

Asynchronous Multimodal Sensor Data Fusion for Label-Scarce Health Indicator Mining in Composite Structure Monitoring

Yang Zhou¹, Chao Zhang², Xiao Xiao³, *

1 School of Artificial Intelligence and Low-Altitude Technology, Hefei College of Finance & Economics, Hefei, China, 230601

2 Institute of Intelligent Manufacturing, Hefei College of Finance & Economics, Hefei, China, 230601

3 School of Smart Finance and Business, Hefei College of Finance & Economics, Hefei, China, 230601

** Correspondence: 1164307064@qq.com*

Abstract

Structural health monitoring (SHM) of composite aerospace panels requires comprehensive health indicators (HIs) that capture multiscale damage evolution under complex fatigue loading. However, constructing reliable HIs is hampered by the asynchronous acquisition rates of passive and active sensing modalities, the scarcity of ground-truth labels, and the heterogeneous information content of complementary sensor streams. This study proposes a multi-level asynchronous multimodal sensor data fusion framework that integrates passive acoustic emission (AE) and active guided wave (GW) sensing to mine high-quality HIs for T-stiffener composite panels without labeled failure data. Three intra-modality AE pipelines and one GW pipeline are developed under an inductive semi-supervised learning paradigm, using physics-informed proxy labels and criteria-driven regularization to enforce monotonicity, prognosability, and trendability. A causal zero-order-hold synchronization scheme aligns the dense AE timeline with sparse GW acquisitions. Late-fusion regression models—including Gaussian process regression, gradient boosting, support vector machines, multilayer perceptrons, and long short-term memory networks—are benchmarked for inter-modality HI fusion. Leave-one-out cross-validation on 12 AE units and 5 GW units demonstrates that the fused HIs consistently exceed a Fitness score of 90%, with the LSTM-based fusion achieving up to $97\% \pm 2\%$ cohort-level performance. The proposed framework offers a scalable and deployable blueprint for next-generation prognostics in safety-critical composite structures.

Keywords: Asynchronous sensor fusion; Health indicator mining; Semi-supervised learning; Acoustic emission; Guided waves; Composite structures; Structural health monitoring

Article History

Received: March 1, 2026

Accepted: May 17, 2026

Published: June 30, 2025

Asynchronous Multimodal Sensor Data Fusion for Label-Scarce Health Indicator Mining in Composite Structure Monitoring

1. Introduction

Composite materials, particularly carbon fiber-reinforced polymer (CFRP) laminates, have become indispensable in aerospace structural applications due to their superior specific strength and stiffness properties (Soutis, 2005). However, the complex and often unpredictable damage mechanisms in composite structures under fatigue loading—including matrix cracking, delamination, fiber breakage, and interfacial debonding—present formidable challenges for structural health monitoring (SHM) and prognostics and health management (PHM) (Worden & Farrar, 2007; Farrar & Worden, 2012). Reliable health indicators (HIs) that track structural degradation from the initial healthy state to failure are essential inputs for remaining useful life (RUL) estimation and maintenance decision-making in safety-critical aerospace operations (Pecht, 2008; Jardine et al., 2006).

Two sensing modalities dominate SHM of composite aerospace panels: passive acoustic emission (AE) and active guided waves (GW). AE captures transient elastic waves released by microscale damage events (Grosse & Ohtsu, 2008; Hamstad, 1994), providing event-rich temporal information about damage initiation and propagation. GW methods actively interrogate the structure at discrete intervals, offering spatially resolved snapshots of the current damage state (Rose, 2014; Giurgiutiu, 2008). While both modalities originate from the same underlying wave physics (Lamb, 1917; Alleyne & Cawley, 1992), they encode fundamentally complementary information: AE excels in temporal sensitivity to ongoing damage processes, whereas GW provides quantitative estimates of damage location and severity at each interrogation (Su et al., 2006; Lissenden, 2021). Exploiting this complementarity through multimodal fusion is therefore a natural strategy for constructing more robust and informative HIs (Lu, 2019; Zhang & Lu, 2021).

Despite the intuitive appeal of multimodal fusion, several significant challenges have hindered its practical implementation. First, AE and GW sensors operate at fundamentally different temporal rates: AE is continuously acquired and windowed at sub-hundred-cycle intervals, whereas GW interrogations are typically spaced thousands of cycles apart. This asynchrony creates a temporal alignment problem that, if not properly addressed, introduces information leakage or spurious correlations in downstream learning (Lee et al., 2014; Saxena & Goebel, 2008). Second, ground-truth HI labels are unavailable for composite structures

under realistic fatigue loading, as no universally accepted physical surrogate of global structural health exists for such systems (Mosallam et al., 2016). Traditional supervised learning therefore cannot be directly applied, and alternative learning paradigms are needed (Van Engelen & Hoos, 2020; Chapelle et al., 2009). Third, the heterogeneous feature spaces of AE (hundreds of time- and frequency-domain descriptors) and GW (high-dimensional spatio-temporal signal tensors) demand modality-specific preprocessing and representation learning pipelines (Virkler et al., 1979; Mujica et al., 2010).

