended for severe domain divergence without any target labels. It presents a novel prototype-guided approach that uses wavelet time scattering to obtain deformation-stable and translation-invariant features, and then performs prototype-guided adversarial domain adaptation in which class prototypes are iteratively refined, target samples are pseudo-labelled under confidence

and similarity constraints, and alignment is encouraged around these prototypes rather than around raw features, which helps to maintain class separation. Interpretability and computational efficiency are treated as primary design requirements rather than afterthoughts. The proposed pipelines provide diagnostic views that include trajectories of samples in the learned manifolds, capsule-based feature activations associated with specific damage states, attention weights that reveal the relative contribution of each sensor, and the evolution of class prototypes during adaptation. All methods are implemented on central processing unit workstations with constrained hyperparameter searches in order to reflect realistic industrial computation budgets. Across two heterogeneous case studies, namely a small-scale glass fibre-reinforced epoxy wind turbine blade tested under controlled temperature variation and a carbon-epoxy composite plate monitored by guided waves for multiple damage severities, the proposed frameworks consistently outperform strong baseline methods in their respective regimes of shift and supervision. In conclusion, this thesis demonstrates that tailored domain adaptation strategies, when combined with physically informed augmentation and explicit interpretability mechanisms, can substantially improve the reliability and deployability of machine learning-based damage detection systems for composite structures under realistic operational variability. Taken together, the augmentation procedures and three domain adaptation frameworks form a practically deployable toolbox that mitigates data scarcity, accommodates changing environmental and operational conditions and delivers auditable damage detection for composite structures.

Statement of Original Authorship

The work contained in this thesis has not been previously submitted to meet requirements for an award at this or any other higher education institution. To the best of my knowledge and belief, the thesis contains no material previously published or written by another person except where due reference is made.

Signature: *Nima Rezaadeh*

Acknowledgements

The research presented in this thesis was funded by Birmingham City University and supported by many individuals, to whom I am deeply grateful.

First and foremost, I would like to express my sincere gratitude to my supervisory team: my Director of Studies, Dr. Mario De Oliveira, and my supervisor, Dr. Fawaz Annaz. Their invaluable insights, detailed feedback, and constant support throughout my PhD journey were instrumental in the completion of this thesis. Their profound knowledge and guidance have significantly enriched both my research and my development as a researcher.

I would also like to extend my special thanks to Dr. Roger Tait for his expert advice on utilising the High-Performance Computing resources. I am grateful to the staff in the Doctoral Research College (DRC), the School of Architecture, Built Environment, Computing and Engineering, the technicians, and the IT help desk at Birmingham City University for their timely assistance, essential resources, and administrative support during my research.

I am profoundly thankful to my family for their endless patience, understanding, and love. Their unwavering support has been the cornerstone of my endurance and success; words cannot fully convey the depth of my gratitude for having them in my life.

I would also like to thank my colleagues and fellow PhD students for making my time at the university both delightful and memorable.

Table of Contents

| | |
|---|-----------|
| Keywords | i |
| Abstract | ii |
| Statement of Original Authorship | iv |
| Acknowledgements | v |
| List of Figures | x |
| List of Algorithms | xiv |
| Nomenclature | xv |
| Chapter 1: Introduction | 1 |
| 1.1 Background | 1 |
| 1.2 Research Aim | 7 |
| 1.3 Research Questions | 8 |
| 1.4 Research Objectives | 8 |
| 1.5 Structure of the Thesis | 9 |
| 1.6 Publications | 11 |
| Chapter 2: Literature Review | 13 |
| 2.1 Overview | 13 |
| 2.2 Signal Structure and Invariance Requirements | 14 |
| 2.3 Feature Extraction in SHM | 14 |
| 2.3.1 Classical Descriptors in Time, Frequency and Time–Frequency Domains | 14 |
| 2.3.2 Wavelet Time Scattering | 15 |
| 2.3.3 Manifold and Embedding Methods | 15 |
| 2.3.4 Neural Feature Learners for Signals | 16 |
| 2.4 Data Augmentation and Dataset Completion for Time Series | 19 |
| 2.4.1 Classical Time–Domain Augmentations | 19 |
| 2.4.2 Frequency and Time–Frequency Domain Augmentations | 19 |
| 2.4.3 Minority Class Balancing with SMOTE Variants | 20 |
| 2.4.4 Cross–Condition Interpolation | 20 |
| 2.4.5 Generative Models for Synthetic Data | 21 |
| 2.4.6 Fidelity and Validation of Synthetic Data | 21 |
| 2.5 EOV and Dataset Shift | 22 |
| 2.5.1 Direct Signal-Level Baseline Compensation | 22 |
| 2.5.1.1 Subtraction and Phase or Time Alignment | 23 |
| 2.5.1.2 Adaptive Filtering | 24 |
| 2.5.1.3 Real-Time and Simulation-Aided Corrections | 24 |
| 2.5.2 Adaptive and Multi-Baseline Strategies | 25 |
| 2.5.2.1 Evolving Baseline Libraries | 25 |
| 2.5.2.2 Virtual Baseline Synthesis | 26 |
| 2.5.2.3 Regression and Latent Space Models | 26 |
| 2.5.3 Reference-Free Approaches | 27 |
| 2.5.3.1 Purely Data Driven Baseline | 28 |
| 2.5.3.2 Physics-Based Invariants and Instantaneous References | 28 |

| | | |
|-------------------|---|-----------|
| 2.5.3.3 | Physics-Informed Learning Hybrids | 29 |
| 2.5.3.4 | Relative and Instantaneous Referencing and Building Block Validation..... | 30 |
| 2.5.3.5 | Transfer Learning and Domain Adaptation for Mechanical Assets | 30 |
| 2.6 | Explainable Artificial Intelligence..... | 32 |
| 2.7 | Summary..... | 35 |
| Chapter 3: | Methodology and Methods | 37 |
| 3.1 | Overview | 37 |
| 3.2 | Data Augmentation..... | 38 |
| 3.2.1 | Spline Interpolation | 39 |
| 3.2.2 | Convolutional Conditional Variational Autoencoder..... | 40 |
| 3.3 | Damage Detection Pipelines..... | 41 |
| 3.3.1 | UCTRF Framework..... | 41 |
| 3.3.1.1 | Main Contribution..... | 42 |
| 3.3.1.2 | Uniform Manifold Approximation and Projection..... | 43 |
| 3.3.1.3 | Capsule Neural Networks..... | 46 |
| 3.3.1.4 | Transfer Adaptive Boosting | 48 |
| 3.3.1.5 | Random Forest | 49 |
| 3.3.1.6 | Weight Adjustments Representation | 50 |
| 3.3.2 | GAT-CAMDA Framework | 53 |
| 3.3.2.1 | Main Contribution..... | 54 |
| 3.3.2.2 | Graph Attention Networks | 55 |
| 3.3.2.3 | Domain-Adversarial Neural Networks in GAT-CAMDA | 57 |
| 3.3.2.3.1 | MULTI-LAYER PERCEPTRON..... | 57 |
| 3.3.2.3.2 | CLASSIFICATION LOSS | 58 |
| 3.3.2.3.3 | DOMAIN DISCRIMINATOR LOSS | 59 |
| 3.3.2.4 | Maximum Mean Discrepancy Loss..... | 59 |
| 3.3.2.5 | Correlation Alignment Loss | 60 |
| 3.3.2.6 | Joint Optimisation | 60 |
| 3.3.2.7 | Optuna-Based Hyperparameter Optimisation | 61 |
| 3.3.2.8 | Computing Sensor Importance..... | 61 |
| 3.3.3 | SPADA Framework | 63 |
| 3.3.3.1 | Main Contribution..... | 64 |
| 3.3.3.2 | Wavelet Time Scattering..... | 65 |
| 3.3.3.3 | Domain Adaptation in SPADA | 66 |
| 3.3.3.3.1 | PROTOTYPE CONSTRUCTION AND ATTENTION-BASED UPDATING..... | 67 |
| 3.3.3.3.2 | TRAINING OBJECTIVES | 68 |
| 3.3.3.3.3 | OVERALL OBJECTIVE FUNCTION..... | 69 |
| 3.3.3.3.4 | SUMMARY OF THE TRAINING PROCESS | 69 |
| 3.3.3.4 | Interpretability of SPADA Internal Activity | 71 |
| 3.3.3.4.1 | TRACKING PROTOTYPE TRAJECTORY EVOLUTION | 73 |
| 3.3.3.4.2 | COMPUTING INSTANCE-PROTOTYPE COSINE SIMILARITY | 74 |
| 3.3.3.4.3 | MONITORING PROTOTYPE ATTENTION WEIGHT DYNAMICS | 75 |
| 3.3.3.4.4 | VISUALISING DECISION BOUNDARIES IN FEATURE SPACE | 76 |
| 3.3.4 | Framework Selection Criteria and Application Contexts..... | 79 |

| | |
|---|-----------|
| Chapter 4: Case Studies | 81 |
| 4.1 Overview | 81 |
| 4.2 Small-scale Wind Turbine Blade Under Varying Climate Conditions (WTB-VibClimate)..... | 81 |
| 4.3 Carbon-Epoxy Composite PlaTe (CONCEPT) | 84 |
| Chapter 5: Results and Discussion | 89 |
| 5.1 Overview | 89 |
| 5.2 Effects of EOV | 89 |
| 5.2.1 WTB-VibClimate | 90 |
| 5.2.2 CONCEPT..... | 92 |
| 5.3 Data Augmentation..... | 93 |
| 5.3.1 SMOTE | 94 |
| 5.3.2 Signal Windowing..... | 96 |
| 5.3.3 Spline Interpolation | 97 |
| 5.3.4 Conv-CVAE | 99 |
| 5.4 Damage Detection | 102 |
| 5.4.1 UCTRF | 102 |
| 5.4.1.1 Reducing the Dimensionality | 102 |
| 5.4.1.2 Feature Extraction through CapsNets..... | 104 |
| 5.4.1.3 Damage Detection without DA | 105 |
| 5.4.1.4 Damage Detection with DA | 106 |
| 5.4.1.5 Comparison Study..... | 109 |
| 5.4.1.5.1 CLASSIFICATION USING DANN..... | 110 |
| 5.4.1.5.2 CLASSIFICATION USING TCA | 110 |
| 5.4.1.6 Weight Adjustments in Mitigating EOVs | 112 |
| 5.4.1.7 Practical Considerations and Applicability | 115 |
| 5.4.2 GAT-CAMDA..... | 115 |
| 5.4.2.1 Damage Detection without DA | 116 |
| 5.4.2.2 Damage Detection with DA | 117 |
| 5.4.2.3 Comparative Study..... | 121 |
| 5.4.2.4 Hyperparameter Importance..... | 124 |
| 5.4.2.5 Sensor Importance..... | 125 |
| 5.4.3 SPADA..... | 127 |
| 5.4.3.1 WTB-VibClimate..... | 127 |
| 5.4.3.1.1 FEATURE EXTRACTION | 127 |
| 5.4.3.1.2 DAMAGE DETECTION WITHOUT DA..... | 131 |
| 5.4.3.1.3 DAMAGE DETECTION WITH DA | 133 |
| 5.4.3.2 CONCEPT | 137 |
| 5.4.3.2.1 FEATURE EXTRACTION | 138 |
| 5.4.3.2.2 DAMAGE DETECTION WITHOUT DA..... | 141 |
| 5.4.3.2.3 DAMAGE DETECTION WITH DA | 142 |
| 5.4.3.3 Internal Mechanism Visualisation and Interpretability Analysis | 145 |
| 5.4.3.3.1 PROTOTYPE TRAJECTORY EVOLUTION..... | 145 |

| | | |
|---------------------|--|------------|
| 5.4.3.3.2 | INSTANCE-PROTOTYPE COSINE SIMILARITY | 147 |
| 5.4.3.3.3 | PROTOTYPE ATTENTION WEIGHT DYNAMICS | 149 |
| 5.4.3.3.4 | DECISION BOUNDARIES IN FEATURE SPACE..... | 151 |
| | 5.4.3.4 Comparison Study..... | 152 |
| | 5.4.3.5 Computation Efficiency | 154 |
| Chapter 6: | Conclusions and Future Work | 156 |
| | Summary of Research and Key Findings..... | 156 |
| 6.1 | Contributions to Knowledge..... | 157 |
| 6.2 | Addressing the Research Questions..... | 159 |
| 6.3 | Limitations..... | 161 |
| 6.4 | Future Work..... | 162 |
| 6.5 | Concluding Remarks | 164 |
| Bibliography | | 167 |

List of Figures

| | |
|--|----|
| Figure 1. Literature review structure and topic coverage..... | 13 |
| Figure 2. A schematic of UCTRF for damage detection..... | 42 |
| Figure 3. A schematic of CapsNets as the feature extraction..... | 47 |
| Figure 4. Schematic overview of the dynamic routing process | 48 |
| Figure 5. Overview of the DA procedure using TrAdaBoost and a random forest classifier | 50 |
| Figure 6. Conceptual diagram of (a) the self-attention operation and (b) the multi-head attention formulation within a GAT layer | 57 |
| Figure 7. Diagrammatic overview of the damage classification module in GAT-CAMDA | 58 |
| Figure 8. A schematic of GAT-CAMDA..... | 61 |
| Figure 9. Schematic representation of the WTS process up to the second order for a single channel | 66 |
| Figure 10. A schematic of SPADA framework..... | 71 |
| Figure 11. The 3.5kW WT made by Sonkyo Energy (Ou et al., 2021) | 81 |
| Figure 12. Workflow of the data acquisition system in WTB-VibClimate (Ou et al., 2021) | 82 |
| Figure 13. Components of the WTB-VibClimate climate-chamber experiment: (a) overall arrangement, (b) climate controller, (c) shaker with an insulation foam box, (d) strain gauges s_1 and s_2 on the low pressure side, (e) strain gauges s_{23} and s_{24} on the high pressure side, (f) climate sensor, (g) cracks 1 and 2, and (h) crack 3 (Ou et al., 2021) | 83 |
| Figure 14. Sensor configuration in WTB-VibClimate | 83 |
| Figure 15. Raw signals from WTB-VibClimate for (a) healthy and (b) cracked conditions..... | 84 |
| Figure 16. The test rig setup in the CONCEPT test..... | 85 |
| Figure 17. Raw signals from CONCEPT for (a) class C0 and (b) C11 | 88 |
| Figure 18. (a) Mode shapes identified at the baseline temperature of 10 °C, based on normalised frequency response functions; (b) FRF–temperature maps for the WTB-VibClimate dataset showing the variation of response magnitude across temperature levels..... | 91 |
| Figure 19. (a) CWI correlation–stretch map and (b) Fractional wave-speed change versus temperature for CONCEPT dataset | 92 |
| Figure 20. Comparison of the original and synthesised data using SMOTE for WTB-VibClimate..... | 95 |
| Figure 21. Experimental and synthesised time-domain responses of the healthy plate recorded by (a) PZT2, (b) PZT3, and (c) PZT4 | 97 |

| | |
|---|-----|
| Figure 22. Agreement between spline-generated and measured signals at Cp30 across PZT2, PZT3, and PZT4, reported using (a) average DTW and (b) average CC | 98 |
| Figure 23. Performance of Conv-CVAE signal synthesis at Cp30 across PZT2, PZT3, and PZT4, summarised by (a) average DTW and (b) average CC | 100 |
| Figure 24. CWI analysis of synthetic scenario C12: (a) correlation–stretch surface; (b) temperature-velocity relationship | 101 |
| Figure 25. Visualisation of the Wp5 embeddings obtained using (a) UMAP and (b) PCA | 103 |
| Figure 26. Comparative evaluation of clustering performance: (a) UMAP silhouette scores and (b) PCA silhouette scores | 103 |
| Figure 27. Scatter plot of after applying CapsNets for (a) Wp5 and (b) Wn15 | 104 |
| Figure 28. Confusion matrices of UCTRF with Wp5 as the source and Wn15 as the target domain without DA for (a) channel 1 and (b) channel 2..... | 105 |
| Figure 29. UCTRF performance across temperature conditions and damage categories for (a) channel 1 and (b) channel 3 | 107 |
| Figure 30. Confusion matrices of UCTRF with Wp5 as the source and (a) Wn15 and (b) Wp35 as the target domain with DA for channel 1 | 108 |
| Figure 31. Confusion matrices of UCTRF with Wp5 as the source and (a) Wn15 and (b) Wp35 as the target domain with DA for channel 3 | 108 |
| Figure 32. Confusion matrices of DANN with Wp5 as the source and Wn15 as the target domain for (a) channel 1 and (b) channel 3 | 110 |
| Figure 33. Confusion matrices of DANN with Wp5 as the source and Wn15 as the target domain for (a) channel 1 and (b) channel 3 | 111 |
| Figure 34. Evolution of scaled instance weights in the source and target domains at (a) epoch 1, (b) epoch 500, and (c) the final epoch..... | 113 |
| Figure 35. Per-class evolution of instance weights for selected Wn15 target-domain samples..... | 114 |
| Figure 36. GAT-CAMDA classification results with DA removed for (a) the source-domain test set and (b) the target-domain test set | 117 |
| Figure 37. Progression of the hyperparameter optimisation process | 118 |
| Figure 38. Hyperparameter slice plot across the optimisation trials | 119 |
| Figure 39. Target test results obtained with the best-performing GAT-CAMDA configuration, presenting (a) the confusion matrix and (b) the classification metrics | 119 |
| Figure 40. Two-dimensional t-SNE embeddings of the extracted feature space, showing (a) the representation before DA and (b) the representation after DA within GAT-CAMDA..... | 120 |
| Figure 41. Damage detection performance obtained with (a) Deep CORAL, (b) PRED, (c) SA, and (d) DANN | 123 |

| | |
|---|-----|
| Figure 42. Estimated hyperparameter importance for damage detection performance | 124 |
| Figure 43. Attention-derived sensor contribution estimates for the best-performing framework configuration | 125 |
| Figure 44. Attention-based sensor contribution estimates obtained using unnormalised signals..... | 126 |
| Figure 45. Scatter plot of the top-2 features ranked by mutual information, based on scattering transforms employed on Wp20; samples are colour-coded by class..... | 129 |
| Figure 46. Scatter plot of the top-2 features ranked using Wp20 data, visualised across 3 temperature conditions (Wn15, Wp20, Wp40); for (a) WTS-based and (b) CNNs-based feature extraction | 131 |
| Figure 47. Damage detection utilising WTS and CNNs without DA when (a) Wn15, (b) Wp20, and (c) Wp40 was assigned as the source domain..... | 132 |
| Figure 48. Damage detection utilising WTS and DA when (a) Wn15, (b) Wp20, and (c) Wp40 was assigned as the source domain | 135 |
| Figure 49. Confusion matrices of SPADA with Wp20 as the source and (a) Wn10, (b) Wp0, (c) W30, and (d) Wp35 as target domains | 137 |
| Figure 50. Scatter plot of the top-2 features ranked by mutual information, based on scattering transforms employed on Cp30; samples are colour-coded by class..... | 139 |
| Figure 51. Scatter plot of the top-2 features ranked using Cp30 data, visualised across 3 temperature conditions (Cp0, Cp30, Cp60); for (a) WTS-based and (b) CNNs-based feature extraction | 140 |
| Figure 52. Damage detection utilising WTS and CNNs without DA when (a) Cp0 and (b) Cp60 was assigned as the source domain | 141 |
| Figure 53. Damage detection utilising WTS and DA when (a) Cp0 and (b) Cp60 was assigned as the source domain | 143 |
| Figure 54. Confusion matrices for SPADA on CONCEPT for (a) source Cp0 and target Cp60, (b) source Cp0 and target Cp40, (c) source Cp60 and target Cp0, and (d) source Cp60 and target Cp10..... | 144 |
| Figure 55. Prototype trajectory evolution through t-SNE embedding for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer | 146 |
| Figure 56. Instance-prototype cosine similarity heatmap for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer..... | 148 |
| Figure 57. Prototype attention and confidence for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer..... | 150 |
| Figure 58. Decision boundaries in t-SNE embedding space for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer..... | 152 |
| Figure 59. Damage detection results of UDA benchmarks against SPADA for (a) source Wp20 and target Wn15, (b) source Wp20 and target Wp40, (c) source Cp0 and target Cp60, and (d) source Cp60 and target Cp10 | 153 |

List of Tables

| | |
|--|-----|
| Table 1. WT blade health conditions..... | 84 |
| Table 2. Summary of the test rig and instruments used in the CONCEPT experiments | 86 |
| Table 3. Damage scenarios and severities for simulated defects in the CONCEPT experiment | 87 |
| Table 4. Hyperparameter values used in CapsNets | 104 |
| Table 5. Damage classification results applying UCTRF through channel 1 | 106 |
| Table 6. Damage classification results applying UCTRF through channel 3 | 106 |
| Table 7. Hyperparameter grid search configuration for TCA implementation | 111 |
| Table 8. Allocation of samples to the source and target domains within the GAT-CAMDA framework..... | 116 |
| Table 9. Hyperparameter settings used to run GAT-CAMDA..... | 118 |
| Table 10. Performance comparison between GAT-CAMDA and alternative feature-based approaches | 122 |
| Table 11. Hyperparameter search space defining the randomised search candidates and values in SPADA | 134 |
| Table 12. System hardware and software specifications for SPADA pipeline | 155 |

List of Algorithms

| | |
|--|----|
| Algorithm 1. UCTRF Framework for SHM | 51 |
| Algorithm 2. GAT-CAMDA Pipeline for SHM | 62 |
| Algorithm 3. SPADA Framework for SHM | 77 |

Nomenclature

Acronyms

| | |
|------------|---|
| ADASYN | Adaptive Synthetic Sampling |
| BRF | Best Reconstruction Frequency |
| CapsNet | Capsule Network |
| CC | Cross-Correlation |
| CNNs | Convolutional Neural Networks |
| CONCEPT | Carbon-Epoxy Composite Plate Dataset |
| CORAL | Correlation Alignment |
| CWI | Coda Wave Interferometry |
| Conv-CVAE | Convolutional Conditional Variational Autoencoders |
| DA | Domain Adaptation |
| DANN | Domain-Adversarial Neural Network |
| DL | Deep Learning |
| DTW | Dynamic Time Warping |
| EEMD | Ensemble Empirical Mode Decomposition |
| EEG | Electroencephalography |
| EOVs | Environmental and Operational Variabilities |
| FRFs | Frequency Response Functions |
| GAN | Generative Adversarial Network |
| GAT | Graph Attention Network |
| GAT-CAMDA | Graph Attention–based Correlation Alignment and Maximum Discrepancy Adaptation framework |
| GNN | Graph Neural Network |
| GPR | Gaussian Process Regression |
| GRL | Gradient Reversal Layer |
| IMWMOTE | Improved Majority Weighted Minority Oversampling Technique |
| ML | Machine Learning |
| MMD | Maximum Mean Discrepancy |
| PCA | Principal Component Analysis |
| PBSHM | Population-Based Structural Health Monitoring |
| PZT | Lead Zirconate Titanate |
| PWAS | Piezoelectric Wafer Active Sensor |
| RF | Random Forest |
| RKHS | Reproducing Kernel Hilbert Space |
| RMSprop | Root Mean Square Propagation |
| RUL | Remaining Useful Life |
| SAE | Stacked Autoencoder |
| SHM | Structural Health Monitoring |
| SMOTE | Synthetic Minority Over-sampling Technique |
| SPADA | Scattering–Prototype–Adversarial Domain Adaptation |
| STFT | Short-Time Fourier Transform |
| TCA | Transfer Component Analysis |
| TL | Transfer Learning |
| TrAdaBoost | Transfer Adaptive Boosting |
| UCTRF | UMAP–CapsNet–TrAdaBoost–Random Forest framework |

| | |
|----------------|--|
| UDA | Unsupervised Domain Adaptation |
| UMAP | Uniform Manifold Approximation and Projection |
| VAE | Variational Autoencoder |
| WT | Wind Turbine |
| WTB | Wind Turbine Blade |
| WTB-VibClimate | Small-scale Wind Turbine Blade under Varying Climate Conditions dataset |
| WTS | Wavelet Time Scattering transform |
| XAI | Explainable Artificial Intelligence |

Symbols

| | |
|------------------------------|--|
| C | Number of Channels |
| $H(\omega)$ | Frequency Response Function |
| $\Delta v/v$ | Fractional Wave-Speed Change |
| ψ_ω | Band-Pass Wavelet Filters |
| Φ_j | Low-Pass (Scaling) Filter |
| S | Scattering Coefficient |
| X | High-Dimensional Input Data |
| U | Low-Dimensional Embedding |
| d | Embedding Dimensionality |
| ρ_i | Local Connectivity Parameter |
| σ_i | Smoothing Bandwidth |
| μ_{ij} | High-Dimensional Fuzzy Membership |
| ν_{ij} | Low-Dimensional Fuzzy Membership |
| s_j | Input Vector to Capsule j |
| v_j | Output Activity Vector of Capsule j |
| c_{ij} | Coupling Coefficient between Capsule i and Capsule j |
| b_{ij} | Log Prior Probability for Linking Capsule i to Capsule j |
| λ_{loss} | Weighting Factor on the Loss |
| N_s | Number of Samples in Source Domain |
| N_t | Number of Samples in Target Domain |
| w_{s_i} | Instance Weights for Source Sample |
| w_{t_i} | Instance Weights for Target Sample |
| a | Learnable Attention Vector for GAT Attention |
| e_{ij} | Unnormalized Attention Score |
| α_{ij} | Normalised Attention Coefficient |
| H | Number of Attention Heads |
| $\sigma(\cdot)$ | Nonlinear Activation Function |
| \mathcal{L}_{cls} | Classification Loss |
| \mathcal{L}_{adv} | Adversarial Domain Loss |
| \mathcal{L}_{MMD} | MMD Discrepancy Loss |
| $\mathcal{L}_{\text{CORAL}}$ | CORAL Covariance Alignment Loss |
| γ | Growth-Rate Parameter |
| p | Normalised Epoch Index |
| $g(\cdot)$ | Shared Feature Extractor |
| z_i | Latent Feature Vector |
| p_s | Source Prototype in Latent Space |
| p_t | Target Prototype in Latent Space |
| τ | Entropy Threshold |
| T_{temp} | Softmax Temperature |
| β | Similarity Weighting |
| w_i | Attention Weight |
| d_{src} | Source Prototype Drift Metric |
| d_{tgt} | Target Prototype Drift Metric |
| \mathcal{L}_{tgt} | Target Pseudo-Label Loss |
| $\mathcal{L}_{\text{proto}}$ | Prototype Alignment Loss |

Chapter 1: Introduction

1.1 Background

The structural integrity of composite materials has become paramount in modern engineering applications, particularly as these materials increasingly replace traditional metallic components in safety-critical infrastructure (Rezazadeh et al., 2024a). Composite structures, characterised by their high strength-to-weight ratios, design flexibility, and corrosion resistance, have proliferated across aerospace, renewable energy, maritime, and civil engineering sectors (Rajak et al., 2019). Wind turbine blades (WTBs) fabricated from glass or carbon fibre-reinforced polymers, aircraft fuselages employing carbon-epoxy laminates, and bridge decks utilising hybrid composite systems exemplify this widespread adoption (Arwood et al., 2023). However, the heterogeneous and anisotropic nature of composites introduces unique failure mechanisms that differ fundamentally from those observed in homogeneous materials, necessitating sophisticated damage detection methodologies capable of identifying subtle structural anomalies before catastrophic failure occurs (Heslehurst, 2014).

Damage in composite structures originates from multiple sources throughout the operational lifecycle. Manufacturing defects, including voids, resin-rich regions, and fibre misalignment, introduce initial imperfections that can propagate under service loads (Mehdikhani et al., 2019). Impact events from bird strikes, hail, or maintenance activities create barely visible impact damage that may extend through multiple plies whilst remaining undetectable through visual inspection (Janeliukstis et al., 2025). Cyclic loading induces fatigue-driven matrix cracking and fibre-matrix debonding, whilst environmental exposure causes moisture ingress, thermal cycling stress, and ultraviolet degradation (Sethi and Ray, 2015). The prevalent damage modes in composites include delamination between plies, transverse matrix cracking, fibre breakage, and interfacial debonding. In rotating systems such as wind turbine blades, additional failure modes emerge through mass imbalance, leading to excessive vibrations and accelerated fatigue. These damage mechanisms often interact synergistically; for instance, matrix cracks serve as initiation sites for delamination,

whilst moisture ingress accelerates crack propagation through freeze-thaw cycling. The complexity and interconnected nature of these failure modes demand monitoring systems capable of detecting, discriminating, and tracking multiple damage types simultaneously (Ge et al., 2024).

Conventional structural health monitoring (SHM) systems for composite structures employ networks of permanently installed or periodically deployed sensors to assess structural integrity through comparative analysis (Wilson et al., 2018). The typical SHM procedure involves baseline data acquisition from the pristine structure, continuous or periodic monitoring during service, signal processing to extract damage-sensitive features, and statistical pattern recognition to identify deviations indicative of damage. Vibration-based methods analyse changes in natural frequencies, mode shapes, and damping ratios, whilst guided wave approaches examine alterations in wave propagation characteristics, including velocity, attenuation, and scattering patterns. These conventional systems offer multiple merits: they provide objective and quantitative assessments of structural health, enable condition-based maintenance strategies that reduce operational costs, and detect internal damage that is inaccessible to visual inspection (Liu et al., 2024b). However, significant limitations constrain their practical deployment. Environmental and operational variations (EOVs) often mask damage signatures, leading to false alarms or missed detections (Mousavi and Gandomi, 2021). The requirement for extensive baseline data under multiple operational conditions creates substantial data management challenges. Moreover, physics-based damage indices often lack sensitivity to incipient damage whilst remaining susceptible to benign structural changes such as boundary condition variations or sensor degradation (Aulakh et al., 2025). These limitations have catalysed the development of more sophisticated, data-driven SHM frameworks that are capable of learning complex damage patterns while adapting to operational variability.

Machine learning (ML) and deep learning (DL) techniques have revolutionised SHM by automating feature extraction, capturing nonlinear damage-environment interactions, and enabling adaptive decision boundaries that evolve with operational conditions (Cha et al., 2024). Unlike conventional threshold-based methods, ML frameworks learn discriminative patterns directly from sensor data through supervised, unsupervised, and semi-supervised paradigms. Convolutional neural networks (CNNs)

automatically extract hierarchical features from time-frequency representations, capturing both local and global damage signatures without manual feature engineering. Recurrent architectures model temporal dependencies in sequential measurements, enabling early-stage damage prognostics. Autoencoders learn compact representations of healthy behaviour, with reconstruction errors indicating anomalous structural states (Tang et al., 2024). These frameworks demonstrate superior performance in laboratory settings, achieving high classification accuracies across multiple damage scenarios. However, their practical deployment faces a critical challenge: the scarcity of labelled damage data, particularly for rare but catastrophic failure modes (Rezazadeh et al., 2023a). Composite structures are expensive to damage deliberately for training purposes, and genuine damage events occur infrequently (Liu et al., 2024c). This data scarcity leads to overfitting, where models memorise training examples rather than learning generalisable damage patterns, or underfitting, where insufficient examples prevent models from capturing the full complexity of damage manifestations. The imbalanced nature of SHM data, with abundant healthy observations but limited damage examples, further exacerbates these challenges (Alzubaidi et al., 2023).

Data augmentation techniques offer one pathway to address data scarcity by synthetically expanding limited datasets whilst preserving physical plausibility. Time-domain augmentation methods include windowing to segment continuous acquisitions, noise injection to simulate measurement uncertainty, and time-stretching to emulate operational speed variations. Frequency-domain approaches apply spectral masking, phase randomisation, and harmonic synthesis to generate realistic variations (Yu et al., 2025). Interpolation-based methods, including spline fitting and synthetic minority oversampling, create intermediate examples between existing observations. Advanced generative models, such as variational autoencoders and generative adversarial networks, learn underlying data distributions to synthesise novel but physically consistent signals (Iglesias et al., 2023). However, augmented data introduce a fundamental challenge: whilst they may accurately represent the source domain conditions under which they were generated, they often fail to generalise to target domains characterised by different environmental or operational states. This domain shift arises primarily from EOVs that systematically alter sensor measurements without indicating actual structural damage (Mumuni and Mumuni, 2022).

EOVs constitute the most significant impediment to robust damage detection in composite structures. Temperature fluctuations modify material properties, with elastic moduli decreasing nonlinearly as temperature rises, thereby altering vibration characteristics and wave propagation velocities (Rezazadeh et al., 2024a). In carbon-epoxy composites, the transverse modulus can decrease by 30% over severe operational temperature ranges, whilst the longitudinal modulus remains relatively stable due to fibre dominance (Li et al., 2020; Yuan et al., 2023). Moisture absorption causes matrix swelling and plasticisation, reducing interlaminar shear strength and shifting damping characteristics (Lijuan et al., 2025). Loading variations change stress distributions and boundary conditions, modifying structural dynamics independent of damage (Guo et al., 2026). Manufacturing tolerances introduce specimen-to-specimen variability even within nominally identical structures (Rupal et al., 2020). These EOVs manifest as systematic shifts in feature distributions between training (source) and deployment (target) conditions. For instance, a classifier trained on room-temperature data may fail catastrophically when deployed in arctic or desert environments, interpreting temperature-induced stiffness changes as structural damage. The challenge intensifies in composite structures due to their anisotropic thermal expansion, hygroscopic behaviour, and viscoelastic properties, which create complex, nonlinear responses to environmental stimuli (Keshmiry et al., 2023). Furthermore, EOVs interact with damage in intricate ways; temperature changes may open or close cracks, moisture may temporarily increase damping that masks damage signatures and loading conditions may activate or deactivate certain failure mechanisms. These interactions necessitate monitoring frameworks that can disentangle damage effects from environmental influences whilst maintaining sensitivity to subtle structural changes (Rezazadeh et al., 2025c).

Multiple strategies have emerged to mitigate EOV effects in SHM, each with distinct advantages and limitations (Wang et al., 2014). Baseline subtraction methods store reference signals under various environmental conditions and select the most appropriate baseline for current measurements, but require extensive data storage and assume that stored baselines adequately span operational conditions (Rezazadeh et al., 2025c). Adaptive filtering techniques, including Kalman filters and recursive least squares, track slowly varying environmental trends to isolate damage-induced changes, though they struggle with abrupt environmental transitions or when damage

coincides with environmental changes. Reference-free methods exploit physical invariants, such as reciprocity in wave propagation or time-reversal focusing, that remain stable under EOV, whilst often require dense sensor arrays and sophisticated signal processing. Temperature compensation approaches explicitly model temperature-damage interactions through physics-based or data-driven regression, but demand accurate temperature measurement and may not capture complex nonlinear dependencies (Rezazadeh et al., 2025c). Multi-baseline strategies maintain libraries of healthy-state measurements across environmental conditions, enabling nearest-neighbour matching, though the curse of dimensionality limits their scalability (Rezazadeh et al., 2025c). Despite these advances, existing methods typically address specific EOV sources in isolation rather than providing comprehensive solutions for multiple, simultaneous variabilities. Moreover, they often require extensive calibration data or make restrictive assumptions about EOV characteristics that may not hold in practice.

Transfer learning (TL) and domain adaptation (DA) have emerged as promising paradigms for addressing both data scarcity and EOV challenges by enabling knowledge transfer from data-rich source domains to data-scarce target domains (Tang et al., 2024). Unlike conventional methods that treat each operational condition independently, TL and DA frameworks learn domain-invariant representations that generalise across environmental states (Yang et al., 2025a). Feature-based approaches project source and target data into a shared latent space where domain discrepancies are minimised whilst preserving class discriminability. Instance-based methods reweight source samples based on their relevance to the target domain, effectively filtering out domain-specific patterns. Model-based techniques fine-tune pre-trained networks using limited target data, leveraging learned representations whilst adapting to new conditions (Rezazadeh et al., 2025a). Unsupervised domain adaptation (UDA), which operates without target labels, remains challenging due to the risk of negative transfer where forced alignment degrades rather than improves performance. Furthermore, current DA methods rarely incorporate physical knowledge about damage mechanisms or material behaviour, treating sensor signals as abstract data rather than measurements governed by mechanical principles (Lu et al., 2024).

The interpretability of ML models represents a critical requirement for industrial adoption of automated SHM systems, particularly in safety-critical applications where diagnostic decisions carry significant consequences. Black-box models that provide accurate but unexplainable predictions face resistance from practitioners who require understanding of failure mechanisms to inform maintenance decisions (Rezazadeh et al., 2025b). Post-hoc interpretation methods, including gradient-based attribution, local surrogate models, and activation maximisation, attempt to explain trained models by identifying input features that most influence predictions. However, these techniques often produce inconsistent or misleading explanations for complex models (Zschech et al., 2025). For SHM applications, interpretability must extend beyond identifying relevant input features to revealing how models adapt across domains, which source examples influence target predictions, and how environmental factors modulate damage signatures. Visualisation of internal adaptation dynamics, including feature alignment trajectories, instance weighting evolution, and prototype movements, provides practitioners with diagnostic insights into model behaviour (Rezazadeh et al., 2025a, 2025b). This transparency enables validation against physical understanding, identification of failure modes, and calibration of confidence thresholds for operational deployment.

Computational efficiency emerges as a paramount consideration for real-time SHM systems deployed in resource-constrained environments (Mardanshahi et al., 2025). Industrial SHM applications demand continuous processing of multi-channel sensor streams with bounded latency for timely damage alerts (Okagawa et al., 2022). Edge computing scenarios, such as embedded systems on offshore wind turbines or remote pipeline networks, impose strict constraints on memory, processing power, and energy consumption (Zhang et al., 2023). Whilst DL models achieve superior accuracy in laboratory settings with powerful GPUs, their deployment on industrial hardware requires careful architectural design (Ngo et al., 2025). Model compression techniques, including pruning, quantisation, and knowledge distillation, reduce computational requirements but may compromise detection sensitivity (Dantas et al., 2024). Efficient feature extraction methods that provide compact, discriminative representations without expensive neural processing offer one solution (Al-Thelaya et al., 2023). Physics-informed architectures that encode domain knowledge reduce model complexity whilst maintaining performance (Jiang et al., 2025b). The trade-off

between model sophistication and computational efficiency must be carefully balanced; overly complex models may be impractical for deployment, whilst oversimplified approaches may miss critical damage signatures (Resck et al., 2024). Cloud-based architectures that partition processing between edge devices and remote servers provide flexibility but introduce latency and reliability concerns (The future of edge computing, 2022). Therefore, SHM frameworks must be designed with deployment constraints as primary considerations rather than post-hoc optimisations.

The selection of appropriate feature extraction techniques critically influences damage detection performance, particularly under domain shift conditions (Buckley et al., 2023). Features must preserve damage-sensitive information whilst exhibiting robustness to environmental variabilities and measurement noise (Ashkarkalaei et al., 2025). Wavelet-based representations capture multi-scale temporal structures characteristic of damage-induced transients whilst providing natural denoising through coefficient thresholding (Akkaya, 2025). Scattering transforms extend wavelet analysis through cascaded convolutions and modulus operations, yielding translation-invariant features with guaranteed stability to deformations (Rezazadeh et al., 2023b). Empirical mode decomposition adaptively decomposes signals into intrinsic mode functions that separate damage-related oscillations from environmental trends (Ahmed et al., 2022). Graph-based features exploit spatial relationships in sensor networks, with attention mechanisms learning which sensor pairs carry diagnostic information under different conditions (Rezazadeh et al., 2025b).

The choice between hand-crafted features that encode physical understanding and learned representations that discover patterns from data depends on available domain knowledge, dataset size, and interpretability requirements (Bento et al., 2022). Hybrid approaches that combine physics-based preprocessing with data-driven refinement offer a pragmatic balance, leveraging mechanical insights whilst maintaining adaptability (Hu and You, 2022). Importantly, features should be designed to facilitate DA; representations that naturally separate damage-related and environment-related variations simplify subsequent alignment procedures.

1.2 Research Aim

This research aims to develop and validate robust DL frameworks for damage detection in SHM of composite materials by leveraging cross-domain adaptation to

address data scarcity, maintain diagnostic accuracy under EOVs, and ensure model interpretability and computational efficiency for practical industrial deployment.

1.3 Research Questions

This research addresses the following primary questions:

1. How can data augmentation techniques be effectively employed to alleviate the scarcity of labelled fault data in SHM of composite materials?
2. How do EOVs impact data distribution and the performance of DL-based damage detection systems, and what strategies can mitigate their effects?
3. How can sophisticated feature extraction techniques distil health-relevant, domain-invariant features?
4. To what extent can DA techniques enable trained models to generalise across varying operating conditions, machines, or environmental states without extensive retraining?
5. How can the internal decision-making processes of DL models be rendered interpretable and explainable to foster trust and support adoption in industrial SHM applications?
6. What computational demands and practical constraints must be addressed to facilitate the real-time deployment of advanced ML and DL pipelines in industrial SHM systems?

1.4 Research Objectives

To address the research questions and fulfil the research aim, the following specific objectives have been established:

1. To investigate and implement time-series data augmentation methods to synthetically expand limited fault datasets whilst preserving the physical characteristics of structural responses.
2. To quantify the impact of EOVs (temperature) on damage detection accuracy and to characterise the resulting distribution shifts in sensor data.

3. To develop advanced feature extraction methodologies that capture characteristics of structural responses whilst exhibiting robustness to environmental variations.
4. To design and implement effective and novel DA frameworks to facilitate knowledge transfer across different operational domains without requiring extensive labelled target data.
5. To integrate explainable AI techniques into the damage detection pipeline to visualise and interpret the model's decision-making process and identify salient signal features.
6. To validate the proposed methodologies through experimental case studies on WTB vibration data under varying climatic conditions and carbon-epoxy composite plate monitoring.
7. To conduct comparative analyses of the proposed frameworks against established reference methods across multiple domain shift scenarios to demonstrate performance improvements and robustness.
8. To evaluate the computational efficiency of the developed frameworks and demonstrate their feasibilities for real-time industrial deployment on standard computing hardware without requiring specialised GPU resources.

1.5 Structure of the Thesis

The structure of this thesis is organised as follows:

1. Chapter 1: Introduction

This chapter introduces the context and motivation for the research, outlining the challenges associated with SHM of composite materials. It highlights the critical issues of data scarcity, environmental and operational variability, and cross-domain generalisation in machine learning-based damage detection. The chapter defines the research aim, presents the research questions and objectives, and provides the structure of the thesis.

2. Chapter 2: Literature Review

This chapter presents a comprehensive review of existing methodologies in structural health monitoring, focusing on signal structure and invariance requirements, feature extraction techniques, data augmentation strategies for time-series data, environmental and operational variability mitigation approaches, transfer learning and domain adaptation methods, and explainable artificial intelligence techniques. The review adopts a hybrid narrative-systematic strategy, providing selective coverage of feature extraction methods whilst offering detailed treatment of domain adaptation and interpretability techniques relevant to the research objectives.

3. Chapter 3: Methodology and Methods

This chapter describes the methodological foundations for the three domain adaptation frameworks developed in this research. It details the data augmentation approaches, including spline interpolation and Convolutional Variational Autoencoders (CONV-CVAE) for temperature-dependent signal synthesis. The chapter presents comprehensive mathematical formulations and algorithmic specifications for the UCTRF framework (integrating UMAP, CapsNet, TrAdaBoost, and Random Forest), the GAT-CAMDA framework (combining graph attention networks with multiple alignment mechanisms), and the SPADA framework (incorporating Wavelet Transform Scattering, prototype-based attention, and adversarial learning with confidence-gated pseudo-labelling).

4. Chapter 4: Case Studies

This chapter describes the two experimental datasets used to validate the proposed methodologies. The first case study examines a small-scale wind turbine blade under varying climate conditions (WTB-VibClimate), encompassing vibration measurements across substantial temperature ranges. The second case study investigates a carbon-epoxy composite plate (CONCEPT) subjected to different loading configurations, utilising guided wave monitoring techniques. The chapter details the experimental setups, data acquisition protocols, damage scenarios, and environmental conditions for both case studies.

5. Chapter 5: Results and Discussion

This chapter presents comprehensive experimental results and analyses across multiple dimensions. It quantifies the effects of environmental and operational variability on damage detection performance, evaluates data augmentation strategies, and validates the three domain adaptation frameworks across challenging source-target transfers. The chapter includes comparative studies against established baseline methods, detailed interpretability analyses revealing internal adaptation mechanisms through prototype trajectories, attention dynamics, and decision boundary visualisations, and computational efficiency assessments demonstrating feasibility for real-time industrial deployment on standard hardware.

6. Chapter 6: Conclusions and Future Work

The final chapter summarises the key findings of the research, highlighting the contributions of the three complementary domain adaptation frameworks such as UCTRF, GAT-CAMDA, and SPADA to structural health monitoring of composite materials. It addresses how each research question was answered through theoretical development and empirical validation, acknowledges limitations of the current work, and discusses promising directions for future research including multi-source domain adaptation, open-set recognition, physics-informed adaptation, continual learning, real-time implementation, uncertainty quantification, and cross-modal generalisation.

1.6 Publications

The following is a list of publications that have resulted from this PhD research project. These publications have contributed to the field and have been disseminated through peer-reviewed journals.

- **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, doi:10.1177/14759217241313350.
- **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, doi:10.3390/s25226847.

- **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(15), 3146, doi:10.3390/electronics14153146.
- **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 Materials and Structures* 2025, 34(7), doi: 10.1088/1361-665X/ade7db.
- **Rezazadeh, N.;** Sharif-Khodaei, Z.; Aliabadidai, M.H.; Annaz, F.; De Oliveira, M. Interpretable Domain-Adaptive Framework for Structural Health Monitoring Under Environmental Variability, *Mechanical Systems and Signal Processing*, Under Review.
- **Rezazadeh, N.;** De Luca, A.; Lamanna, G.; Annaz, F.; De Oliveira, M. A Novel Synthesis Framework to Mitigate Data Scarcity in Mechanical Fault Diagnosis Using Sparse Identification of Nonlinear Dynamic Systems and Empirical Wavelet, *Engineering with Computers*. Under Review.

Chapter 2: Literature Review

2.1 Overview

SHM systems detect and diagnose damage in engineering structures through sensor data analysis. A persistent challenge is data scarcity, particularly for labelled fault conditions, as deliberate damage is costly and many damage mechanisms only manifest after prolonged service. Recent research addresses this by re-using data from similar structures or synthetic signals from numerical models. However, sensor responses from different assets or simulations rarely follow identical distributions due to variations in fabrication, boundary conditions, and operational loading, causing domain shifts that degrade ML model performance when transferred between assets.

This chapter adopts a hybrid narrative-systematic review strategy, with selective coverage of feature extraction methods aligned with the thesis methodology, whilst EOV mitigation, TL and DA, and interpretability receive detailed treatment. The review is organised into sections covering signal structure and invariance requirements, data augmentation and dataset completion for time series, EOV and dataset shift, and XAI. Figure 1 represents the structure of the literature review.




Figure 1. Literature review structure and topic coverage

2.2 Signal Structure and Invariance Requirements

Signals recorded in mechanical and civil SHM are non-stationary and multiscale. In thin plates and shells, ultrasonic Lamb waves are subject to dispersion and mode conversion, which produces overlapping components in time–frequency representations. Rotating machinery responses contain periodic bearing and gear excitations superimposed on broadband noise, and all responses vary with environmental and operational conditions such as temperature, rotational speed, prestress and loading. Accordingly, features are expected to exhibit translation invariance, scale invariance, and either permutation invariance to sensor ordering or graph equivariance to array topology, while preserving damage information and suppressing variability unrelated to damage.

2.3 Feature Extraction in SHM

Damage detection in mechanical systems is reliant on robust, discriminative features extracted from measured responses (Zong et al., 2023). It has been shown that when carefully designed feature extractors are combined with resilient ML models, they can yield transferable representations that generalise across different structures and conditions. Meanwhile, automatic feature extraction layers in advanced DL architectures have been adopted to learn hierarchical representations directly from raw data, thereby reducing the need for manual feature engineering and improving adaptability to varying operational conditions (Soomro et al., 2024). In the following subsections, both classical feature engineering approaches and modern DL-based feature extraction methods are detailed, examining how each addresses the challenges of damage detection under variable operational conditions.

2.3.1 Classical Descriptors in Time, Frequency and Time–Frequency Domains

Early SHM methods were based on engineered features derived from time and frequency analyses. Examples include peak amplitudes, energy, root–mean–square values, kurtosis, skewness, resonant frequencies and mode shapes. Time–frequency representations such as short–time Fourier transforms, and continuous wavelet transforms were employed to permit localisation of transients (Perfetto et al., 2023). Although these descriptors are straightforward to compute and interpret, generalisation across varying operational conditions is often not achieved because of sensitivity to

changes in signal energy and frequency content. For TL and DA, such descriptors are typically brittle under EOV and are best used as transparent baselines or as inputs to subsequent alignment.

2.3.2 Wavelet Time Scattering

WTS is a multiresolution transform built from cascades of fixed Morlet wavelets followed by modulus and averaging operators. Deformation–stable and translation–invariant representations are produced without learning filter coefficients. In mechanical fault diagnosis, scattering features have been shown to capture multiscale modulation effects due to defects while being more robust to nuisance variability than conventional time–frequency indices (Susu et al., 2020).

In composite laminates, the use of WTS with probabilistic regression to localise impacts was reported by the authors in (Ojha et al., 2024), with improved robustness to variability compared with conventional scalogram features, which supports adoption of wavelet scattering as a low–variance front end for both guided wave imaging and data–driven classifiers. For rotating machinery, scattering transform features coupled with a least squares proximal twin support vector machine (SVM) were proposed in (Ma et al., 2018) and robust fault diagnosis under noise was demonstrated, which highlights the advantage of deformation–tolerant descriptors over classical time–frequency indices. Similarly, fault diagnosis in rotor systems under uncertainty caused by operational variabilities was addressed in (Rezazadeh et al., 2022) through the employment of WTS in combination with Long Short–Term Memory and SVM, and damage–sensitive features were effectively extracted even in the presence of operational variations. Focusing on compound faults, WTS with improved soft threshold denoising was combined by Guo et al. (Guo et al., 2022) to separate interacting bearing defects from vibration data, further evidencing robustness in realistic operating conditions. For TL and DA, WTS can reduce source–target discrepancy as a low–variance, physics–informed front end, while retaining damage modulations needed for transfer.

2.3.3 Manifold and Embedding Methods

When the intrinsic dimensionality of measured signals is low, manifold–learning techniques can be used to reveal underlying structures. Principal component analysis (PCA) is a linear method that projects high–dimensional data onto directions of

maximal variance. Kernel PCA, locally linear embedding (LLE), t -distributed stochastic neighbour embedding (t -SNE) and uniform manifold approximation and projection (UMAP) are non-linear methods that aim to preserve neighbourhood relationships while revealing clusters. In vibration-based SHM, such embeddings have been used for visualising damage clustering and for dimensionality reduction prior to classification (Rezazadeh et al., 2024b).

For example, the reduction of high-dimensional vibration data to identify structural damage under harmonic excitation was undertaken by Golinval (Golinval, 2017). Damage detection was achieved by evaluating the subspace angle between eigenvector spaces of undamaged and potentially damaged states, with validation on a real truss structure. A visualised classification method via t -SNE for planetary gearbox fault identification from raw mechanical data was developed by Chen et al. (Chen et al., 2018), where a separable low-dimensional structure supporting accurate recognition was revealed. A robust LLE algorithm for machinery fault diagnosis was proposed by Zhang et al. (Zhang et al., 2018) and improved clustering and identification of bearing faults under noise were shown, evidencing utility for non-linear feature extraction in vibration-based monitoring. The projection of raw Lamb wave signals to a low-dimensional manifold followed by an inductive mapping to cluster damage scenarios in complex stiffened composite structures more effectively than hand-crafted indices was demonstrated by Rahbari et al. (Rahbari et al., 2021), with UMAP motivated as a competitive non-linear map. The practical benefit of an explicit out-of-sample extension to project previously unseen signals was also highlighted, which is essential for online deployment. An advantage of UMAP over t -SNE is the inductive mapping learned from training data that allows new data to be projected without retraining, enabling online deployment (McInnes and Healy, 2018). For TL and DA, embeddings are useful for exploration and visual diagnostics, although instability across runs and distortion of global structure should be noted; parametric or inductive variants are preferable when downstream transfer is intended.

2.3.4 Neural Feature Learners for Signals

Neural feature learners are used to extract informative patterns from raw signals, thereby reducing reliance on manual feature engineering. CNNs efficiently capture local temporal and spectral structures through shared convolutional filters, while

attention mechanisms have been used to enhance focus on fault–relevant components. Graph–based models such as Graph Attention Networks (GATs) are used to model spatial–temporal dependencies among sensors, improving generalisation under varying conditions (Lu and Chen, 2025). Capsule networks (CapsNets) further preserve hierarchical and spatial relationships, offering robust representations of damage–sensitive features across complex operational settings (Marchisio et al., 2022).

CNNs are the most widely adopted architecture in the SHM of mechanical systems; convolutional layers effectively process one–dimensional and time–frequency representations to capture local signal patterns, with weights shared across time to achieve translation equivariance and with pooling used to promote limited invariance and scalability for large datasets (Rezazadeh et al., 2021). In recent years, deep CNNs combined with complementary strategies have been employed to address the complexity of composite structures where standalone CNNs can fall short (Mirkhalaf and Rocha, 2024). For instance, a deep CNN combined with wavelet transforms to detect delamination in carbon fibre reinforced polymer plates was employed by Wu et al. (Wu et al., 2021), while a fusion of CNNs with a stacked autoencoder to localise and quantify defects in composite plates was reported by Gao and Hua (Gao and Hua, 2022). A physics–guided CNN framework that enhanced generalisation and interpretability under EOV was proposed by Xu and Liu (Xu and Liu, 2024). These hybrid approaches illustrate that CNNs, when synergised with complementary techniques, can deliver robust and precise damage detection in sophisticated composite systems.

Limitations related to the loss of salient features can arise in CNNs, primarily due to pooling operations that retain only maximum or average values, thereby potentially discarding critical information. To alleviate these shortcomings, CapsNets were introduced in (Hinton et al., 2011), in which capsules, namely clusters of neurons that process vector–valued information, were used to encode properties of a feature including orientation and spatial relationships. Intricate inter–feature dependencies are preserved as a result, and recognition performance can be enhanced when compared with conventional CNNs (Kwabena Patrick et al., 2022).

Building upon this framework, a CapsNet–driven approach for the diagnosis of faults in wind turbine gearboxes was proposed by Liang et al. (Liang et al., 2023). A weight–sharing capsule network aimed at diagnosing the health of automotive transmissions under EOV was introduced by Huang et al. (Huang et al., 2020). A domain–separated capsule architecture for guided wave damage detection under EOV was proposed by Jiang et al. (Jiang et al., 2025a) and stronger generalisation than plain CNNs was shown, which supports the idea that capsule features can better preserve physically meaningful configuration information under condition changes. Despite the implementation of numerous feature extraction methodologies for damage detection in composite panels, the application of CapsNets in this domain has been reported less frequently than CNNs.

Blades and large composite plates are commonly instrumented with sparse arrays, in which the information content lies not only in each channel but also in the topology of actuator–sensor paths and their acoustic reciprocity (Michaels, 2016). Graph neural networks (GNNs) and, in particular, graph attention networks assign trainable weights to edges, thereby allowing focus to be placed on paths that are most informative for damage localisation (Vrahatis et al., 2024).

GATs and their variants have been employed in the SHM of engineering systems in the present decade. Attention mechanisms within GATs were employed by Niu et al. (Niu et al., 2022) to manage incomplete SHM datasets by capturing spatial–temporal dependencies among sensor nodes. A virtual sensing approach for predicting bearing loads was developed by Zhao et al. (Zhao et al., 2024), in which heterogeneous temporal GNNs were used to model inter–sensor relationships explicitly, and superior performance in load estimation tasks compared with conventional CNNs was achieved. For WTB, acoustic emission and guided response localisation were formulated as a graph problem by Zhao and Chen (Zhao and Chen, 2024) and an exponential smoothing multi–head graph attention network was introduced to stabilise temporal variation while highlighting spatially coherent damage zones, with outperformance of baselines that ignore graph structure reported. A companion study from the same group showed that even a single–sensor blade set–up can be recast as a graph over time–frequency slices and localised with a graph convolutional surrogate, further endorsing relational encoders for sparse sensing on complex blades.

For TL and DA, neural encoders provide adaptable representations that can be aligned across domains while preserving damage cues. CNN front ends offer efficient local pattern capture, CapsNets can retain configuration information relevant under EOV, and graph attention mechanisms can prioritise informative actuator–sensor paths in sparse arrays, all of which support transfer under practical sensing constraints.

2.4 Data Augmentation and Dataset Completion for Time Series

Data augmentation and dataset completion techniques have become essential in time–series analysis, particularly for SHM, where the collection of sufficient data across diverse conditions is often impractical. These methods are used to enrich limited datasets, enhance model generalisation, and fill gaps at missing or unobserved operating states. By generating realistic signal variations, augmentation is used to mitigate overfitting and to balance class distributions. The following subsections outline major approaches across time, frequency, and latent domains, as well as interpolation and generative strategies for robust dataset completion.

2.4.1 Classical Time–Domain Augmentations

In time–domain signals, basic operations include the addition of Gaussian noise, the scaling of amplitudes, resampling at slightly different rates, time warping, the cropping of windows from longer signals, and the shuffling of segments. Energy consumption time series were augmented by Demir et al. (Demir et al., 2021) with jittering and scaling to improve regression accuracy under varying loads. Voice pathology detection was enhanced by Javanmardi et al. (Javanmardi et al., 2024) through the application of time stretching and noise injection to speech recordings. EEG–based mental state classification was improved through amplitude scaling and random segment permutation by Basheer et al. (Basheer et al., 2024). These studies indicate that classical augmentations can be used to diversify temporal data and to increase robustness across different systems. However, for complex engineering systems such as composite structures, simple augmentations may fail to capture the physical dependencies and mode interactions inherent in the signals.

2.4.2 Frequency and Time–Frequency Domain Augmentations

Beyond waveform edits, augmentations can be applied in the spectral domain by perturbing magnitudes or phases. Random time–frequency regions of leak–noise

spectrograms were masked by Wu et al. (Wu et al., 2024) and paired with phase-randomised iterative amplitude-adjusted Fourier transform surrogates to boost CNN detection in water-distribution pipelines. STFT-based time-frequency masking with adaptive thresholds that suppress nuisance components was introduced by Matthew et al. (Ejiofor Matthew et al., 2024), functioning as analysis and as spectral occlusion augmentation to improve robustness on non-stationary vibration signals. Mini-batch spectrogram augmentations, including frequency and time masking, were evaluated by Kivaisi and Zhao (Kivaisi and Zhao, 2024) for classification, and it was shown that the combination of multiple masks within a batch yielded stronger generalisation, an approach that is transferable to machinery spectrograms. A potential challenge with these methods is that spectral perturbations can unintentionally distort key frequency-mode relationships or introduce artefacts, which can lead to misinterpretation of damage-sensitive features in complex mechanical systems.

2.4.3 Minority Class Balancing with SMOTE Variants

The SMOTE generates synthetic minority samples by interpolating between existing neighbours. Variants such as Borderline-SMOTE and ADASYN focus sampling near class boundaries or harder-to-learn regions. In mechanical fault diagnosis, scarce rotating-machinery faults in nuclear power plants were balanced by Yin et al. (Yin et al., 2024) using ADASYN with EEMD and continuous wavelet transform features feeding a ResNet-style classifier, and performance under imbalanced samples was improved. Feature-level SMOTE was introduced by Liu et al. (Liu et al., 2024a), with interpolation in a learned latent space for gas-turbine fault diagnosis used to preserve temporal coherence better than raw-space SMOTE. An importance-weighted oversampling method, IMWMOTE, was proposed by Wang et al. (Wang et al., 2024) to target informative minority regions and to boost accuracy on heterogeneous machinery-fault datasets. While effective, these oversampling strategies can distort sequence dynamics or physics-consistent class boundaries if applied naively, so sequence-aware or latent-space interpolation is preferable for complex engineering systems.

2.4.4 Cross-Condition Interpolation

In SHM, data are often missing at certain operating conditions; spline or regression interpolation across temperature or loading can be used to estimate

responses at unobserved points and to provide a continuous baseline. Frequency–temperature relationships on the Ting Kau Bridge were modelled by Ni et al. (Ni et al., 2005), enabling interpolation of natural frequencies across the thermal range for baseline normalisation. Environmental and operational effects, notably temperature, on a footbridge were quantified and regressed out by Moser and Moaveni (Moser and Moaveni, 2011), with the learned trends used to interpolate modal properties and to reduce false alarms. EOV of modal features was characterised by Wah et al. (Soo Lon Wah et al., 2021) using regression, allowing condition–consistent interpolation before damage assessment. While effective, such interpolation can over–smooth damage cues or mislead under extrapolation, and care is needed to respect physics, including mode ordering and thermal dependencies, when constructing continuous baselines.

2.4.5 Generative Models for Synthetic Data

Deep generative models are trained to learn the underlying distribution of time-series and can be used to synthesise realistic signals to combat data scarcity. In SHM, a normalised conditional variational autoencoder (VAE) for bearing–rotor systems was introduced by Zhao et al. (Zhao et al., 2022), with class–conditioned latent sampling used to generate minority–class vibration patterns under imbalanced data. A physics-informed cycle–consistent generative adversarial network (GAN) was adapted by Ge and Sadhu (Ge and Sadhu, 2024) to translate signals between source and target domains, thereby generating target–like responses that reduce domain shift in structural vibration monitoring. A time–series diffusion framework to generate high–fidelity vibration signals for data augmentation was proposed by Yi et al. (Yi et al., 2024), and improved fault diagnosis with limited samples was reported. These examples show that VAEs, GANs, and diffusion models can be used to generate condition–aware data while maintaining realism through physics–based constraints. However, improper tuning may lead to unrealistic dynamics, mode distortion, or overfitting to artefacts that compromise the physical interpretability of generated signals.

2.4.6 Fidelity and Validation of Synthetic Data

Although higher precision in data synthesis has been reported in SHM applications, the quality of the produced data should be evaluated through appropriate tests. Augmented data should be validated to ensure that plausible behaviours are

represented. Statistical similarity tests, including two-sample tests such as the Kolmogorov–Smirnov test, dynamic time warping (WTP) distances, cross-correlation measures, and maximum mean discrepancy, can be used to assess similarity between synthetic and real data. Validation should also consider whether augmentation preserves label information and does not introduce misleading patterns, as highlighted by Vallevik et al. (Vallevik et al., 2024) and Bahrpeyma et al. (Bahrpeyma et al., 2021), who proposed frameworks to ensure fidelity, diversity, and representativeness in synthetic data generation. In addition, augmentation should be confined to the source domain during DA to avoid target leakage, and reports should include with-augmentation and without-augmentation baselines to quantify the effect of the chosen strategy.

2.5 EOV and Dataset Shift

One of the primary challenges in any SHM framework lies in extracting damage-related features from the acquired signals. Threshold values for various conventional features are often defined in international standards, such as ISO 10816, which provides guidelines for the measurement and evaluation of mechanical vibration in machines. Comparing real-time measurements with these standard values, abnormalities in structural or mechanical behaviour can be identified, and an experienced practitioner may even infer the damage type or its underlying cause. However, multiple factors hinder accurate SHM performance, including ambient noise, EOV, and potential issues with measurement instruments, all of which introduce uncertainty into sensor readings. Extensive research has been undertaken to mitigate these effects, for instance through advanced filtering techniques for noise reduction.

In the context of EOV, before the emergence of ML-based SHM frameworks, direct compensation and baseline alignment were the primary strategies for neutralising its influence, whereas modern intelligent approaches increasingly employ adaptive and multi-baseline and reference-free techniques. The following section elaborates on these previous studies.

2.5.1 Direct Signal-Level Baseline Compensation

Direct compensation acts at waveform level by co-registering the current record with a healthy baseline in time, phase, and where needed amplitude, before subtraction

yields a difference primarily driven by damage rather than by EOV. The appeal of this family is its simplicity, speed, and suitability for embedded deployment in composite panels, aluminium structures, pipelines, aircraft skins, and WT blades; however, performance typically degrades when the stored baseline becomes obsolete through ageing or when the environment shifts far beyond the range seen during commissioning (Croxford et al., 2008; Mariani et al., 2020). Within mechanical systems, 3 strands repeatedly appear, namely subtraction with phase or time alignment, adaptive filtering, and real-time and simulation aided corrections, which together delineate how signal level compensation is extended from short-term and bounded conditions to wider operating regimes (Rezazadeh et al., 2025c).

2.5.1.1 Subtraction and Phase or Time Alignment

The central idea is that much of the environmentally induced mismatch manifests as arrival time shifts and phase rotation because temperature changes wave speed and alters coupling; therefore, careful co registration in time and phase prior to subtraction suppresses a large fraction of the drift and leaves a residual enriched with damage related scattering. Building on this rationale, Nandyala and Darpe et al. (Nandyala et al., 2020) introduced a modified delay sum imaging with a damage residue measure that compares geometrically similar propagation paths and subtracts a calibrated baseline residue, thereby maintaining accurate localisation of impact damage in cross ply laminates under temperature change. In a complementary aluminium context, Mariani et al. (Mariani et al., 2020) formulated an explicit phase compensation that addresses temperature dependent velocity and transducer drift; once phase is corrected before baseline subtraction, spurious residuals decrease markedly under moderate thermal swings in pipe and plate like geometries. Extending this line, Salmanpour et al. (Salmanpour et al., 2017b) conducted a wide-ranging comparative analysis of guided wave temperature correction procedures and proposed a minimum residual alignment that selects a correction minimising misfit while preserving sensitivity to defect scattering, with demonstrations on both composite and aluminium plates across temperature windows. Overall, these studies show that subtraction can remain effective provided the baseline is representative and the co-registration accurately removes time and phase offsets, while the reported drift of baselines with service age explains why alignment alone cannot protect performance indefinitely.

2.5.1.2 Adaptive Filtering

When slow thermal drift cannot be fully removed by a fixed alignment, signal level learning can track the mismatch between baseline and current record in real-time. An adaptive linear neuron, adaptive linear neuron ADALINE, updates its coefficients so that the predicted baseline component follows temperature driven changes and the residual highlights novelty consistent with damage. In composite plate experiments with piezoelectric transducers, Wang and Gao et al. (Wang et al., 2014) demonstrated that ADALINE based compensation tracks velocity shifts with rising temperature and stabilises residuals, which reduces the need to archive multiple temperature specific baselines; independently, Huang and Yuan et al. confirmed the feasibility of ADALINE temperature compensation in a different configuration and excitation regime, again yielding residuals that emphasise defect scattering rather than environmental drift (Huang et al., 2014). These results indicate that lightweight online adaptation is feasible for embedded SHM in mechanical assets, while also showing that very large or abrupt excursions can exceed the tracking range of such low-order predictors, which motivates coupling with improved baseline selection or model supported correction when EOVS varies widely.

2.5.1.3 Real-Time and Simulation-Aided Corrections

A natural extension replaces the static reference with a baseline that is consistent with today's environment and operation, which can be achieved by reconstructing the expected healthy response from a measurement informed model or by synthesising a temperature matched reference from limited anchors. In an aluminium plate study, Aryan et al. (Aryan et al., 2016) combined scanning laser vibrometry with high fidelity finite element modelling (FEM) to reconstruct the healthy time trace at the current temperature; with that reconstructed baseline in place, subtraction operates reliably across large thermal variations without requiring dense baseline libraries. For composite laminates with piezoelectric sensing, it was demonstrated that a single baseline temperature correction learns to map a laboratory reference to the present thermal state, which reduces archival demands while retaining detectability under substantial temperature change (Salmanpour et al., 2017a). In field operation, Simon et al. (Simon et al., 2024) applied radar measurements with trend decomposition to remove temperature induced drift from the baseline for WT blade monitoring,

improving indicator stability under daily and seasonal variability while maintaining sensitivity to genuine changes. Collectively, these studies show that direct compensation can be pushed well beyond small perturbations when a current condition baseline is reconstructed or when environmental variability is removed as a trend, albeit at the cost of additional modelling effort, parameter identification, and continued maintenance of the model to reflect structural ageing.

2.5.2 Adaptive and Multi-Baseline Strategies

Adaptive and multi-baseline strategies reduce dependence on any single static reference by curating, updating, or synthesising baselines that better reflect the environment and operation at the time of measurement (Goodwin and Olazabal, 2025). The central idea is to treat EOV as a slowly evolving context that can be learned or represented by a small set of informative exemplars, rather than exhaustively storing every condition. In mechanical systems, this family extends the useful range of baseline dependent methods for composite laminates, aluminium plates and shells, pipelines, aircraft components, and WTB, while introducing moderate computational and curation overheads that remain practical for long term monitoring (Rezazadeh et al., 2025c).

2.5.2.1 Evolving Baseline Libraries

The concept idea in these approaches is to grow or refine a reference library during service so that each new healthy measurement either augments or reweights the set of admissible baselines. By doing so, the monitoring system continually reduces mismatch between the current record and the selected reference without enumerating dense condition grids, which is valuable when temperature bands widen over seasons and when sensors experience gradual ageing. In a seminal contribution on guided waves for composite plates, Putkis and Croxford (Putkis and Croxford, 2013) proposed continuous baseline growth, whereby healthy data are appended incrementally, and selection rules choose the most compatible baseline at runtime; the study improved sensitivity relative to relying on a single commissioning reference while avoiding the storage burden of exhaustive condition catalogues. In parallel, Abbas et al. (Abbas et al., 2020) investigated optimal baseline selection for a stainless-steel element and demonstrated that reselecting or updating baselines to minimise residual mismatch under changing temperature and operation leads to more stable indicators than fixed

references; the work formalised selection as an optimisation step that is repeated as the environment drifts. Together these studies illustrate a pragmatic pattern for mechanical assets. Rather than attempting to anticipate every environmental state offline, evolving libraries allow a manageable set of baselines to be curated from service data, provided that stewardship rules ensure new entries are verified as healthy and that outdated entries are retired in a controlled way.

2.5.2.2 Virtual Baseline Synthesis

A complementary route is to synthesise a temperature consistent baseline from a small number of anchors, thereby compressing storage without sacrificing fidelity within the trained range. The governing principle is to learn a low-dimensional mapping from condition to waveform so that amplitude and phase adjustments predicted for today's environment can be applied to one or two stored exemplars. Yue and Aliabadi (Yue and Aliabadi, 2020) devised such a scalable approach for anisotropic carbon fibre reinforced polymer structures by identifying dimensionless amplitude and phase compensation factors from a minimal set of references; extrapolating those factors yielded accurate reconstructions of baselines at unmeasured temperatures, which is especially attractive for embedded guided wave SHM where memory is constrained. In a related development for composite structures, Fendzi et al. (Fendzi et al., 2016) employed Hilbert domain representations with linear parameter fitting to reconstruct baselines at new temperatures with small error, demonstrating that carefully chosen analytic descriptors can carry temperature dependence while leaving damage related content largely invariant. In practice, these synthesis methods provide strong compression of baseline libraries however retain admissible accuracy inside the envelope spanned by the anchors, although performance naturally depends on how well the learned mapping captures the true environmental dependence of the measured paths.

2.5.2.3 Regression and Latent Space Models

This category uses regression and latent variable modelling to predict and subtract EOV trends, often blending data driven inference with physics to improve extrapolation and quantify uncertainty. The working idea is to model how temperature, load, and other context variables reshape the healthy response, then remove that predicted component so that the residual emphasises damage. Amer et al. (Amer et al.,

2021) introduced a physics assisted Gaussian process regression (GPR) that merges parametric guided wave models with non-parametric learning; validated on an aluminium coupon across varying load and damage conditions, the framework compensated for load effects and provided probabilistic confidence on the corrected signal. Addressing field operation over long horizons, Qu et al. (Qu et al., 2025) employed GPR to track gradual environmental drift in an offshore WT structure undergoing structural evolution; the study showed that a multi-phase adaptive methodology can capture long-term changes in operational states and stabilise indicators in practice. Turning to vibration-based diagnostics, Font Moré et al. (Font-Moré et al., 2024) used PCA as a latent variable model to factor out environmental variability in power spectral density features measured on a small scale WT blade; by projecting to a low-dimensional subspace that is insensitive to temperature and operation, the approach isolated damage sensitive components more reliably. At full scale, Roberts et al. (Roberts et al., 2024) combined multivariate nonlinear regression with robust estimation to remove confounding operational factors in a Vestas WT blade case, thereby improving the signal to noise ratio of residual-based damage indicators under wide ranging EOV. Totally, these examples show that regression and latent modelling provide flexible and predictive compensation for mechanical systems, with physics assistance and probabilistic outputs enhancing trustworthiness; nevertheless, the models must be trained on data that adequately span the expected environmental and operational space to avoid biased residuals.

2.5.3 Reference-Free Approaches

Reference-free approaches remove the dependency on a stored pristine baseline by either learning self-referenced decision boundaries from data or exploiting invariants of wave physics that remain stable under EOV (Huynh et al., 2025). The common intention is to ensure that the indicator depends primarily on the presence of damage rather than on temperature, humidity, or duty cycle, which can be especially valuable for mechanical systems such as composite laminates, aluminium plates and shells, pipelines, aircraft components, and WT blades. These methods often trade storage and curation burdens for modelling effort and, in learning-based cases, for greater attention to interpretability and uncertainty.

2.5.3.1 Purely Data Driven Baseline

Purely data driven diagnostics construct features and classifiers that separate healthy and damaged responses while treating EOVS as a nuisance factor that the model learns to ignore. The underlying mechanism is representation learning. A network compresses the input to a latent space that preserves damage related structure and suppresses nuisance variation, or a nonparametric learner discovers discriminative patterns directly in the measurements. Because there is no reliance on a stored healthy trace, the approach is agnostic to baseline drift and can be deployed on devices with modest storage if inference is efficient (Zhao et al., 2025).

Kashyap et al. (Kashyap et al., 2024) proposed a convolutional autoencoder that operates on guided waves in honeycomb composite sandwich components and demonstrated that tiny ML deployment on an edge device can separate temperature effects from damage without access to a stored reference. Zhu et al. (Zhu et al., 2025) combined unsupervised shapelets with shift invariant dictionary learning on complex anisotropic composites and achieved localisation without baselines, illustrating how compact motifs stabilise against EOVS and still carry fault information.

2.5.3.2 Physics-Based Invariants and Instantaneous References

Physics-based reference-free methods build indicators from invariants of propagation, reciprocity, or material nonlinearity that hold across a range of environmental states. The mechanism is to work at frequencies, path combinations, or signal transforms where the environmental imprint is either weak or cancels out by construction, so that any change in the indicator originates predominantly in damage related scattering, stress, or modulation.

Time reversal is a leading example, Sharma et al. (Sharma et al., 2021) showed that certain excitations termed best reconstruction frequencies (BRFs) preserve time reversibility even as temperature varies in plate like aluminium, and Zhu et al. (Zhu et al., 2024a) examined related behaviour for composites, so locating and operating at the BRF yields a damage index that remains stable when velocity drifts due to temperature. Reference-free practice can also avoid history entirely by comparing simultaneously acquired paths. Anton et al. (Anton et al., 2009) introduced instantaneous baselines that rely on concurrent path comparisons in aluminium plates, while Mesnil et al. (Mesnil et al., 2023) developed truly self-referenced guided wave

detection for woven composite parts of complex shape, and Zhu et al. (Zhu et al., 2024b) appraised linear baseline-free techniques in anisotropic composites with an emphasis on reciprocity and careful path design. Beyond reversibility, Shi et al. (Shi et al., 2023) used acoustoelastic Lamb waves to infer stress without a stored baseline in aluminium components, Lee et al. (Lee et al., 2012) realised baseline-free ultrasonic monitoring in pipelines at temperatures up to 300 °C using optical fibre guided laser generation, Lim et al. (Lim et al., 2014) detected fatigue related sidebands through nonlinear ultrasonic modulation on aluminium and aircraft fittings, Qiu et al. (Qiu et al., 2019) compensated distance to neutralise temperature induced velocity changes in aluminium plates, Loshelder et al. (Loshelder et al., 2023) removed ambient temperature, pressure, and humidity fluctuations by means of apex shifted transforms across metal and polymer specimens, and Sohn et al. (Sohn et al., 2006) employed consecutive outlier analysis to stabilise condition screening in composites under changing operations. Collectively, these studies indicate that when geometry and sensing remain stable, physics driven indices can deliver robust detection without curating large libraries of baselines.

2.5.3.3 Physics-Informed Learning Hybrids

Hybrid methods encode governing physics inside a learning pipeline so that the model needs less data to generalise beyond trained conditions and can work without a stored baseline. The mechanism is to constrain the hypothesis space using equations, priors, or interpretable structures, and to combine these with flexible learners that capture residual complexity, which is particularly suitable for metals and beam like components where simple parameterisations of wave speed, attenuation, or stiffness already exist.

Miele et al. (Miele et al., 2023) developed multi-fidelity physics-informed learning for probabilistic damage diagnosis, which reduces baseline dependence by tying the latent representation to a mechanistic description and by propagating uncertainty to the final decision. Koune and Cicirello (Koune and Cicirello, 2024) proposed a physics-informed variational autoencoder with disentangled factors that separate environmental effects from damage related content on a Euler Bernoulli beam model, which is representative of many slender mechanical components. Amer and Kopsaftopoulos (Amer and Kopsaftopoulos, 2023) embedded structural damage

indices inside GPR to support baseline-free inference with probabilistic outputs that quantify confidence, a desirable property when past data are sparse or unrepresentative. El Mountassir et al. (El Mountassir et al., 2021) combined PCA with adaptive modelling to detect damage in highly attenuating pipeline media without baselines, illustrating the practical reach of hybrids in industrial piping. These results show that physics-informed architectures frequently outperform purely empirical schemes when extrapolating across EOV, while being more deployable than full order simulations.

2.5.3.4 Relative and Instantaneous Referencing and Building Block Validation

Relative referencing compares paths to each other rather than to an historical record, and instantaneous mapping searches a small path specific transformation that aligns measurements at the time of acquisition, both of which reduce global environmental drift without maintaining a large archive. Yue et al. (Yue et al., 2021) also advocate a building block philosophy in which compensation schemes are validated progressively from coupons to larger sub-assemblies to reduce recalibration overheads on complex airframe parts.

Lee et al. (Lee et al., 2011) introduced path to path relative baseline features for guided waves in composite and aluminium plates, thereby nullifying common environmental shifts at feature level; Salmanpour et al. (Salmanpour et al., 2017c) proposed an instantaneous baseline mapping that uses sensor mapping for real-time localisation when temperature varies; Yue et al. (Yue et al., 2021) set out a staged validation that moves from small coupons to fuselage scale elements, which streamlines transitions to complex assets and clarifies the limits of compensation at each scale.

2.5.3.5 Transfer Learning and Domain Adaptation for Mechanical Assets

The essential mechanism in such techniques is to align distributions or features between a source and a target so that damage related structure is preserved and nuisance variation is suppressed. This is well suited to mechanical fleets, for example WTs and aircraft components, where data are plentiful in aggregate, but labels are scarce for any single asset.

Population based SHM (PBSHM) shares information across nominally similar assets and adjusts to turbine to turbine or component to component variability through feature alignment and mixture modelling. The approach reduces the labelling burden and helps enforce consistency of indicators across a fleet by treating EOV as a distributional shift that can be compensated in the learned representation (Gardner et al., 2022).

Bull et al. (Bull et al., 2021b) reported PBSHM concepts that motivate knowledge sharing at population level; Bull et al. (Bull et al., 2021a) demonstrated transferable damage detectors between structures with rigorous experimental validation, which is a cornerstone for fleet deployment. In the presence of unknown or emerging fault classes, Zhou and Lai (Zhou and Lai, 2024) proposed an open set DA framework with time frequency modules that maintained classification performance under varying environmental conditions for WTBs. When the temperature of operation is unknown, Yang et al. (Yang et al., 2025b) introduced a multi-source dynamic adaptive generalisation network for composite crack detection that improved robustness to temperature shifts without requiring labels in the target data. Beyond blades and aerostructures, Michau and Fink (Michau and Fink, 2021) formulated unsupervised transfer for anomaly detection under complementary operating conditions, which captures the logic of transferring between operating regimes that do not share labels, a situation that occurs frequently in bearings and engines. Together these studies show that careful feature alignment and population aware learning are practical whenever structures are similar but not identical and labels are limited.

When experiments cannot cover severe EOV or when damage cases are rare, simulation augments data so that models learn to align simulated and real signals. The mechanism is either adversarial adaptation, where a discriminator encourages feature indistinguishability between simulation and experiment, or explicit statistical alignment in a common subspace. Physics constraints regularise the mapping so that the network respects conservation and material behaviour.

Sawant et al. (Sawant et al., 2023) used TL with unsupervised objectives to compensate temperature effects in guided wave monitoring of composites, demonstrating that synthetic data can improve generalisation when labels are scarce. Wang et al. (Wang et al., 2023) performed adversarial DA between FE simulation and

experiment for fatigue crack detection using Lamb waves in aluminium, reducing the gap sufficiently for reliable inference on real measurements. Da Silva et al. (Silva et al., 2021) employed transfer component analysis (TCA) to stabilise impedance-based diagnostics under temperature changes in aluminium beams, showing that subspace alignment can be as effective as adversarial strategies when the physics is relatively stationary. Physics-informed adversarial transfer has also been proposed for composites, where constraints enforce plausible mappings between synthetic and experimental domains for different damage modes (Liu et al., 2023), while attention enhanced CycleGANs have been used to align simulated and field signals under temperature changes so that transferred models respect the structure of the data (Ge and Sadhu, 2024). For liquid cooled plate like components, manifold alignment has been used to bridge differences between simulation and experiment under combined temperature and load variability, providing a path to robust transfer when coupled physics dominate the response (Liu et al., 2025). These advances improve data efficiency and extrapolation across EOv provided modelling assumptions remain credible and the difference between domains is not dominated by unmodelled boundary conditions.

2.6 Explainable Artificial Intelligence

While DL algorithms achieve impressive accuracy in SHM of mechanical systems, their decision mechanisms remain opaque, i.e., a limitation often called the black-box problem. XAI methods have recently been applied to structural damage classification from guided-wave and vibration sensor data (Luckey et al., 2022).

Conventional post-hoc attribution techniques, e.g., SHapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), Gradient-weighted Class Activation Mapping (Grad-CAM), and layer-wise relevance propagation are widely used to interpret deep models in this field (Rezazadeh et al., 2025d). Grad-CAM was adapted for one-dimensional vibration data by transforming time-series into spectrograms. A recent study converted bearing vibration signals to short-time Fourier transform spectrograms and used Grad-CAM to highlight frequency-time regions that most influenced fault classifications (Kim and Kim, 2024). The process involved inverse transformation to inspect highlighted frequencies and linked outlier frequencies to severity of damage, demonstrating that saliency maps

provide intuitive visual evidence of where the model is focusing. Based on this, Brito et al. (Brilo et al., 2023) utilised a TL approach for rotating machinery created synthetic vibration signals to pretrain a one-dimensional convolutional network and employed Grad-CAM to visualise which segments were responsible for each fault class.

SHAP and Lime quantify each feature's contribution to a prediction and were originally designed for tabular data. In mechanical fault diagnosis, SHAP has been used to turn physics-based descriptors into interpretable predictors. For instance, Yan et al. (Yan et al., 2024) treated fault characteristic frequencies as physics-informed features and used SHAP to identify those that most influenced bearing-fault decisions under varying conditions, thereby linking learned importance back to established vibration signatures. In a different application, Yao et al. (Yao et al., 2025) combined diffusion modelling for WT condition monitoring with a Shapley-value analysis to attribute anomalies and fault identification to specific variables, which provided transparent post-hoc justification for the model outputs.

Lime has likewise been adapted to vibration-based monitoring. Hanchate et al. (Hanchate et al., 2023) applied Lime in a vibration sensing scheme to highlight time–frequency regions that locally drove condition-monitoring decisions, which supported practitioner validation. Beyond component-level studies, Gawde et al. (Gawde et al., 2024) used Lime in an explainable predictive-maintenance strategy for multi-fault equipment, where the local surrogate analyses were reported to clarify which signal attributes governed each predicted fault state.