To address these challenges, this study proposes a multi-level asynchronous multimodal sensor data fusion framework for label-scarce HI mining in composite structure monitoring. The principal contributions of this work are as follows: (1) three complementary intra-modality AE pipelines—covering semi-supervised LSTM, PCA-guided two-stage machine learning, and CEEMDAN-based deep ensemble learning—are systematically developed and benchmarked; (2) a physics-informed semi-supervised learning paradigm using nonlinear proxy targets and criteria-based regularization is formulated to enforce monotonicity, prognosability, and trendability without ground-truth labels; (3) a causal zero-order-hold synchronization scheme reconciles the asynchronous AE and GW timelines without introducing future information leakage; and (4) a comprehensive inter-modality late-fusion study benchmarks eight regression architectures, identifying LSTM as the most effective fusion model. The proposed framework is validated on a publicly available dataset of run-to-failure composite T-stiffener specimens under compression-compression fatigue, achieving fused HI Fitness scores consistently above 90% under strict leave-one-out cross-validation.

The remainder of this paper is organized as follows. Section 2 surveys related work on HI construction and sensor fusion for SHM. Section 3 describes the research design and methodology, including the semi-supervised learning framework and fusion architectures. Section 4 presents the experimental dataset and the analytical framework with detailed algorithmic specifications. Section 5 reports quantitative results. Section 6 discusses the findings in the context of the research questions. Section 7 addresses theoretical and practical implications, followed by limitations and future research directions in Section 8. Section 9 concludes the paper.

2. Literature Review

The construction of reliable HIs from SHM sensor data has evolved considerably over the past two decades. Early approaches relied on physics-based HIs (PHIs) derived from well-understood material properties—such as bearing RMS and relative RMS—for which physical interpretation was straightforward (Li et al., 2012; Yan et al., 2004). For composite structures under fatigue, however, no universally accepted PHI exists, motivating the

development of virtual health indicators (VHIs) that are engineered to satisfy prognostic objectives rather than reflect a specific physical quantity (Nuhaa et al., 2022; Si et al., 2011).

Among VHI construction methods, self-organizing map (SOM)-based approaches were among the earliest data-driven contributions (Jardine et al., 2006; Staszewski, 2002). Recurrent neural network-based HIs subsequently demonstrated improvements in monotonicity and prognosability over SOM baselines for rolling element bearings (Heimes, 2008; Zhao et al., 2017). Convolutional and convolutional-recurrent architectures later extended these approaches to exploit both spatial and temporal information from vibration data (Zhu et al., 2018; Liu et al., 2019a; Zhao et al., 2019). Autoencoder variants—including contractive, variational, and sparse autoencoders—have been widely applied for unsupervised feature compression in rotating machinery prognostics (Bengio et al., 2013; Zhang et al., 2019; Liu et al., 2019b). Despite their success in rotating machinery, these methods face significant obstacles when applied to composite structures, where damage is distributed, multimodal, and quasi-statically evolving (Talreja & Singh, 2012; Haghani et al., 2012).

Sensor fusion for SHM has been approached at multiple levels of abstraction. Low-level (data) fusion directly combines raw sensor signals or processed representations (Staszewski et al., 2004; Boller et al., 2009). Feature-level fusion aggregates descriptors extracted from individual modalities (Zhao et al., 2020; Migot et al., 2022). Decision-level or late fusion combines the outputs of independently trained unimodal models, offering greater modularity and robustness to modality-specific failures (Kim & Bhattacharyya, 2012). In the context of PHM, late fusion is particularly attractive because unimodal HI models can be trained and validated independently before their outputs are combined by a lightweight meta-learner (Vachtsevanos et al., 2006). Bayesian fusion frameworks have also been explored for combining AE and GW measurements in damage classification (Lu, 2017; Fromme et al., 2006), though their extension to continuous HI construction under label scarcity remains underexplored.