In addition to post-hoc attribution, some researchers incorporate interpretability into the model architecture. Sparse Temporal Logic Network (STLN) is an ante-hoc architecture that conceptualises neurons as logical propositions. It employs a word network with wavelet kernels to extract interpretable features, a transformer encoder with sparse attention to locate informative segments and a logic network that synthesises a formal temporal logic description (Chen and Dong, 2025). Experiments on bearing data showed that STLN achieved competitive accuracy while producing logical rules describing the decision process, demonstrating the feasibility of combining signal processing, attention mechanisms and temporal logic for transparent diagnosis.

The Variational Attention-based Transformer Network is another interpretable architecture designed for rotary machine fault diagnosis (Li et al., 2024). It starts from a transformer encoder but introduces a sparse Dirichlet prior on attention weights to force the network to highlight a few critical segments of the vibration signal. Variational inference and Laplace approximation are used to sample attention weights, and heat maps of the resulting weights illustrate causal associations between signal patterns and fault types. This approach addresses the fact that standard transformers capture associations across entire sequences but do not reveal which patterns drive classification.

A multi-scale CNN combined with a transformer encoder has been developed to diagnose reciprocating piston pump faults under noisy underground conditions. The network uses a multi-scale CNN encoder to capture local features and a transformer encoder to capture global context; patch segmentation aligned with the crankshaft phase reduces random noise and improves the interpretability of global features (Lai et al., 2025). A feature-fusion module enhances interactions between local and global features, enabling the model to achieve high diagnostic accuracy while revealing which signal segments contribute to predictions.

Other work integrates interpretability techniques into CNN variants. An interpretable quadratic convolutional residual network combines quadratic neurons with residual connections to capture bearing fault features and applies layer-wise relevance propagation to produce attention maps (Luo et al., 2025). Visualising contributions of input signals makes the diagnostic process transparent and helps users understand why particular vibration patterns indicate specific faults.

Another aspect of changing the black-box nature of DL methods is revealing the internal activities of DA frameworks to show how a pre-trained model is adjusting to work well in the new domain where training or labelled data is scarce. Rezazadeh et al. (Rezazadeh et al., 2024a), presented the activation patterns in terms of heat maps in dense layers of a fine-tuning framework employed for the DA and the subsequent fault diagnosis of rotor systems.

Despite these advances, limitations of existing explainable artificial intelligence techniques remained evident for time-series signals from mechanical systems. Feature-attribution methods were typically predicated on the independence of input features

and on a stable mapping to physical quantities, whereas vibration and guided-wave representations were highly correlated and deep models often operated in latent spaces without direct physical meaning. Standard transformer architectures captured associations across long sequences, but the drivers of classification were not exposed unless additional sparsity or structural constraints were imposed. Most importantly, conventional explainable artificial intelligence tools focused on outputs of trained models and did not reveal how internal representations evolved during DA or TL. Insight into changes such as instance weighting, sensor-level sensitivities and movements of class prototypes between source and target domains was required for domain-adaptive diagnosis, although no standard approach made these internal dynamics visible. The limited attempts to impose interpretability through auxiliary modelling, including contrastive or equation-derived formulations, were task specific and did not generalise across architectures, which left a clear gap for techniques that expose intermediate representational changes in a user-readable form.

2.7 Summary

The chapter opened with the motivation that labelled fault data are scarce, and feature distributions differ between assets, and a hybrid narrative–systematic strategy was stated in which feature coverage is selective while interpretability is treated in depth.

Signal characteristics were outlined as non-stationary and multiscale, with desired properties for features set as translation and scale invariance, and either permutation invariance to sensor ordering or graph equivariance to array topology, so that damage information is retained while nuisance variability is suppressed.

Within feature extraction, 3 families were profiled. Classical time, frequency and time–frequency descriptors were noted as computationally light whilst brittle under EOV. WTS was positioned as a physics–informed, low–variance front-end offering deformation stability and translation invariance. Manifold and embedding methods, including PCA and non–linear embeddings, were highlighted for visualisation and dimensionality reduction, with inductive or parametric variants preferred for deployment. Neural feature learners were summarised as follows: CNNs for efficient local pattern capture, CapsNets for preserving configuration information through

vector-valued representations, and GATs for weighting actuator-sensor paths in sparse arrays.

Data augmentation and dataset completion were reviewed across time, spectral and latent domains. Classical time-domain operations were presented as useful for diversity and regularisation; frequency and time-frequency masking as helpful against nuisance components, with a caution about mode distortion; synthetic oversampling as effective when sequence-aware or latent-space interpolation is used; cross-condition interpolation as practical when physics constraints are respected; and deep generative models as capable of synthesising condition-aware signals when guided by physics, provided tuning avoids artefacts. A validation protocol for synthetic data was outlined, including statistical similarity tests, temporal alignment measures and checks that labels are preserved.

XAI methods were then surveyed; post-hoc attribution, ante-hoc and hybrid designs have been applied to guided-wave and vibration data to improve transparency, though challenges remain with correlated inputs, latent representations without direct physical meaning, and limited diagnostics of representation change during adaptation.

Stepping back, these sections establish the signal properties, feature choices, augmentation strategies and interpretability tools required for robust SHM under EOV, and they provide the rationale for the front ends and validation procedures adopted later in the thesis.

Chapter 3: Methodology and Methods

3.1 Overview

This chapter presents the methodological foundations for the damage detection frameworks developed in this research. The chapter comprises four principal sections: Data Synthesising, UCTRF Framework, GAT-CAMDA Framework, and SPADA Framework.

The methodological development is guided by three requirements: first, the capability to learn from severely limited labelled data, particularly for damaged structural states; second, transferability across domains such that models trained under one set of conditions generalise to unseen or minimally labelled target conditions; third, interpretable outputs enabling practitioners to identify which sensors or features contribute to diagnostic decisions.

Each framework addresses the DA challenge through different paradigms. UCTRF employs instance-based adaptation with minimal target supervision to reweight training samples for cross-domain knowledge transfer. GAT-CAMDA implements a semi-supervised through adversarial training and feature distribution alignment, requiring no labelled target data during adaptation. SPADA operates in the fully UDA regime, leveraging pseudo-labelling with confidence gating to adapt without any target domain labels.

Each framework is presented with comprehensive mathematical formulations and algorithmic specifications. All notation is standardised throughout the chapter, with symbols defined upon first use and maintained consistently. The mathematical exposition for UCTRF includes UMAP fuzzy membership formulation, CapsNet squashing and routing mechanisms, TrAdaBoost weight updates, and Random Forest aggregation. For GAT-CAMDA, the formulations encompass GAT attention computation, MMD kernel-based measures, CORAL covariance alignment, adversarial training objectives, Optuna optimisation, and sensor importance quantification. For SPADA, the development covers wavelet scattering coefficients,

DANN architectures, prototype construction and updating, confidence gating, attention weighting, multi-component loss functions, and interpretability metrics.

This chapter focuses exclusively on methodological exposition. Empirical validation, performance benchmarking, ablation studies, comparative analyses, hyperparameter sensitivity investigations, implementation details (learning rates, batch sizes, convergence criteria), and hardware specifications are documented in the Results and Discussion chapter.

3.2 Data Augmentation

Four distinct data augmentation approaches are explored in this research, each offering different theoretical foundations and operational characteristics suited to specific aspects of the SHM problem. Spline interpolation provides a deterministic, closed-form solution for generating temperature-dependent responses by modelling smooth functional relationships between known temperature points, making it suitable when the underlying physical behaviour exhibits continuous, differentiable temperature dependence. Conv-CVAE represents a probabilistic deep generative approach that learns complex, potentially nonlinear distributions in latent space, enabling the synthesis of realistic signal variations when explicit functional forms are unknown or intractable. SMOTE addresses class imbalance through feature-space interpolation between minority class samples, operating independently of physical or temporal structure. Signal windowing enables data augmentation through segmentation of continuous acquisitions into nonoverlapping fixed-length observations, increasing effective sample size whilst maintaining temporal coherence within windows.

The selection of which methods to apply is determined by the specific characteristics of each experimental scenario, as detailed in the Results and Discussion chapter. The following subsections detail spline interpolation and Conv-CVAE, which required substantial development or adaptation for SHM applications. SMOTE is cited with reference to its original publication as well as the details of its implementation are discussed in the Results and Discussion chapter; similarly signal windowing is discussed in the same chapter.

3.2.1 Spline Interpolation

Spline interpolation was adopted as a deterministic data synthesising method to model the temperature-dependent behaviour of healthy and damaged structural signals. It enables the generation of intermediate responses by estimating smooth transitions between known temperature points, thereby preserving the physical continuity of structural behaviour. The method effectively mitigates discontinuities caused by abrupt temperature variations and maintains the physical plausibility of the synthesised data. Furthermore, it captures complex non-linear temperature–response relationships that frequently arise in SHM applications. Its adaptability to irregular temperature intervals and resilience to noise make it particularly suitable for datasets with limited observations. Interpolating healthy baseline responses and appropriately scaling the damaged signals, the method aims to promote coherence between datasets acquired under different operational conditions. This forms a simple at the same time reliable foundation for subsequent ML and DA stages, where data balance and representativeness are essential for model generalisation. The technical formulation of this approach is described below.

In this approach, first, the mean healthy response at temperature T , denoted $\bar{x}_{\text{healthy}}(T)$, was calculated using spline interpolation; a scaling factor was then computed as:

$$F_{\text{scal}}(T) = \frac{\bar{x}_{\text{healthy}}(T)}{\bar{x}_{\text{healthy}}(T_{\text{Baseline}})} \quad (1)$$

Damaged signals at temperature T were generated by scaling the baseline damaged responses (F_{scal}) and adding small Gaussian noise (ϵ):

$$\bar{x}_{\text{damaged}}(T) = F_{\text{scal}}(T) \cdot x_{\text{damaged}}(T_{\text{Baseline}}) + \epsilon \quad (2)$$

This approach enables the extrapolation of damaged behaviours across temperature ranges while preserving realistic variability and statistical coherence with measured data. It maintains the relative amplitude and temporal characteristics of the original responses, facilitating that the synthesised signals remain physically meaningful and analytically consistent for subsequent modelling and evaluation.

3.2.2 Convolutional Conditional Variational Autoencoder

The second technique called Conv-CVAE was developed to synthesise temperature-dependent damaged guided-wave signals. This deep generative framework integrates convolutional feature extraction with conditional variational inference to model the joint distribution of signals and their temperature conditions. Leveraging latent-space offsets and amplitude masks, it produces realistic signal variations that reflect the influence of temperature and structural damage. The approach provides a data-driven mechanism for extending limited experimental datasets, enabling the generation of plausible synthetic responses across a wide range of operational environments.

Let $\mathcal{D}_s = \{(x^{(n)}, y^{(n)})\}_{n=1}^{N_s}$ and $\mathcal{D}_t = \{\tilde{x}^{(m)}\}_{m=1}^{N_t}$ represent the labelled source and unlabelled target datasets, with $x^{(n)} \in \mathbb{R}^{N_{time} \times C}$ where C is the number of sensors. For temperature-dependent synthesis, the baseline condition at T_{Baseline} was used as the reference. The Conv-CVAE encoder maps an input signal $x^{(n)}$ and its associated temperature T into a latent representation z :

$$q_\phi(z|x^{(n)}, T) = \mathcal{N}(\mu_\phi(x^{(n)}, T), \sigma_\phi^2(x^{(n)}, T)I) \quad (3)$$

where μ_ϕ and σ_ϕ are functions parameterised by convolutional layers and stand for the mean vector and variance of the latent distribution (per dimension), respectively; $q_\phi(\cdot)$ is posterior distribution (encoder only), $\mathcal{N}(\cdot)$ is the Gaussian distribution, and I represents the identity matrix that makes the covariance matrix a diagonal. The decoder reconstructs signals from latent variables and temperature conditions:

$$\hat{x}^{(n)} = f_\theta(z, T) \quad (4)$$

where $f_\theta(\cdot)$ is the decoder. For each damage level l , a latent offset Δz_l is computed relative to the healthy baseline at T_{Baseline} :

$$\Delta z_l = \frac{1}{N} \sum_{i=1}^N z_{l,i}(T_{\text{Baseline}}) - \frac{1}{N} \sum_{i=1}^N z_{\text{healthy},i}(T_{\text{Baseline}}) \quad (5)$$

This offset characterises the displacement in the latent space that arises from structural damage and serves as a transferable feature across temperature conditions. To further enhance the realism of the generated signals, per-channel and per-time

amplitude masks were derived from the envelope ratios between damaged and healthy signals at T_{Baseline} . These masks introduce both deterministic scaling (through the mean of the mask) and controlled stochasticity (through the mask’s standard deviation combined with random jitter), thereby improving the variability and fidelity of the synthesised samples.

The training objective of the Conv-CVAE combines multiple reconstruction and regularisation criteria. The total loss includes mean squared error (MSE) for temporal-domain accuracy, spectral loss in the Fourier domain to preserve frequency content, wavelet-domain similarity for capturing multi-resolution structures, and amplitude envelope consistency to maintain realistic modulation patterns. The Kullback-Leibler (KL) divergence term regularises the latent distribution, encouraging smooth transitions and preventing overfitting to specific temperature conditions. The final synthesised damaged signals at temperature T are expressed as:

$$\bar{x}_{\text{damaged}}(T) = f_{\theta}(z_{\text{healthy}}(T) + \Delta z_l, T) \odot \mathcal{M}_l + \epsilon \quad (6)$$

in which \mathcal{M}_l is the learned amplitude mask for damage level l , \odot denotes elementwise multiplication, and ϵ is a small Gaussian noise term for stochasticity.

This method enables the synthesis of realistic temperature-dependent damaged signals using only baseline damaged data obtained at T_{Baseline} . Modelling latent-space transformations and temperature-conditioned decoding, the Conv-CVAE generates statistically coherent and physically consistent responses across different thermal conditions. The resulting dataset provides an enriched basis for the development and validation of fault detection algorithms, improving both robustness and generalisation under EOVs.

3.3 Damage Detection Pipelines

In this section, the three designed hybrid damage detection pipelines called UCTRF, GAT-CAMDA, and SPADA within their components are elaborated separately.

3.3.1 UCTRF Framework

UCTRF addresses DA scenarios characterised by severe data scarcity and distribution shift between training (source) and deployment (target) conditions. The

framework integrates 4 components in a sequential pipeline: UMAP reduces dimensionality whilst preserving manifold structure, CapsNet extracts hierarchical features through vector-valued capsules, TrAdaBoost adapts features across domains through instance reweighting with limited target supervision, and Random Forest performs final classification with ensemble robustness.

UMAP compresses high-dimensional signals into low-dimensional representations, reducing computational burden whilst maintaining topological structure. CapsNet, with its classification layer removed, functions as a feature encoder producing structured capsule vectors that capture part-whole relationships in vibration signals. TrAdaBoost then adjusts these features for domain shift by iteratively reweighting samples, reducing weights of misleading source instances whilst emphasising those that align with target data patterns. Random Forest provides the final discriminative mapping, leveraging ensemble voting for robust classification. Unlike conventional supervised approaches requiring extensive labelled data in both domains, UCTRF operates effectively with severely limited target labels by leveraging source domain knowledge through weighted instance transfer. An overview of the framework using a WTB as a case study is presented in Figure 2.




Figure 2. A schematic of UCTRF for damage detection

3.3.1.1 Main Contribution

As a unified SHM framework for mechanical systems, UCTRF is designed, in which UMAP and CapsNet feature extraction, together with TrAdaBoost and random forest, are combined to address data scarcity and temperature-driven EOVs.

The effectiveness of the proposed framework is reflected in multiple key contributions:

- Robust temperature adaptability is demonstrated, with consistent damage detection performance maintained under varying thermal conditions.
- Domain adaptation under limited data availability is enabled through the integration of advanced dimensionality reduction, hierarchical feature extraction, and transfer learning methods.
- Enhanced feature interpretability is achieved, as UMAP facilitates the separation of complex vibration signals into distinct and physically meaningful clusters, supporting stable adaptation through TrAdaBoost.
- Superior detection accuracy is reported on temperature-affected datasets, surpassing the performance of the alternative techniques.

The novelty of UCTRF is attributed to its purposeful integration of UMAP for structure-preserving dimensionality reduction, CapsNet for capturing spatial hierarchies in the data, and TrAdaBoost with random forest for adaptive reweighting in DA. Through this combination, longstanding challenges related to distribution shifts and data insufficiency are effectively addressed, and improved generalisability to unseen temperature conditions is attained. This framework constitutes a substantial advancement in supervised DA within the field of SHM.

The following subsections detail each component: UMAP's cross-entropy optimisation in fuzzy topological space, the dynamic routing algorithm of CapsNet, the iterative weight update procedure for TrAdaBoost, random forest's ensemble decision mechanism, and weight adjustments representation.

3.3.1.2 Uniform Manifold Approximation and Projection

Dimensionality reduction aids in identifying informative features within vibration data, improving both interpretability and computational efficiency. UMAP is a non-linear DR method that maps complex, high-dimensional signals into a simpler, low-dimensional Euclidean space. This transformation preserves both local and global topological structures, facilitating the analysis and visualisation of data patterns.

Grounded in manifold learning theory, UMAP assumes that data points reside on locally uniform manifolds embedded in the high-dimensional space.

The UMAP algorithm comprises two fundamental stages: (1) constructing a fuzzy topological representation of the data in the original high-dimensional space, and (2) optimising a low-dimensional embedding that faithfully preserves this topological structure. Let \mathbf{X} represent the vibration input data in a D_{orig} -dimensional space, which is to be embedded into \mathbf{U} in D_{emb} dimensions. The algorithm requires specification of a distance metric $d(x_i, x_j)$ between any two points, where $i, j = 1, 2, \dots, n$, and the number of nearest neighbours K_{nn} .

The first stage constructs a weighted graph representing the topological structure of the high-dimensional data. For each point x_i , UMAP identifies its K_{nn} nearest neighbours, identified as the set N_i , using the specified distance metric $d(x_i, x_j)$. To establish local connectivity while adapting to varying data densities, two critical parameters are computed for each point.

The first parameter, ρ_i , represents the distance to the nearest neighbour and ensures that each point maintains at least one strong connection within the manifold:

$$\rho_i = \min_{j \in N_i} \{d(x_i, x_j) | d(x_i, x_j) > 0\}, \quad 1 \leq j \leq K_{nn} \quad (7)$$

The second parameter, ϱ_i , is a smoothing bandwidth that controls the rate at which similarity decays with distance beyond ρ_i . This parameter is determined through a binary search procedure to satisfy the following constraint:

$$\log_2 K_{nn} = \sum_{j \in N_i} \exp \left[-\frac{\max(0, d(x_i, x_j) - \rho_i)}{\varrho_i} \right] \quad (8)$$

$$1 \leq j \leq K_{nn}$$

$$1 \leq i \leq n$$

This constraint facilitates that the sum of fuzzy memberships within each neighbourhood approximates the expected perplexity, thereby achieving a target effective number of neighbours. Using these parameters, UMAP defines asymmetric conditional fuzzy memberships $p_j | i$ that quantify the probability that point x_j belongs to the local neighbourhood of x_i :

$$p_{j|i} = \begin{cases} \exp \left[-\frac{\max(0, d(x_i, x_j) - \rho_i)}{q_i} \right] & , j \neq i \\ 0 & , j = i \end{cases} \quad (9)$$

To obtain a symmetric representation suitable for graph construction, these asymmetric memberships are combined using a probabilistic union operation:

$$p_{ij} = p_{j|i} + p_{i|j} - p_{j|i} p_{i|j} \quad (10)$$

By construction, $p_{ij} = p_{ji}$ (symmetric) and $p_{ii} = 0$. These symmetric fuzzy memberships define a weighted graph that encapsulates the neighbourhood relationships in the high-dimensional space.

The second stage aims to construct a low-dimensional representation that keeps the topological structure encoded in the high-dimensional fuzzy graph. Let u_i denote the embedded coordinates of point x_i in the low-dimensional space U . UMAP defines analogous conditional fuzzy memberships $p_{j|i}$ in the embedding space using a smooth approximation to a step function:

$$\frac{1}{[1 + \alpha \|u_i - u_j\|^{2\beta_{kern}}]} \approx \begin{cases} 1 & , \quad \|u_i - u_j\| \leq \Gamma \\ \exp[-(\|u_i - u_j\|) - \Gamma] & , \quad \|u_i - u_j\| > \Gamma \end{cases} \quad (11)$$

Symmetrised fuzzy memberships in the embedding space are then obtained through the same probabilistic union:

$$q_{ij} = q_{j|i} + q_{i|j} - q_{j|i} q_{i|j} \quad (12)$$

UMAP refines the embedding by minimising the cross-entropy between fuzzy set representations in the high-dimensional and low-dimensional spaces:

$$CE = \sum_{i \neq j} \left[p_{ij} \log \left(\frac{p_{ij}}{q_{ij}} \right) - (1 - p_{ij}) \log \left(\frac{1 - p_{ij}}{1 - q_{ij}} \right) \right] \quad (13)$$

This objective function reconciles 2 competing effects: the first term promotes proximity in the embedding for points that are close in the high-dimensional space, whereas the second term discourages dissimilar points from being positioned too closely (repulsive force). The gradient of the cross-entropy with respect to the embedded coordinates is:

$$\frac{\partial CE}{\partial u_i} = \sum_j \left[\frac{2\alpha\beta_{kern}[d(u_i, u_j)]^{2(\beta_{kern}-1)}}{1 + \alpha[d(u_i, u_j)]^{2\beta_{kern}}} p_{ij} - \frac{2\beta_{kern}}{[d(u_i, u_j)]^2 (1 + \alpha[d(u_i, u_j)]^{2\beta_{kern}})} (1 - p_{ij}) \right] (u_i, u_j) \quad (14)$$

This gradient is employed in stochastic gradient descent with negative sampling to efficiently optimise the embedding coordinates. Upon convergence, UMAP yields a low-dimensional embedding U_{emb} for each original signal X_n :

$$\{U_{\text{emb}}\} = f(X_n; \theta_{\text{hyper}}) \quad (15)$$

where $f(\cdot)$ denotes the UMAP transformation function and θ_{hyper} represents the set of hyperparameters governing the algorithm's behaviour. Key hyperparameters include the distance metric (e.g., Euclidean, Chebyshev, or cosine), the minimum distance between embedded points (`min_dist`), and the number of nearest neighbours K_{nn} . These parameters jointly determine the balance between preserving local neighbourhood structures and maintaining the global manifold topology, thereby controlling the fidelity and interpretability of the resulting low-dimensional representation.

3.3.1.3 Capsule Neural Networks

The final stage of the feature extraction module in UCTRF employs CapsNets, introduced by Sabour *et al.* (Sabour et al., 2017). The architecture begins with convolutional layers followed by normalisation and activation (ReLU) to extract low-level features from vibration signals. These outputs are then forwarded to a primary capsule layer comprising multiple capsules, which convert scalar outputs into vector representations to capture richer relationships among signal components.

CapsNets also incorporate a digit capsule layer that encodes higher-level entities and their attributes, as shown in Figure 3. First, the feature vector for each observation is fed into the convolutional layers with dimensions $1 \times D_{\text{emb}}$, where D_{emb} denotes the reduced dimension produced by the UMAP stage. After convolution, the resulting activation matrix has dimensions $N_{\text{cft}} \times D_{\text{emb}}$, where N_{cft} denotes the number of convolutional filters. The primary capsule layer then maps this matrix to an output of size $N_{\text{prim}} \times D_{\text{cap}}$, with each of the N_{prim} primary capsules producing a D_{cap} -

dimensional vector. Subsequently, the digit capsules, supported by dynamic routing in the Agreement Routing module, aggregate these lower-level features into more discriminative representations associated with specific conditions or faults. The digit capsule output has dimensions $K \times L_{vec}$, where K denotes the number of classes and L_{vec} denotes the output vector length.




Figure 3. A schematic of CapsNets as the feature extraction

CapsNets employ a non-linear squashing function to normalise capsule output magnitudes while preserving their directions:

$$F_j = \frac{\|s_j\|^2}{1 + \|s_j\|^2} \cdot \frac{s_j}{\|s_j\|} \quad (16)$$

where s_j is the input vector to capsule j , $\|s_j\|$ its Euclidean norm, and F_j the output activity vector. This facilitates that vector lengths remain within $[0, 1]$, reflecting the probability of feature presence.

Rather than using max-pooling, CapsNets apply a dynamic routing algorithm that iteratively updates the coupling coefficients between lower-level and higher-level capsules based on agreement. The coefficients are computed as:

$$c_{ij} = \frac{\exp(b_{ij})}{\sum_k \exp(b_{ik})} \quad (17)$$

where b_{ij} is the log prior probability that capsule i should be linked to capsule j . Predicted outputs $\hat{F}_{i|j} = w_{ij}F_i$ are aggregated as $s_j = \sum_i c_{ij}\hat{F}_{i|j}$, and coefficients are refined over successive routing iterations based on equations 15 and 16. The complete routing mechanism and its iterative refinement process are illustrated in Figure 4.

CapsNet training typically employs a margin loss function controlled by three parameters: m_{plus} , m_{minus} , and λ_{loss} . m_{plus} (approximately 0.9) defines the minimum

acceptable length for correctly classified vectors, m_{minus} (near to 0.1) defines the threshold for incorrect ones, and λ_{loss} adjusts the weighting of misclassified samples. Within the UCTRF framework, the final softmax classification layer is omitted, allowing the digit capsule outputs to be used directly as feature representations for each observation.




Figure 4. Schematic overview of the dynamic routing process

3.3.1.4 Transfer Adaptive Boosting

TrAdaBoost is implemented to mitigate the effects of EOVs by enabling effective knowledge transfer between source and target domains when labelled target data are scarce. As an extension of AdaBoost, it employs a reverse boosting strategy that iteratively decreases the weights of misclassified source samples while increasing those of correctly classified target samples, enhancing cross-domain generalisation (Dai et al., 2007).

The TrAdaBoost algorithm proceeds through the following steps:

- (a) Standardise the importance weights for the source and target domains:

$$\sum w_S + w_T = 1 \quad (18)$$

where w_S and w_T denote the weight values associated with the source and target domains, respectively.

- (b) Fit a classifier, identified by g , to the labelled samples from the source and target domains (F_S, y_S) and (F_T, y_T) , incorporating the corresponding importance weights w_S and w_T .

(c) Evaluate the error vectors for the training samples:

$$\begin{aligned}\epsilon_S &= L_{01}(g(F_S), y_S) \\ \epsilon_T &= L_{01}(g(F_T), y_T)\end{aligned}\tag{19}$$

where L_{01} refers to the 0-1 loss function. When $g(F_S)$ assigns the correct label to a sample $\epsilon_S = 0$; otherwise, $\epsilon_S = 1$.

(d) Compute the weighted error over the target samples:

$$E_T = \frac{w_T^T \epsilon_T}{m_T}\tag{20}$$

(e) Update the weights of source and target instances:

$$\begin{aligned}w_S &= w_S \delta_S^{\epsilon_S}, & \delta_S &= \left(1 + \sqrt{\frac{2 \ln m_S}{N_{iter}}}\right)^{-1} \\ w_T &= w_T \delta_T^{-\epsilon_T}, & \delta_T &= E_T / (1 - E_T)\end{aligned}\tag{21}$$

where m_S and m_T correspond to the sample counts in the source and target domains, respectively.

(f) Repeat the procedure from step (a) until the specified number of boosting iterations N_{iter} is completed.

The final output is derived from the most recent $N_{iter}/2$ estimators, each contributing to the ultimate prediction in accordance with its weight factor δ_T ; this encourages that each estimator's impact on the final outcome reflects its relevance and reliability.

A key advantage of TrAdaBoost is its classifier-agnostic nature, meaning that it can be integrated with a broad range of classification algorithms, including ridge classifiers, random forests, support vector machines, neural networks, and others. This flexibility enables the choice of a base classifier that is best suited to the specific task (He et al., 2020).

3.3.1.5 Random Forest

The classification stage constitutes the final step of the UCTRF framework. In this study, a random forest classifier was adopted. Random forest is an ensemble learning approach that is widely used for classification and regression, and it is also

applied to feature selection and anomaly detection. During training, the method constructs multiple decision trees. For classification, the predicted label is determined either by majority voting or by averaging the predictions across the trees. This ensemble structure reduces the likelihood of overfitting and supports effective modelling of diverse data types. In addition, the method is resilient to noise and outliers, which strengthens its generalisation performance (Breiman, 2001; Srivastava et al., 2022).

When implementing a random forest, several hyperparameters must be specified, including the maximum tree depth, the minimum number of samples required to split a node, the minimum number of samples required at a leaf node, and the maximum number of features considered at each split. One of the most influential hyperparameters is the number of estimators, which determines how many decision trees are included in the ensemble. This setting affects prediction diversity and stability; although increasing it can improve accuracy, it also increases computational cost.

Figure 5 presents the procedure for performing DA using the TrAdaBoost algorithm, with random forest employed as the base classifier.




Figure 5. Overview of the DA procedure using TrAdaBoost and a random forest classifier

3.3.1.6 Weight Adjustments Representation

The iterative reweighting process of TrAdaBoost provides insight into the adaptation dynamics by tracking how instance importance evolves throughout training. At the initial stage, equal weights are assigned to all instances from both the source and target domains. As adaptation progresses through successive boosting

iterations, these weights are updated according to classification performance: source domain instances resembling target domain patterns receive increased weights, whilst those dissimilar to the target domain are down-weighted; conversely, target domain instances that prove difficult to classify correctly receive increased emphasis, whereas those correctly classified early in adaptation see their weights reduced.

To visualise this dynamic process, weight values at specific epochs (initial, intermediate, and final) are extracted and plotted as scatter diagrams in two-dimensional feature space. Feature dimensions exhibiting substantial overlap between source and target domain observations are deliberately selected for visualisation, as these regions most clearly illustrate the adaptation mechanism. To enhance visual interpretability, raw weight values are scaled logarithmically: initial weights are mapped to moderate scatter point sizes, with all other weights scaled proportionally, whilst minimum weights approaching zero are clipped to a floor value to maintain visibility. This scaling transforms weight magnitudes into visually distinguishable scatter point sizes, enabling qualitative assessment of instance importance evolution. The resulting visualisations reveal which source domain instances contribute constructively to target domain classification (indicated by sustained or increased point sizes) and which target domain instances remain challenging throughout adaptation (indicated by growing point sizes), thereby providing interpretability into the instance-based transfer learning process.

In Algorithm 3 a pseudocode for UCTRF was presented.

Algorithm 1. UCTRF Framework for SHM

Input: Source-domain data and labels, and target-domain data and labels.

Output: Trained UCTRF model, instance weight dynamics, damage classification performance.

1. Configuration and Setup

1.1 Set random seeds for reproducibility.

1.2 Configure GPU device if available, otherwise utilise CPU.

1.3 Define global hyperparameters for UMAP, capsule network, and training process.

2. Data Preprocessing

2.1. Load source and target data and convert to appropriate numerical formats.

2.2. Apply UMAP dimensionality reduction: fit on combined data, transform both domains to a low-dimensional space.

2.3. Reshape data for capsule network input as sequential features.

3. Capsule Network Architecture

-
- 3.1 Define initial convolutional layer for low-level feature extraction.
 - 3.2 Construct primary capsule layer to generate multiple capsule types representing different feature aspects.
 - 3.3 Implement digit capsule layer with dynamic routing to iteratively refine capsule agreements.
 - 3.4 Define margin loss function to penalise incorrect classifications whilst down-weighting absent classes.
- 4. CapsNet Training on Source Domain**
- 4.1 Initialise Adam optimiser and train capsule network on source domain using margin loss.
 - 4.2 Iterate over mini-batches: forward propagation, loss computation, backpropagation, parameter updates.
- 5. Feature Extraction**
- 5.1 Extract hierarchical capsule representations from both domains by passing data through trained network.
 - 5.2 Partition target features into adaptation and testing subsets using stratified splitting.
- 6. TrAdaBoost Domain Adaptation**
- 6.1 Initialise TrAdaBoost ensemble with Random Forest base learners and uniform instance weights.
 - 6.2 For each boosting iteration:
 - 6.2.1 Train base classifier on combined weighted data and evaluate on target adaptation set.
 - 6.2.2 Update weights: reduce for correctly classified source instances, increase for misclassified target instances.
 - 6.2.3 Store classifier with confidence weight and record instance weight trajectory.
 - 6.3 Generate predictions via weighted majority voting using latter half of ensemble.
- 7. Model Evaluation**
- 7.1 Apply trained model to predict labels on held-out target test set.
 - 7.2 Compute confusion matrix, overall accuracy, and generate heatmap visualisation.
- 8. Interpretability Analysis**
- 8.1 Visualise instance weight evolution for source and target domains over boosting iterations.
 - 8.2 Export complete weight histories to CSV files for detailed analysis.
 - 8.3 Interpret patterns: decreasing source weights indicate negative transfer, increasing target weights indicate challenging instances.
- 9. Final Outputs**
- 9.1 Trained UCTRF framework comprising UMAP reducer, capsule network extractor, and TrAdaBoost ensemble.
 - 9.2 Classification metrics including accuracy, confusion matrix, and per-class performance statistics.
 - 9.3 Interpretability visualisations showing weight dynamics and complete history records.
 - 9.4 Insights into domain adaptation effectiveness and hierarchical capsule feature learning.
-

3.3.2 GAT-CAMDA Framework

GAT-CAMDA advances a semi-supervised DA by explicitly modelling inter-sensor relationships through graph-structured representations and by simultaneously employing multiple complementary domain alignment strategies, without the need for labelled target data during the DA stage. In this framework, the labelled validation phase of the target domain was used for optimal model selection, which is why the pipeline is referred to as semi-supervised. This pipeline aligns distributions through adversarial learning, mean embedding matching, and covariance alignment operating concurrently in a label-free manner.

The graph-based formulation treats each sensor as a node, with attention mechanisms learning edge weights dynamically based on neighbour importance. Multi-head attention captures information across representation subspaces, enhancing robustness. Domain alignment combines three strategies operating without target labels: Gradient Reversal Layer for adversarial invariance (forcing the feature extractor to produce representations that cannot distinguish source from target), MMD for mean alignment in reproducing kernel Hilbert space (minimising distributional divergence through kernel embeddings), and CORAL for second-order statistics matching (aligning covariance matrices). This hybrid approach addresses distributional divergence more comprehensively than single-strategy methods, enabling adaptation without any labelled target observations.

The semi-supervised nature of GAT-CAMDA makes it particularly suited for scenarios where target domain labels are partially unavailable that is common in real-world deployments where structures operate under new environmental conditions and labelling damage states requires expensive inspections or destructive testing. Eliminating the need for target labels entirely, GAT-CAMDA reduces deployment barriers and enables continuous adaptation as environmental conditions evolve.

Additionally, automated hyperparameter optimisation via Optuna with Tree-structured Parzen Estimators identifies optimal configurations efficiently. Sensor importance quantification leverages learnt attention weights to identify critical sensor contributions.

3.3.2.1 Main Contribution

This study presents a novel graph-based semi-supervised DA framework, GAT-CAMDA, to mitigate key challenges in the SHM of composite plates under temperature-induced EOVs. The main contributions are summarised as follows:

- Improved generalisability is achieved through temperature domain merging, whereby multiple temperature domains are combined into a single target domain. This strategy better represents real-world operational variability and improves the robustness and practical applicability of the SHM process.
- Cross-temperature feature alignment is facilitated through the integration of CORAL and MMD losses, enabling feature distributions to be aligned across varying thermal conditions. The effectiveness of CORAL is examined within the SHM context, while GATs are utilised to model complex spatial-temporal dependencies present in ultrasonic guided wave data, which contributes to improved damage detection accuracy.
- Model transparency and sensor relevance analysis are supported by the use of GAT attention weights, which allow for the visualisation and quantification of the contributions of individual sensors. This improves interpretability and informs sensor placement and analysis in both research and applied SHM scenarios.
- Robust semi-supervised DA with unlabelled target training data is achieved by using labelled target validation data for model selection, while adaptation during training relies entirely on unlabelled data. Hyperparameter tuning through Optuna results in a generalisable solution that outperforms existing methods under cross-domain temperature variations.

The originality of this pipeline stems from its ability to generalise across varied temperature conditions by integrating multiple domains within a semi-supervised DA framework. Unlike prior studies that are limited to isolated or narrowly specified scenarios, the proposed approach more faithfully represents real operational variability. Furthermore, combining CORAL loss, which is still relatively underused in SHM, with the feature extraction strengths of GATs substantially improves the alignment of domain-specific feature spaces. By integrating attention-based visualisation for explainability and employing optimisation during training, the

proposed framework achieves improved damage detection while remaining transparent, scalable, and suitable for practical SHM applications under complex environmental conditions.

The following subsections present the GAT architecture and attention mechanisms, the MMD formulation, CORAL alignment, adversarial training with gradient reversal scheduling, MLP classifier and discriminator designs, the Optuna optimisation procedure, and sensor importance computation.

3.3.2.2 Graph Attention Networks

GATs were employed as the feature extractor in GAT-CAMDA to capture both local and global relational structures in the input data. Conventional feature extraction approaches can be effective at identifying localised patterns; however, they often struggle to represent complex interdependencies, particularly in multi-sensor datasets typical of SHM systems. By contrast, GATs use attention mechanisms to adaptively weight connections according to the relative importance of neighbouring features (Li et al., 2023), thereby producing a more context-sensitive and informative representation.

Let the source and target datasets be denoted as D_s and D_t , where each contains signals from C sensors. A sensor graph was formed where node i corresponds to sensor i , and \mathcal{N}_i denotes the neighbourhood of node i . The operation of a GAT layer can be summarised in the following steps (Abdel-Basset et al., 2023):

Feature transformation where each node's signal vector, x_i , and a signal from its neighbourhood, x_j , undergo a shared linear transformation using a learnable weight matrix W :

$$h_i = Wx_i, h_j = Wx_j \quad (22)$$

These transformed features, h_i and h_j , serve as inputs for the subsequent attention mechanism.

Computation of attention scores, in which a self-attention mechanism evaluates the unnormalised attention score Ie_{ij} to quantify the relevance of node j 's features to node i :

$$Ie_{ij} = a(h_i, h_j) \quad (23)$$

where a is a learned function parameterised by a weight vector.

To obtain the attention coefficients α_{ij} , the unnormalised scores are passed through a softmax function applied over the neighbourhood $\mathcal{N}(i)$ of node i :

$$\alpha_{ij} = \frac{\exp(Ie_{ij})}{\sum_{k \in \mathcal{N}(i)} \exp(Ie_{ik})} \quad (24)$$

If a is modelled as a single-layer feedforward neural network with a LeakyReLU activation function to preserve gradient flow for negative inputs, the attention coefficient becomes:

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(a^\top [h_i \| h_j]))}{\sum_{k \in \mathcal{N}(i)} \exp(\text{LeakyReLU}(a^\top [h_i \| h_k]))} \quad (25)$$

Here, a^\top is the transpose of the learnable attention vector, and $\|$ denotes concatenation.

Feature aggregation in which each node aggregates the features from its neighbours via a weighted sum:

$$z_i = \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij} h_j \right) \quad (26)$$

where z_i is the refined feature of node i , and σ denotes an activation function, such as ReLU.

To enhance stability and increase representational capacity, GATs employ multi-head attention. In this setting, each head k computes its own attention coefficients and performs feature aggregation independently:

$$z_i = \parallel_{k=1}^K \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(k)} h_j \right) \quad (27)$$

where K specifies the total number of attention heads, and the operator \parallel indicates concatenation. At the output layer, concatenation is commonly replaced by an averaging operation:

$$z_i = \sigma \left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(k)} h_j \right) \quad (28)$$

Figures 6-a and 6-b illustrate the single-head and multi-head attention mechanisms (with $M_{\text{head}} = 4$), respectively. Each line colour and style in Figure 6-b represents one of the 4 independent attention calculations.




Figure 6. Conceptual diagram of (a) the self-attention operation and (b) the multi-head attention formulation within a GAT layer

3.3.2.3 Domain-Adversarial Neural Networks in GAT-CAMDA

The DANN unit in GAT-CAMDA framework incorporates two interdependent prediction components operating in a competitive at the same time complementary manner to promote domain-invariant feature extraction. The damage classifier infers fault labels $y_{(n)}$ from the latent embeddings, whereas the domain discriminator assesses whether the embeddings are drawn from the source domain \mathcal{D}_s or the target domain \mathcal{D}_t . Adversarial learning is enabled by a Gradient Reversal Layer (GRL), which behaves as an identity mapping during forward propagation but inverts the gradients during backpropagation. This mechanism prompts the feature extractor to learn representations that reduce the discriminator's capacity to differentiate between the 2 domains. This adversarial interplay facilitates that extracted features remain discriminative for damage classification whilst invariant to domain-specific variations.

3.3.2.3.1 Multi-layer Perceptron

The damage classifier in GAT-CAMDA is a multi-layer perceptron (MLP) comprising three blocks: input, hidden, and addition. Each block includes linear layers, batch normalisation, and dropout (Figure 7). This compact design balances

classification power and generalisation ability; overly complex classifiers often overfit the source data and fail to generalise effectively across domains (Huang et al., 2021). The streamlined MLP structure supports efficient use of extracted features without biasing them toward the source distribution.




Figure 7. Diagrammatic overview of the damage classification module in GAT-CAMDA

By contrast, the domain discriminator is constructed from 5 sequential blocks. In each block, a linear transformation is followed by a LeakyReLU activation, which introduces nonlinearity while maintaining gradient flow for negative inputs and thereby mitigating the risk associated with standard ReLU units. Dropout is applied to regularise the network and limit overfitting. The final output layer then converts the learned representation into domain-specific logits, allowing the model to distinguish between samples from the source and target domains. A comparatively deeper discriminator is adopted to ensure adequate capacity for modelling subtle domain shifts and for generating an informative adversarial signal for the feature extractor. This configuration maintains a functional equilibrium between the modules: the classifier preserves generalisability, whereas the discriminator remains sufficiently expressive to provide a challenging objective, thereby encouraging the learning of robust, domain-invariant features.

3.3.2.3.2 Classification Loss

The damage classifier operates on source domain samples where labelled training data are available. The classification loss \mathcal{L}_{cls} employs cross-entropy:

$$\mathcal{L}_{\text{cls}} = -\frac{1}{N_s} \sum_{n=1}^{N_s} \sum_{k=1}^K \mathbb{1}[y_{n_s} = k] \log(p(y_{n_s} = k | z_{n_s})) \quad (29)$$

in which N_s denotes the number of source samples, K represents the number of classes, z_{n_s} presents the extracted feature vector of the n -th observation belonging to the source domain, and $\mathbb{1}[y_{n_s} = k]$ is an indicator function that is 1 if the true label of sample n is k and 0 otherwise. $p(y_{n_s} = k|z_{n_s})$ is the predicted probability of sample n belonging to class k .

3.3.2.3.3 Domain Discriminator Loss

The domain discriminator operates on latent embeddings z extracted by the GAT layers. The adversarial loss \mathcal{L}_{adv} implements binary cross-entropy to distinguish domain membership:

$$\mathcal{L}_{\text{adv}} = -\frac{1}{N_s + N_t} \left[\sum_{n_s} \log(p(D_{n_s} = 0|z_{n_s})) + \sum_{n_t} \log(p(D_{n_t} = 1|z_{n_t})) \right] \quad (30)$$

where N_t denotes the number of target samples, D_n indicates domain membership (0 for source, 1 for target), and $p(\cdot)$ represents the discriminator's probability that z_n originates from which domain. Through the GRL, the feature extractor learns domain-invariant representations by maximising discriminator confusion during backpropagation, whilst the discriminator itself is trained to accurately distinguish domains. The strength of adversarial training is progressively adjusted per epoch:

$$\gamma_e = \alpha_{\text{kern}} \left(\frac{2}{1 + \exp(-\mu_{\text{GRL}} \cdot e/E)} - 1 \right) \quad (31)$$

where μ_{GRL} controls the growth rate, e is the current epoch, and E is the total number of epochs.

3.3.2.4 Maximum Mean Discrepancy Loss

The MMD loss quantifies distributional discrepancy by comparing the mean embeddings of the source and target features after mapping them into a reproducing kernel Hilbert space (RKHS). It evaluates how well feature distributions align across domains; the MMD loss function is:

$$\mathcal{L}_{\text{MMD}} = \left\| \frac{1}{N_s} \sum_{n=1}^{N_s} \Phi(z_{n_s}) - \frac{1}{N_t} \sum_{m=1}^{N_t} \Phi(z_{m_t}) \right\|^2 \quad (32)$$

where $\Phi(\cdot)$ maps the features into a higher-dimensional RKHS. Mapping into RKHS enables MMD to capture complex, nonlinear discrepancies between distributions. A smaller \mathcal{L}_{MMD} indicates stronger feature alignment, signifying effective knowledge transfer across domains. When applying MMD, the kernel type (linear or radial basis function) and its parameter γ_{RBF} (which controls sensitivity to distributional differences) must be carefully tuned.

3.3.2.5 Correlation Alignment Loss

The CORAL loss complements MMD by aligning second-order statistics between source and target feature sets. Whereas MMD focuses on mean alignment, CORAL contributes to covariance consistency, thereby reducing residual domain shifts. CORAL loss can be computed as:

$$\mathcal{L}_{\text{CORAL}} = \frac{1}{4D_{\text{feat}}^2} \|\text{Cov}(Z_s) - \text{Cov}(Z_t)\|_F^2 \quad (33)$$

where Z_s and Z_t represent the feature sets of source and target domains, D_{feat} is the dimension of the feature space, and $\|\cdot\|_F$ denotes the Frobenius norm; Cov present the covariance function. CORAL can be applied either on a per-class basis or over the entire feature set. Minimising $\mathcal{L}_{\text{CORAL}}$, the model enforces similarity in both the mean and covariance structures, thereby enhancing cross-domain generalisation.

3.3.2.6 Joint Optimisation

The framework optimises a composite loss function that balances classification performance on labelled source data:

$$\mathcal{L} = \lambda_{\text{cls}}\mathcal{L}_{\text{cls}} + \lambda_{\text{adv}}\mathcal{L}_{\text{adv}} + \lambda_{\text{MMD}}\mathcal{L}_{\text{MMD}} + \lambda_{\text{CORAL}}\mathcal{L}_{\text{CORAL}} \quad (34)$$

where $\lambda_{(\cdot)}$ are weighting coefficients controlling the relative contribution of each component. This joint optimisation combining discrepancy-based strategies (MMD, CORAL) and adversarial training (GRL) facilitates robust, transferable features across source and target domains.

Figure 8 displays a schematic of GAT-CAMDA framework for both the DA and classification stages (on the condition of 4 GATs layers).

\mathcal{L}_{cls} and \mathcal{L}_{adv} are optimised via separate backpropagation routes that correspond to their dedicated trainable modules. In contrast, \mathcal{L}_{MMD} and $\mathcal{L}_{\text{CORAL}}$ function as

regularisation objectives imposed on the common feature space. Their gradients therefore act solely on the feature extractor, shaping shared representations to reduce domain mismatch rather than propagating through additional trainable branches.




Figure 8. A schematic of GAT-CAMDA

3.3.2.7 Optuna-Based Hyperparameter Optimisation

Hyperparameter optimisation in GAT-CAMDA was performed using Optuna, which applies a TPE to efficiently explore the search space. The optimisation goal is to achieve the highest classification accuracy on the target-domain validation set:

$$\theta_{\text{hyper}}^* = \arg \max_{\theta_{\text{hyper}}} \text{ACC}_{D_t}(\theta_{\text{hyper}}) \quad (35)$$

This automated optimisation encourages a balanced configuration that enhances both model accuracy and generalisation under diverse operational conditions.

3.3.2.8 Computing Sensor Importance

In GAT-CAMDA and to quantify each sensor's contribution, the attention weights obtained from GAT were collected and analysed. For every edge $(i, j) \in \varepsilon$ with coefficient α_{ij} , contributions were equally divided between the two connected sensors. The overall importance score for sensor R can be calculated as:

$$S_R = \frac{\sum_{gl=1}^{GL} \sum_{(i,j) \in \varepsilon} \mathbb{1}\{R \in \{i, j\}\} \alpha_{ij}^{(gl)}}{\sum_{gl=1}^{GL} \sum_{(i,j) \in \varepsilon} \alpha_{ij}^{(gl)}} \quad (36)$$

in which GL is the total number of GAT layers and $\mathbb{1}\{\cdot\}$ is an indicator function; here $\mathbb{1}\{R \in \{i, j\}\} = 1$ when sensor R is one of the endpoints of edge (i, j) and 0 otherwise. Aggregating these scores across layers highlights the sensors most influential to the decision process, thereby enhancing the interpretability and transparency of the pipeline.

The pseudocode describing GAT-CAMDA is provided in Algorithm 2.

Algorithm 2. GAT-CAMDA Pipeline 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.

3.3.3 SPADA Framework

SPADA addresses fully UDA where no target domain labels exist, representing the most challenging adaptation scenario. The framework employs WTS as the signal processing stage to obtain inherently stable features before neural processing. WTS coefficients, derived through cascaded wavelet convolutions and modulus operations, provide translation-invariant, noise-robust representations that preserve discriminative information whilst suppressing nuisance variability.

A prototype-based attention mechanism provides class level structure and weighting, with cross domain alignment encouraged indirectly via adversarial training and prototype updates. Source prototypes use true labels; target prototypes are constructed from confident pseudo-labels generated by the model itself and filtered through entropy-based gating. Attention weights, computed through cosine similarity to prototypes, enhance representative samples dominate training whilst uncertain samples are down-weighted. Prototypes are updated online, which enables dynamic refinement as adaptation proceeds, and pseudo label quality improves.

The training step in SPADA integrates four losses: source classification (weighted cross entropy with label smoothing), target pseudo labelling (cross entropy on confident target samples only), adversarial domain discrimination (using a gradient reversal schedule to promote domain invariant features), and prototype compactness (encouraging source features to cluster around their source prototypes, with target prototypes updated from confident pseudo labels and influencing sample weights, so cross domain alignment is indirect rather than explicit). This multi objective design balances source performance, self-training-based target adaptation, domain invariance, and class structure.

The pseudo-labelling mechanism distinguishes SPADA from purely adversarial approaches. This pipeline actively generates and refines pseudo-labels to guide adaptation. The confidence gating and prototype dynamics of this framework provide

explicit mechanisms for managing pseudo label noise, which is critical when ground truth is unavailable.

Interpretability mechanisms track prototype trajectories, similarity structure, attention dynamics, and latent space separation throughout training, enabling diagnosis of adaptation quality and identification of failure modes.

3.3.3.1 Main Contribution

As a fully UDA framework for mechanical systems, SPADA is introduced, in which WTS feature extraction, prototype-based attention, entropy-gated pseudo-labelling, and adversarial alignment are combined to enable cross-domain damage diagnosis without target labels. Its effectiveness and scope are evidenced by the following contributions:

- Robust temperature adaptability is demonstrated, with reliable detection maintained under high EOVs.
- Generalisability across systems and modalities is evidenced through validation on two distinct systems employing different signal types, indicating resilience to sensing and domain variability.
- Stable and noise-robust representations are obtained via WTS, providing translation-invariant features that preserve discriminative structure prior to learning.
- Prototype-guided self-training is achieved, as cosine-similarity attention and online prototype updates emphasise representative samples and suppress uncertain ones, thereby improving class cohesion and mitigating pseudo-label noise.
- Fully unsupervised adaptation is supported through the joint use of source classification, confident target learning, adversarial domain discrimination, and prototype compactness, yielding strong performance without any target labels.

The contribution beyond component integration is provided by three specific advances: first, a confidence-gated prototype update rule that couples entropy filtering with similarity-weighted attention to control pseudo-label noise; second, an indirect alignment strategy in which prototype compactness preserves inter-class margins while adversarial training encourages domain invariance, thereby reducing collapse and over-alignment; third, an interpretability suite that quantifies adaptation quality

through prototype trajectories, similarity structure, attention dynamics, and latent-space separation, enabling diagnosis of failure modes and informed model selection in fully UDA settings. Validation on two heterogeneous systems substantiates the robustness and generalisability of these advances.

The following subsections detail the WTS and scattering coefficients, the DANN architecture with a shared feature extractor, prototype construction and online updating, entropy-based confidence gating, attention weight computation, the 4 components loss formulation, and the interpretability procedures.

3.3.3.2 Wavelet Time Scattering

To obtain stable and interpretable features from raw signals, SPADA employs the WTS transform. WTS processes each input signal through a hierarchical sequence of operations, where the output of one stage serves as the input to the next. At each stage, three fundamental operations are performed: (i) convolution through wavelet transformation to capture localised frequency content, (ii) nonlinearity extraction using the modulus operation to enhance stability, and (iii) averaging through scaling functions to support translation invariance. For a single channel $x_i(t)$, the zero-order scattering coefficient is obtained by local averaging of the input signal as (Susu et al., 2020; Rezazadeh et al., 2023b):

$$S_{(\text{scat},0)}(x_i) = x_i * \phi_J \quad (37)$$

where $*$ denotes convolution and ϕ_J represents the low-pass filter at scale J . This coefficient retains coarse, low-frequency components of the signal.

The first-order scattering coefficient is derived by convolving the signal with a band-pass wavelet filter $\psi_{(\omega_1)}$, taking the modulus of the result, and subsequently applying a low-pass filter:

$$S_{(\text{scat},1)}(x_i, \omega_1) = |x_i * \psi_{\omega_1}| * \phi_J \quad (38)$$

Here, ψ_{ω_1} is a band-pass wavelet filter indexed by the frequency parameter ω_1 and $|\cdot|$ is the complex modulus. The modulus operation captures amplitude-modulated oscillations, while low-pass filtering provides translation invariance.

Higher order coefficients are built by iterating wavelet–modulus–averaging along a path $(\omega_1, \dots, \omega_\eta)$. The η -th order scattering coefficient is formulated as:

$$S_{(\text{scat},\eta)}(x_i, \omega_1, \dots, \omega_\eta) = | \dots | | x_i * \psi_{\omega_1} | * \psi_{\omega_2} | \dots * \psi_{\omega_\eta} | * \phi_J \quad (39)$$

where $\eta \geq 0$ is the scattering order. Increasing η captures progressively complex temporal dependencies and multiscale interactions.

Concatenating all scattering coefficients up to a maximum order η , the scattering representation of the i -th channel is expressed as:

$$s_i^{\text{scat}} = [S_{(\text{scat},0)}(x_i), S_{(\text{scat},1)}(x_i, \omega_1), \dots, S_{(\text{scat},\eta)}(x_i, \omega_1, \dots, \omega_\eta)] \in \mathbb{R}^{D_{\text{scat}}} \quad (40)$$

Finally, the multi-channel scattering feature vector is constructed by concatenating the representations from all C sensor channels:

$$z = [s_1^{\text{scat}}; s_2^{\text{scat}}; \dots; s_C^{\text{scat}}] \in \mathbb{R}^D \quad (41)$$

where $D = C \cdot D_{\text{scat}}$ denotes the total feature dimension. Figure 9 presents a schematic of WTS implemented on a signal $x(t)$ up to second order ($\eta = 2$).




Figure 9. Schematic representation of the WTS process up to the second order for a single channel

The resulting features are invariant to small temporal shifts, robust to noise, and maintain discriminative properties essential for cross-domain fault diagnosis.

3.3.3.3 Domain Adaptation in SPADA

The DA block in SPADA consists of four loss functions: source classification loss, target pseudo-label loss, domain adversarial loss, and prototype alignment loss. A DANN block was designed in SPADA that has a shared feature extraction unit for projecting the extracted features from WTS into a latent space that is shared between the source and target domains. This block is responsible for calculating the source classification and domain adversarial losses. In the following, first the prototype block

of SPADA is elaborated, then the mathematics for calculation of the loss functions along with the training procedure are detailed.

3.3.3.3.1 Prototype Construction and Attention-based Updating

To achieve class-level alignment between domains, class prototypes are constructed in the latent space h . This latent space is produced from projecting the scattering features z through a shared feature extractor g_θ in DANN as:

$$h = g_\theta(z) \quad (42)$$

Each prototype represents the mean latent embedding of samples belonging to a specific class. For the source domain, prototypes are calculated from features associated with true labels as:

$$\mu_k^{(s)} = \frac{1}{|\mathcal{D}_s^k|} \sum_{h \in \mathcal{D}_s^k} h \quad (43)$$

where \mathcal{D}_s^k is the subset of source features belonging to class label k . For the target domain, prototypes are formed using pseudo-labelled target samples, which are iteratively updated during training:

$$\mu_k^{(t)} = \frac{1}{|\mathcal{D}_t^k|} \sum_{h \in \mathcal{D}_t^k} h \quad (44)$$

Pseudo-labels for target samples are generated from the classifier output through temperature-scaled softmax probabilities:

$$p^t(h) = \text{softmax}\left(\frac{\text{logits}(h)}{\tau_{temp}}\right) \quad (45)$$

The temperature parameter τ_{temp} controls the sharpness of the probability distribution. A normalised entropy $H(p^t(h))$ is computed to measure prediction confidence, and a binary confidence gate is applied such that only samples with entropy below a threshold τ are used for prototype updating and target loss computation.

To modulate the influence of each sample, a prototype-based attention weight is assigned. For source samples, the weight is based on the cosine similarity between each feature and its corresponding source prototype:

$$w^{(s)} = (1 - \beta_w) + \beta_w \cos(h, \mu_k^{(s)}) \quad (46)$$

where β_w controls the contribution of similarity. For target samples, the weight is additionally modulated by the entropy-based confidence:

$$w^{(t)} = \left[(1 - \beta_w) + \beta_w \cos(h, \mu_k^{(t)}) \right] (1 - H(p^t(h))) \quad (47)$$

This mechanism fosters that highly representative and confident samples contribute more strongly to the adaptation process. Prototypes in both domains are initialised at zero and updated online throughout training.

3.3.3.2 Training Objectives

The learning process in SPADA is governed by four complementary loss functions, each addressing a specific aspect of the adaptation problem.

(a) Source Classification Loss (\mathcal{L}_{cls})

A weighted cross-entropy objective is applied to the labelled source samples using their prototype-based attention weights as:

$$\mathcal{L}_{\text{cls}} = -\frac{1}{N_s} \sum_{n=1}^{N_s} w_n^{(s)} \sum_{k=1}^K \mathbb{1}[y^{(n)} = k] \log(p(y_{(n_s)} = k | h_{(n_s)})) \quad (48)$$

where $p^k(h^{(n)})$ denotes the predicted probability of class k for the n -th sample, and $\mathbb{1}[\cdot]$ is the indicator function.

(b) Target Pseudo-label Loss (\mathcal{L}_{tgt})

This loss is used to guide the learning on unlabelled target data through confident pseudo-labels. Only target samples passing the confidence gate ($H(p^t(h)) \leq \tau$) are used. The loss is expressed as:

$$\mathcal{L}_{\text{tgt}} = \frac{1}{\sum_{m=1}^{N_t} m^{(m)}} \sum_{m=1}^{N_t} m^{(m)} w_{(m)}^{(t)} \log(p(y_{(m_t)} = \hat{y}_{(m_t)} | h_{(m_t)})) \quad (49)$$

where $\hat{y}^{(m)} = \arg \max_k p(y_{(m_t)} = k | h_{(m_t)})$.

(c) Domain Adversarial Loss (\mathcal{L}_{adv})

To promote domain invariance, the domain discriminator is trained to distinguish between source and target representations, while the feature extractor is trained in an adversarial manner to prevent such discrimination. The loss is defined as:

$$\mathcal{L}_{\text{adv}} = -\frac{1}{N_s + N_t} \left[\sum_{n=1}^{N_s} \log \left(p(D_{(n_s)} = 0 | h_{(n_s)}) \right) + \sum_{m=1}^{N_t} \log \left(p(D_{(m_t)} = 1 | h_{(m_t)}) \right) \right] \quad (50)$$

(d) Prototype Alignment Loss ($\mathcal{L}_{\text{proto}}$)

To improve intra-class compactness within the source domain, the features are encouraged to align with their corresponding source prototypes:

$$\mathcal{L}_{\text{proto}} = \frac{1}{N_s} \sum_{n=1}^{N_s} \left(1 - \cos \left(h_{(n_s)}, \mu_{y_{(n_s)}}^{(s)} \right) \right) \quad (51)$$

This term enforces class cohesion in the latent space and stabilises the prototype-based weighting mechanism.

3.3.3.3 Overall Objective Function

The complete training objective combines the 4 losses with weighting coefficients that balance their relative contributions:

$$\mathcal{L} = \lambda_{\text{cls}} \mathcal{L}_{\text{cls}} + \lambda_{\text{tgt}} \mathcal{L}_{\text{tgt}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}} + \lambda_{\text{proto}} \mathcal{L}_{\text{proto}} \quad (52)$$

in which $\lambda_{(\cdot)}$ is the weight value related to the corresponding loss function.

This multi-objective optimisation promotes that the model learns to (i) remain discriminative for labelled source data, (ii) adapt effectively to the unlabelled target domain, (iii) preserve domain invariance, and (iv) maintain class-level structural consistency.

3.3.3.4 Summary of the Training Process

During training, the WTS features from both domains are first extracted and standardised. Mini-batch training is performed with source batches (labelled) and target batches (unlabelled). The model alternates between forward passes through the shared feature extractor, the classifier, and the domain discriminator. Prototypes for

both domains are updated online at each iteration. Source prototypes are updated using true labels, while target prototypes are updated using confident pseudo-labels filtered by the confidence gate. The four loss components are computed for each mini-batch and combined according to the total loss function. The model parameters are then updated via backpropagation using the combined gradients, with the GRL advancing adversarial domain learning. Throughout training, model selection is performed using an unsupervised validation score that combines prediction entropy, diversity, and prototype compactness (encouraging source features to cluster around their source prototypes, with target prototypes updated from confident pseudo labels and influencing sample weights, so cross domain alignment is indirect rather than explicit). The best model checkpoint, determined by this score, is retained for final evaluation.

To ensure robust model selection without relying on target domain labels, the unlabelled target adaptation set is randomly partitioned into two equal subsets, A and B. At each training epoch, the model is evaluated independently on both subsets using an unsupervised scoring metric that quantifies adaptation quality through 4 complementary criteria: prediction diversity (entropy of the averaged class probability distribution across samples), prediction confidence (inverse of mean per-sample entropy), prototype compactness (mean cosine similarity between high-confidence samples and their assigned prototypes), and coverage (proportion of samples exceeding the confidence threshold). The unsupervised score is computed as the sum of diversity minus mean entropy, plus weighted contributions from compactness and coverage. Evaluating on both A and B splits and summing their respective scores, the selection process mitigates potential bias from a single data partition. The model checkpoint achieving the highest cumulative unsupervised score across all training epochs is retained as the optimal model, ensuring selection is based purely on intrinsic properties of the target domain predictions rather than supervised accuracy metrics. Figure 10 represents a graphical schematic of SPADA.




Figure 10. A schematic of SPADA framework

3.3.3.4 Interpretability of SPADA Internal Activity

SPADA incorporates four interpretability views that monitor internal activity during training to verify that source and target latent distributions align whilst class discriminability is preserved. These views are consistent with WTS features, the DANN latent space, the prototype-based attention mechanism, and the unsupervised model selection strategy. Crucially, these views do not alter training dynamics; they serve purely as diagnostic tools for post-hoc analysis and validation.

Let $h = g_{\theta}(z) \in \mathbb{R}^{D_{\text{lat}}}$ denote the latent feature representation of scattering coefficients z extracted by the feature encoder and K be the number of classes. The source class subset for class k is fixed throughout training and defined as $D_s^k = \{i \in D_s \mid y_{(i_s)} = k\}$, where $y_{(i_s)}$ is the true label of source instance i . The source prototype for class k at epoch e is computed as the arithmetic average of all source instances belonging to that class:

$$\mu_k^S(e) = \frac{1}{|D_s^k|} \sum_{i \in D_s^k} h_i(e) \quad (53)$$

Since source labels are available, this computation remains stable across epochs, though the feature representations $h_i(e)$ evolve as the network parameters update.

For the target domain, class membership is determined by pseudo-labels and filtered by prediction confidence. The target class subset at epoch e is dynamically constructed as:

$$D_t^{(k,\tau)}(e) = \left\{ j \in D_t \mid \hat{y}_{j_t}(e) = k, H\left(p^T\left(h_{(j_t)}(e)\right)\right) \leq \tau \right\} \quad (54)$$

where $\hat{y}_{j_t}(e)$ denotes the pseudo-label assigned to target instance j at epoch e , $H(\cdot)$ is the Shannon entropy of the temperature-scaled softmax output $p^T(h)$, and τ is the entropy threshold for confidence gating. Only target instances with pseudo-label k and entropy below τ contribute to the target prototype. The target prototype is then computed as:

$$\mu_k^T(e) = \frac{1}{\left|D_t^{(k,\tau)}(e)\right|} \sum_{j \in D_t^{(k,\tau)}(e)} h_j(e) \quad (55)$$

This formulation ensures that target prototypes are updated using only confident predictions, avoiding contamination from ambiguous or misclassified samples. The coverage metric quantifies the proportion of target samples meeting the confidence criterion at epoch e :

$$C(e) = \frac{1}{N_t} \sum_{j \in D_t} \mathbb{1}\left\{H\left(p^T\left(h_{(j_t)}(e)\right)\right) \leq \tau\right\} \quad (56)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function and N_t is the total number of target samples. Coverage provides a direct measure of how many target samples the model considers sufficiently confident for pseudo-labelling.

The cosine similarity between a feature vector h and a prototype μ measures their angular alignment in the latent space, defined as:

$$\text{sim}(h, \mu) = \frac{h^\top \mu}{\|h\|_2 \|\mu\|_2} \quad (57)$$

This normalised similarity ranges from -1 to +1, with values near +1 introducing strong alignment. Instance-level attention weights leverage this similarity to modulate the contribution of each sample to the loss. For source instances, attention at epoch e is computed as:

$$a_i^S(e) \propto \exp \left\{ \alpha \cdot \text{sim} \left(h_i(e), \mu_{y_{(i_s)}}^S(e) \right) \right\} \quad (58)$$

where α is a temperature parameter used to adjust the sharpness of the attention distribution, and $y_{(i_s)}$ is the true class label. For target instances, attention additionally incorporates prediction confidence:

$$a_j^T(e) \propto \left(1 - H \left(p^T \left(h_{(j_t)}(e) \right) \right) \right) \cdot \exp \left\{ \alpha \cdot \text{sim} \left(h_j(e), \mu_{\hat{y}_{(j_t)}(e)}^T(e) \right) \right\} \quad (59)$$

The term $\left(1 - H \left(p^T \left(h_{(j_t)}(e) \right) \right) \right)$ downweights instances with high entropy (low confidence), whilst the exponential term upweights instances with features aligned to their assigned prototype. These attention weights are normalised within each mini-batch to form probability distributions before being applied to the loss computation.

3.3.3.4.1 Tracking Prototype Trajectory Evolution

To characterise how class structure evolves in the shared latent space and whether target centres migrate toward corresponding source centres whilst preserving inter-class distinctness, SPADA tracks per-class drift and alignment metrics across epochs. The source prototype drift quantifies the Euclidean displacement of the source prototype for class k between consecutive epochs:

$$\Delta_k^S = \|\mu_k^S(e) - \mu_k^S(e-1)\|_2 \quad (60)$$

Similarly, the target prototype drift measures the movement of the target prototype:

$$\Delta_k^T = \|\mu_k^T(e) - \mu_k^T(e-1)\|_2 \quad (61)$$

The alignment gap quantifies the Euclidean distance between source and target prototypes for the same class:

$$A_k(e) = \|\mu_k^S(e) - \mu_k^T(e)\|_2 \quad (62)$$

These metrics are computed independently for each class k and tracked throughout training. A diminishing trend in $\Delta_k^T(e)$ coupled with a reduction in $A_k(e)$ indicates that target prototypes are stabilising and converging toward their source

counterparts, signalling successful domain alignment. Conversely, persistent large values of $\Delta_k^t(e)$ or increasing $A_k(e)$ suggest unreliable pseudo-labels, excessive pressure to remove domain-specific information, or insufficient adaptation. Trajectory crossings wherein prototype for different classes move closer together may indicate collapsing class boundaries and warrant intervention. Early-stage prototype movement is expected as the GRL schedule strengthens, and the confidence gate admits more target samples. The desired outcome is convergence with stability: prototypes settle into distinct limiting locations across classes, preserving discriminative margins whilst achieving cross-domain alignment for corresponding classes.

3.3.3.4.2 Computing Instance–Prototype Cosine Similarity

To examine within-class compactness and between-class confusion patterns, SPADA constructs a similarity matrix that quantifies how individual instances relate to all class prototypes at a selected epoch. For a dataset with N instances and K classes, the similarity matrix $M \in \mathbb{R}^{N \times K}$ is computed with entries:

$$M_{i,k} = \text{sim}(h_i(e), \mu_k(e)) \quad (63)$$

where $h_i(e)$ is the latent feature of instance i at epoch e , $\mu_k(e)$ denotes either $\mu_k^s(e)$ or $\mu_k^t(e)$ depending on whether source or target domain inspection is performed, and $\text{sim}(u, v) = \frac{u^\top v}{\|u\|_2 \|v\|_2}$ is the cosine similarity. Each row i of M represents the similarity profile of instance i across all K prototypes, whilst each column k reflects how all instances relate to prototype k .

The computation proceeds as follows: (1) select the epoch e for analysis (typically the best epoch determined by unsupervised validation); (2) extract latent features $h_i(e)$ for all instances in the domain of interest; (3) retrieve prototypes $\mu_k(e)$ for all K classes; (4) compute pairwise cosine similarities to populate M . The resulting matrix is visualised as a heatmap, with rows typically sorted by domain (source first, then target) and class label to facilitate interpretation.

A sharply diagonal block structure in the source similarity matrix indicates that source instances exhibit high similarity to prototypes of their true class and low similarity to prototypes of other classes, confirming compact and well-separated clusters. When adaptation is effective, confident target instances exhibit a similar

pattern: high on-diagonal values for their assigned pseudo-label class and low off-diagonal values elsewhere. Off-diagonal bands or elevated values identify pairs of confusable classes and can flag imbalances in attention allocation or confidence distribution. Progressive enhancement of diagonal dominance across epochs signals sharpening class structure and improving discriminability. Persistent off-diagonal mass concentrated on a particular class indicates that the corresponding prototype acts as an attractor, potentially absorbing instances from neighbouring classes, and merits closer examination for prototype collapse or class overlap.

3.3.3.4.3 Monitoring Prototype Attention Weight Dynamics

To assess whether the model appropriately weights target samples it deems confident and whether reliance on target data grows in a controlled manner, SPADA tracks epoch-wise summary statistics of attention weights and coverage. These summaries are computed using robust statistics to avoid sensitivity to outliers. For a set of attention weights $\{a_j(e)\}$ at epoch e , the median and interquartile range are defined as:

$$\begin{aligned}\tilde{a}(e) &= \text{median}_j\{a_j(e)\}, \\ \text{IQR}(a)(e) &= \text{IQR}_j\{a_j(e)\}, \\ \mathcal{C}(e) &= (\text{as defined in Equation 56})\end{aligned}\tag{64}$$

in which IQR denotes the interquartile range, computed as the difference between the 75th and 25th percentiles of the attention weight distribution. This measure quantifies dispersion robustly: a small $\text{IQR}(a)(e)$ indicates attention is distributed evenly across instances, whilst a large $\text{IQR}(a)(e)$ signifies that attention is concentrated on a subset of samples. The median $\tilde{a}(e)$ provides a central tendency measure resistant to extreme values.

The computation procedure is as follows: (1) at each epoch e , collect attention weights $a_j^t(e)$ for all target instances $j \in D_t$; (2) compute the median $\tilde{a}(e) = \text{median}(\{a_j^t(e)\})$; (3) compute the 25th percentile $Q_1(e)$ and 75th percentile $Q_3(e)$, then $\text{IQR}(a)(e) = Q_3(e) - Q_1(e)$; (4) simultaneously compute coverage $\mathcal{C}(e)$ as the fraction of target instances with entropy below τ . These statistics are tracked separately for source and target domains and plotted across training epochs.

A healthy adaptation trajectory is indicated by a rising or stable $\tilde{a}(e)$ concurrent with increasing $C(e)$, suggesting that as more target samples are admitted via the confidence gate, the model continues to assign meaningful attention weights. In contrast, a decline in $\tilde{a}(e)$ whilst $C(e)$ rises indicates that newly admitted samples receive low attention, possibly due to unreliable pseudo-labels or noisy feature representations. A large $\text{IQR}(a)(e)$ showing attention concentrated on a small subset suggests that prototype updates are driven by a narrow group of instances, potentially leading to instability or bias. Comparing $\tilde{a}(e)$ with $C(e)$, co-movement (both rising together) implies that additional target samples admitted by the entropy gate are also judged informative by the attention mechanism; divergence (e.g., $C(e)$ increasing but $\tilde{a}(e)$ flat or declining) flags a mismatch, such as admission of many low-quality or ambiguous samples that do not meaningfully contribute to adaptation.

3.3.3.4.4 Visualising Decision Boundaries in Feature Space

SPADA provides a qualitative visualisation of class separation and source-target overlap in the latent space at selected epochs by applying dimensionality reduction. Specifically, a projection $\Pi: \mathbb{R}^{D_{lat}} \rightarrow \mathbb{R}^2$ is fitted on a reference set comprising both source and target features at a chosen epoch (typically the best epoch), using t-SNE or PCA. The procedure is: (1) collect latent features $h_i(e)$ for all source and target instances at epoch e ; (2) concatenate these features into a single matrix; (3) fit the projection Π to this matrix; (4) transform both instance features and prototypes into 2D: $u_i = \Pi(h_i(e))$ and $v_k = \Pi(\mu_k(e))$.

The resulting 2D scatter plot displays source and target instances with distinct markers (e.g., circles for source, squares for target), colour-coded by class label or pseudo-label. Source prototypes $\Pi\mu_k^s(e)$ and target prototypes $\Pi\mu_k^t(e)$ are overlaid with distinct markers (e.g., 'X' for source prototypes, '+' for target prototypes), also colour-coded by class. The ideal geometric structure exhibits class-wise overlap between source and target samples, i.e., instances from the same class but different domains intermingle with clear separation between different classes. This pattern confirms domain-invariant features whilst preserving discriminability.

Excessive invariance manifests as reduced inter-class separation, where clusters for different classes merge or lose distinctness, indicating over-suppression of task-relevant information. Insufficient adaptation results in disjoint source and target

clusters within the same class, signalling that domain shift has not been adequately mitigated. Alignment between this visualisation and the quantitative trends in $A_k(e)$ (alignment gap) and $\Delta_k^t(e)$ (target prototype drift) reinforces interpretation: converging prototypes (small $A_k(e)$) should correspond visually to overlapping source-target clusters, whilst stable prototypes (small $\Delta_k^t(e)$) should appear as fixed cluster centres.

Algorithm 3 presents the algorithm for SPADA.

It will a good idea if you add a table that summarise all the special features of the different approaches/methods/etc (under 3.3.3) have. Similarly for Section 3.3.2 and other sections.

Algorithm 3. SPADA Framework for SHM

Input: Source domain signals and labels, target domain adaptation signals (unlabelled), target domain test signals and labels (held out for final evaluation)

Output: Domain-adaptive model, performance metrics, interpretability visualisations

1. Data Loading

- 1.1 Load raw time-series signals from source domain and corresponding labels.
- 1.2 Load raw time-series signals from target domain adaptation set (unlabelled).
- 1.3 Load raw time-series signals from target domain test set and corresponding labels (held out for final evaluation).

2. Feature Extraction

- 2.1 Apply per-sample, per-channel standardisation to all loaded signals.
- 2.2 Extract wavelet scattering transform features using Kymatio with configurable scale and order parameters.
- 2.3 Apply log-stabilisation and concatenate multi-channel features.

3. Data Splitting and Preprocessing

3.1 Source Domain:

- 3.1.1 Split extracted source features into training, validation, and test sets (stratified by labels).
- 3.1.2 Create labelled data loaders for supervised training and evaluation.

3.2 Target Domain:

- 3.2.1 Split extracted target adaptation features into two subsets: DA set (for domain adaptation training) and validation set (for unsupervised model selection), both unlabelled.
- 3.2.2 Further partition target validation set into equal A and B subsets for unsupervised cross-validation scoring.
- 3.2.3 Retain target test features separately (held out, used only for final evaluation).

3.3 Standardisation:

- 3.3.1 Fit StandardScaler on combined source training features and target DA features (unlabelled).
- 3.3.2 Transform all source and target feature splits using the fitted scaler.

4. Model Architecture

- 4.1 Feature extractor: Two-layer network with batch normalisation and dropout.
 - 4.2 Label classifier: Two-layer network for class prediction.
 - 4.3 Domain discriminator: Two-layer network with gradient reversal layer.
-

4.4 Prototype module: Maintain source and target class prototypes, updated via exponential moving average.

5. Training Loop

5.1 For each epoch:

5.1.1 Sample mini-batches from source (labelled) and target DA (unlabelled) data.

5.1.2 Forward pass through feature extractor, classifier, and discriminator.

5.1.3 Generate pseudo-labels for target samples via temperature-scaled softmax.

5.1.4 Update prototypes using source true labels and confident target pseudo-labels.

5.1.5 Compute instance weights based on cosine similarity between features and prototypes, modulated by prediction entropy.

5.1.6 Calculate combined loss from four components: source classification, target pseudo-labelling, domain adversarial, and prototype alignment.

5.1.7 Backpropagate and update parameters.

5.2 Unsupervised model selection at each epoch:

5.2.1 Evaluate model separately on target validation subsets A and B (both unlabelled).

5.2.2 Compute unsupervised score combining prediction diversity, confidence, prototype compactness, and coverage.

5.2.3 Sum scores from both subsets and retain checkpoint if score improves.

5.3 Load best checkpoint selected by unsupervised criterion.

6. Hyperparameter Optimisation

6.1 Define search space over adaptation parameters, prototype parameters, training parameters, and scattering parameters.

6.2 Sample configurations from the full grid space.

6.3 For each sampled configuration:

6.3.1 Train model across multiple random seeds following the training loop in Section 5.

6.3.2 Compute mean unsupervised score across all seeds.

6.3.3 Record per-seed unsupervised scores, test accuracies (for analysis only), and best epochs.

6.4 Select the configuration with the highest mean unsupervised score.

6.5 Retrain the selected best configuration across all seeds for final test evaluation.

7. Evaluation and Reporting

7.1 Evaluate best model on target test set and compute classification accuracy, confusion matrix, and per-class metrics.

8. Interpretability Analysis

8.1 Prototype Trajectory Evolution:

8.1.1 Apply t-SNE to concatenated source and target prototypes across all logged epochs.

8.1.2 Visualise trajectories with solid lines (source) and dashed lines (target), marking starting and ending positions.

8.2 Instance-Prototype Cosine Similarity:

8.2.1 Compute cosine similarity matrix between all instances and prototypes at the best epoch.

8.2.2 Generate heatmap organised by domain and class, highlighting diagonal block structure.

8.3 Prototype Attention Weight Dynamics:

8.3.1 Plot median source and target attention weights with interquartile range across epochs.

8.3.2 Overlay pseudo-label coverage ($\text{entropy} \leq \tau$) and median target weight on third panel.

8.3.3 Mark best epoch determined by unsupervised validation with vertical line.

8.4 Decision Boundaries in Feature Space:

8.4.1 Apply t-SNE jointly to instances and prototypes at the best epoch.

8.4.2 Visualise source/target instances with distinct markers and overlay source and target prototypes, all colour-coded by class.

3.3.4 Framework Selection Criteria and Application Contexts

The three frameworks developed in this research named as UCTRF, GAT-CAMDA, and SPADA address complementary aspects of the DA problem in SHM under varying levels of data availability and supervision. Their deployment in subsequent chapters is guided by the nature of the datasets, the presence or absence of labelled target information, and the degree of domain shift encountered between environmental and operational conditions.

UCTRF is applied in scenarios where a small quantity of labelled target domain data is available. Its instance-based weighting mechanism enables effective adaptation under limited supervision, making it suitable for experimental cases in which partial labelling is feasible, such as early-stage monitoring or controlled laboratory campaigns. Leveraging minimal target information, UCTRF provides interpretable adaptation through the evolution of sample weights and feature embeddings.

GAT-CAMDA is employed in scenarios where target domain data are unlabelled during training, while inter-sensor relationships remain critical for achieving robust adaptation. By leveraging graph-attention mechanisms alongside concurrent adversarial, MMD, and CORAL alignment strategies, GAT-CAMDA enables comprehensive distributional matching across domains. This configuration is particularly well suited to dense sensor networks, such as those used on large structural components, where relational dependencies among sensors have a significant impact on diagnostic accuracy.

SPADA is designed for conditions exhibiting substantial domain divergence and no available target labels. Combining WTS features, prototype-based attention, and confidence-gated pseudo-labelling, SPADA autonomously adapts across domains with high EOVs. Its emphasis on interpretable internal dynamics, i.e., prototype trajectories, attention weights, and entropy evolution enables diagnostic transparency even in fully unsupervised regimes.

Overall, these frameworks provide a systematic progression across levels of supervision and domain complexity, supporting that the methodological strategy

remains adaptable to a broad spectrum of real-world SHM scenarios. Their complementary characteristics, i.e., instance-level, feature-level, and prototype-level alignment for UCTRF, GAT-CAMDA, and SPADA, respectively offer a unified framework capable of adapting across varying availability of target supervision.

Chapter 4: Case Studies

4.1 Overview

This study evaluates the proposed damage detection frameworks using 2 publicly available datasets. Both datasets were selected because their experiments share a comparable configuration in which a temperature chamber was used to control the ambient conditions during data acquisition. In addition, the datasets were derived from two different types of signals such as guided wave and vibration signals, thereby enabling a comprehensive evaluation across diverse sensing approaches. The datasets are described in detail in the following sections.

4.2 Small-scale Wind Turbine Blade Under Varying Climate Conditions (WTB-VibClimate)

The first dataset comprises experimental measurements collected from a small-scale WTB associated with the blade of the Sonkyo Energy Windspot 3.5 kW WT model (Figure 11), which was made publicly available in (Ou et al., 2021). The blade adopts a 3-layer sandwich composite construction and has a length of 1.75 m and a mass of 5.0 kg.



Figure 11. The 3.5kW WT made by Sonkyo Energy (Ou et al., 2021)

Data acquisition was conducted at 12 temperature levels, spanning from $-15\text{ }^{\circ}\text{C}$ to $40\text{ }^{\circ}\text{C}$ in $5\text{ }^{\circ}\text{C}$ increments (Wn15, Wn10, Wn5, Wp0, Wp5, Wp10, Wp15, Wp20,

Wp25, Wp30, Wp35, Wp40), while humidity was maintained at 60%. Two excitation modes were applied separately: white noise over 0 to 400 Hz and a sine sweep over 1 to 300 Hz. Each excitation was applied for approximately 120 s, although the recorded signal duration exhibited minor variation. The sampling frequency was kept constant at 1666 Hz, and excitation was introduced at a fixed location on the blade surface throughout. Responses were measured using 2 sensor groups, accelerometers and strain gauges, each deployed with a distinct configuration. Further details on the sensor specifications are reported in (Ou et al., 2021). Figure 12 summarises the WTBVib Climate data acquisition workflow. Figure 13 provides an overview of the experimental arrangement, including the climate controller (Figure 13-a), the shaker with an insulation foam box (Figure 13-b), strain gauges S1 and S2 on the low pressure side (Figure 13-c), strain gauges S23 and S24 on the high pressure side (Figure 13-d), the climate sensor (Figure 13-e), cracks 1 and 2 (Figure 13-f), and crack 3 of WTB-VibClimate (Figure 13-g and Figure 13-h).




Figure 12. Workflow of the data acquisition system in WTB-VibClimate (Ou et al., 2021)




Figure 13. Components of the WTB-VibClimate climate-chamber experiment: (a) overall arrangement, (b) climate controller, (c) shaker with an insulation foam box, (d) strain gauges s_1 and s_2 on the low pressure side, (e) strain gauges s_{23} and s_{24} on the high pressure side, (f) climate sensor, (g) cracks 1 and 2, and (h) crack 3 (Ou et al., 2021)

In Figure 14 the sensor configuration embedded on WTB in the WTB-VibClimate case study are showed; for the acceleration sensors only one configuration was utilised whilst for strain gauges two various configurations were implemented.




Figure 14. Sensor configuration in WTB-VibClimate

For the WTB, 13 distinct health states were defined. These comprised a single undamaged baseline, 9 cracked configurations in which both the number of cracks, ranging from 1 to 3, and crack length were varied, and 3 icing accretion cases implemented by attaching unbalanced masses of 44 g, using 1 to 3 masses. Table 1 reports the full set of conditions together with the corresponding fault quantities and severities, and the assigned index is used later for classification.

Table 1. WT blade health conditions

| Health scenario | No. of observations | Label | Commentary (Crack length in cm) |
|-----------------|---------------------|-------|---|
| Intact | 20 | 0 | Healthy |
| Unbalanced | 5 | 1 | 1 added mass |
| | 5 | 2 | 2 added masses |
| | 5 | 3 | 3 added masses |
| Cracked | 5 | 4 | 1 crack, $L_{c1} = 5$ |
| | 5 | 5 | 2 cracks, $L_{c1} = 5, L_{c2} = 5$ |
| | 5 | 6 | 3 cracks, $L_{c1} = 5, L_{c2} = 5, L_{c3} = 5$ |
| | 5 | 7 | 3 cracks, $L_{c1} = 10, L_{c2} = 5, L_{c3} = 5$ |
| | 5 | 8 | 3 cracks, $L_{c1} = 10, L_{c2} = 10, L_{c3} = 5$ |
| | 5 | 9 | 3 cracks, $L_{c1} = 10, L_{c2} = 10, L_{c3} = 10$ |
| | 5 | 10 | 3 cracks, $L_{c1} = 15, L_{c2} = 10, L_{c3} = 10$ |
| | 5 | 11 | 3 cracks, $L_{c1} = 15, L_{c2} = 15, L_{c3} = 10$ |
| | 5 | 12 | 3 cracks, $L_{c1} = 15, L_{c2} = 15, L_{c3} = 15$ |

In Figure 15-a and 15-b, two representative observations belonging to the healthy and cracked (with label 0 and 12) are presented from WTB-VibClimate through exciting the blade with white noise; both signals were recorded through channel 1 (sensor a_1 in Figure 14) and at temperature 20 °C (Wp20).




Figure 15. Raw signals from WTB-VibClimate for (a) healthy and (b) cracked conditions

4.3 Carbon-Epoxy Composite PLATE (CONCEPT)

The second dataset, CONCEPT: CarbON-epoxy Composite PlaTe (Ferreira et al., 2024), consists of Lamb wave recordings obtained from a unidirectional carbon-

epoxy laminate plate in both undamaged and damaged states. The experiments were structured to examine how temperature variation and progressive damage affect the laminate's structural condition.

The plate was instrumented with 4 Lead Zirconate Titanate (PZT) transducers supplied by Acellent Technologies. PZT1 was designated as the actuator, whereas PZT2, PZT3, and PZT4 operated as sensing elements. Testing was conducted under free-free boundary conditions to minimise support-induced restrictions on wave propagation and to better reflect the intrinsic dynamic response of the plate. Figure 16 depicts the test rig and measurement configuration, and Table 2 summarises the principal equipment.




Figure 16. The test rig setup in the CONCEPT test

Table 2. 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 | 1 Actuator, 3 Sensors |
| Mounting Conditions | Boundary Condition | Free-Free |
| Temperature Control | Range | 0°C to 60°C |
| | Increment | 10°C |
| Excitation Signal | Type | 5-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 |

Temperature regulation was provided by a Thermotron thermal chamber, which enabled controlled and repeatable environmental conditions. The chamber uses a cascade refrigeration architecture, and the 0 °C setting was achieved through mechanical cooling rather than reliance on ambient sub-zero conditions. Excitation was applied using a sinusoidal tone burst, after which the corresponding responses were recorded at the specified sampling rate. Data generation and acquisition were performed with dedicated hardware controlled through LabVIEW.

For the undamaged plate, measurements were collected at 7 temperatures between 0 °C and 60 °C in 10 °C steps, labelled Cp0, Cp10, Cp20, Cp30, Cp40, Cp50, and Cp60, with 100 recordings obtained at each level. For the damaged cases, data were acquired only at 30 °C, which was treated as the reference condition, and no damaged measurements were taken at the remaining temperatures.

Damage was introduced by placing industrial adhesive putty on the plate surface to emulate delamination-like defects. The affected area was expanded progressively within the region between PZT1 and PZT2 to evaluate how defect growth influences wave attenuation and propagation characteristics. Table 3 summarises the defined health states, their severities, and brief descriptions, and it also includes the temperature labels. The dataset has been used in several studies to assess SHM

methods designed to mitigate EOVs (da Silva et al., 2021; Almeida et al., 2024; Ferreira et al., 2024).

Table 3. Damage scenarios and severities for simulated defects in the CONCEPT experiment

| Damage Scenario | Severity | Label | Description | Temperature | Temperature Label |
|-----------------|----------|-------|-------------------------------|-------------|-------------------|
| Healthy | 0% | C0 | No damage | 0 | 0 |
| | | | | 10 | 1 |
| | | | | 20 | 2 |
| | | | | 30 | 3 |
| | | | | 40 | 4 |
| | | | | 50 | 5 |
| | | | | 60 | 6 |
| Damaged D1 | 0.196% | C1 | Industrial putty | 30 | 3 |
| Damaged D2 | 0.282% | C2 | Increased coverage of putty | | |
| 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 | | |

Figure 17-a, 17-b show two signals recorded at 30 °C (Cp30) corresponding to the healthy and the most severely damaged conditions, labelled C0 and C11 in the CONCEPT dataset. Both signals were recorded by channel 1 (PZT2 in Figure 16). The juxtaposition of these two extreme damage states at a controlled temperature isolates damage-induced signal variations from thermal effects, providing a baseline

comparison for subsequent analyses of temperature-dependent distribution shifts. The selection of C0 and C11 represents the maximum contrast available within the dataset, establishing the bounds of damage-related signal variation against which intermediate damage states and cross-temperature adaptations can be evaluated. This pairing serves to illustrate the fundamental challenge addressed by the proposed methods: discriminating between damage states whilst accounting for confounding environmental factors that may produce signal alterations of comparable or greater magnitude than the damage signatures themselves.




Figure 17. Raw signals from CONCEPT for (a) class C0 and (b) C11

From Figure 17, it can be observed that there is a decrease in the recorded maximum amplitude due to the presence of the damage (i.e., C11). The main wave packet in the damaged case is more attenuated, which is consistent with additional energy loss and scattering introduced by the putty-based defect.

However, this “healthy vs. most severe damage” comparison represents an easy visual case. In practice, the differences between neighbouring damage severities can be much smaller, and temperature-driven variability can introduce changes in amplitude, phase, and waveform shape that may partially mimic (or hide) damage-related signatures. In addition, each recording is a high-dimensional time series, so directly feeding raw samples to a learning algorithm is often inefficient and sensitive to noise and operational variability.

Therefore, beyond the visually apparent amplitude reduction, a more sophisticated approach is required to capture both the explicit and the hidden waveform differences and convert them into compact numerical representations (features) that are readable and analysable by machine learning frameworks, enabling reliable comparison across health states and operating conditions.

Chapter 5: Results and Discussion

5.1 Overview

In this chapter, the impact of the designed damage detection frameworks on the case studies presented in the previous chapter is demonstrated and analysed. Prior to this, the influence of environmental and operational variabilities on the collected signals is examined. The augmented datasets are then validated using standard evaluation metrics.

5.2 Effects of EOV

EOV can alter both the underlying system physics and the distribution of observations across source and target domains, which can degrade downstream classification performance (Kamariotis et al., 2025). Characterising how EOV influences the system, one can design appropriate feature-extraction pipelines and DA methods to neutralise its effects as far as possible. In many cases, raw time-domain signals do not reveal EOV, because such variability typically modifies stiffness, damping, and wave-propagation speed in the monitored structure. More diagnostic representations are therefore required, for example frequency-response functions (FRFs) (Lee and Shin, 2002), mode-shape analysis, and Coda Wave Interferometry (CWI) (Chen et al., 2025).

In FRF analysis, the response of the structure to a known input is represented in the frequency domain as the ratio of output to input, $H(\omega) = Y(\omega)/F(\omega)$. From this, shifts in resonance frequencies, changes in modal damping, and variations in modal residues can be estimated to quantify EOV effects. CWI is a technique that estimates small fractional changes in wave speed, dv/v , by correlating the late, multiply scattered coda of repeat waveforms; positive dv/v indicates faster propagation, negative dv/v indicates slower propagation.

The practicality of each method, and possible alternatives, depends on factors such as the availability and calibration of actuation, sensor bandwidth and signal-to-noise ratio, the aperture and density of the sensing array, dispersion and multimodality of the wavefield, the degree of linearity and time invariance over the measurement

window, access constraints, computational budget, and the availability of a suitable baseline or labels.

For instance, FRF-based approaches are often impractical for thin composite plates excited with guided waves using PZT transducers, because the input from a bonded PZT is a distributed, frequency-dependent traction rather than a calibrated point force; the guided-wave field is dispersive and multi-modal, so a single-input single-output FRF mixes modes and obscures modal parameters; thin plates exhibit closely spaced, lightly damped modes that require dense spatial sampling to separate; and small drifts in boundary conditions or temperature violate the approximate time-invariance assumed when comparing transfer functions. Consequently, CWI and other wavefield-based methods can be more effective in such settings.

In this section, the EOV effects on the physics of the mechanical systems were analysed in two case studies, WTB-VibClimate and CONCEPT, which were subsequently used to validate the proposed SHM frameworks. For WTB-VibClimate, which comprised low-frequency vibration responses under shaker excitation, mode-shape and FRF analyses were applied. For CONCEPT, which comprised guided-wave measurements, CWI was employed. Collectively, these characterisations of EOV were used to guide feature selection and the DA strategy adopted in the ensuing experiments; results were presented in the following sections.

5.2.1 WTB-VibClimate

The purpose of this analysis is to quantify how temperature affects mode shapes and FRFs of the WTB-VibClimate system. To examine this effect, FRFs were estimated in MATLAB® using the H1 estimator from healthy measurements at all available temperatures labelled Wn15 to Wp40. Twenty runs per temperature were processed with channel 1 as the response and the force channel as the input. Signals were trimmed by discarding the first 10000 samples and the last 20000 samples, detrended, and band-pass filtered between 0.5 and 380 Hz using a zero-phase finite impulse response filter of order 800. Spectral estimates for FRF computation employed Welch's method with a sampling rate of 1666 Hz, a Hann window of 4 seconds, 50% overlap, and a fast Fourier transform length of 8192.

Building on this acquisition and preprocessing, FRFs were pooled across repeated runs using weights proportional to the squared coherence, retaining only

frequency bins with squared coherence at or above 0.8 for averaging, while plots were masked below 0.7. Modal candidates were identified at the baseline temperature of 10 °C using peak picking with a maximum of 4 modes, a minimum peak prominence of 3% of the baseline maximum, a minimum inter peak spacing of 1 Hz, and a matching tolerance of 1 Hz for ridge initialisation. Modal ridges were then tracked across temperatures within non overlapping frequency bands defined around the baseline peaks, and half power bandwidth calculations were set up to provide damping estimates for each mode. When multiple response channels are available, mode shapes are obtained from the complex FRFs at each modal peak and normalised across channels at the baseline, with an optional output only frequency domain decomposition check at the baseline based on the first singular value of the response spectral density matrix. The corresponding mode shape plots and FRF maps are presented in Figure 18-a and Figure 18-b, respectively.




Figure 18. (a) Mode shapes identified at the baseline temperature of 10 °C, based on normalised frequency response functions; (b) FRF–temperature maps for the WTB-VibClimate dataset showing the variation of response magnitude across temperature levels

Figure 18-a shows 4 tracked modal frequencies that are highest at -15 °C and decrease approximately linearly as temperature rises to $+40$ °C, with the largest absolute shift in the mode near 300 Hz, a smaller but clear shift around 215–220 Hz, and more modest changes in the lower bands near 140 Hz and 120 Hz. The dashed horizontal markers indicate the 10 °C baseline, with the curves lying above the baseline at sub-zero temperatures and below it at warmer conditions. Figure 18-b corroborates these trends: bright FRF ridges remain in the same modal bands but migrate towards lower frequency with increasing temperature, and the slight thickening and reduced intensity of the ridges at higher temperatures suggest modest peak broadening consistent with increased damping. No mode crossing is evident, so mode ordering is preserved across the range. Moreover, parallel to an increase in the temperature, a decline in the natural frequency is visible (the ridges are moving to the left) more

specifically for the last two highest frequencies. Collectively, the figures demonstrate a systematic temperature dependence of the dynamic response that is consistent with thermal softening and that motivates temperature compensation when using these data for classification.

5.2.2 CONCEPT

For the CONCEPT guided-wave data, temperature-induced EOV was quantified using CWI. A baseline at 30 °C was formed by median-stacking the kept runs. For each temperature level (0–60 °C), every observation that passed force-channel quality screening was pre-processed (detrrending and zero-phase band-pass filtering around the burst) and compared against the baseline through a stretch-correlation search over $\epsilon \in [\pm 0.5\%]$. Estimation was carried out using a moving-window approach (40 μ s windows, 10 μ s hop) with a window-quality gate (median $\rho \geq 0.90$). For each temperature, run-wise dv/v estimates were aggregated using the median, with the mean also reported for completeness.

The analysis was performed only on the healthy plate and Channel 1, as this sensor was later shown to have the highest importance (owing to its augmented amplitude). At each temperature, all 100 observations were processed, and the curves reported in the figures correspond to the median and mean aggregated across those observations. The results of this analysis are shown in Figure 19-a and 19-b.




Figure 19. (a) CWI correlation–stretch map and (b) Fractional wave-speed change versus temperature for CONCEPT dataset

Analysing Figure 19-a, one can observe that a clear, near-monotonic decrease in dv/v with temperature was observed, crossing approximately 0 at the 30 °C baseline. Approximate values were +0.43% (0 °C), +0.22% (10 °C), –0.20% (20 °C), close to 0% (30 °C), about –0.47% (40 °C), and roughly –0.50% (50-60 °C). The median and mean curves were essentially coincident across all temperatures, indicating low

between-run dispersion and a stable, repeatable temperature effect. The sign pattern implied faster propagation at sub-baseline temperatures and slower propagation at elevated temperatures, with a mild saturation of the negative dv/v beyond 40 °C.

From Figure 19-b, similarly it can be understood that the ridge of maximum correlation, $\epsilon(T)$, moved from negative stretch at low temperatures to positive stretch at higher temperatures, consistent with the dv/v trend (since $dv/v = -\epsilon$). Correlation along the ridge remained high (0.9 or above), suggesting that temperature primarily induced a quasi-uniform phase dilation rather than substantial waveform distortion. This behaviour indicated that the dominant EOV mechanism in the guided-wave setting was a small fractional change in wave speed.

5.3 Data Augmentation

To evaluate the performance of the designed SHM frameworks, reliance was placed on datasets that captured signals from multiple health scenarios under varying environmental and operational conditions, with temperature being the primary factor. As outlined in the case studies, limitations were noted in both WTB-VibClimate and CONCEPT. In WTB-VibClimate, approximately 20 observations were available for the healthy condition and only 5 for damaged conditions across different temperatures. In CONCEPT, 100 healthy observations were available across temperatures, whereas damaged observations were present only at the baseline of 30 °C and were absent at other temperatures.

To address these gaps, dataset-specific procedures were adopted. For WTB-VibClimate, two methodologies were implemented. First, SMOTE was applied within a compact feature space derived from the vibration signals, the augmented dataset was saved, and it was used within one SHM framework. Second, a short-time windowing procedure was implemented, whereby records were segmented into fixed-length, tapered windows with controlled overlap under a local stationarity assumption, and the resulting dataset was used in an alternative SHM framework. For CONCEPT, two methodologies were also implemented. Spline interpolation was used along the temperature dimension within feature or latent representations to estimate damaged-condition characteristics at unmeasured temperatures, and the resulting dataset was used in a dedicated framework. In a separate pathway, a Conv-CVAE was deployed,

conditioned on temperature and damage labels, to synthesise guided-wave responses, and the generated dataset was used independently in another framework.

These choices were motivated by the data type and underlying physics. Windowing was suited to broadband vibration behaviour when local stationarity could be assumed, whereas in guided waves the coexistence of A0 and S0 modes, boundary reflections, dispersion and mode conversion produced overlapping arrivals that rendered naive windowing impractical, so model-based synthesis was preferred.

Implementation details were provided in the subsequent sections for each case study. The outcomes of the augmentation procedures were examined through domain-relevant visualisations and comparative diagnostics.

5.3.1 SMOTE

To enlarge the WTB-VibClimate dataset in the first instance, the SMOTE framework was applied, increasing the number of damaged observations per scenario from 5 to 20. SMOTE expands minority classes by creating artificial samples through interpolation between existing instances. In practice, it selects a minority-class sample together with k nearest neighbours and then generates new points along the line segments linking these samples in the feature space.

The value of k was set to 2 after empirical testing in which several candidate values were evaluated using MMD. The choice $k = 2$ produced the most favourable outcome, producing synthetic data that remained close to the original distribution while still providing meaningful variability. Following augmentation, each health condition in both the source and target domains comprised 20 samples.

Effective augmentation requires a careful trade-off between diversity and fidelity. If synthetic samples are too similar to the originals, augmentation adds little value and does not improve generalisability. Conversely, if the generated samples vary excessively, they may drift away from the defining characteristics of the class and introduce misleading patterns.

To evaluate the impact of SMOTE augmentation, MMD was computed for the healthy WTB case and for 12 faulty conditions. For the healthy class, the comparison was constructed by randomly drawing 5 samples from the original dataset and 15 samples to represent the synthetic set. MMD quantifies how far the original and

generated samples depart from one another by mapping both sets into a high-dimensional space through a kernel and then comparing their mean embeddings. In this study, values above 0.1 were treated as indicating substantial distributional change, which can improve generalisation but may weaken class distinctiveness. Scores in the range 0.01 to 0.1 were interpreted as reflecting moderate alignment, whereas values below 0.01 were regarded as overly similar and therefore less beneficial for augmentation.

An RBF kernel was selected for these calculations because it is well suited to high-dimensional data and can capture non-linear structure. The kernel parameter was set to $\gamma = 1.0$ to balance sensitivity to distributional differences with computational cost. Using this configuration, the resulting MMD scores indicated that SMOTE injected useful variability into the faulty classes while preserving close agreement with the underlying original distributions. Figure 20 reports the MMD outcomes across all 13 classes, with emphasis on data collected at Wn10 and Wp35 from channel 1. All computations were performed in Python within a Jupyter environment. The augmentation and evaluation were applied only to channels 1 and 3 because, as discussed later, the UCTRF approach used with this dataset was configured to operate on signals from a single channel.




Figure 20. Comparison of the original and synthesised data using SMOTE for WTB-VibClimate

Across the 13 classes, the MMD results for Wn10 and Wp35 indicate that SMOTE augmentation preserves the underlying data distribution while still adding

appreciable variation. In the healthy case, where 5 original samples were compared with 15 synthetic samples selected from the experimental measurements, the MMD scores were 0.04 and 0.05. Both values lie within the 0.01 to 0.1 interval, suggesting that the generated samples are neither near-duplicates of the originals nor excessively displaced from them. This supports the view that the augmentation introduces useful diversity without compromising distributional consistency. As such, the augmentation process supports improved model generalisation without compromising the fidelity of the original class characteristics. This analysis thus reinforces the suitability of SMOTE for enhancing dataset variability in a controlled manner, thereby contributing to the robustness and adaptability of the proposed model.

5.3.2 Signal Windowing

In the second approach to expand the WTB-VibClimate dataset, a windowing strategy was employed to produce additional samples for training and testing, facilitated by the uniform sampling frequency maintained during the experiment.

Each original observation was divided into 5 equal segments of 39200 data points. To ensure accurate segmentation and avoid edge effects or noise artefacts at signal boundaries, the first 4000 data points and the last 1 data point were excluded, producing a clean segment of 196000 data points evenly divided into 5 non-overlapping windows. Although each observation contains readings from 8 sensors, only data from 3 accelerometers (channels 1, 4, and 8) were retained to reduce computational complexity and assess framework performance under sensor-limited conditions. Any future reference to channels 1, 2, and 3 in the WTB-VibClimate assumption within SPADA corresponds to channels 1, 4, and 8 of the original dataset.

Following windowing, 25 observations were derived for each damaged condition. The same procedure was applied to the healthy condition (using the first original observation), resulting in a dataset of dimensions $325 \times 3 \times 39200$ for each temperature level given 13 health scenarios. To prevent data leakage between training, DA, and testing phases, 15 target domain observations were allocated to DA, whilst the remaining 10 were withheld for testing. The 10 test observations were generated from 2 original recordings, whilst the 15 DA observations were obtained from 3 separate original observations.

5.3.3 Spline Interpolation

Spline interpolation was used to produce synthetic signals for damaged cases at temperatures that were not experimentally recorded, specifically 0 °C, 10 °C, 20 °C, 40 °C, 50 °C, and 60 °C. The procedure was implemented in MATLAB® on a per-sensor and per-temperature basis and involved both interpolation and extrapolation, with 30 °C treated as the reference condition. Signal synthesis relied on temperature-dependent scaling factors estimated from healthy-state measurements across the full temperature range, which were then applied to the reference damaged signals to emulate the expected thermal influence. To introduce realistic sample-to-sample variation, a bounded perturbation of plus or minus 1% was added to the scaled outputs (Rezazadeh et al., 2025b). Using this strategy, 100 observations were generated for each damaged state at each target temperature and for each sensor.

The synthesis method was then evaluated using data at the source temperature Cp30. In this assessment, 50 healthy observations were randomly selected to construct the synthetic signals, and the remaining 50 observations were reserved for validation. Figure 21 compares a measured healthy example from Cp0 at the baseline temperature with its synthesised counterpart across the 3 sensor channels.




Figure 21. Experimental and synthesised time-domain responses of the healthy plate recorded by (a) PZT2, (b) PZT3, and (c) PZT4

Figure 21 indicates strong agreement between the measured and generated time-domain responses for all 3 sensors. To quantify how well the synthesised signals

reproduce the experimental behaviour, 2 evaluation measures were considered: DTW (Senin, 2008) and Cross-Correlation (CC) (Habermehl et al., 2021). The assessment was performed at the reference temperature of 30 °C, for which recordings were available for every health scenario. For each class, 50 observations were used to construct synthetic counterparts, and the remaining 50 observations were reserved as a hold-out set for direct comparison. This split was applied consistently across classes to maintain a uniform basis for evaluation.

DTW evaluates similarity by identifying the minimum-cost alignment between 2 time series, with values approaching 0 indicating closer correspondence between experimental and synthesised signals. CC measures their linear association, where coefficients near -1 or 1 reflect strong agreement in waveform shape, including amplitude and phase trends. The validation results are summarised in Figure 22-a and Figure 22-b, which report mean DTW and mean CC values averaged across the 3 sensors.




Figure 22. Agreement between spline-generated and measured signals at Cp30 across PZT2, PZT3, and PZT4, reported using (a) average DTW and (b) average CC

The validation metrics indicate that the synthesis strategy performs well: DTW remains low and CC is consistently near 1 for all sensors, implying close correspondence between generated and measured responses. PZT2 exhibits a higher DTW value (0.122) than PZT3 (0.034) and PZT4 (0.0321), which is plausibly linked to the larger signal amplitudes observed for PZT2. Despite this, the CC coefficients are uniformly high, with values of 0.99994 for PZT2, 0.99977 for PZT3, and 0.99979 for PZT4. Collectively, these results suggest that the synthesised signals preserve both

waveform timing and amplitude behaviour, and that the method remains reliable even in the presence of sensor-dependent amplitude differences.

5.3.4 Conv-CVAE

Conv-CVAE as a temperature-conditioned generative augmentation method was implemented using PyTorch and Jupyter Notebook to complete CONCEPT dataset. The Conv-CVAE model was trained using 150 epochs with a batch size of 128, a latent dimension of 64, and a learning rate of $3 * 10^{-4}$, optimised using the AdamW optimiser. Training data consisted of 100 healthy observations across each of the 7 temperature levels (Cp0 to Cp60, in 10 °C intervals), normalised using global statistics computed from the full healthy dataset.

The encoder structure involved a sequence of 3 one-dimensional convolutional layers with filter sizes of 7, 5, and 5 and stride 2, each followed by batch normalisation and GELU activations; the decoder mirrored this with transposed convolutions for signal reconstruction. Temperature was encoded as a normalised scalar and concatenated at the fully connected layer in both encoder and decoder paths. For each damage level (12 levels including healthy), a latent offset vector Δz was computed at the reference temperature (Cp30). Additionally, per-sensor and per-time amplitude masks were derived from the Hilbert envelopes of healthy and damaged signals at Cp30; median envelope ratios were Gaussian-smoothed ($\sigma = 6$) and their dispersion estimated as 15% of the interquartile range. The loss combined mean-squared error, a spectral term on the log-magnitude FFT, wavelet-domain similarity (4-level Daubechies-4), envelope consistency, and KL divergence, with weights chosen to balance reconstruction fidelity and latent regularisation.

After training, damaged signals at any temperature CpT were generated by (i) encoding a healthy seed at CpT , (ii) adding the corresponding damage offset Δz_d , and (iii) decoding with the same temperature condition. The decoded waveform was then modulated by the learned amplitude mask for damage level d ; stochastic variability was introduced by sampling a zero-mean Gaussian scaled by the mask's standard-deviation map. This produces temperature-consistent, damage-conditioned waveforms derived from healthy seeds. For each damage condition and temperature, 30 synthesised observations were generated. To ensure a balanced dataset, 30 observations were randomly selected from the healthy class at the same temperature.

Consequently, each saved dataset from Cp0 to Cp60 contains a total of 360 signals with dimensions $360 \times 3 \times 1000$. On the condition of Cp30, from each class 30 observations were selected randomly from the original set, and its synthesised dataset was discarded.

In the same manner as spline interpolation (i.e., for the Cp30), DTW and CC were calculated to evaluate the effectiveness of the deployed Conv-CVAE in data generation; the outcomes are presented in Figure 23-a and 23-b.




Figure 23. Performance of Conv-CVAE signal synthesis at Cp30 across PZT2, PZT3, and PZT4, summarised by (a) average DTW and (b) average CC

Figures 23-a and 23-b demonstrate that Conv-CVAE reproduces Cp30 responses with high fidelity. Mean DTW values are very low for PZT2 (0.000362), PZT3 (0.000103), and PZT4 (0.000094), whilst mean CC values remain high at 0.990627, 0.983301, and 0.984319, respectively, indicating close matching in waveform shape and amplitude. Compared to spline interpolation, DTW scores decrease by more than 330 times across all sensors whilst CC remains at levels indicating reliable similarity, confirming that Conv-CVAE captures the principal temporal and structural characteristics of measured signals.

The superior DTW performance highlights Conv-CVAE's ability to capture complex, non-linear timing variations that spline-based methods cannot represent, stemming from its temperature-conditioned latent representations and learned amplitude masks. Whilst CC is marginally reduced compared to spline interpolation, values remain high, indicating maintained fidelity without input duplication.

For Conv-CVAE-generated damaged scenario C12, CWI was applied to verify whether temperature-driven EOV patterns match those in the experimental healthy plate, using the same setup as the healthy CONCEPT analysis. Force-channel-based quality screening was not possible due to unavailable excitation traces. At each temperature (0–60 °C), 30 synthesised observations per PZT were compared against the 30 °C damaged baseline. Run-wise dv/v estimates were aggregated using median (with mean reported for consistency), and correlation–stretch curves were similarly averaged. Results are shown in Figure 24-a (dv/v –temperature trends) and Figure 24-b (correlation–stretch maps referenced to 30 °C damaged baseline).




Figure 24. CWI analysis of synthetic scenario C12: (a) correlation–stretch surface; (b) temperature-velocity relationship

In Figure 24-a, the dv/v –temperature trends, and in Figure 24-b, the ϵ^* ridge trajectories, mirror the qualitative patterns seen for the experimental healthy plate. Specifically, dv/v is positive at the colder conditions, approaches and passes through zero around the 30 °C reference, and becomes negative as temperature increases, with

broadly comparable amplitudes for all 3 PZTs. Consistent with this, the location of the peak correlation migrates smoothly from negative to positive stretch as temperature rises, while the correlation level remains uniformly high. Together, these observations suggest that the Conv-CVAE outputs primarily reflect a coherent, temperature-driven time-stretching mechanism, rather than non-physical or erratic waveform deformation. Accordingly, the Conv-CVAE can be considered to reproduce the intrinsic temperature sensitivity of guided-wave propagation in the composite plate, enabling a physics-informed semi-synthetic augmentation of the experimental data.

5.4 Damage Detection

In the following sections, the results of damage detection using the 3 proposed frameworks are presented and discussed. The WTB-VibClimate case study was examined for the UCTRF framework. The CONCEPT case study was employed under the GAT-CAMDA framework. Both case studies were analysed in the context of the SPADA framework.

5.4.1 UCTRF

The performance of the UCTRF damage detection framework was investigated through sequential assessments covering DR, CapsNet feature extraction, damage classification without and with DA (full UCTRF pipeline), benchmarking against 2 alternative TL approaches (DANN and TCA), and examination of instance weight evolution during adaptation. Validation was conducted using the WTB-VibClimate case study with the SMOTE-augmented dataset. Single accelerometer channels were used, with channels 1 and 3 evaluated independently to characterise framework capability. Across all experiments, Wp5 served as the source domain, whilst Wn15, Wn10, Wn5, Wp15, Wp25, and Wp35 were separate target domains.

5.4.1.1 Reducing the Dimensionality

After enlarging the dataset with SMOTE, UMAP was applied to compress each vibration signal to a 200-dimensional representation. This step was intended to simplify the high-dimensional input, retain the underlying topological structure, and support improved class separation prior to feature extraction with CapsNets. The suitability of UMAP was examined by benchmarking it against PCA, a commonly adopted DR baseline. Figure 25-a and Figure 25-b display the corresponding scatter

plots for Wp5 from channel 1, allowing a direct visual comparison between the 2 techniques.




Figure 25. Visualisation of the Wp5 embeddings obtained using (a) UMAP and (b) PCA

The projections suggest that UMAP yields more distinct class groupings than PCA. In Figure 25-a, the clusters are more clearly partitioned, whereas Figure 25-b shows a more congested embedding in which several classes concentrate in comparatively compact regions. To move beyond visual inspection, clustering quality was evaluated using the silhouette score. This index ranges from -1 to 1 and compares within-cluster compactness with between-cluster separation, with larger values indicating tighter and more isolated clusters. Such separation is important when DR is used to support class discrimination and subsequent analysis.

Using the mean silhouette score across the 13 classes, UMAP produced 0.5793 , while PCA produced -0.3067 . These values indicate that UMAP forms substantially more coherent and separable groupings for this dataset. Figure 26 further reports per-class silhouette scores in bar-chart form.




Figure 26. Comparative evaluation of clustering performance: (a) UMAP silhouette scores and (b) PCA silhouette scores

Despite UMAP outperforming PCA, Figure 26-a still shows that certain classes are not fully separable, with noticeable inter-class mixing in parts of the embedding. This residual overlap indicates that dimensionality reduction alone is insufficient and that more discriminative feature learning is required. For this reason, CapsNets were applied to the UMAP-compressed representations to learn richer features and improve class distinguishability.

5.4.1.2 Feature Extraction through CapsNets

In UCTRF, the last feature engineering step uses CapsNets to process the embeddings produced by UMAP. This stage required specification of the model settings listed in Table 4. The resulting feature distributions for Wp5 and Wn15 from channel 1 are shown as scatter plots in Figure 27-a and Figure 27-b, respectively. The hyperparameter values were selected through iterative manual tuning to obtain the best observed performance.

Table 4. Hyperparameter values used in CapsNets

| Parameter | Value | Parameter | Value |
|------------------------------|-------|-------------------|-------|
| Number of epochs | 200 | m_plus | 0.9 |
| Batch size | 15 | m_minus | 0.1 |
| Learning rate | 0.001 | Loss_lambda | 0.5 |
| Number of routing iterations | 3 | Type of optimiser | Adam |




Figure 27. Scatter plot of after applying CapsNets for (a) Wp5 and (b) Wn15

Figure 27 reveals that the spatial distribution of class observations differs notably between Wp5 and Wn15 as will be considered as the source and target domains, respectively. For instance, the region associated with label 8 observations (indicated by yellow-coloured shapes enclosed within the red dashed circle) has shifted between domains. This shift illustrates a domain discrepancy, which hinders the ability of the

network trained exclusively on source domain samples to accurately classify observations from the target domain.

5.4.1.3 Damage Detection without DA

The UCTRF framework was initially implemented without the DA stage (TrAdaBoost), by training the classifier (random forest) using Wp5 and evaluating it on Wn15 under two assumptions, i.e., signals from channel 1 and channel 3 separately, following the UMAP and CapsNets stages. From Wn15, 5 observations per class were randomly selected for the test phase. The number of estimators in the random forest was set to 20, and the learning rate was fixed at 0.1. These values were determined through manual trial and error. The damage classification results for the channel 1 (Ch1) and channel 3 (Ch3) scenarios are presented in Figure 28-a and 28-b, respectively.




Figure 28. Confusion matrices of UCTRF with Wp5 as the source and Wn15 as the target domain without DA for (a) channel 1 and (b) channel 2

The confusion matrices indicate that, in the absence of DA, the classification model encountered significant difficulty in correctly identifying observations across different domains with respect to both fault severity and type. Specifically, the model achieved an overall accuracy of approximately 11% when using channel 1, and around 9% when utilising channel 3. These results highlight the considerable challenge posed by domain shifts and underscore the necessity of incorporating DA techniques to improve classification performance.

5.4.1.4 Damage Detection with DA

This section presents the performance of the UCTRF framework when it was applied on its whole format using the augmented dataset of WTB-VibClimate through SMOTE. To this end, Wp5 was considered as the source domain, with samples from 6 selected temperatures, i.e., Wn15, Wn10, Wn5, Wp15, Wp25 and Wp35 serving as individual target domains. It is important to note that this process was conducted separately for channel 1 and channel 3. The average performance metrics for each channel are reported in Table 5 (channel 1) and Table 6 (channel 3).

Overall damage detection performance is summarised using the F1 score, aggregated across all damage states and evaluated for each target-domain condition. The results are shown in Figure 29-a for channel 1 and Figure 29-b for channel 3. For every target domain, 15 samples per class were used during the DA stage, covering both training and validation, and the remaining 5 samples per class were held out for testing.

Table 5. Damage classification results applying UCTRF through channel 1

| Target | Macro average of precision | Macro average of recall | Macro average of F1-score | Macro average of accuracy |
|--------|----------------------------|-------------------------|---------------------------|---------------------------|
| Wn15 | 0.99 | 0.98 | 0.98 | 0.98 |
| Wn10 | 0.99 | 0.98 | 0.98 | 0.98 |
| Wn5 | 0.97 | 0.97 | 0.97 | 0.97 |
| Wp15 | 1.00 | 1.00 | 1.00 | 1.00 |
| Wp25 | 0.97 | 0.97 | 0.97 | 0.97 |
| Wp35 | 0.99 | 0.98 | 0.98 | 0.98 |

Table 6. Damage classification results applying UCTRF through channel 3

| Target | Macro average of precision | Macro average of recall | Macro average of F1-score | Macro average of accuracy |
|--------|----------------------------|-------------------------|---------------------------|---------------------------|
| Wn15 | 0.96 | 0.95 | 0.95 | 0.95 |
| Wn10 | 1.00 | 1.00 | 1.00 | 1.00 |
| Wn5 | 0.98 | 0.97 | 0.97 | 0.97 |
| Wp15 | 1.00 | 1.00 | 1.00 | 1.00 |
| Wp25 | 0.98 | 0.97 | 0.97 | 0.97 |
| Wp35 | 0.94 | 0.92 | 0.92 | 0.92 |




Figure 29. UCTRF performance across temperature conditions and damage categories for (a) channel 1 and (b) channel 3

Figure 29 together with Tables 5 and 6 shows that channel 1 yields the strongest results across the evaluated metrics. Across all temperature conditions, the macro-averaged precision, recall, F1-score, and accuracy remain above 0.95. At $-15\text{ }^{\circ}\text{C}$ and $-10\text{ }^{\circ}\text{C}$, precision and recall are greater than 0.99, and both F1-score and accuracy reach 0.98. Comparable performance is observed at $-5\text{ }^{\circ}\text{C}$ and $+25\text{ }^{\circ}\text{C}$, where all metrics exceed 0.97. Moreover, the scores reach 1.00 for every metric at $+15\text{ }^{\circ}\text{C}$ and $+35\text{ }^{\circ}\text{C}$. These results are notable given that only 20 samples per class were available for training.

By contrast, channel 3, while still achieving strong overall results, shows a modest reduction in temperature classification accuracy, most noticeably at the extreme settings of $-15\text{ }^{\circ}\text{C}$ and $+35\text{ }^{\circ}\text{C}$. This pattern suggests that differences in sensor location or responsiveness can affect predictive performance, and it reinforces the suitability of channel 1 for extracting the features required for reliable inference across changing environmental conditions.

The confusion matrices for two representative target domains, i.e., Wn15 and Wp35, are presented in Figures 30-a and 30-b for channel 1. Out of 65 test observations, only 1 misclassification occurred, where a sample from the unbalanced scenario was incorrectly identified as representing a cracked WTB condition. For channel 3, the overall accuracy reached 98.46%, with a single error in which a severely

unbalanced WTB observation (class 3) was incorrectly assigned to the slightly unbalanced category (class 1).




Figure 30. Confusion matrices of UCTRF with Wp5 as the source and (a) Wn15 and (b) Wp35 as the target domain with DA for channel 1

For channel 3, Figures 31-a and 31-b present the confusion matrices for the -15°C and $+35^{\circ}\text{C}$ conditions, respectively. At -15°C , all samples were assigned to their correct classes, corresponding to an accuracy of 100%. At $+35^{\circ}\text{C}$, performance decreased to 92.31%, with 5 samples from class 2 incorrectly predicted, indicating reduced reliability for this class under high-temperature conditions.




Figure 31. Confusion matrices of UCTRF with Wp5 as the source and (a) Wn15 and (b) Wp35 as the target domain with DA for channel 3

The performance gap observed between channel 1 and channel 3 is largely explained by how their DA models were developed. Separate models were trained for

each channel using the same source domain but with independent adaptation and tuning procedures, and they were then assessed on target-domain data spanning multiple temperature settings. As a result, channel-specific differences in the training data used during adaptation and in the subsequent parameter updates likely contributed to the differences in classification accuracy.

Channel 1, positioned closer to the base of the WTB, likely captured vibration signals that contained a greater proportion of structurally relevant information, even within a controlled experimental environment. As a result, the model trained on data from channel 1 learned features and patterns that were particularly sensitive to temperature-induced variations, contributing to its consistently strong classification performance.

In contrast, channel 3, located further from the base, recorded signals influenced by different sections of the blade and the associated material characteristics. This may have hindered the model's ability to generalise effectively across temperature conditions. The combination of variation in the training data and the sensor's physical placement relative to the blade structure likely contributed to the slightly reduced classification accuracy observed for channel 3, especially under extreme temperature scenarios.

In summary, the disparities in classification outcomes can be attributed to the unique training conditions experienced by each sensor. These findings highlight the importance of sensor placement and training data characteristics in determining a model's capacity to accurately classify temperature conditions within a controlled experimental setting.

5.4.1.5 Comparison Study

The effectiveness of UCTRF was compared with two DA benchmarks, i.e., DANN and TCA. For DANN, the same combined feature extraction approach, comprising DR using UMAP followed by CapsNets, was applied with Wp5 as the source domain and Wn15 as the target domain, using channel 1 and channel 3 separately. In the case of TCA, the raw vibration data (augmented using SMOTE) were directly input into its internal feature extraction process, as TCA is a feature projection technique that can serve as a feature extraction method. This also enabled an indirect comparison between the hybrid feature extraction stage used in UCTRF and the

approach embedded within TCA. The results obtained from these two DA alternatives are presented in the following sections.

5.4.1.5.1 Classification using DANN

The DANN model was optimised using Adam with a learning rate of 0.001 and a batch size of 64, while the gradient reversal layer was assigned a weight of 1.0. The domain adaptation outcomes are reported using confusion matrices in Figure 32-a for channel 1 and Figure 32-b for channel 3.




Figure 32. Confusion matrices of DANN with Wp5 as the source and Wn15 as the target domain for (a) channel 1 and (b) channel 3

The confusion matrices presented in Figures 32-a and 32-b show that the application of DANN resulted in classification accuracies of 83% for channel 1 and 85% for channel 3. In contrast, the proposed pipeline, i.e., UCTRF, demonstrated superior performance, as shown in the confusion matrices in Figures 30-a and 31-a. Specifically, UCTRF achieved an improvement of around 15% in classification accuracy for both sensor datasets, highlighting its enhanced capability in adapting to domain shifts and accurately identifying fault conditions under varying temperature scenarios.

5.4.1.5.2 Classification using TCA

After data augmentation, TCA was applied to the original vibration signals without an additional feature extraction step. This is because TCA performs its own transformation, mapping source and target samples into a shared latent space to reduce distributional mismatch.

An RBF kernel was adopted to accommodate non-linear structure in the signals. Model settings were selected via a grid search in which 3 candidate values were tested for each hyperparameter. The explored ranges and the chosen values are reported in Table 7.

Table 7. Hyperparameter grid search configuration for TCA implementation

| Parameter | Tested values |
|---------------------|---------------|
| Kernel bandwidth | 0.1, 1, 10 |
| Regularisation term | 0.01, 0.1, 1 |
| MMD weight | 0.5, 1, 2 |

Following the training, validation, and testing phases, the optimised TCA model was evaluated using Wp5 as the source domain for channel 1 and channel 3, resulting in two independently trained networks. As the target domain, Wn15 was chosen for each sensor to have a consistence comparison with the UCTRF and DANN. With the kernel bandwidth, regularisation term, MMD weight set to 1, 0.1, and 0.5 respectively, the TCA model achieved classification accuracies of 63% and 62% for channel 1 and channel 3, as illustrated in Figure 33-a and 33-b.




Figure 33. Confusion matrices of DANN with Wp5 as the source and Wn15 as the target domain for (a) channel 1 and (b) channel 3

Figure 33 highlights that, even in its optimised configuration, TCA fails to effectively project the source and target domain data into a shared feature space. The limited performance of TCA can be attributed to its static feature projection mechanism, which does not accommodate the dynamic nature of domain shifts. Unlike the developed framework (UCTRF), TCA lacks the capacity for hierarchical feature learning and adaptive representation, both of which are crucial for capturing complex

variations across domains. Consequently, its inability to adjust to differing data distributions results in suboptimal classification accuracy.

5.4.1.6 Weight Adjustments in Mitigating EOVs

Instance-based DA approaches such as TrAdaBoost address source–target mismatch by reweighting training samples over successive iterations, allowing the model to selectively emphasise informative instances whilst suppressing those that introduce negative transfer. Here, all source and target observations began with the same weight of 0.0022, calculated as the reciprocal of the total number of training samples. As boosting progressed, source samples that better matched the target distribution were emphasised through increased weights, whereas poorly aligned source samples were progressively downweighted. Within the target set, the weighting scheme prioritised difficult cases: misclassified samples received larger weights to focus learning on challenging decision boundaries, whilst samples consistently classified correctly were assigned smaller weights in later rounds.

To illustrate this adaptive weighting mechanism, source domain data from Wp5 and target domain data from Wn15, both from channel 1, were selected as representative examples. For visualising EOv mitigation, feature spaces with overlaps between different class observations and between source and target domain data were selected, as such overlaps demonstrate how the weighting scheme resolves ambiguous regions. Accordingly, features 26 and 84, which are of lower importance than the first and second features but retain discriminative information, were chosen.

Weight evolution across the adaptation process is depicted in Figures 34-a to 32-c, corresponding to the first, 500th, and final epochs, respectively. Weights were scaled to scatter point sizes for clarity: the initial weight of 0.0022 was scaled to point size 300, whilst other weights were adjusted proportionally, with minimum weights scaled to point size 50 to ensure visibility. These visualisations provide insight into how TrAdaBoost progressively emphasises informative instances whilst diminishing less relevant samples, thereby facilitating effective DA in the presence of EOVs.




Figure 34. Evolution of scaled instance weights in the source and target domains at (a) epoch 1, (b) epoch 500, and (c) the final epoch

The scatter plots show that source-domain samples were progressively downweighted, indicating their distribution differed from the target domain and many contributed little to accurate target classification. In contrast, target-domain weights generally declined because a large proportion were classified correctly from early stages, reflecting the effectiveness of the hybrid feature extraction stage. Nevertheless, a smaller group of target observations became increasingly influential, with rising weights when the model struggled to classify them. As shown in Figure 34, these higher-weight samples typically appear near boundaries of neighbouring class clusters. For instance, a light-blue sample from class 2 was strongly upweighted because it lay close to the orange cluster corresponding to class 9, and near a small set of black points from class 4.

Figure 35 presents the weight trajectories for multiple target domain instances at $-15\text{ }^{\circ}\text{C}$ across all epochs, plotted in their original magnitudes prior to scaling. Each instance represents one of the 13 classes and was selected because its weight exhibited greater fluctuation relative to its class peers. The mismatch between instance indices and class labels arises from the shuffled partitioning of the target domain data into training and testing subsets.




Figure 35. Per-class evolution of instance weights for selected Wn15 target-domain samples

It is worth noting that, for some fault scenarios, the weight trajectories of the selected instances are only slightly distinguishable from those of the remaining damage conditions. Even so, the amplitudes of almost all weights fell below 0.1, and

only one case, i.e., label 2, showed a sustained increase in its allocated weight up to the final epoch. Beyond the influence of proximity to clusters from other categories, any instance that behaves as an outlier tends to receive a higher weight during instance-based DA.

Across each temperature case, the DA process repeatedly updates the displayed instance weights over successive epochs until the predefined training limit is met, with the aim of improving separation and classification of the health states. Faster convergence is likely if the feature extraction stage produces more discriminative representations. Similarly, adopting a projection or feature learning approach that embeds source and target data in a common latent space could increase the contribution of source samples through higher weights and may further reduce the number of epochs required for convergence.

5.4.1.7 Practical Considerations and Applicability

The proposed framework was run on a high-performance computing cluster with 72 CPU cores and 200 GB of RAM. The UMAP-based DR stage took approximately 1.5 minutes per health scenario, CapsNet training required about 45 minutes for 200 epochs, and TrAdaBoost-based DA completed in around 5 minutes per boosting iteration.

Owing to its modular construction, the framework can be deployed on cloud infrastructure, including Microsoft Azure, where compute-intensive stages may be scheduled across virtual machines or containerised environments via Azure Batch. Working directly with vibration signal data also simplifies integration into established monitoring workflows, since it does not depend on specialised data formats. In addition, the relatively lightweight CapsNet and TrAdaBoost components support scalable execution and straightforward adaptation to different wind turbine blade configurations and EOVs.

5.4.2 GAT-CAMDA

This section investigates the capability of the GAT-CAMDA framework to detect damage under the EOv assumption, with the CONCEPT dataset used as the case study. The source domain is formed from measurements covering all health conditions at Cp30. Initially, the data were synthesised using the spline interpolation

approach introduced earlier. The end-to-end hybrid model, which combines DA with classification, was then evaluated on the basis that the test set was constructed by pooling the remaining augmented datasets, namely Cp0, Cp10, Cp20, Cp40, Cp50 and Cp60. A comparative study was undertaken to demonstrate the advantage of GAT-CAMDA; this analysis also determined which receiver sensors contributed most to the detection task and identified the hyperparameters that had a pronounced effect on the outcomes. Since GAT is a graph neural network method employed for feature extraction, features were obtained from all 3 channels and subsequently concatenated.

5.4.2.1 Damage Detection without DA

Before reporting the damage detection outcomes obtained with the full GAT-CAMDA configuration, a preliminary assessment was carried out to isolate the contribution of the feature extraction stage and to appraise the framework’s classification performance. In this assessment, the trained GAT-CAMDA was first evaluated using the source-domain test subset. The DA module was then removed, and the same trained network was tested on the target-domain test subset. All experiments were developed in Python within a Jupyter Notebook environment.

For this purpose, Cp30 was partitioned into training, validation, and testing subsets comprising 63%, 27%, and 10% of the data, respectively. Within the aggregated target dataset, 70% of the samples were assigned to DA, while preserving equal representation across the target temperatures via the ‘stratify’ option. A further 15% was used for validation, and the final 15% was retained as the target test set. As a result, the target test set contains 90 observations for each of the 12 health scenarios, C0 to C11, consisting of 15 samples at each of the 6 target temperatures. Table 8 summarises the composition of the source and target datasets.

Table 8. Allocation of samples to the source and target domains within the GAT-CAMDA framework

| Domain | Number of Observations (per class) | | |
|--------|------------------------------------|------------|---------|
| | Training | Validation | Testing |
| Source | 63 | 27 | 10 |
| Target | 420 | 90 | 90 |

The confusion matrices presented in Figure 36-a and Figure 36-b provide an overview of damage detection performance in 2 settings: the Cp30 test phase and the test subset drawn from the combined target dataset.




Figure 36. GAT-CAMDA classification results with DA removed for (a) the source-domain test set and (b) the target-domain test set

The confusion matrix in Figure 36-a demonstrates that the GAT employed is proficient at distinguishing between the different damage conditions. Furthermore, the classifier, trained on the source domain dataset, achieved perfect accuracy (100%) on its corresponding test set, highlighting the effectiveness of the training process. In contrast, Figure 36-b shows that when the same classifier is applied to the target test set, its performance drops significantly, achieving only around 36% accuracy. This reduction in performance is attributed to shifts in the feature distribution due to EOVs. These results underline the importance of incorporating a DA phase to maintain the framework’s effectiveness when transferring from the source to the target domain.

5.4.2.2 Damage Detection with DA

In this section, the limitations observed when applying GAT-CAMDA without the DA stage to damage identification in the composite plate are addressed by deploying the complete GAT-CAMDA framework, including its DA module. The model was run for 100 trials under an Optuna optimisation schedule to identify the most suitable framework setting, informed by the 13 hyperparameters listed in Table 9. The training duration was held constant at 50 epochs, and the optimisation objective was defined as the highest classification accuracy achieved on the target validation set.

Within the DA stage, class labels were not used so that the procedure operated in an unsupervised manner, consistent with industrial practice where labelled measurements are frequently scarce or expensive to acquire. A fixed random seed of 42 was adopted to support reproducibility, and GAT-CAMDA was developed and

executed in Python using Jupyter Notebook. Under the optimisation strategy described above, the best damage detection performance occurred at trial 56, as shown by the optimisation history in Figure 37.

Table 9. Hyperparameter settings used to run GAT-CAMDA

| Hyperparameter | Value | Hyperparameter | Value |
|--------------------|------------------------------|----------------------|----------------|
| Learning Rate | $[1 * 10^{-4}, 1 * 10^{-2}]$ | Number of GAT Heads | 1, 2, 4 |
| Weight Decay | $[1 * 10^{-5}, 1 * 10^{-3}]$ | 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 |




Figure 37. Progression of the hyperparameter optimisation process

Figure 37 indicates that hyperparameter choice strongly influences overall performance. For instance, trial 20 yielded an accuracy of approximately 20%, whereas the best configuration at trial 56 increased accuracy to 95.83%. Figure 38 presents a slice plot summarising the explored ranges and categorical options, including the pooling function, the normalisation method, and the selected MMD kernel. The trial delivering the highest objective value is explicitly identified, and its corresponding hyperparameter settings are highlighted within a red rectangle.




Figure 38. Hyperparameter slice plot across the optimisation trials

The main observations in Figure 38 show that introducing a normalisation step, indicated by the ‘use norm’ option on the horizontal axis, reduced GAT-CAMDA performance relative to using the raw signals during signal processing. Moreover, smaller learning rates were generally linked to better results. Mean pooling was also not well suited to GAT-CAMDA, as the best accuracy achieved under this setting was approximately 75%. The target test performance of the best-performing framework variant is reported in Figure 39-a and Figure 39-b. Figure 39-b summarises the associated classification metrics, namely precision, recall, and F1-score.




Figure 39. Target test results obtained with the best-performing GAT-CAMDA configuration, presenting (a) the confusion matrix and (b) the classification metrics

The reported metrics show that GAT-CAMDA separates the 12 structural health scenarios effectively, delivering high precision and recall for most classes and an overall accuracy of 95.83%. Classes with a precision of 1.0 indicate that the model assigns those damage conditions without generating false positives, supporting its sensitivity to subtle changes in wave propagation behaviour. In addition, the uniformly strong recall values across classes suggest that missed detections are rare, which is critical in SHM where undetected damage can threaten safety and structural integrity. Although class C10 exhibits a modest reduction in precision, its recall remains high, implying that the model occasionally assigns other samples to C10, but rarely overlooks true C10 cases.

To examine how GAT-CAMDA affects the distribution of learned representations, t-SNE was performed on the extracted features before and after the hybrid DA stage. The corresponding embeddings are presented in Figure 40-a and Figure 40-b, where each point denotes a single observation linked to a particular temperature and damage condition.




Figure 40. Two-dimensional t-SNE embeddings of the extracted feature space, showing (a) the representation before DA and (b) the representation after DA within GAT-CAMDA

Comparing the t-SNE visualisations in Figure 40 before and after the DA stage shows that GAT-CAMDA produces a more coherent and better-aligned feature space. After adaptation, samples sharing the same temperature and damage condition group

into tighter clusters, suggesting that the framework mitigates the distribution shifts induced by temperature changes. Consequently, points from the same class concentrate into denser regions with clearer boundaries between classes, indicating that class-discriminative structure is retained while previously overlapping representations are reduced.

Overall, these results suggest that GAT-CAMDA can represent subtle signal differences and distinguish healthy from damaged states with high reliability, which is essential for robust and transferable SHM. By reducing temperature as a confounding factor, the framework increases inter-class separation while improving within-class consistency, supporting its suitability for scalable SHM deployments.

5.4.2.3 Comparative Study

The proposed GAT-CAMDA framework was benchmarked against a set of widely used feature-based DA techniques to assess its performance under semi-supervised adaptation conditions. All approaches were evaluated using identical training and test splits, with 70% of the pooled target-domain data consistently assigned to the feature alignment stage, whilst the remaining 30% was reserved for testing. This allocation reflects a realistic scenario where limited labelled target data are available for alignment, and unlabelled target samples must be classified in deployment. For each baseline, hyperparameters were adjusted via trial-and-error to obtain the best-performing setting, ensuring fair comparison within the constraints of their original formulations.

Overall, 10 DA strategies were tested on the augmented CONCEPT dataset: Transfer Component Analysis (TCA), Feature Selection with MMD (fMMD), CORAL, Deep CORAL, SrcOnly Prediction (PRED), Subspace Alignment (SA), Domain-Adversarial Training of Neural Networks (DANN), Dynamic Unsupervised Adaptation (DUA), and Balanced Adaptation Regularisation based Transfer Learning (BAR). These methods represent a spectrum of adaptation paradigms, including subspace projection methods (TCA, SA), statistical moment matching approaches (CORAL, Deep CORAL, fMMD), adversarial alignment techniques (DANN, DUA), and regularisation-based transfer strategies (BAR). Within this set, 2 methods (PRED and BAR) are supervised and require target labels during alignment, whilst the

remaining 8 operate under fully unsupervised domain adaptation (UDA) conditions, making the comparison relevant to varying degrees of supervision availability.

Optuna-driven optimisation was not applied to the baseline methods to remain consistent with their originally reported formulations and to avoid introducing optimisation-specific effects that are not part of their intended design. This decision ensures that performance differences reflect inherent algorithmic capabilities rather than disparities in hyperparameter tuning effort. Similarly, GAT-based feature extraction was not introduced into these baselines, since doing so would fundamentally alter their architectures and compromise comparability. Instead, all baselines operate on raw time-domain signals or conventional feature representations as originally intended, whereas GAT-CAMDA leverages graph-based spatial feature learning as an integral component of its design. Table 10 lists the classification metrics for all baselines alongside the results obtained with GAT-CAMDA, enabling direct comparison of accuracy, F1-score, and computational efficiency across methods.

Table 10. Performance comparison between GAT-CAMDA and alternative feature-based approaches

| Model | Accuracy (%) | Precision | F1-Score |
|-------------------------------------|---------------------|------------------|-----------------|
| TCA (Pan et al., 2011) | 52 | 0.63 | 0.48 |
| fMMD (Uguroglu and Carbonell, 2011) | 58 | 0.67 | 0.54 |
| CORAL (Sun et al., 2016) | 59 | 0.66 | 0.56 |
| Deep CORAL (Sun and Saenko, 2016) | 63 | 0.63 | 0.62 |
| PRED (III, 2009) | 79 | 0.80 | 0.79 |
| SA (Fernando et al., 2013) | 79 | 0.81 | 0.79 |
| DUA (Mirza et al., 2022) | 82 | 0.81 | 0.79 |
| DANN (Ganin et al., 2016) | 82 | 0.83 | 0.82 |
| BAR (Hu et al., 2022) | 91.02 | 0.92 | 0.92 |
| GAT-CAMDA | 95.83 | 0.96 | 0.96 |

Table 10 shows that, even with the added difficulty introduced by combining disjoint target domains, GAT-CAMDA delivered markedly stronger performance than almost all comparator methods. The framework reached 95.83% accuracy, with both precision and F1-score equal to 0.96. This represents a clear gain over BAR, which was the strongest baseline, achieving 91.02% accuracy and precision and F1-score values of 0.92, yet still falling short of GAT-CAMDA.

Further class-level evidence is provided by the confusion matrices in Figure 41-a through Figure 41-d, which report the target test predictions for Deep CORAL, PRED, SA, and DANN, respectively.




Figure 41. Damage detection performance obtained with (a) Deep CORAL, (b) PRED, (c) SA, and (d) DANN

Figure 41 illustrates that GAT-CAMDA demonstrates superior precision in classifying the intact plate condition (C0) as well as specific damage states such as C5 and C11, in comparison to other methods including SA and DANN, which frequently misclassify closely related classes, such as C2 and C3. Nonetheless, it is important to note that class C6 posed a significant classification challenge for all evaluated approaches. Despite this, GAT-CAMDA addressed the challenge more effectively, further supporting its robustness in handling complex structural health scenarios.

5.4.2.4 Hyperparameter Importance

To clarify how each of the 13 optimised hyperparameters contributes to GAT-CAMDA performance, Optuna was used to estimate their relative importances, which are reported in Figure 42. This visual summary indicates which settings most strongly affect the outcome, allowing future refinement efforts to focus on the most influential parameters to improve both predictive accuracy and computational efficiency.




Figure 42. Estimated hyperparameter importance for damage detection performance

Figure 42 indicates that the learning rate dominates the optimisation landscape, with an importance of 0.59, which points to its central role in shaping model performance. The weighting applied to the MMD loss and the number of attention heads in the GAT are the next most consequential factors, with importance values of 0.10 and 0.08, respectively, highlighting their strong contribution to the framework’s effectiveness. While the hidden-layer dimensionality has a smaller importance score of 0.04, it still plays a meaningful part in refining the model’s behaviour for damage detection. Overall, the results suggest that optimisation should concentrate first on these high-impact settings to improve both accuracy and efficiency in GAT-CAMDA.

It should be noted that the hyperparameter importance values in Figure 42 represent global, aggregated effects estimated by Optuna’s TPE over all 100 optimisation trials, rather than the configuration-specific fluctuations in performance shown in Figure 38. Consequently, parameters that exert a stable influence across many trials, such as the learning rate, receive higher importance scores, even when some categorical choices, such as enabling normalisation or selecting particular MMD

kernels, produce pronounced but localised changes in accuracy. This explains why the learning rate ranks highest in the overall importance ordering, while other factors, including the MMD kernel, pooling strategy, and normalisation, show meaningful yet less consistent contributions. Accordingly, the 2 figures should be interpreted together: Figure 38 captures trial-level sensitivity, whereas Figure 42 summarises influence across the full optimisation space.

5.4.2.5 Sensor Importance

Within GAT-CAMDA, the sensor-specific attention scores were normalised to quantify each sensor’s relative contribution to damage detection. Under the framework’s optimal configuration, these scores were computed from the degree to which each sensor supported the identification of structural anomalies. The resulting values were scaled to the 0 to 1 interval, providing a uniform and readily interpretable measure of sensor importance. Figure 43 displays the normalised scores using a gradient bar chart, which visualises the comparative influence of each sensor on overall detection performance.




Figure 43. Attention-derived sensor contribution estimates for the best-performing framework configuration

As noted earlier, the best-performing setting was obtained when the signals were processed without normalisation. Under these conditions, the measurements from PZT2 exhibit an amplitude that is approximately 8 times larger than those recorded by the other 2 receivers, PZT3 and PZT4. The response from PZT4 is also marginally

greater than that of PZT3. These amplitude differences are reflected in the estimated sensor contributions, where importance decreases in the order PZT2, PZT4, and PZT3.

The same optimal model was then assessed again with all hyperparameters held fixed and with signal normalisation still disabled, and the resulting sensor contribution scores are reported in Figure 44.




Figure 44. Attention-based sensor contribution estimates obtained using unnormalised signals

The comparatively high contribution attributed to PZT3 in the proposed SHM framework, when unnormalised signals were used throughout signal processing and model training, is consistent with earlier observations reported for this dataset (Ferreira et al., 2024). In particular, PZT3 was found to exhibit a wider confidence interval than PZT2, which indicates greater variability in its recordings. This variability is largely driven by temperature changes that modify the composite’s mechanical properties. As temperature rises, the carbon epoxy laminate becomes less stiff, with the shear modulus G_{12} decreasing by 7%, whereas the elastic modulus E_1 shows a smaller reduction of 1.3%. This contrast is especially relevant because PZT3 is positioned at 45 degrees relative to the excitation source, which increases its sensitivity to shear-modulus changes due to its orientation with respect to the dominant wave propagation direction.

The larger spread observed in the PZT3 measurements can be linked to its placement and to the viscoelastic behaviour of the epoxy resin, which differs from that

of the carbon fibres. Consequently, although PZT3 exhibits a lower signal-to-noise ratio, it can be more responsive to subtle perturbations in wave propagation that may indicate damage. Figure 44 reflects this interaction between sensor location, temperature-dependent material behaviour, and guided-wave response in composite structures.

5.4.3 SPADA

To assess the effectiveness of SPADA in detecting damage in mechanical systems under EOVs, the two case studies WTB-VibClimate and CONCEPT were analysed separately. This section presents results from both intermediate and final stages of the evaluation. It includes damage detection outcomes without applying the DA section of SPADA, as well as the results obtained using the complete SPADA framework. Furthermore, ablation and comparative studies were conducted, and the corresponding findings for both case studies are reported. Finally, the internal activities of SPADA were visualised to demonstrate how the framework adjusts the source and target domains to mitigate EOVs effects in the damage detection; the computational efficiency of the pipeline was also discussed to evaluate its capability in the real-time deployment.

5.4.3.1 WTB-VibClimate

To evaluate the performance of SPADA on damage detection for small-scale WTBs using vibration signals, the augmented WTB-VibClimate dataset was employed following the windowing procedure described previously. To prevent data leakage from the test set during the DA stage, the dataset was partitioned as follows. From the windowed observations, 10 samples (derived from 2 original observations) were randomly selected and reserved exclusively for final testing. The remaining 15 samples (windowed from 3 original observations) were allocated for training and validation, with 11 samples assigned to the DA phase and 4 samples designated for the unsupervised validation phase. This stratified allocation ensured class balance across all subsets whilst maintaining complete separation between adaptation and testing data. The following sections present the results and discussions for each implementation stage.

5.4.3.1.1 Feature Extraction

In SPADA, a high-level WTS was utilised to extract discriminative features from the raw vibration signals. Each channel's signal was first standardised to remove mean offsets and normalise the variance, thereby ensuring consistency across the dataset.

The scattering transform was implemented using the Morlet wavelet as the mother wavelet. This choice is particularly suitable for vibration analysis, as the Morlet wavelet combines good frequency localisation with adequate time resolution, allowing both transient and oscillatory behaviours to be effectively captured (Łuczak, 2023). The transform was configured with a maximum scale corresponding to 2^{10} , ensuring that long-range temporal dependencies were represented, and with two wavelets per octave, which provided a balanced frequency resolution across the spectrum. Furthermore, coefficients were retained up to the second scattering order, thereby capturing not only primary spectral content but also nonlinear interactions between frequency bands.

The resulting scattering coefficients underwent logarithmic compression to reduce the impact of large fluctuations and enhance stability. Subsequently, temporal averaging was applied to produce fixed-length descriptors that are invariant to small shifts in time while remaining sensitive to subtle differences in oscillatory structure. Finally, the features from the 3 sensor channels were concatenated to form a unified representation. Through this procedure, the WTS yielded compact, translation-invariant, and noise-robust features. These descriptors preserved essential information about fine-scale transients as well as broader structural variations in the signals, thereby providing a reliable input space for the subsequent DA and classification modules in the SPADA framework. However, it must be acknowledged that the selection of two wavelets per octave and second-order scattering represents a design compromise between feature expressiveness and computational efficiency, potentially overlooking higher-order modulations that might characterise more subtle fault signatures.

Moreover, the temporal averaging window, whilst conferring shift invariance, inherently discards fine temporal dynamics that could prove diagnostically significant in certain fault evolution patterns. The logarithmic compression parameter, though empirically justified, introduces a degree of arbitrariness that warrants sensitivity analysis to ensure robustness across diverse operating conditions. Consequently,

whilst the chosen configuration provides a theoretically sound and practically viable feature extraction framework, alternative parameterisations merit systematic exploration to ascertain whether performance gains can be realised for specific fault types or environmental conditions.

To evaluate the efficacy of the implemented feature extraction stage in the context of WTB-VibClimate, the scattering transform was applied to the 3-channel vibration signals collected under a representative temperature condition, specifically Wp20, which reflects a typical operational scenario. Scattering features were extracted separately from each channel, then normalised, averaged over time, and concatenated to form comprehensive feature vectors. These features were ranked based on their mutual information with the class labels, and the top two features were visualised using a scatter plot; Figure 45 illustrates the resulting plot.




Figure 45. Scatter plot of the top-2 features ranked by mutual information, based on scattering transforms employed on Wp20; samples are colour-coded by class

Figure 45 shows the feature extraction method effectively separates observations across classes, though notable overlap persists between crack-related classes (labels 5 and 6) and classes 9 and 10. These overlaps likely arise from subtle differences in crack characteristics, such as length and severity, which are harder to distinguish in low-dimensional projections but potentially separable in the full 64-dimensional

feature space where higher-order feature interactions may provide additional discriminative information.

To examine EOV impact on data distribution and assess WTS's capture of damage-related and domain-invariant features, scatter analysis was performed on 3 scenarios under different temperatures (Wn15, Wp20 and Wp40), selected to span a wide thermal range and represent below-ambient, near-ambient, and elevated temperature conditions. The two most discriminative features, identified by mutual information ranking, were visualised in Figure 46-a to maximise class separability in the two-dimensional projection.

For comparison, WTS was substituted with a CNN-based framework comprising 4 one-dimensional convolutional layers with batch normalisation and ReLU activation, concluded by global average pooling and linear projection to a 64-dimensional embedding space. This architecture was chosen to provide comparable representational capacity to WTS whilst relying on learned rather than hand-crafted features. The CNN was trained using source domain data, and the same 3 datasets were processed with corresponding scatter plots in Figure 46-b. Features in both figures were ranked based on Wp20 for consistency with Figure 45, with 5 observations per domain plotted per class for visibility.

In Figure 46-a (WTS), the class clusters exhibit a clear domain-wise ordering when features are ranked on Wp20, with samples from Wn15 consistently located to the left, Wp20 centred, and Wp40 to the right along the first MI feature. This left–centre–right pattern repeats across classes whilst preserving intra-class compactness (as identified with the red-dashed rectangular), indicating that WTS encodes temperature-driven shifts as an approximately monotonic displacement in feature space whilst maintaining class structure, a property that is conducive to alignment across domains and facilitates subsequent DA. In contrast, Figure 46-b (CNNs) shows weaker domain regularity and greater intermingling of Wn15, Wp20, and Wp40 samples within many class groups, despite comparable overall separability. This suggests that the CNN trained on Wp20 captures class-discriminative cues but with reduced domain awareness and poorer alignment under EOV, making it more susceptible to distribution shifts when deployed on target domains with different thermal characteristics.




Figure 46. Scatter plot of the top-2 features ranked using Wp20 data, visualised across 3 temperature conditions (Wn15, Wp20, Wp40); for (a) WTS-based and (b) CNNs-based feature extraction

5.4.3.1.2 Damage Detection without DA

In this section, the damage detection task was conducted under the assumption that no DA stage was implemented. For this purpose, the datasets recorded at 3 different temperatures (Wn15, Wp20 and Wp40) were each independently designated

as the source domain; the remaining datasets were used as the target domains. The loss weights associated with the DA stage (that is, λ_t for the target self-training objective and λ_d for the domain-adversarial objective) were set to zero. A grid search across 200 feature spaces was performed to identify the optimal classification accuracy. The graphs in Figure 47-a, 47-b and 47-c show the damage detection results obtained without applying DA on the assumption Wn15, Wp20, and Wp40 were considered as the source domains, respectively. These results were achieved after feature extraction using WTS, training the classifier on the labelled source data, and evaluating it on the test set of the target domains. Furthermore, to demonstrate the effectiveness of the proposed WTS-based feature extraction step, the same CNNs framework described in the previous section was applied as an alternative to WTS under the 3 source domain assumptions. The corresponding results obtained when tested on the various target domains are also presented in the same figures for comparison.




Figure 47. Damage detection utilising WTS and CNNs without DA when (a) Wn15, (b) Wp20, and (c) Wp40 was assigned as the source domain

Figure 47 shows a clear temperature proximity effect for both approaches. Using WTS, average accuracies across targets were 76.57% with Wn15 as source, 81.54% with Wp20, and 60.0% with Wp40. The corresponding CNN averages were markedly lower at 35.87%, 46.85%, and 25.24%. Wp20 consistently provided the strongest generalisation for both methods.

WTS degraded more gracefully as the source–target temperature gap widened, whereas CNNs were highly sensitive to mismatch. For example, with Wn15 as source, WTS reached 98.46% on nearby Wn5 but fell to 49.23% on Wp35, whilst CNNs dropped from 68.46% to 13.08%. With Wp40 as source, WTS reached 95.38% on Wp35 but only 36.15% on Wn15, whilst CNNs moved from 85.38% to 15.38%. Overall, WTS outperformed CNNs by margins of approximately 41, 35, and 35 percentage points for the 3 source settings, indicating that WTS features retain more damage-relevant information under temperature variation. The persistently lower accuracies at larger temperature differences indicate a systematic generalisation gap, providing strong justification for incorporating a DA stage to mitigate temperature-induced covariate shift and stabilise performance across dissimilar operating conditions.

5.4.3.1.3 Damage Detection with DA

To address the performance gap under large temperature differences, the full SPADA pipeline with DA was deployed. Hyperparameters were tuned via random search over 200 configurations, each run with 5 independent seeds. Epochs were selected by the unsupervised two-part hold-out score on a validation split, and the optimal configuration was chosen by highest mean validation accuracy across seeds (ties broken by lower s.d., then earlier index). Test performance was reported as mean \pm s.d. Reproducibility was ensured through consistent seed control across Python, NumPy, PyTorch, and scikit-learn. The class-balanced target-domain data allocation was 60% for DA and 40% (10 observations per class) for testing.

In WTS, the maximum scale was constrained to 8 per the scattering condition $2^8 \leq T$ where $T = 39200$ (augmented signal length). Remaining WTS hyperparameters were specified as small search grids. Table 11 summarises all hyperparameters used in the random search.

Table 11. Hyperparameter search space defining the randomised search candidates and values in SPADA

| Component | Value | Count | Note |
|------------------------------------|--|-------|--|
| Epochs per trial | 300, 400 | 2 | Training epochs for each grid trial |
| Feature width | 64, 128 | 2 | Hidden dimension of the feature extractor |
| Learning rate | $1.5 * 10^{-5}$, $2 * 10^{-5}$, $5 * 10^{-5}$ | 3 | RMSprop optimiser |
| Domain loss weight | 0.5, 1 | 2 | Weight for adversarial domain loss |
| Prototype weighting strength | 0.05, 0.1, 0.2 | 3 | Strength of prototype-based sample weighting |
| Entropy gate | 0.5, 0.6, 0.7 | 3 | Confidence gate used during training |
| RMSprop momentum | 0.5, 0.9 | 2 | Optimiser momentum |
| RMSprop weight decay | $1 * 10^{-6}$, $1 * 10^{-5}$ | 2 | Optimiser weight decay |
| Label smoothing | 0.1, 0.2 | 2 | Applied in the source cross-entropy loss |
| Prototype-alignment loss weight | 0.15, 0.2 | 2 | Coefficient on the prototype-alignment term |
| Confidence threshold | 0.7, 0.8 | 2 | Used in the unsupervised two-part hold-out scoring |
| Prototype blend | 0.8, 0.95 | 2 | Used in the unsupervised two-part hold-out scoring |
| Weight on target pseudo-label loss | 0.02, 0.03, 0.05, 0.07, 0.1, 0.12 | 6 | Coefficient on target pseudo-label loss |
| Softmax temperature | 0.9, 1.0, 1.2, 1.5, 2, 2.5, 3, 3.5 | 8 | Temperature for softmax calibration |
| Schedule steepness | 5, 10, 12, 15, 20 | 5 | Slope parameter of the training schedule |
| Grid batch size | 4, 8, 16, 24, 32 | 5 | Mini-batch size for model training |
| Global seed (per-trial) | 1, 2, 3, 4, 5 | 5 | Random seed for initiation and splits |
| Maximum scale (in WTS) | 8, 9, 10, 15 | 4 | Log2 of the largest averaging scale |
| Quality factor (in WTS) | 1, 2 | 2 | Filters per octave in the wavelet bank |
| Maximum scattering order (in WTS) | 1, 2 | 2 | Highest scattering order included |
| Log epsilon (in WTS) | $1 * 10^{-6}$, $1 * 10^{-5}$ | 2 | Stabiliser inside logarithmic compression |
| Batch size (in WTS) | 24, 64 | 2 | Mini-batch size for feature extraction |

In accordance with the without DA scenario, 3 source-to-target settings were considered, with Wn15, Wp20, and Wp40 each designated as the source domain in turn, whilst the remaining datasets were assigned to target domains using the same partitions as in the without DA condition. The highest target accuracies obtained are presented in Figures 48-a to 48-c, with CNN+DA results also plotted for consistent comparison.




Figure 48. Damage detection utilising WTS and DA when (a) Wn15, (b) Wp20, and (c) Wp40 was assigned as the source domain

When Wp20 was used as the source (Figure 48-b), accuracies were uniformly high across all targets, giving the best average of 97.58%. When Wn15 was the source (Figure 48-a), perfect results were obtained for nearby targets, but performance decreased for the hottest targets (Wp30 66.92%, Wp35 61.54%), yielding 86.29% on average. With Wp40 as the source (Figure 48-c), the coldest targets were most challenging (Wn10 56.92%, Wn15 59.23%), while warm targets remained strong, for an average of 79.65%. A brief look back at Figure 47 (without DA) showed much

lower averages, i.e., 35.87% for Wn15 as source, 46.85% for Wp20, and 25.24% for Wp40. Under the present domain-adapted setting, these were raised to 86.29%, 97.62%, and 79.65%, which indicates that the earlier low-accuracy gaps were largely closed. For instance, the case from Wn15 to Wp35 increased from about 13.08% in Figure 47 to 61.54% here.

The CNN with DA resulted in lower accuracies overall, and in multiple target domains performance fell below the setting without DA. This pattern is consistent with negative transfer in UDA, where the alignment objective can erode class margins when source class-wise separability is high while the source to target discrepancy is large. The DA stage likely enforced global alignment without sufficient class awareness, producing misaligned features and weaker generalisation. In this context, the WTS pipeline proved more resilient, indicating that class-aware or temperature-aware adaptation, stronger regularisation, or instance weighting may be required to mitigate negative transfer in the CNN-based approach.

These results emphasise the importance of choosing a suitable source domain; a mid-range source such as Wp20 kept temperature differences moderate to both cooler and warmer targets, which stabilised alignment and aided prototype formation. In practice, anchoring adaptation on a central operating condition, or creating a balanced source mix around it, is likely to yield the most reliable transfer.

To gain a clearer understanding of which classes the SPADA framework (with DA) struggles to classify accurately, confusion matrices are presented for 4 scenarios considered as the target domain, i.e., Wn10, Wp0, W30 and Wp35. These are illustrated in Figures 49-a to 49-d, respectively, under the assumption that Wp20 serves as the source domain. These target domains were selected because SPADA did not achieve full classification accuracy on them, allowing for a more detailed examination of its limitations.

Considering the plotted confusion matrices in Figure 49, it can be understood that, across the selected target domains, SPADA's residual errors are concentrated almost entirely in the pairwise confusions between labels 5 and 6; the misclassifications are symmetric and limited in number, indicating a tight decision boundary rather than widespread class drift. These two labels correspond to closely related crack configurations (two 5 cm cracks vs. three 5 cm cracks), which are

expected to induce similar stiffness reductions and mode-shape perturbations. Under temperature shift, the induced spectral signatures become even more alike due to thermal softening and modest peak broadening, narrowing the margin between their class prototypes. The effect is amplified by (i) sensor-limited operation (3 channels), which reduces spatial sensitivity to crack multiplicity; (ii) windowed segmentation, which preserves local transients but weakens global geometry cues; and (iii) conservative pseudo-labelling during adaptation, which can slightly relax class margins for near-neighbour classes.




Figure 49. Confusion matrices of SPADA with Wp20 as the source and (a) Wn10, (b) Wp0, (c) W30, and (d) Wp35 as target domains

5.4.3.2 CONCEPT

SPADA underwent dual validation by assessing its effectiveness in detecting damage in composite plates, employing the augmented dataset generated for

CONCEPT through Conv-CVAE. To avoid potential bias during the DA process, the dataset corresponding to the baseline temperature used for completing the others, i.e., Cp30, was excluded and only used to show feature distribution in the scatter plots. Consequently, only Cp0, Cp10, Cp20, Cp40, Cp50, and Cp60 were considered. Two distinct scenarios were examined. In the first, the dataset corresponding to the lowest temperature, Cp0, was designated as the source domain, while each of the remaining temperatures served individually as target domains. In the second scenario, Cp60, representing the highest temperature, was used as the source domain, with the others again serving as separate targets.

During the DA stage, 50% of the relevant target domain was allocated for training, while the remaining 50% was retained for testing, with no prior exposure to the model. Both subsets were randomly selected and class-balanced. Accordingly, in the testing phase, 15 observations per class were classified. The following sections present a detailed analysis of the performance of SPADA when evaluated on the CONCEPT dataset.

5.4.3.2.1 Feature Extraction

To implement WTS on the CONCEPT case study, the same feature extraction block used in WTB-VibClimate was employed, with the maximum scale limited to or below due to the shorter signal length in the CONCEPT data, where each channel comprises 1000 data points. This constraint ensures that the scattering coefficients remain well-defined across all decomposition levels, as wavelet scales exceeding the signal duration would produce unreliable or undefined transform outputs.

To assess the efficacy of the feature extraction stage, the scattering transform was applied to 3-channel guided wave signals recorded under the baseline temperature condition (Cp30) assigning the maximum scale to 2^7 . Similar to the previous case, scattering features from each channel were normalised, averaged over time, and concatenated into feature vectors. The features were then ranked by mutual information with the class labels, and the top two were visualised in a scatter plot (Figure 50). The selection of Cp30 as the baseline condition provides a reference distribution against which temperature-induced shifts can subsequently be characterised, whilst the mutual information ranking prioritises features that maximise class discriminability under nominal conditions. This visualisation serves to verify that

the scattering transform produces a feature space wherein damage states exhibit measurable separation prior to the introduction of cross-temperature domain shifts.




Figure 50. Scatter plot of the top-2 features ranked by mutual information, based on scattering transforms employed on Cp30; samples are colour-coded by class

Figure 50 demonstrates that the implemented WTS successfully separates the clusters corresponding to different classes. However, as previously discussed, this is not the only criterion for selecting a feature extraction method within the SPADA framework, as a meaningful hierarchy between class clusters should also be maintained when temperature varies. To illustrate this using the CONCEPT dataset, Figure 51-a and 51-b present the two-dimensional scatter plots of the extracted features ranked by mutual information for 3 domains, Cp0, Cp30, and Cp60, obtained using WTS and CNNs, respectively. For CNNs, the same supervised framework applied in WTB-VibClimate was employed and only 5 observations' scatter were plotted for each health scenario and each domain.

Analysing Figure 51, it can be observed that CNNs achieved reasonable distinguishability in separating domain clusters, with visible class-wise separations (the apparent overlaps result from the figure's density). Similarly, WTS demonstrated strong capability in both domain and class-level separability. The domain separation pattern observed in the previous case study, WTB-VibClimate (highlighted with a red box in Figure 46-a), is again evident here but in a different positioning. As shown in

Figure 51-a, the Cp30 clusters (baseline temperature) are positioned in the lower left region rather than at the centre of their class areas. Moreover, Figure 51-b indicates that the CNN-extracted features from different classes exhibit a large domain gap, which poses challenges for the subsequent UDA adjustment stage.




Figure 51. Scatter plot of the top-2 features ranked using Cp30 data, visualised across 3 temperature conditions (Cp0, Cp30, Cp60); for (a) WTS-based and (b) CNNs-based feature extraction

5.4.3.2.2 Damage Detection without DA

In the CONCEPT case study, the DA stage in SPADA was initially omitted to examine the extent to which WTS could compensate for the domain gap caused by variations in background temperature. As noted earlier, Cp0 and Cp60 were designated as independent source domains, while the remaining datasets were used as separate target domains in each scenario, with λ_t and λ_d set to zero. A grid search across 200 feature spaces was conducted to determine the optimal classification accuracy. Figure 52-a and 52-b present the damage detection results obtained without DA, assuming Cp0 and Cp60 as the source domains, respectively. These results were achieved following feature extraction using WTS, training the classifier on labelled source data, and evaluating it on the test sets of the target domains. To further assess the effectiveness of the proposed WTS-based feature extraction stage, the same CNNs framework described previously was employed as an alternative under the 3 source domain assumptions. The corresponding results for the various target domains are also shown in the same figures for comparison.




Figure 52. Damage detection utilising WTS and CNNs without DA when (a) Cp0 and (b) Cp60 was assigned as the source domain

Figure 52 shows that WTS significantly outperformed CNNs in extracting domain invariant features. In both cases that used Cp0 and Cp60 as the source domains, target observations with temperatures closer to the chosen source were classified with high accuracy. For example, when the classifier was trained with Cp0 as the source domain, samples from Cp10 were classified with 100% accuracy, whereas under the same source and target selection the classifier trained on features extracted with CNNs

achieved about 52%. Performance deteriorated further for the Cp0 to Cp40 pairing, where accuracy fell to 0%.

The same pattern is visible in Figure 52-b. With Cp60 as the source domain, observations from Cp50 were classified with 100% accuracy using features from WTS, while CNNs achieved approximately 78%. The lower accuracy that appears when the temperature gap is wider, for example Cp60 as the target with Cp0 as the source, or Cp0 as the target with Cp60 as the source, indicates the need for an additional stage, i.e., DA. The systematic degradation with increasing temperature separation suggests that both feature extraction methods retain residual temperature-dependent characteristics, though WTS exhibits substantially greater robustness within narrow thermal ranges. Critically, the complete classification failure of CNN-based features at moderate temperature shifts (0% accuracy for Cp0-to-Cp40) raises questions regarding the generalisability of end-to-end learned representations when faced with distribution shifts that fall outside the training envelope, whereas WTS maintains partial discrimination capability even under challenging conditions.

5.4.3.2.3 Damage Detection with DA

In this section the whole pipeline of SPADA (including DA stage) was employed for the damage detection of CONCEPT case study and for the augmented dataset utilising CONV-CVAE framework. The same code applied for WTB-VibClimate and with the listed hyperparameters in Table 11 was developed in this case by searching 200 random scenarios, only by restricting the maximum scale was constrained to 2^9 ($J = 9$), since in CONCEPT dataset each channel's signal has a length of 1000, and the maximum scale should not exceed the signal length ($2^9 = 512$). As a result, 4 values for J were considered in the search space, i.e., 6, 7, 8, 9.

Another notable change concerns the target-domain sample allocation. Each target dataset was split into two per-class balanced subsets, with 15 samples per class in the DA phase and 15 per class in the test phase. Figure 53-a and 53-b represent the damage-detection results on the CONCEPT dataset with Cp0 and Cp60 used as the source domains.




Figure 53. Damage detection utilising WTS and DA when (a) Cp0 and (b) Cp60 was assigned as the source domain

Figure 53 shows that enabling the DA stage in SPADA increases damage detection accuracy across both source domain selections, with the largest gains observed when the temperature gap between source and target domains is wider. For example, with Cp0 as the source and Cp40 as the target, accuracy rose from 48.33% without DA to 94.44% with DA. Under a wider gap, Cp0 as the source and Cp60 as the target, accuracy increased from 45% to 93.33% after activating DA.

When Cp60 was used as the source domain, similar improvements were obtained. Accuracy increased from 41.67%, 51.67%, and 66.11% to 91.67%, 94.44%, and 100% when Cp0, Cp10, and Cp20 were the target domains, respectively. Cases with small source to target differences, such as Cp0 to Cp10, already achieved 100% without DA and remained at 100% with DA.

Figure 53-a and 53-b also indicate that there is still room to improve accuracy under the widest gaps. For instance, Cp60 as the source and Cp0 as the target yielded 91.67% with DA. The same pattern was observed when the full SPADA framework was applied to WTB-VibClimate.

To determine residual misclassification patterns of SPADA on CONCEPT where accuracy did not reach 100%, Figure 54-a to 54-d present the test phase confusion matrices for 4 source and target settings: Cp0 as the source with Cp60 as the target, Cp0 as the source with Cp40 as the target, Cp60 as the source with Cp0 as the target, and Cp60 as the source with Cp10 as the target.




Figure 54. Confusion matrices for SPADA on CONCEPT for (a) source Cp0 and target Cp60, (b) source Cp0 and target Cp40, (c) source Cp60 and target Cp0, and (d) source Cp60 and target Cp10

In Cp0 as source and Cp60 as target (Figure 54-a), the matrix is mostly diagonal, with the remaining errors concentrated between adjacent severity classes. The largest exchange occurs between C2 and C3, with smaller leakages from C4 into C3, from C1 into C2, and from C7 into C8. When the gap narrows to Cp0 as source and Cp40 as target (Figure 54-b), almost all classes are classified correctly; the main deviations are a split of C9 into C10 and a minor C3 into C4 spill. These patterns are consistent with the physics of the CONCEPT set up. Signals are Lamb waves, so neighbouring severities can result in highly similar dispersion and attenuation signatures along a fixed path, especially under larger temperature separation across 0 to 60 °C in 10 °C increments labelled Cp0 to Cp60.

With Cp60 as source and Cp0 as target (Figure 54-c), the residual errors are again local, concentrated within the C2, C3, and C4 block, indicating limited margins between adjacent prototypes under the largest shift. Bringing the target closer to Cp10 (Figure 54-d) tightens the diagonal further. The largest residuals occur for class C2, with 2 samples predicted as C1 and 3 as C3. Minor errors remain for class C3, with 2 samples predicted as C4 and 1 as C5, and for class C8, with 2 samples predicted as C9; all other classes are correctly classified. Interpreting classes as healthy C0 and increasing simulated delamination severities C1 to C11 created by incremental putty coverage, these local swaps are plausible because adjacent severities perturb the wavefield in similar ways between the actuator and receivers. They suggest value in class conditional alignment or explicit pairwise margin penalties for the hard pairs, rather than global alignment alone.

5.4.3.3 Internal Mechanism Visualisation and Interpretability Analysis

To enhance transparency and enable mechanistic understanding of the DA process, SPADA was instrumented to log internal states at every 10 training epochs during the DA stage. The logged data comprised prototype vectors, instance-level features, pseudo-labels, entropy values, and attention weights for both domains. Four complementary t-SNE visualisations were generated to reveal progressive domain alignment whilst preserving class separability: prototype trajectory evolution (Figure 55), instance-prototype cosine similarity (Figure 56), prototype attention weight dynamics (Figure 57), and decision boundaries in feature space (Figure 58). The analysis is demonstrated through two representative CONCEPT scenarios: C0-to-C40 transfer (94.44% accuracy) and C0-to-C50 transfer (100% accuracy). Epoch numbers on horizontal axes represent logged snapshots and must be multiplied by 10 to obtain actual training epochs. The entropy threshold for pseudo-label confidence gating was set to $\tau = 0.6$ throughout.

5.4.3.3.1 Prototype Trajectory Evolution

Prototype trajectories were visualised by projecting the logged high-dimensional prototype vectors into two-dimensional space using t-SNE. The t-SNE transformation was fitted on the concatenated set of all source and target prototypes across the logged epochs, ensuring consistent projection axes. Trajectories were sampled at every 10th logged epoch, with source prototypes marked by solid lines with circular markers and

target prototypes marked by dashed lines with square markers. Starting positions are distinguished by larger filled markers with black edges, whilst ending positions are indicated by star symbols with black edges. Figures 55-a and 55-b present these trajectories for C0-to-C40 and C0-to-C50, respectively.





Figure 55. Prototype trajectory evolution through t-SNE embedding for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer


In both scenarios, the trajectory visualisations reveal distinct alignment patterns that correlate with classification performance. For C0-to-C50 (Figure 55-b), which achieved 100% accuracy, source and target prototypes for each class converge to nearly coincident final positions, indicating complete domain alignment. This tight convergence is particularly evident for classes C6, C7, and C11, where prototypes originating from disparate initial positions successfully migrate toward shared representations. In contrast, C0-to-C40 (Figure 55-a), which achieved 94.44% accuracy, exhibits notably imperfect alignment for class C8, where source and target prototypes began at approximately overlapping positions but diverged substantially by the final epoch, ending at opposite regions of the embedding space. This divergence likely contributes to the residual classification errors observed for this class. Additionally, classes C3 and C4 in the C0-to-C40 scenario demonstrate partial convergence, with final positions closer than initial states but not fully coincident, whilst classes such as C6 and C7 achieve satisfactory alignment. The superior prototype coalescence in C0-to-C50 across all classes validates the framework's capacity to achieve complete domain invariance under this particular temperature shift, whereas the persistent C8 divergence in C0-to-C40 highlights the increased adaptation challenge posed by the larger thermal gradient.

5.4.3.3.2 Instance-Prototype Cosine Similarity

The similarity heatmap displays cosine similarities between individual instances and all prototypes at the final epoch selected by the unsupervised validation criterion. The heatmap is organised with rows representing instances (source instances in the upper block, target instances in the lower block) and columns representing prototypes (source prototypes on the left, target prototypes on the right), with each prototype corresponding to one class. White lines delineate the domain boundaries. This visualisation enables direct assessment of whether instances align more strongly with prototypes of their correct class and appropriate domain, thereby revealing the quality of learned representations and the effectiveness of cross-domain alignment. High-intensity values along the diagonal blocks indicate successful prototype-based classification, whilst off-diagonal activations suggest residual misalignment or confusion between classes. Figure 56-a and 56-b present the similarity matrices for C0-to-C40 and C0-to-C50, respectively.



(a)



(b)

Figure 56. Instance-prototype cosine similarity heatmap for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer

Both heatmaps exhibit pronounced diagonal structures, with elevated similarity values (bright yellow) concentrated along the main diagonals of the 4 quadrants,


confirming that instances align most strongly with prototypes of their corresponding class and domain. The C0-to-C50 case (Figure 56-b) demonstrates sharper and more continuous diagonal patterns with minimal off-diagonal activations, consistent with its 100% classification accuracy. In contrast, C0-to-C40 (Figure 56-a) shows comparatively weaker diagonal intensity in certain regions and moderate off-diagonal similarity values, particularly in the source-to-target and target-to-source quadrants, reflecting the residual 5.56% error rate. The superior block structure in C0-to-C50 indicates that both source and target prototypes serve as highly reliable class representatives under this temperature shift, whereas the diffuse patterns in C0-to-C40 suggest incomplete prototype convergence for specific classes, corroborating the trajectory analysis findings. Notably, the off-diagonal activations in C0-to-C40 are not uniformly distributed but exhibit localised clusters, suggesting that specific class pairs are more susceptible to cross-domain confusion under moderate temperature shifts. This structured degradation pattern implies that the 10°C difference between conditions may exceed critical thresholds for certain classes whilst remaining well within tolerance for others.

5.4.3.3.3 Prototype Attention Weight Dynamics

Three synchronised panels track attention dynamics across training epochs. The left panel displays median source instance attention weights with interquartile range (IQR) shading, the middle panel shows median target instance attention weights with IQR shading, and the right panel overlays pseudo-label coverage (the proportion of target samples with entropy below threshold $\tau = 0.6$) as a solid line against median target weight as a dashed line. A vertical dashed line marks the epoch selected by unsupervised validation. Perfect domain alignment is characterised by stable high source weights, progressively increasing target weights converging toward unity, and coverage approaching 1.0, indicating confident pseudo-label assignment across the target domain. Conversely, stagnant or declining target weights alongside low coverage signal inadequate alignment, reflecting persistent uncertainty in target predictions and weak prototype correspondence. Figures 57-a and 57-b present the attention dynamics for C0-to-C40 and C0-to-C50, respectively. The C0-to-C50 condition exhibits near-monotonic growth in target weights accompanied by coverage exceeding 0.95 by the selected epoch, demonstrating successful knowledge transfer.


In contrast, C0-to-C40 shows premature plateauing of target weights and coverage stabilising below 0.85, suggesting that a subset of target instances resist confident pseudo-labelling throughout training. This divergence aligns with the heatmap observations, wherein incomplete prototype convergence in C0-to-C40 manifests as persistent entropy in the attention mechanism's target assignments.

Prototype Attention & Confidence Dynamics C0-C40 ($\tau = 0.6$; best@epoch 29)



(a)

Prototype Attention & Confidence Dynamics C0-C50 ($\tau = 0.6$; best@epoch 12)



(b)

Figure 57. Prototype attention and confidence for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer


Source attention weights remain at unity throughout training in both scenarios, confirming stable confidence in labelled source data. For C0-to-C50 (Figure 57-b), the framework achieves rapid convergence, with the best epoch identified at epoch 12 and target weights stabilising around 0.3 with moderate IQR throughout. Coverage rises swiftly to approximately 0.5 within the first 10 epochs and plateaus, whilst median target weight exhibits modest fluctuations. In contrast, C0-to-C40 (Figure 57-a) demonstrates delayed convergence, reaching its best epoch at epoch 29. Target attention weights display greater variability, as evidenced by wider IQR bands, and fluctuate between 0.2 and 0.5 across training. Coverage follows a more erratic trajectory, peaking near 0.6 around epoch 33 before declining slightly. The synchronous rise of coverage and target weights in both cases validates the entropy-

based gating mechanism; however, the earlier stabilisation and tighter variance in C0-to-C50 corroborate its superior 100% accuracy, whereas the prolonged adaptation period and higher variability in C0-to-C40 reflect the additional challenge posed by the larger temperature gradient.


5.4.3.3.4 Decision Boundaries in Feature Space

Learned representations were visualised by projecting both instances and prototypes into two-dimensional t-SNE space at the best epoch determined by unsupervised validation. Source instances are plotted as circles, target instances as squares, source prototypes as 'X' markers, and target prototypes as filled plus symbols, all colour-coded by class. Successful DA is characterised by instances clustering according to class rather than domain, with source and target instances intermixed within shared class-specific regions, and by source and target prototypes positioned in proximity at cluster centroids, confirming convergent class representations across domains. Figures 58-a and 58-b present the decision boundaries for C0-to-C40 and C0-to-C50, respectively.

In C0-to-C50 (Figure 58-b), instances form tight, well-separated clusters organised by class identity with minimal inter-class overlap. Source and target prototypes for each class are positioned either coincidentally or in immediate proximity, validating complete prototype convergence. The uniform intermixing of source and target instances within each cluster confirms successful domain-invariant feature learning across all classes. Conversely, C0-to-C40 (Figure 58-a) exhibits looser clustering and greater spatial dispersion. Whilst most classes demonstrate reasonable alignment, several exhibit noticeable gaps between source and target prototypes, indicating incomplete convergence. The wider spread of instances and prototype displacement in C0-to-C40 visually corroborates the 5.56% accuracy deficit relative to C0-to-C50, confirming that tighter prototype coalescence and more compact class-specific clustering directly correspond to superior classification performance.



(a)



(b)

Figure 58. Decision boundaries in t-SNE embedding space for (a) C0-to-C40 transfer and (b) C0-to-C50 transfer

5.4.3.4 Comparison Study

SPADA was evaluated against six benchmarks that employed WTS for feature extraction and incorporated different UDA modules, including Adversarial Domain Adaptation with Prototypes (ADA-Proto), Multi-class Adversarial Domain Adaptation with Prototype Attention (MADA-Proto), Generative Adversarial Network with Prototype Weighting (GAN-Proto), Central Moment Discrepancy (CMD), CORAL, and Neighbour Refinement Consistency with Virtual Adversarial Training (NRC-VAT). Experiments were conducted on the WTB-VibClimate and CONCEPT

datasets, which were augmented through windowing and CONV-CVAE, respectively. Challenging domain shifts were considered, where Wp20 was used as the source domain with Wn15 and Wp40 as target domains for WTB-VibClimate. In the CONCEPT case study, Cp0 was used as the source domain with Cp60 as the target, and Cp60 was used as the source domain with Cp10 as the target. For each of these codes, 200 grid search was performed to find the most optimum reported test classification. Figure 59-a to 59-d represent the classification output in terms of the accuracy for each of the aforementioned source-target combinations. The selected benchmarks span three distinct adaptation paradigms: adversarial alignment, statistical moment matching, and consistency regularisation. This provides a comprehensive assessment of SPADA's performance relative to established methodological families. The chosen domain shifts encompass both symmetric (Wp20-to-Wp40) and asymmetric (Cp60-to-Cp10) temperature transitions, enabling evaluation of bidirectional adaptation robustness under varying magnitudes of distribution mismatch.




Figure 59. Damage detection results of UDA benchmarks against SPADA for (a) source Wp20 and target Wn15, (b) source Wp20 and target Wp40, (c) source Cp0 and target Cp60, and (d) source Cp60 and target Cp10

Comparisons in Figure 59 reveal that SPADA delivers the highest accuracy across all source-to-target transfers, with the largest advantage observed on the more severe WTB-VibClimate shift from Wp20 to Wp40, where temperature-induced domain divergence is most pronounced. On WTB-VibClimate, SPADA attains 96.15% for Wp20 to Wn15, exceeding the strongest baseline, ADA-Proto, by 1.06 percentage points, and reaches 90% for Wp20 to Wp40, outperforming the next best method by 11.47 percentage points. This substantial margin on the challenging Wp20 to Wp40 transfer underscores SPADA's capacity to handle severe domain shifts where other methods exhibit significant degradation. Prototype-based adversarial baselines are generally stronger than CMD, CORAL and NRC-VAT on these transfers, confirming the importance of prototype-guided alignment for maintaining class structure, although they still trail SPADA, indicating that the combination of wavelet-based features with confidence-gated pseudo-labelling and attention-weighted prototype refinement provides more effective preservation of class structure under EOV. On CONCEPT, where augmentation used CONV-CVAE to generate synthetic damage scenarios, overall accuracies are high for all methods due to the controlled experimental conditions and high-quality sensor data, but SPADA retains a consistent lead across all transfers. It achieves 91.67% for Cp0 to Cp60, ahead of GAN-Proto by 1.67 percentage points, and 94.44% for Cp60 to Cp10, ahead of ADA-Proto by 1.11 percentage points, demonstrating robust performance across both positive and negative temperature shifts. Considered across all 4 shifts, SPADA also shows the smallest performance spread at 6.15 percentage points, compared with markedly larger variability for the strongest baselines, which evidence both improved peak performance and superior stability across distinct domain shifts. This consistency is particularly valuable for industrial deployment, where reliable performance across diverse operating conditions is essential for maintaining trust in automated damage detection systems.

5.4.3.5 Computation Efficiency

This study placed particular emphasis on the computational efficiency and practical implementation of the SPADA framework. All training, evaluation, and the 1000-combination grid search (200 * 5 for 5 seed scenarios) were conducted in Python using Jupyter Notebook on a standard workstation with CPU only. Despite the

extensive grid search, the process was completed within 198 minutes for CONCEPT and 269 minutes for WTB-VibClimate. The framework therefore demonstrates suitability for real-time industrial deployment. The system specifications are presented in Table 12.

Table 12. System hardware and software specifications for SPADA pipeline

| Component | Specification | Component | Specification |
|------------------|--|-----------------------|--|
| Operating System | Microsoft Windows 11 Pro for Workstation (Build 26200) | Main libraries | PyTorch 2.0.1, Scikit-learn, NumPy, Matplotlib |
| Processor | Intel(R) Xeon(R) Gold 6248R CPU @ 3.00GHz, 2993 MHz, 24 Core(s), 48 Logical Processor(s) | Execution environment | Jupyter Notebook (Python 3.10) |
| Physical Memory | 192 GB | Virtual memory | 250 GB |

Chapter 6: Conclusions and Future Work

Summary of Research and Key Findings

This research addressed the critical challenges of data scarcity, EOV, and cross-domain generalisation in SHM of composite materials. The overarching aim was to develop robust ML frameworks that maintain diagnostic accuracy under varying conditions whilst ensuring computational efficiency and model interpretability for practical industrial deployment. The principal novel contribution of this work lies in the development and validation of 3 complementary DA frameworks, each specifically tailored to different regimes of domain shift severity and supervision availability, unified by a consistent emphasis on interpretability and computational tractability. UCTRF introduces a novel combination of instance-level reweighting with manifold-based feature learning to address mild domain shifts with limited target supervision. GAT-CAMDA proposes a novel architecture that integrates graph attention networks with multi-modal adversarial alignment to tackle moderate shifts by exploiting multi-sensor spatial relationships. SPADA presents a novel prototype-guided adversarial learning strategy that handles severe shifts under fully unsupervised conditions through confidence-gated pseudo-labelling and class-wise alignment.

The research successfully demonstrated that appropriately designed ML frameworks can generalise across substantially different environmental and operational conditions without requiring extensive retraining or abundant target domain labels in SHM systems. Experimental validation on two heterogeneous case studies, i.e., wind turbine blade vibration monitoring under varying climatic conditions (WTB-VibClimate) and carbon-epoxy composite plate monitoring under different loading configurations (CONCEPT) confirmed the practical applicability and robustness of the proposed methodologies. Each framework demonstrated distinct strengths aligned with its design philosophy: UCTRF effectively leveraged minimal target supervision through transfer boosting, GAT-CAMDA exploited sensor topology through attention-based graph neural networks whilst aligning feature distributions, and SPADA achieved superior performance across the most challenging domain shifts through confidence-gated pseudo-labelling. Critically, all frameworks operated

efficiently on standard CPU-only workstations, demonstrating computational feasibility for real-time industrial deployment without specialised hardware requirements.

Multiple key findings emerged from this research. First, a novel data augmentation pipeline combining deterministic methods (windowing, spline interpolation) with probabilistic generative models (CONV-CVAE) effectively addresses limited labelled fault data whilst preserving physical plausibility and discriminative damage signatures, a contribution that extends beyond existing augmentation approaches by explicitly conditioning generation on both damage state and environmental factors. Second, physics-informed feature extraction providing translation invariance and noise robustness proved essential for effective cross-domain generalisation, with hierarchical wavelet-based representations outperforming raw signals or conventional time-frequency features. Third, a novel methodological insight emerged: the complementary nature of the 3 frameworks demonstrated that no single adaptation strategy dominates across all scenarios; rather, the severity of domain shift, availability of target supervision, and characteristics of sensor arrays determine optimal framework selection. This finding contributes a decision framework for practitioners to match adaptation strategies to deployment contexts. Fourth, novel integrated interpretability mechanisms tracking adaptation dynamics transformed opaque processes into diagnosable procedures, addressing the critical trust barrier for industrial adoption through visualisations of manifold trajectories, attention weights, and prototype evolution. Finally, careful architectural design prioritising computational efficiency enabled sophisticated DA on standard hardware, demonstrating that advanced domain adaptation can be deployed without specialised computational infrastructure, a novel practical contribution that removes a significant deployment barrier in resource-constrained industrial environments.

6.1 Contributions to Knowledge

This research makes several significant contributions to the field of SHM and ML-based damage detection. The development and validation of three complementary DA frameworks, i.e., UCTRF, GAT-CAMDA, and SPADA advances the state-of-the-art by addressing different adaptation scenarios encountered in practical SHM applications.

The UCTRF framework contributes a systematic integration of manifold learning, capsule networks, and transfer boosting for scenarios where limited target supervision is available. By combining UMAP dimensionality reduction with CapsNet feature extraction and TrAdaBoost instance reweighting, the framework demonstrates how sequential pipeline architectures can effectively manage severe data scarcity. The capsule network architecture, when used as a feature encoder rather than end-to-end classifier, provides structured vector representations capturing part-whole relationships in vibration signals, whilst transfer boosting iteratively adjusts sample weights to emphasise source instances that align well with target patterns. This contribution demonstrates that classical ML components, when appropriately integrated, remain competitive with purely DL approaches particularly under extreme data scarcity.

The GAT-CAMDA framework contributes an innovative application of graph attention networks to multi-sensor SHM, where sensor spatial relationships carry diagnostic information. By representing sensor arrays as graphs and learning attention weights over sensor connections, the framework identifies which sensors and sensor pairs contribute most to damage classification under varying conditions. The integration of multiple alignment mechanisms such as maximum mean discrepancy for moment matching, correlation alignment for second-order statistics, and adversarial training for domain invariance with graph-based feature extraction provides a comprehensive approach to feature-level DA. The sensor importance quantification derived from attention scores offers interpretability specifically relevant to multi-sensor monitoring systems, enabling practitioners to identify critical sensors and potentially optimise sensor placement strategies. This contribution demonstrates the value of incorporating structural knowledge about sensor topology into DA architectures.

The SPADA framework advances UDA for mechanical systems through three novel mechanisms. The confidence-gated prototype update rule couples' entropy-based filtering with similarity-weighted attention to control pseudo-label noise during self-training, enabling selective incorporation of high-confidence predictions whilst down-weighting uncertain samples. The indirect alignment strategy preserves inter-class margins through prototype compactness whilst encouraging domain invariance

through adversarial training, reducing feature collapse risks associated with explicit distribution matching. The comprehensive interpretability suite quantifies adaptation quality through multiple complementary visualisations prototype trajectories, instance-prototype similarity analysis, attention weight dynamics, and decision boundary evolution enabling practitioners to diagnose potential failure modes and build justified confidence in automated systems. This contribution demonstrates that prototype-based approaches with careful confidence management can achieve effective adaptation without any target domain labels, addressing the most challenging practical scenario.

Beyond the three frameworks, this research contributes validated methodologies spanning the complete diagnostic pipeline from data augmentation through feature extraction to domain-adapted classification. The CONV-CVAE framework for temperature-dependent damage synthesis introduces latent-space offsets and amplitude masks with multi-domain loss functions, enabling realistic signal generation across environmental conditions using only baseline damaged data. The comparative evaluation of the 3 frameworks across multiple domain shift scenarios establishes that framework selection should be guided by adaptation scenario characteristics such as supervision availability, shift severity, and sensor topology rather than seeking a universal solution. The demonstrated computational efficiency across all frameworks challenges the perception that advanced DA necessarily demands specialised hardware, broadening potential deployment scenarios to resource-constrained industrial environments.

6.2 Addressing the Research Questions

The research systematically addressed 6 research questions through theoretical development, methodological innovation, and empirical validation. Regarding data augmentation for scarce fault data, the investigation of 4 complementary strategies, e.g., windowing, spline interpolation, CONV-CVAE, and SMOTE demonstrated that physics-constrained synthesis successfully expands limited datasets whilst maintaining discriminative damage signatures. The empirical validation confirmed that models trained on augmented data generalised effectively to real test cases, validating the approach for practical applications. The success of different augmentation strategies across the two case studies, i.e., windowing for WTB-

VibClimate and CONV-CVAE for CONCEPT demonstrated that augmentation method selection should align with dataset characteristics and available prior knowledge about environmental dependencies.

Concerning EOVs effects, the research quantified substantial performance degradation under temperature variations without adaptation, characterised through distribution shift analysis revealing systematic changes in feature statistics and class separability. Mitigation was achieved through multiple complementary mechanisms deployed across the 3 frameworks: physics-informed feature extraction providing inherent stability, statistical and adversarial alignment techniques reducing distributional differences, and adaptive weighting mechanisms emphasising relevant training samples. The combination of these approaches successfully maintained diagnostic reliability under severe environmental variations (temperature), with each framework demonstrating effective mitigation strategies appropriate to its adaptation paradigm.

For cross-domain generalisation capabilities, the empirical validation demonstrated effective knowledge transfer across substantial environmental variations with varying levels of target supervision. UCTRF successfully leveraged minimal target labels to guide adaptation, GAT-CAMDA achieved semi-supervised adaptation for moderate shifts through comprehensive feature alignment, and SPADA maintained high performance on severe shifts without any target labels through confidence-gated pseudo-labelling. The frameworks successfully adapted across temperature ranges spanning tens of degrees and different structural configurations with varying loading conditions, confirming that appropriately designed adaptation mechanisms enable model deployment under conditions significantly different from training environments. The comparative results established that adaptation capability scales with framework sophistication, with more complex mechanisms justified for more severe distribution shifts.

Regarding model interpretability, each framework integrated interpretability mechanisms appropriate to its architecture. UCTRF provided interpretability through visualisation of manifold embeddings and instance weights revealing which training samples contributed most to target predictions. GAT-CAMDA offered sensor importance scores quantifying contributions of individual sensors and sensor pairs to

diagnostic decisions, particularly valuable for multi-sensor array optimisation. SPADA incorporated comprehensive visualisation mechanisms, i.e., prototype trajectories, instance-prototype similarity analysis, attention weight dynamics, and decision boundary evolution enabling detailed diagnosis of adaptation dynamics. These complementary approaches demonstrated that interpretability mechanisms must be tailored to framework architecture, with different visualisation strategies appropriate for instance-based, feature-based, and prototype-based adaptation paradigms.

Concerning computational requirements and practical constraints, the research demonstrated that all 3 frameworks can operate within realistic constraints. Completion of comprehensive training and hyperparameter optimisation on standard CPU workstations eliminates specialised hardware requirements across diverse architectural choices. UCTRF achieved efficiency through dimensionality reduction prior to neural processing, GAT-CAMDA through sparse graph representations of sensor topology, and SPADA through pre-computed stable features and efficient prototype operations. These design choices positioned all frameworks as viable for deployment in resource-constrained operational environments typical of industrial SHM.

6.3 Limitations

A number of limitations warrant acknowledgement and suggest directions for future investigation. The experimental validation, whilst comprehensive within its scope, was conducted on laboratory-scale structures under controlled conditions. Field deployment on full-scale operational assets would introduce additional complexities including sensor degradation, installation variability, and uncontrolled environmental factors beyond temperature. The datasets comprised limited damage classes and relatively pristine signals compared to industrial environments where multiple simultaneous faults, progressive damage evolution, and contaminating noise sources are prevalent.

All 3 frameworks assume that source and target domains share identical damage classes, which may not hold in practice where previously unseen failure modes can emerge during operation. Open-set DA represents an important extension not addressed in this work. Additionally, the frameworks primarily addressed single-

source to single-target adaptation scenarios, whereas operational fleets consist of multiple assets with varying health states and operational histories. Multi-source DA approaches that can aggregate knowledge from heterogeneous sources warrant further investigation.

The computational efficiency demonstrations validated feasibility on standard hardware but did not address real-time inference constraints for continuous monitoring applications. The frameworks were evaluated in batch processing mode with offline hyperparameter optimisation, whereas operational deployment requires streaming data processing with bounded latency guarantees. Whilst the interpretability mechanisms provide valuable insights, they require expert knowledge to interpret effectively. Standardised interpretation protocols and automated diagnostic criteria that can flag potential adaptation failures without requiring specialist expertise would enhance practical utility.

6.4 Future Work

The findings and limitations of this research suggest multitude promising directions for future investigation. Multi-source DA would enable aggregating knowledge from multiple structural assets simultaneously, potentially improving detection of subtle or uncommon faults underrepresented in any single source. This extension could integrate population-based SHM concepts with the prototype-based, graph-based, and instance-based adaptation mechanisms demonstrated in this research, providing principled approaches to fleet-wide learning whilst respecting individual asset characteristics. Research challenges include developing appropriate weighting schemes to emphasise relevant sources for each target asset and addressing potential negative transfer when sources exhibit conflicting patterns.

Open-set recognition and novel damage detection represent critical extensions for operational safety. Real-world monitoring inevitably encounters damage modes not represented in training data, requiring mechanisms to quantify prediction uncertainty, establish confidence thresholds for known classes, and detect outliers falling outside the manifold of training classes. Integration of out-of-distribution detection methods with the 3 frameworks could enable systems that recognise when target samples belong to unknown classes, preventing misclassification of unseen faults as known benign conditions. The confidence-gating mechanisms in SPADA

provide a foundation for such extensions, as entropy-based filtering already identifies uncertain predictions. Extending this principle to distinguish between uncertain predictions within known classes versus truly unseen damage patterns would enhance operational safety.

Physics-informed DA offers opportunities to further improve robustness and interpretability across all frameworks. Dispersion curve relationships for guided waves, modal properties for vibration monitoring, and constitutive material behaviours for composite structures provide powerful constraints that could regularise adaptation and prevent physically implausible mappings. Research directions include developing physics-informed losses penalising violations of conservation laws or material property bounds, incorporating finite element models to generate synthetic training data with controlled damage parameters, and integrating symbolic regression to identify interpretable, physics-consistent transformation rules. Such extensions could be incorporated into UCTRF through physics-based feature selection, into GAT-CAMDA through physics-informed graph construction, and into SPADA through physics-constrained prototype initialisation.

Continual learning frameworks that can incrementally incorporate new information without catastrophic forgetting would enable monitoring systems that adapt alongside progressive asset degradation and evolving operational conditions. Research challenges include developing rehearsal strategies retaining knowledge of previous environmental conditions whilst learning new operating regimes, investigating architectural approaches accommodating new domains without interfering with existing capabilities, and establishing criteria for when adaptation should be performed versus when retraining is necessary. The instance-weighting mechanisms in UCTRF could be extended to weight historical data during continual updates, whilst the prototype-based approach in SPADA naturally supports incremental prototype refinement as new conditions are encountered.

Real-time implementation and edge deployment require addressing latency constraints, memory limitations, and energy efficiency considerations. Research directions include investigating model compression techniques such as pruning, quantisation, and knowledge distillation that reduce computational requirements without compromising adaptation performance, developing incremental adaptation

algorithms updating parameters online as new data arrive, and designing hierarchical architectures where intensive adaptation occurs periodically whilst lightweight inference operates continuously. The modular architectures of UCTRF and SPADA with distinct feature extraction and classification stages may facilitate such decompositions, whilst GAT-CAMDA's graph sparsity provides natural opportunities for computational reduction.

Uncertainty quantification and risk-based decision-making represent important extensions for safety-critical applications. Developing well-calibrated uncertainty estimates could enable risk-based frameworks balancing detection performance against inspection costs and failure consequences. Research directions include investigating Bayesian approaches providing posterior distributions over predictions, developing conformal prediction methods providing valid coverage guarantees under distribution shift, and integrating structural reliability models with ML diagnostics to propagate prediction uncertainties into structural safety assessments. The ensemble nature of Random Forest in UCTRF naturally provides uncertainty estimates, whilst the attention mechanisms in GAT-CAMDA and prototype distances in SPADA could be calibrated to provide confidence scores suitable for risk-based decision frameworks.

Cross-modal and cross-system generalisation would enable more comprehensive monitoring leveraging complementary information sources. Developing frameworks that transfer knowledge across fundamentally different signal types such as vibration, acoustic emission, guided waves, strain measurements could enhance diagnostic capabilities. Research directions include investigating cross-modal representation learning identifying abstract damage concepts independent of specific sensing modalities, developing fusion architectures combining predictions from adapted models operating on different signal types, and exploring whether the adaptation strategies demonstrated here generalise beyond rotating machinery and composite plates to other structural systems. The graph-based representation in GAT-CAMDA appears particularly promising for cross-modal fusion, as heterogeneous sensor types could be represented as different node types within a unified graph structure.

6.5 Concluding Remarks

This research has demonstrated that advanced ML frameworks integrating principled data augmentation, physics-informed feature extraction, and sophisticated

DA techniques can address fundamental challenges in SHM of composite materials. The development of three complementary frameworks provides practitioners with adaptation strategies appropriate for different operational scenarios characterised by varying levels of supervision availability, domain shift severity, and sensor topology characteristics. The methodologies enable reliable cross-domain damage diagnosis under severe environmental variations without requiring extensive target domain labels or prohibitive computational resources, thereby advancing the practical deployability of automated monitoring systems for safety-critical applications.

Each framework demonstrated distinct strengths aligned with its design philosophy and target scenario. UCTRF effectively managed severe data scarcity through instance reweighting with minimal target supervision, GAT-CAMDA exploited sensor spatial relationships through attention-based graph learning whilst achieving semi-supervised adaptation for moderate shifts, and SPADA achieved superior performance on severe shifts through confidence-gated prototype-based learning without any target labels. The integration of interpretability mechanisms tailored to each framework's architecture addresses the critical gap between achieving high predictive accuracy and building justified confidence in automated diagnostic systems, which is essential for adoption in safety-critical applications.

The experimental validation on heterogeneous structural systems demonstrated generalisability across different composite materials, sensing modalities, and damage mechanisms. The robust performance under substantial environmental variations confirms that the proposed methodologies can accommodate the variability characteristic of real-world structural monitoring scenarios. The demonstrated computational efficiency across all frameworks removes a significant barrier to practical adoption, enabling deployment in resource-constrained industrial environments without requiring specialised computing infrastructure.

As composite materials continue to proliferate in safety-critical applications spanning aerospace, renewable energy, and transportation infrastructure, the need for robust, adaptable, and interpretable automated monitoring systems will intensify. This research provides a foundation for next-generation SHM systems capable of learning from limited data, generalising across diverse conditions, explaining their reasoning, and operating within practical constraints. The identified future work directions offer

promising opportunities to further enhance capabilities, robustness, and applicability, advancing the field towards truly autonomous and trustworthy monitoring systems capable of safeguarding critical infrastructure throughout its operational lifetime.

Bibliography

1. 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, doi:10.1177/14759217241295951.
2. Rajak, D.K.; Pagar, D.D.; Kumar, R.; Pruncu, C.I. Recent Progress of Reinforcement Materials: A Comprehensive Overview of Composite Materials. *Journal of Materials Research and Technology* 2019, 8, 6354–6374, doi:10.1016/j.jmrt.2019.09.068.
3. Arwood, Z.; Cousins, D.S.; Young, S.; Stebner, A.P.; Penumadu, D. Infusible Thermoplastic Composites for Wind Turbine Blade Manufacturing: Static Characterization of Thermoplastic Laminates under Ambient Conditions. *Composites Part C: Open Access* 2023, 11, 100365, doi:10.1016/j.jcomc.2023.100365.
4. Heslehurst, R.B. *Defects and Damage in Composite Materials and Structures*; CRC Press, 2014; ISBN 9781466580480.
5. Mehdikhani, M.; Gorbatiikh, L.; Verpoest, I.; Lomov, S. V Voids in Fiber-Reinforced Polymer Composites: A Review on Their Formation, Characteristics, and Effects on Mechanical Performance. *J Compos Mater* 2019, 53, 1579–1669, doi:10.1177/0021998318772152.
6. Janeliukstis, R.; Baranovskis, D.; Katunin, A.; Zorin, I.; Burgholzer, P.; Lopes, H.; Dragan, K.; Rucevskis, S.; Gaile, L.; Chen, X. Nondestructive Evaluation of Barely Visible Impact Damage in Composite Structures – a Review. *Compos Struct* 2025, 373, 119661, doi:10.1016/j.compstruct.2025.119661.
7. Sethi, S.; Ray, B.C. Environmental Effects on Fibre Reinforced Polymeric Composites: Evolving Reasons and Remarks on Interfacial Strength and Stability. *Adv Colloid Interface Sci* 2015, 217, 43–67, doi:10.1016/j.cis.2014.12.005.
8. Ge, J.; Catalanotti, G.; Falzon, B.G.; Higgins, C.; McClory, C.; Thiebot, J.-A.; Zhang, L.; He, M.; Jin, Y.; Sun, D. Process Characteristics, Damage Mechanisms and Challenges in Machining of Fibre Reinforced Thermoplastic Polymer (FRTP) Composites: A Review. *Compos B Eng* 2024, 273, 111247, doi:10.1016/j.compositesb.2024.111247.
9. Wilson, C.L.; Lonkar, K.; Roy, S.; Kopsaftopoulos, F.; Chang, F.-K. 7.20 Structural Health Monitoring of Composites. In *Comprehensive Composite Materials II*; Elsevier, 2018; pp. 382–407.
10. Liu, J.; Li, Q.; Li, L.; An, S. Structural Damage Detection and Localization via an Unsupervised Anomaly Detection Method. *Reliab Eng Syst Saf* 2024, 252, 110465, doi:10.1016/j.ress.2024.110465.
11. Mousavi, M.; Gandomi, A.H. Structural Health Monitoring under Environmental and Operational Variations Using MCD Prediction Error. *J Sound Vib* 2021, 512, 116370, doi:10.1016/j.jsv.2021.116370.
12. Aulakh, D.S.; Rawat, D.; Bhalla, S. A Novel Integrated Framework for Damage Detection in Steel Structures Using Magnetic Piezoelectric Sensors and Hybrid Damage Index. *Sci Rep* 2025, 15, 36083, doi:10.1038/s41598-025-19936-1.

13. Cha, Y.-J.; Ali, R.; Lewis, J.; Büyüköztürk, O. Deep Learning-Based Structural Health Monitoring. *Autom Constr* 2024, *161*, 105328, doi:10.1016/j.autcon.2024.105328.
14. Tang, S.; Ma, J.; Yan, Z.; Zhu, Y.; Khoo, B.C. Deep Transfer Learning Strategy in Intelligent Fault Diagnosis of Rotating Machinery. *Eng Appl Artif Intell* 2024, *134*, 108678, doi:10.1016/j.engappai.2024.108678.
15. Rezazadeh, N.; Felaco, A.; Fallahy, S.; Lamanna, G. Application of Supervised and Unsupervised Machine Learning to the Classification of Damaged Rotor-Bearing Systems. *Macromol Symp* 2023, *411*, doi:10.1002/masy.202200219.
16. Liu, Y.; Jiang, H.; Yao, R.; Zeng, T. Counterfactual-Augmented Few-Shot Contrastive Learning for Machinery Intelligent Fault Diagnosis with Limited Samples. *Mech Syst Signal Process* 2024, *216*, 111507, doi:10.1016/j.ymsp.2024.111507.
17. Alzubaidi, L.; Bai, J.; Al-Sabaawi, A.; Santamaría, J.; Albahri, A.S.; Al-dabbagh, B.S.N.; Fadhel, M.A.; Manoufali, M.; Zhang, J.; Al-Timemy, A.H.; et al. A Survey on Deep Learning Tools Dealing with Data Scarcity: Definitions, Challenges, Solutions, Tips, and Applications. *J Big Data* 2023, *10*, 46, doi:10.1186/s40537-023-00727-2.
18. Yu, Y.; Ma, R.; Gu, W.; Ma, Z. DiffAT: Effective Data Augmentation with Diffusion Models for Time Series Forecasting. *Eng Appl Artif Intell* 2025, *161*, 112091, doi:10.1016/j.engappai.2025.112091.
19. Iglesias, G.; Talavera, E.; González-Prieto, Á.; Mozo, A.; Gómez-Canaval, S. Data Augmentation Techniques in Time Series Domain: A Survey and Taxonomy. *Neural Comput Appl* 2023, *35*, 10123–10145, doi:10.1007/s00521-023-08459-3.
20. Mumuni, A.; Mumuni, F. Data Augmentation: A Comprehensive Survey of Modern Approaches. *Array* 2022, *16*, 100258, doi:10.1016/j.array.2022.100258.
21. Li, F.; Leng, J.; Liu, Y.; Remillat, C.; Scarpa, F. Temperature Dependence of Elastic Constants in Unidirectional Carbon Fiber Reinforced Shape Memory Polymer Composites. *Mechanics of Materials* 2020, *148*, 103518, doi:10.1016/j.mechmat.2020.103518.
22. Yuan, X.W.; Li, W.G.; Xiao, Z.M.; Zhang, Y.M. Prediction of Temperature-Dependent Transverse Strength of Carbon Fiber Reinforced Polymer Composites by a Modified Cohesive Zone Model. *Compos Struct* 2023, *304*, 116310, doi:10.1016/j.compstruct.2022.116310.
23. Lijuan, J.; Xin, W.; Xiao, C.; Yinzi, Z.; Yongzhan, Y. A Review of Moisture Absorption and Corrosion Resistance Performance of FRTP in Marine Environments. *Polym Test* 2025, *145*, 108760, doi:10.1016/j.polymertesting.2025.108760.
24. Guo, W.; Sun, W.; Guo, Q.; Ding, R.; Liu, C.; Liu, Y. Effect of Stress Loading on Microstructural Evolution and Precipitate Behavior during High-Temperature Creep of Transient Liquid Phase Bonded RAFM Steel Joints. *Materials Science and Engineering: A* 2026, *949*, 149403, doi:10.1016/j.msea.2025.149403.
25. Rupal, B.S.; Anwer, N.; Secanell, M.; Qureshi, A.J. Geometric Tolerance and Manufacturing Assemblability Estimation of Metal Additive Manufacturing (AM) Processes. *Mater Des* 2020, *194*, 108842, doi:10.1016/j.matdes.2020.108842.

26. Keshmiry, A.; Hassani, S.; Mousavi, M.; Dackermann, U. Effects of Environmental and Operational Conditions on Structural Health Monitoring and Non-Destructive Testing: A Systematic Review. *Buildings* 2023, *13*, 918, doi:10.3390/buildings13040918.
27. 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, doi:10.1088/1361-665X/ade7db.
28. Wang, Y.; Gao, L.; Yuan, S.; Qiu, L.; Qing, X. An Adaptive Filter-Based Temperature Compensation Technique for Structural Health Monitoring. *J Intell Mater Syst Struct* 2014, *25*, doi:10.1177/1045389X13519001.
29. Yang, C.; Taniguchi, K.; Akashi, Y. Transfer Learning with Unsupervised Domain Adaptation for Personal Thermal Comfort Prediction. *Energy Build* 2025, *332*, 115449, doi:10.1016/j.enbuild.2025.115449.
30. 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, doi:10.1177/14759217241313350.
31. Lu, Y.; Wong, W.K.; Yuan, C.; Lai, Z.; Li, X. Low-Rank Correlation Learning for Unsupervised Domain Adaptation. *IEEE Trans Multimedia* 2024, *26*, 4153–4167, doi:10.1109/TMM.2023.3321430.
32. 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, doi:10.3390/s25226847.
33. Zschech, P.; Weinzierl, S.; Kraus, M. Inherently Interpretable Machine Learning: A Contrasting Paradigm to Post-Hoc Explainable AI. *Business & Information Systems Engineering* 2025, doi:10.1007/s12599-025-00964-0.
34. Mardanshahi, A.; Sreekumar, A.; Yang, X.; Barman, S.K.; Chronopoulos, D. Sensing Techniques for Structural Health Monitoring: A State-of-the-Art Review on Performance Criteria and New-Generation Technologies. *Sensors* 2025, *25*, 1424, doi:10.3390/s25051424.
35. Okagawa, S.; Bernus, P.; Noran, O. Realtime Health Monitoring of Composite Structures Using FBG Sensors. *IFAC-PapersOnLine* 2022, *55*, 157–162, doi:10.1016/j.ifacol.2022.09.200.
36. Zhang, P.; He, Z.; Cui, C.; Xu, C.; Ren, L. An Edge-Computing Framework for Operational Modal Analysis of Offshore Wind-Turbine Tower. *Ocean Engineering* 2023, *287*, 115720, doi:10.1016/j.oceaneng.2023.115720.
37. Ngo, D.; Park, H.-C.; Kang, B. Edge Intelligence: A Review of Deep Neural Network Inference in Resource-Limited Environments. *Electronics (Basel)* 2025, *14*, 2495, doi:10.3390/electronics14122495.
38. Dantas, P.V.; Sabino da Silva, W.; Cordeiro, L.C.; Carvalho, C.B. A Comprehensive Review of Model Compression Techniques in Machine Learning. *Applied Intelligence* 2024, *54*, 11804–11844, doi:10.1007/s10489-024-05747-w.
39. Al-Thelaya, K.; Gilal, N.U.; Alzubaidi, M.; Majeed, F.; Agus, M.; Schneider, J.; Househ, M. Applications of Discriminative and Deep Learning Feature

- Extraction Methods for Whole Slide Image Analysis: A Survey. *J Pathol Inform* 2023, 14, 100335, doi:10.1016/j.jpi.2023.100335.
40. Jiang, Z.; Wang, X.; Li, H.; Hong, T.; You, F.; Dragoña, J.; Vrabie, D.; Dong, B. Physics-Informed Machine Learning for Building Performance Simulation-A Review of a Nascent Field. *Advances in Applied Energy* 2025, 18, 100223, doi:10.1016/j.adapen.2025.100223.
 41. Resck, L.; M. Raimundo, M.; Poco, J. Exploring the Trade-off Between Model Performance and Explanation Plausibility of Text Classifiers Using Human Rationales. In Proceedings of the Findings of the Association for Computational Linguistics: NAACL 2024; Association for Computational Linguistics: Stroudsburg, PA, USA, 2024; pp. 4190–4216.
 42. The Future of Edge Computing. In; 2022; pp. 333–358.
 43. Buckley, T.; Ghosh, B.; Pakrashi, V. A Feature Extraction & Selection Benchmark for Structural Health Monitoring. *Struct Health Monit* 2023, 22, doi:10.1177/14759217221111141.
 44. Ashkarkalaei, M.; Ghiasi, R.; Pakrashi, V.; Malekjafarian, A. Feature Selection for Unsupervised Defect Detection of a Wind Turbine Blade Considering Operational and Environmental Conditions. *Mech Syst Signal Process* 2025, 230, 112568, doi:10.1016/j.ymsp.2025.112568.
 45. Akkaya, S. Wavelet-Based Denoising Strategies for Non-Stationary Signals in Electrical Power Systems: An Optimization Perspective. *Electronics (Basel)* 2025, 14, 3190, doi:10.3390/electronics14163190.
 46. Rezazadeh, N.; de Oliveira, M.; Perfetto, D.; De Luca, A.; Caputo, F. Classification of Unbalanced and Bowed Rotors under Uncertainty Using Wavelet Time Scattering, LSTM, and SVM. *Applied Sciences (Switzerland)* 2023, 13, doi:10.3390/app13126861.
 47. Ahmed, A.; Serrestou, Y.; Raouf, K.; Diouris, J.-F. Empirical Mode Decomposition-Based Feature Extraction for Environmental Sound Classification. *Sensors* 2022, 22, 7717, doi:10.3390/s22207717.
 48. Bento, N.; Rebelo, J.; Barandas, M.; Carreiro, A. V.; Campagner, A.; Cabitza, F.; Gamboa, H. Comparing Handcrafted Features and Deep Neural Representations for Domain Generalization in Human Activity Recognition. *Sensors* 2022, 22, 7324, doi:10.3390/s22197324.
 49. Hu, G.; You, F. Hybrid Physics-Based and Data-Driven Model Predictive Control for Multi-Zone Building's Thermal Comfort Under Disjunctive Uncertainty. *IFAC-PapersOnLine* 2022, 55, 352–357, doi:10.1016/j.ifacol.2022.07.469.
 50. Zong, S.; Wang, S.; Luo, Z.; Wu, X.; Zhang, H.; Ni, Z. Robust Damage Detection and Localization Under Complex Environmental Conditions Using Singular Value Decomposition-Based Feature Extraction and One-Dimensional Convolutional Neural Network. *Chinese Journal of Mechanical Engineering* 2023, 36, 61, doi:10.1186/s10033-023-00889-3.
 51. Soomro, A.A.; Muhammad, M.B.; Mokhtar, A.A.; Md Saad, M.H.; Lashari, N.; Hussain, M.; Sarwar, U.; Palli, A.S. Insights into Modern Machine Learning Approaches for Bearing Fault Classification: A Systematic Literature Review. *Results in Engineering* 2024, 23, 102700, doi:10.1016/j.rineng.2024.102700.
 52. Perfetto, D.; Rezazadeh, N.; Aversano, A.; De Luca, A.; Lamanna, G. Composite Panel Damage Classification Based on Guided Waves and Machine

- Learning: An Experimental Approach. *Applied Sciences* 2023, 13, 10017, doi:10.3390/app131810017.
53. Susu, A.A.; Agboola, H.A.; Solebo, C.; Lesi, F.E.A.; Aribike, D.S. Wavelet Time Scattering Based Classification of Interictal and Preictal EEG Signals. *Journal of Brain Research* 2020, 3.
 54. Ojha, S.; Jangid, N.; Shelke, A.; Habib, A. Probabilistic Impact Localization in Composites Using Wavelet Scattering Transform and Multi-Output Gaussian Process Regression. *Measurement* 2024, 236, 115078, doi:10.1016/j.measurement.2024.115078.
 55. Ma, S.; Cheng, B.; Shang, Z.; Liu, G. Scattering Transform and LSPTSVM Based Fault Diagnosis of Rotating Machinery. *Mech Syst Signal Process* 2018, 104, 155–170, doi:10.1016/j.ymssp.2017.10.026.
 56. Rezazadeh, N.; De Luca, A.; Perfetto, D. Unbalanced, Cracked, and Misaligned Rotating Machines: A Comparison between Classification Procedures throughout the Steady-State Operation. *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 2022, 44, doi:10.1007/s40430-022-03750-1.
 57. Guo, J.; Si, Z.; Xiang, J. A Compound Fault Diagnosis Method of Rolling Bearing Based on Wavelet Scattering Transform and Improved Soft Threshold Denoising Algorithm. *Measurement* 2022, 196, 111276, doi:10.1016/j.measurement.2022.111276.
 58. Rezazadeh, N.; Polverino, A.; Perfetto, D.; De Luca, A. Dimensionality Reduction in Structural Health Monitoring: A Case Study on Damaged Wind Turbine Blades. *Macromol Symp* 2024, 413, doi:10.1002/masy.202400044.
 59. Golinval, J.-C. Damage Detection in Structures Based on Principal Component Analysis of Forced Harmonic Responses. *Procedia Eng* 2017, 199, 1912–1918, doi:10.1016/j.proeng.2017.09.449.
 60. Chen, J.; Wang, C.; Wang, B.; Zhou, Z. A Visualized Classification Method via T-Distributed Stochastic Neighbor Embedding and Various Diagnostic Parameters for Planetary Gearbox Fault Identification from Raw Mechanical Data. *Sens Actuators A Phys* 2018, 284, 52–65, doi:10.1016/j.sna.2018.10.021.
 61. Zhang, Y.; Ye, D.; Liu, Y. Robust Locally Linear Embedding Algorithm for Machinery Fault Diagnosis. *Neurocomputing* 2018, 273, 323–332, doi:10.1016/j.neucom.2017.07.048.
 62. Rahbari, A.; Rébillat, M.; Mechbal, N.; Canu, S. Unsupervised Damage Clustering in Complex Aeronautical Composite Structures Monitored by Lamb Waves: An Inductive Approach. *Eng Appl Artif Intell* 2021, 97, doi:10.1016/j.engappai.2020.104099.
 63. McInnes, L.; Healy, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. 2018, doi:https://doi.org/10.48550/arXiv.1802.03426.
 64. Lu, J.-H.; Chen, Z.-W. Spatio-Temporal Graph Attention Network and Graph-Based Transformer Architecture for Distributed Urban Wind Sequence Reconstruction and Forecasting. *Measurement* 2025, 252, 117400, doi:10.1016/j.measurement.2025.117400.
 65. Marchisio, A.; Bussolino, B.; Salvati, E.; Martina, M.; Masera, G.; Shafique, M. Enabling Capsule Networks at the Edge through Approximate Softmax and Squash Operations.; 2022; pp. 1–6.

66. Rezazadeh, N.; Ashory, M.-R.; Fallahy, S. Identification of Shallow Cracks in Rotating Systems by Utilizing Convolutional Neural Networks and Persistence Spectrum under Constant Speed Condition. *Journal of Mechanical Engineering, Automation and Control Systems* 2021, 2, 135–147, doi:10.21595/jmeacs.2021.22221.
67. Mirkhalaf, M.; Rocha, I. Micromechanics-Based Deep-Learning for Composites: Challenges and Future Perspectives. *European Journal of Mechanics - A/Solids* 2024, 105, 105242, doi:10.1016/j.euromechsol.2024.105242.
68. Wu, J.; Xu, X.; Liu, C.; Deng, C.; Shao, X. Lamb Wave-Based Damage Detection of Composite Structures Using Deep Convolutional Neural Network and Continuous Wavelet Transform. *Compos Struct* 2021, 276, doi:10.1016/j.compstruct.2021.114590.
69. Gao, F.; Hua, J. Damage Characterization Using CNN and SAE of Broadband Lamb Waves. *Ultrasonics* 2022, 119, 106592, doi:10.1016/j.ultras.2021.106592.
70. Xu, X.; Liu, C. Physics-Guided Deep Learning for Damage Detection in CFRP Composite Structures. *Compos Struct* 2024, 331, 117889, doi:10.1016/j.compstruct.2024.117889.
71. Hinton, G.E.; Krizhevsky, A.; Wang, S.D. Transforming Auto-Encoders. In Proceedings of the Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); 2011; Vol. 6791 LNCS.
72. Kwabena Patrick, M.; Felix Adekoya, A.; Abra Mighty, A.; Edward, B.Y. Capsule Networks – A Survey. *Journal of King Saud University - Computer and Information Sciences* 2022, 34.
73. Liang, P.; Deng, C.; Yuan, X.; Zhang, L. A Deep Capsule Neural Network with Data Augmentation Generative Adversarial Networks for Single and Simultaneous Fault Diagnosis of Wind Turbine Gearbox. *ISA Trans* 2023, 135, doi:10.1016/j.isatra.2022.10.008.
74. Huang, R.; Li, J.; Wang, S.; Li, G.; Li, W. A Robust Weight-Shared Capsule Network for Intelligent Machinery Fault Diagnosis. *IEEE Trans Industr Inform* 2020, 16, doi:10.1109/TII.2020.2964117.
75. Jiang, Q.; Huang, X.; Qu, W.; Xiao, L.; Lu, Y. Domain-Separated Capsule Network for Damage Detection in Aluminum Plates under Varying Vibration Conditions. *Ultrasonics* 2025, 154, 107688, doi:10.1016/j.ultras.2025.107688.
76. Michaels, J.E. Sparse Array Imaging with Guided Waves under Variable Environmental Conditions. In *Structural Health Monitoring (SHM) in Aerospace Structures*; Elsevier, 2016; pp. 255–284.
77. Vrahatis, A.G.; Lazaros, K.; Kotsiantis, S. Graph Attention Networks: A Comprehensive Review of Methods and Applications. *Future Internet* 2024, 16, 318, doi:10.3390/fi16090318.
78. Niu, J.; Li, S.; Li, Z. Restoration of Missing Structural Health Monitoring Data Using Spatiotemporal Graph Attention Networks. *Struct Health Monit* 2022, 21, doi:10.1177/14759217211056832.
79. Zhao, M.; Taal, C.; Baggerohr, S.; Fink, O. Virtual Sensor for Real-Time Bearing Load Prediction Using Heterogeneous Temporal Graph Neural Networks. *PHM Society European Conference* 2024, 8, 8, doi:10.36001/phme.2024.v8i1.3998.

80. Zhao, Z.; Chen, N.-Z. An Exponential Smoothing Multi-Head Graph Attention Network (ESMGAT) Method for Damage Zone Localization on Wind Turbine Blades. *Compos Struct* 2024, 342, 118224, doi:10.1016/j.compstruct.2024.118224.
81. Demir, S.; Mincev, K.; Kok, K.; Paterakis, N.G. Data Augmentation for Time Series Regression: Applying Transformations, Autoencoders and Adversarial Networks to Electricity Price Forecasting. *Appl Energy* 2021, 304, 117695, doi:10.1016/j.apenergy.2021.117695.
82. Javanmardi, F.; Kadiri, S.R.; Alku, P. A Comparison of Data Augmentation Methods in Voice Pathology Detection. *Comput Speech Lang* 2024, 83, 101552, doi:10.1016/j.csl.2023.101552.
83. Basheer, S.; Aldehim, G.; Alluhaidan, A.S.; Sakri, S. Improving Mental Dysfunction Detection from EEG Signals: Self-Contrastive Learning and Multitask Learning with Transformers. *Alexandria Engineering Journal* 2024, 106, 52–59, doi:10.1016/j.aej.2024.06.058.
84. Wu, Y.; Ma, X.; Guo, G.; Jia, T.; Huang, Y.; Liu, S.; Fan, J.; Wu, X. Advancing Deep Learning-Based Acoustic Leak Detection Methods towards Application for Water Distribution Systems from a Data-Centric Perspective. *Water Res* 2024, 261, 121999, doi:10.1016/j.watres.2024.121999.
85. Ejiolor Matthew, D.; Shi, J.; Hou, M.; Cao, H. Improved STFT Analysis Using Time-Frequency Masking for Chatter Detection in the Milling Process. *Measurement* 2024, 225, 113899, doi:10.1016/j.measurement.2023.113899.
86. Kivaisi, A.R.; Zhao, Q. Improved Mini-Batch Multiple Augmentation for Low-Resource Spoken Word Recognition. *Expert Syst Appl* 2024, 252, 124157, doi:10.1016/j.eswa.2024.124157.
87. Yin, W.; Xia, H.; Huang, X.; Wang, Z. A Fault Diagnosis Method for Nuclear Power Plants Rotating Machinery Based on Deep Learning under Imbalanced Samples. *Ann Nucl Energy* 2024, 199, 110340, doi:10.1016/j.anucene.2024.110340.
88. Liu, D.; Zhong, S.; Lin, L.; Zhao, M.; Fu, X.; Liu, X. Feature-Level SMOTE: Augmenting Fault Samples in Learnable Feature Space for Imbalanced Fault Diagnosis of Gas Turbines. *Expert Syst Appl* 2024, 238, 122023, doi:10.1016/j.eswa.2023.122023.
89. Wang, J.; Wei, J.; Huang, H.; Wen, L.; Yuan, Y.; Chen, H.; Wu, R.; Wu, J. IMWMOTE: A Novel Oversampling Technique for Fault Diagnosis in Heterogeneous Imbalanced Data. *Expert Syst Appl* 2024, 251, 123987, doi:10.1016/j.eswa.2024.123987.
90. Ni, Y.Q.; Hua, X.G.; Fan, K.Q.; Ko, J.M. Correlating Modal Properties with Temperature Using Long-Term Monitoring Data and Support Vector Machine Technique. *Eng Struct* 2005, 27, 1762–1773, doi:10.1016/j.engstruct.2005.02.020.
91. Moser, P.; Moaveni, B. Environmental Effects on the Identified Natural Frequencies of the Dowling Hall Footbridge. *Mech Syst Signal Process* 2011, 25, 2336–2357, doi:10.1016/j.ymsp.2011.03.005.
92. Soo Lon Wah, W.; Chen, Y.-T.; Owen, J.S. A Regression-Based Damage Detection Method for Structures Subjected to Changing Environmental and Operational Conditions. *Eng Struct* 2021, 228, 111462, doi:10.1016/j.engstruct.2020.111462.

93. Zhao, X.; Yao, J.; Deng, W.; Jia, M.; Liu, Z. Normalized Conditional Variational Auto-Encoder with Adaptive Focal Loss for Imbalanced Fault Diagnosis of Bearing-Rotor System. *Mech Syst Signal Process* 2022, *170*, 108826, doi:10.1016/j.ymssp.2022.108826.
94. Ge, L.; Sadhu, A. Domain Adaptation for Structural Health Monitoring via Physics-Informed and Self-Attention-Enhanced Generative Adversarial Learning. *Mech Syst Signal Process* 2024, *211*, 111236, doi:10.1016/j.ymssp.2024.111236.
95. Yi, H.; Hou, L.; Jin, Y.; Saeed, N.A.; Kandil, A.; Duan, H. Time Series Diffusion Method: A Denoising Diffusion Probabilistic Model for Vibration Signal Generation. *Mech Syst Signal Process* 2024, *216*, 111481, doi:10.1016/j.ymssp.2024.111481.
96. Vallevik, V.B.; Babic, A.; Marshall, S.E.; Elvatun, S.; Brøgger, H.M.B.; Alagaratnam, S.; Edwin, B.; Veeraragavan, N.R.; Befring, A.K.; Nygård, J.F. Can I Trust My Fake Data – A Comprehensive Quality Assessment Framework for Synthetic Tabular Data in Healthcare. *Int J Med Inform* 2024, *185*, 105413, doi:10.1016/j.ijmedinf.2024.105413.
97. Bahrpeyma, F.; Roantree, M.; Cappellari, P.; Scriney, M.; McCarren, A. A Methodology for Validating Diversity in Synthetic Time Series Generation. *MethodsX* 2021, *8*, 101459, doi:10.1016/j.mex.2021.101459.
98. Croxford, A.J.; Wilcox, P.D.; Lu, Y.; Michaels, J.; Drinkwater, B.W. Quantification of Environmental Compensation Strategies for Guided Wave Structural Health Monitoring.; Kundu, T., Ed.; March 27 2008; p. 69350H.
99. Mariani, S.; Heinlein, S.; Cawley, P. Compensation for Temperature-Dependent Phase and Velocity of Guided Wave Signals in Baseline Subtraction for Structural Health Monitoring. *Struct Health Monit* 2020, *19*, doi:10.1177/1475921719835155.
100. Nandyala, A.R.; Darpe, A.K.; Singh, S.P. Damage Localization in Cross-Ply Laminated Composite Plates under Varying Temperature Conditions Using Lamb Waves. *Meas Sci Technol* 2020, *31*, doi:10.1088/1361-6501/ab6eca.
101. Salmanpour, M.S.; Sharif Khodaei, Z.; Aliabadi, M.H. Guided Wave Temperature Correction Methods in Structural Health Monitoring. *J Intell Mater Syst Struct* 2017, *28*, doi:10.1177/1045389X16651155.
102. Huang, T.; Yuan, S.; Qiu, L. Adaline Network-Based Temperature Compensation Method for SHM Method. In Proceedings of the 7th European Workshop on Structural Health Monitoring, EWSHM 2014 - 2nd European Conference of the Prognostics and Health Management (PHM) Society; 2014.
103. Aryan, P.; Kotousov, A.; Ng, C.T.; Wildy, S. Reconstruction of Baseline Time-Trace under Changing Environmental and Operational Conditions. *Smart Mater Struct* 2016, *25*, doi:10.1088/0964-1726/25/3/035018.
104. Salmanpour, M.S.; Khodaei, Z.S.; Ferri Aliabadi, M.H. Impact Damage Localisation with Piezoelectric Sensors under Operational and Environmental Conditions. *Sensors (Switzerland)* 2017, *17*, doi:10.3390/s17051178.
105. Simon, J.; Moll, J.; Krozer, V. Trend Decomposition for Temperature Compensation in a Radar-Based Structural Health Monitoring System of Wind Turbine Blades. *Sensors* 2024, *24*, 800, doi:10.3390/s24030800.
106. Goodwin, S.; Olazabal, M. Revisiting the Challenges to Monitoring, Evaluation, Reporting, and Learning for Climate Adaptation. *Environ Sci Policy* 2025, *172*, 104199, doi:10.1016/j.envsci.2025.104199.

107. Putkis, O.; Croxford, A.J. Continuous Baseline Growth and Monitoring for Guided Wave SHM. *Smart Mater Struct* 2013, 22, doi:10.1088/0964-1726/22/5/055029.
108. Abbas, S.; Li, F.; Zhu, Y.; Tu, X. Experimental Investigation of Impact of Environmental Temperature and Optimal Baseline for Thermal Attenuation in Structural Health Monitoring Based on Ultrasonic Guided Waves. *Wave Motion* 2020, 93, doi:10.1016/j.wavemoti.2019.102474.
109. Yue, N.; Aliabadi, M.H. A Scalable Data-Driven Approach to Temperature Baseline Reconstruction for Guided Wave Structural Health Monitoring of Anisotropic Carbon-Fibre-Reinforced Polymer Structures. *Struct Health Monit* 2020, 19, doi:10.1177/1475921719887109.
110. Fendzi, C.; Rébillat, M.; Mechbal, N.; Guskov, M.; Coffignal, G. A Data-Driven Temperature Compensation Approach for Structural Health Monitoring Using Lamb Waves. *Struct Health Monit* 2016, 15, doi:10.1177/1475921716650997.
111. Amer, A.; Roy, S.; Kopsaftopoulos, F. Probabilistic Shm under Varying Loads via the Integration of Gaussian Process Regression and Physics-Based Guided-Wave Propagation Models. In Proceedings of the AIAA Scitech 2021 Forum; 2021.
112. Qu, K.; Logan, A.; Miller, E.; Garcia Cava, D. Multi-Phase Adaptive Methodology for Mitigating Environmental and Operational Variability in Slowly Changing Time-Variant Engineering Structures. *Mech Syst Signal Process* 2025, 229, 112494, doi:10.1016/j.ymsp.2025.112494.
113. Font-Moré, J.; Garcia Cava, D.; Perez, M.A.; Avendano-Valencia, L.D. A Latent Variable Approach for Mitigation of Environmental and Operational Variability in Vibrationbased SHM – A Linear Approach. *Research and Review Journal of Nondestructive Testing* 2024, 2, doi:10.58286/30512.
114. Roberts, C.; Avendaño-Valencia, L.D.; García Cava, D. Robust Mitigation of EOVs Using Multivariate Nonlinear Regression within a Vibration-Based SHM Methodology. *Mech Syst Signal Process* 2024, 208, doi:10.1016/j.ymsp.2023.111028.
115. Huynh, T.-C.; Koh, C.G.; Low, Y.M. Reference-Free Impedance-Based Crack Detection for Asymmetric Structures: Application to Tidal Turbine Blades. *J Sound Vib* 2025, 618, 119304, doi:10.1016/j.jsv.2025.119304.
116. Zhao, H.; Sui, D.; Wang, Y.; Ma, L.; Wang, L. Privacy-Preserving Federated Learning Framework for Multi-Source Electronic Health Records Prognosis Prediction. *Sensors* 2025, 25, 2374, doi:10.3390/s25082374.
117. Kashyap, P.; Shivgan, K.; Patil, S.; Raja, B.R.; Mahajan, S.; Banerjee, S.; Tallur, S. Unsupervised Deep Learning Framework for Temperature-Compensated Damage Assessment Using Ultrasonic Guided Waves on Edge Device. *Sci Rep* 2024, 14, doi:10.1038/s41598-024-54418-w.
118. Zhu, H.; Khodaei, Z.S.; Aliabadi, F.M.H. Baseline-Free Damage Detection and Localization on Complex Composite Structures Using Unsupervised Shapelets and Shift-Invariant Dictionary Learning. *Mech Syst Signal Process* 2025, 224, 112035, doi:10.1016/j.ymsp.2024.112035.
119. Sharma, B.N.; Kapuria, S.; Arockiarajan, A. Time Reversibility of Lamb Waves in Thin Plates with Surface-Bonded Piezoelectric Transducers Is Temperature Invariant at the Best Reconstruction Frequency. *Struct Health Monit* 2021, 20, doi:10.1177/1475921720965122.

120. Zhu, H.; Khodaei, Z.S.; Aliabadi, M.H. Investigation of Baseline-Free Techniques for Damage Localisation on Anisotropic Composite Structures. In *Proceedings of the Procedia Structural Integrity*; 2024; Vol. 52.
121. Anton, S.R.; Inman, D.J.; Park, G. Reference-Free Damage Detection Using Instantaneous Baseline Measurements. *AIAA Journal* 2009, 47, doi:10.2514/1.43252.
122. Mesnil, O.; Recoquillay, A.; Fisher, C.; Serey, V.; Sharma, S.; d'Almeida, O. Self-Referenced Robust Guided Wave Based Defect Detection: Application to Woven Composite Parts of Complex Shape. *Mech Syst Signal Process* 2023, 188, doi:10.1016/j.ymsp.2022.109948.
123. Zhu, H.; Sharif Khodaei, Z.; M. H. Aliabadi, F. Appraisal of Linear Baseline-Free Techniques for Guided Wave Based Structural Health Monitoring. *Ultrasonics* 2024, 144, 107445, doi:10.1016/j.ultras.2024.107445.
124. Shi, W.; Li, J.; Zhao, B.; Tan, J. A Baseline-Free Stress Monitoring Strategy Based on Acoustoelastic Lamb Waves Using PWAS Array. *Struct Health Monit* 2023, 22, doi:10.1177/14759217221115849.
125. Lee, H.; Yang, J.; Sohn, H. Baseline-Free Pipeline Monitoring Using Optical Fiber-Guided Laser Ultrasonics. *Struct Health Monit* 2012, 11, 684–695, doi:10.1177/1475921712455682.
126. Lim, H.J.; Sohn, H.; Desimio, M.P.; Brown, K. Reference-Free Fatigue Crack Detection Using Nonlinear Ultrasonic Modulation under Various Temperature and Loading Conditions. *Mech Syst Signal Process* 2014, 45, doi:10.1016/j.ymsp.2013.12.001.
127. Qiu, J.; Li, F.; Abbas, S.; Zhu, Y. A Baseline-Free Damage Detection Approach Based on Distance Compensation of Guided Waves. *Journal of Low Frequency Noise Vibration and Active Control* 2019, 38, doi:10.1177/1461348418813699.
128. Loshelder, A.E.; He, J.; Aktharuzzaman, M.; Harb, M.S.; Rao, J. Apex-Shifted Radon Transform for Baseline-Subtraction-Free (BSF) Damage Scattered Wave Extraction. *Struct Health Monit* 2023, 22, doi:10.1177/14759217231156364.
129. Sohn, H.; Woo Park, H.; Law, K.; Farrar, C.R. Combination of a Time Reversal Process and a Consecutive Outlier Analysis for Baseline-Free Damage Diagnosis. *J Intell Mater Syst Struct* 2006, doi:10.1177/1045389x06066291.
130. Miele, S.; Karve, P.; Mahadevan, S. Multi-Fidelity Physics-Informed Machine Learning for Probabilistic Damage Diagnosis. *Reliab Eng Syst Saf* 2023, 235, doi:10.1016/j.res.2023.109243.
131. Koune, I.; Cicirello, A. Disentangled Representation Learning with Physics-Informed Variational Autoencoder for Structural Health Monitoring. *e-Journal of Nondestructive Testing* 2024, 29, doi:10.58286/29862.
132. Amer, A.; Kopsaftopoulos, F. Gaussian Process Regression for Active Sensing Probabilistic Structural Health Monitoring: Experimental Assessment across Multiple Damage and Loading Scenarios. *Struct Health Monit* 2023, 22, doi:10.1177/14759217221098715.
133. El Mountassir, M.; Yaacoubi, S.; Mourot, G.; Maquin, D. An Adaptive PCA-Based Method for More Reliable Ultrasonic Guided Waves SHM: Data-Driven Modeling and Experimental Validation in High Attenuating Medium. *Struct Control Health Monit* 2021, 28, doi:10.1002/stc.2634.

134. Yue, N.; Khodaei, Z.S.; Aliabadi, M.H. Damage Detection in Large Composite Stiffened Panels Based on a Novel SHM Building Block Philosophy. *Smart Mater Struct* 2021, *30*, doi:10.1088/1361-665X/abe4b4.
135. Lee, J.S.; Park, G.; Kim, C.G.; Farrar, C.R. Use of Relative Baseline Features of Guided Waves for in Situ Structural Health Monitoring. *J Intell Mater Syst Struct* 2011, *22*, doi:10.1177/1045389X10395643.
136. Salmanpour, M.S.; Sharif Khodaei, Z.; Aliabadi, M.H. Instantaneous Baseline Damage Localization Using Sensor Mapping. *IEEE Sens J* 2017, *17*, doi:10.1109/JSEN.2016.2629279.
137. Gardner, P.; Bull, L.A.; Gosliga, J.; Dervilis, N.; Cross, E.J.; Papatheou, E.; Worden, K. Population-Based Structural Health Monitoring. In: 2022; pp. 413–435.
138. Bull, L.A.; Gardner, P.A.; Gosliga, J.; Rogers, T.J.; Dervilis, N.; Cross, E.J.; Papatheou, E.; Maguire, A.E.; Campos, C.; Worden, K. Foundations of Population-Based SHM, Part I: Homogeneous Populations and Forms. *Mech Syst Signal Process* 2021, *148*, 107141, doi:10.1016/j.ymsp.2020.107141.
139. Bull, L.A.; Gardner, P.A.; Dervilis, N.; Papatheou, E.; Haywood-Alexander, M.; Mills, R.S.; Worden, K. On the Transfer of Damage Detectors between Structures: An Experimental Case Study. *J Sound Vib* 2021, *501*, 116072, doi:10.1016/j.jsv.2021.116072.
140. Zhou, M.; Lai, Z. Structural Damage Classification under Varying Environmental Conditions and Unknown Classes via Open Set Domain Adaptation. *Mech Syst Signal Process* 2024, *218*, 111561, doi:10.1016/j.ymsp.2024.111561.
141. Yang, J.; Gan, Z.; Wang, T.; Xie, J.; Pan, T.; He, J.; Wang, Z. Multi-Source Dynamic Adaptive Domain Generalization Network for Crack Detection under Unknown Temperature Environment. *Measurement* 2025, *240*, 115588, doi:10.1016/j.measurement.2024.115588.
142. Michau, G.; Fink, O. Unsupervised Transfer Learning for Anomaly Detection: Application to Complementary Operating Condition Transfer. *Knowl Based Syst* 2021, *216*, 106816, doi:10.1016/j.knosys.2021.106816.
143. Sawant, S.; Sethi, A.; Banerjee, S.; Tallur, S. Unsupervised Learning Framework for Temperature Compensated Damage Identification and Localization in Ultrasonic Guided Wave SHM with Transfer Learning. *Ultrasonics* 2023, *130*, 106931, doi:10.1016/j.ultras.2023.106931.
144. Wang, L.; Liu, G.; Zhang, C.; Yang, Y.; Qiu, J. FEM Simulation-Based Adversarial Domain Adaptation for Fatigue Crack Detection Using Lamb Wave. *Sensors* 2023, *23*, 1943, doi:10.3390/s23041943.
145. Silva, S. da; Yano, M.O.; Gonzalez-Bueno, C.G. Transfer Component Analysis for Compensation of Temperature Effects on the Impedance-Based Structural Health Monitoring. *J Nondestr Eval* 2021, *40*, 64, doi:10.1007/s10921-021-00794-6.
146. Liu, C.; Xu, X.; Wu, J.; Zhu, H.; Wang, C. Deep Transfer Learning-Based Damage Detection of Composite Structures by Fusing Monitoring Data with Physical Mechanism. *Eng Appl Artif Intell* 2023, *123*, 106245, doi:10.1016/j.engappai.2023.106245.
147. Liu, X.; Shang, C.; Wang, W.; Wu, M.; Bao, H. An Enhanced Manifold Alignment Method for Fault Diagnosis of Liquid-Cooled Plate Based on

- Manifold Distance Features. *Measurement* 2025, 242, 116303, doi:10.1016/j.measurement.2024.116303.
148. Luckey, D.; Fritz, H.; Legatiuk, D.; Peralta Abadía, J.J.; Walther, C.; Smarsly, K. Explainable Artificial Intelligence to Advance Structural Health Monitoring. In; 2022; pp. 331–346.
 149. Rezazadeh, N.; Perfetto, D.; Caputo, F.; De Luca, A. Enhancing Air Compressor Fault Diagnosis: A Comparative Study of GPT-2 and Traditional Machine Learning Models. *Macromol Symp* 2025, 414, doi:10.1002/masy.70057.
 150. Kim, K.; Kim, Y.S. Vibration Spectrogram Analysis for Bearing Fault Diagnosis Based on Grad-Cam for Feature Selection and Statistical Approach. *Journal of Mechanical Science and Technology* 2024, 38, 5885–5898, doi:10.1007/s12206-024-1010-3.
 151. Brito, L.C.; Susto, G.A.; Brito, J.N.; Duarte, M.A.V. Fault Diagnosis Using EXplainable AI: A Transfer Learning-Based Approach for Rotating Machinery Exploiting Augmented Synthetic Data. *Expert Syst Appl* 2023, 232, 120860, doi:10.1016/j.eswa.2023.120860.
 152. Yan, T.; Xing, X.; Xia, T.; Wang, D. Relation between Fault Characteristic Frequencies and Local Interpretability Shapley Additive Explanations for Continuous Machine Health Monitoring. *Eng Appl Artif Intell* 2024, 136, 109046, doi:10.1016/j.engappai.2024.109046.
 153. Yao, Q.; Chen, B.; Hu, A.; Zhen, D.; Xiang, L. A Novel Diffusion Model with Shapley Value Analysis for Anomaly Detection and Identification of Wind Turbine. *Expert Syst Appl* 2025, 284, 127925, doi:10.1016/j.eswa.2025.127925.
 154. Hanchate, A.; Bukkapatnam, S.T.S.; Lee, K.H.; Srivastava, A.; Kumara, S. Explainable AI (XAI)-Driven Vibration Sensing Scheme for Surface Quality Monitoring in a Smart Surface Grinding Process. *J Manuf Process* 2023, 99, 184–194, doi:10.1016/j.jmapro.2023.05.016.
 155. Gawde, S.; Patil, S.; Kumar, S.; Kamat, P.; Kotecha, K. An Explainable Predictive Maintenance Strategy for Multi-Fault Diagnosis of Rotating Machines Using Multi-Sensor Data Fusion. *Decision Analytics Journal* 2024, 10, 100425, doi:10.1016/j.dajour.2024.100425.
 156. Chen, G.; Dong, G. Temporal Logic Inference for Interpretable Fault Diagnosis of Bearings via Sparse and Structured Neural Attention. *ISA Trans* 2025, 158, 256–271, doi:10.1016/j.isatra.2025.01.013.
 157. Li, Y.; Zhou, Z.; Sun, C.; Chen, X.; Yan, R. Variational Attention-Based Interpretable Transformer Network for Rotary Machine Fault Diagnosis. *IEEE Trans Neural Netw Learn Syst* 2024, 35, 6180–6193, doi:10.1109/TNNLS.2022.3202234.
 158. Lai, Y.; Li, R.; Ye, Z.; He, Y. A New Fault Diagnosis Method for Reciprocating Piston Pump Based on Feature Fusion of CNN and Transformer Encoder. *Sci Prog* 2025, 108, doi:10.1177/00368504251330003.
 159. Luo, Z.; Pan, S.; Dong, X.; Zhang, X. Interpretable Quadratic Convolutional Residual Neural Network for Bearing Fault Diagnosis. *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 2025, 47, 158, doi:10.1007/s40430-025-05457-5.
 160. Sabour, S.; Frosst, N.; Hinton, G.E. Dynamic Routing between Capsules. In Proceedings of the Advances in Neural Information Processing Systems; 2017; Vol. 2017-December.

161. Dai, W.; Yang, Q.; Xue, G.R.; Yu, Y. Boosting for Transfer Learning. In Proceedings of the ACM International Conference Proceeding Series; 2007; Vol. 227.
162. He, H.; Khoshelham, K.; Fraser, C. A Multiclass TrAdaBoost Transfer Learning Algorithm for the Classification of Mobile Lidar Data. *ISPRS Journal of Photogrammetry and Remote Sensing* 2020, 166, doi:10.1016/j.isprsjprs.2020.05.010.
163. Breiman, L. Random Forests. *Mach Learn* 2001, 45, doi:10.1023/A:1010933404324.
164. Srivastava, R.; Kumar, S.; Kumar, B. Classification Model of Machine Learning for Medical Data Analysis. In *Statistical Modeling in Machine Learning: Concepts and Applications*; 2022.
165. 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, doi:10.3390/s23239327.
166. Abdel-Basset, M.; Moustafa, N.; Hawash, H.; Tari, Z. *Responsible Graph Neural Networks*; 2023;
167. 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); IEEE, October 2021; pp. 8968–8979.
168. Ou, Y.; Tatsis, K.E.; Dertimanis, V.K.; Spiridonakos, M.D.; Chatzi, E.N. Vibration-Based Monitoring of a Small-Scale Wind Turbine Blade under Varying Climate Conditions. Part I: An Experimental Benchmark. *Struct Control Health Monit* 2021, 28, doi:10.1002/stc.2660.
169. Ferreira, L. de P.S.; Teloli, R. de 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, doi:10.1016/j.ymsp.2023.111011.
170. 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, doi:10.1177/1045389X20963171.
171. Almeida, E.F. de; 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. *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 2024, 46, 717, doi:10.1007/s40430-024-05251-9.
172. Kamariotis, A.; Vlachas, K.; Ntertimanis, V.; Koune, I.; Cicirello, A.; Chatzi, E. On the Consistent Classification and Treatment of Uncertainties in Structural Health Monitoring Applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering* 2025, 11, doi:10.1115/1.4067140.
173. Lee, U.; Shin, J. A Frequency Response Function-Based Structural Damage Identification Method. *Comput Struct* 2002, 80, 117–132, doi:10.1016/S0045-7949(01)00170-5.
174. Chen, J.; Zhu, C.; Pu, Y.; Rui, Y.; Liu, B.; Apel, D.B. A Systematic Review of Coda Wave Interferometry Technique for Evaluating Rock Behavior

- Properties: From Single to Multiple Perturbations. *Earth Energy Science* 2025, *1*, 180–192, doi:10.1016/j.ees.2025.03.002.
175. Senin, P. Dynamic Time Warping Algorithm Review. *Science (1979)* 2008, 2007.
 176. 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.
 177. Pan, S.J.; Tsang, I.W.; Kwok, J.T.; Yang, Q. Domain Adaptation via Transfer Component Analysis. *IEEE Trans Neural Netw* 2011, 22, doi:10.1109/TNN.2010.2091281.
 178. Uguroglu, S.; Carbonell, J. Feature Selection for Transfer Learning. In; 2011; pp. 430–442.
 179. Sun, B.; Feng, J.; Saenko, K. Return of Frustratingly Easy Domain Adaptation. In Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI 2016; 2016.
 180. 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); 2016; Vol. 9915 LNCS.
 181. III, H.D. Frustratingly Easy Domain Adaptation. *CoRR* 2009, *abs/0907.1815*.
 182. 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; 2013.
 183. 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; 2022; Vol. 2022-June.
 184. Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; Lempitsky, V. Domain-Adversarial Training of Neural Networks. *Journal of Machine Learning Research* 2016, 17.
 185. 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, doi:10.1109/JSEN.2022.3174396.
 186. Łuczak, D. Mechanical Vibrations Analysis in Direct Drive Using CWT with Complex Morlet Wavelet. *Power Electronics and Drives* 2023, 8, 65–73, doi:10.2478/pead-2023-0005.