Semi-supervised learning (SSL) has emerged as a powerful paradigm for exploiting unlabeled data when ground-truth labels are costly or unavailable (Van Engelen & Hoos, 2020; Chapelle et al., 2009; Zhu, 2005). In PHM, SSL approaches have been applied to fault detection (Zhang et al., 2021a), degradation modeling (Listou Ellefsen et al., 2019), and HI construction (Nitta et al., 2022). Physics-informed constraints—such as monotonicity regularizers, boundary conditions, and degradation priors—have been incorporated into SSL frameworks to improve physical plausibility and sample efficiency (Nascimento & Viana, 2019; Raissi et al., 2019; Karniadakis et al., 2021). However, the simultaneous integration of multiple sensing modalities with asynchronous acquisition rates within a physics-guided SSL paradigm has not been previously addressed for composite aerospace structures, representing the primary gap that this work targets.

Ensemble learning (EL) has been widely adopted in PHM to improve stability and reduce variance in data-driven models (Dietterich, 2000; Polikar, 2006). Simple and weighted averaging ensembles, random forests, boosting methods, and stacked generalization have all demonstrated utility in prognostic applications (Seera & Lim, 2014; Zhang et al., 2020). In the context of HI construction, ensemble strategies that combine predictions from diverse random initializations convert stochastic variation into beneficial diversity, yielding smoother and more monotonic HI trajectories (Heimes, 2008; Liu et al., 2020). The integration of ensemble strategies at both the intra-modality and inter-modality fusion levels is a novel aspect of the present work that systematically exploits diversity across architectures, random seeds, and sensor configurations.

3. Research Design and Methodology

The proposed framework adopts a hierarchical multi-level fusion architecture that integrates passive AE and active GW modalities through sequential intra-modality and inter-modality fusion stages. Figure 1 illustrates the complete system architecture. At the intra-modality level, three AE-specific pipelines and one GW-specific pipeline independently construct unimodal HIs using the inductive semi-supervised learning paradigm. At the inter-modality level, the synchronized unimodal HIs serve as inputs to a late-fusion regression meta-learner that produces a single fused HI. The overarching design principle is that of causal deployability: no future information is used at any stage, and all models are evaluated on strictly held-out test units.




Figure 1. Overall architecture of the proposed asynchronous multimodal sensor data fusion framework for HI mining in composite structures.

The semi-supervised learning (SSL) paradigm, central to all unimodal pipelines, addresses the absence of ground-truth HI labels by substituting them with physics-informed proxy targets derived from known degradation physics. Specifically, a quadratic time-to-failure profile is used as the proxy: the HI target for unit j at cycle t_i is defined as $T_j(t_i) = (t_i / t_{\{EOL,j\}})^2$, which encodes the expectation of accelerating degradation as the end of life approaches—a behavior consistent with observations in composite fatigue (Talreja & Singh, 2012; Christensen, 2013). This proxy target is combined with a criteria-based regularization term that penalizes deviations from monotonicity, prognosability, and trendability (Mo, Pr, Tr) during model training.

Three evaluation criteria are employed throughout this work. Monotonicity (Mo) is computed using the modified Mann-Kendall (MMK) statistic, which accounts for non-uniform time gaps and is more robust to noise than the standard Kendall-tau measure (Coble & Hines, 2011). Prognosability (Pr) quantifies the concentration of end-of-life HI values across units, reflecting the consistency of failure signatures across nominally identical specimens (Saxena et al., 2008). Trendability (Tr) measures the minimum absolute pairwise correlation among unit HI trajectories, capturing the degree to which specimens share a common degradation pattern. The combined cohort-level Fitness is defined as the sum $Mo + Pr + Tr$, with a theoretical maximum of 3.0. Test-only variants of Mo and Pr are separately computed on the held-out unit to assess out-of-cohort generalizability.

Three AE frameworks exploit distinct trade-offs between supervision, dimensionality, and model capacity. AE-Framework 1 (FFT-SSLSTM) employs a compact sequence model combining fully connected and long short-term memory (LSTM) layers (Hochreiter & Schmidhuber, 1997), trained with proxy targets on 201 time- and frequency-domain features extracted via fast Fourier transform (FFT). AE-Framework 2 (FFT-PCA-2S-ML-PBO) separates HI learning into a time-independent model (TIM) for spatial feature compression using PCA and a time-dependent model (TDM) for temporal refinement via LSTM, with hyperparameter tuning by physics-based Bayesian optimization (PBO). AE-Framework 3 (CEEMDAN-SSEDL) replaces FFT with complete ensemble empirical mode decomposition with adaptive noise (CEEMDAN) to extract 504 intrinsic mode function (IMF)-based features, followed by ensemble deep learning blenders (Torres et al., 2011; Wu & Huang, 2009).

The GW framework (HT-SSCNN) processes GW signals with the Hilbert transform (HT) to extract amplitude envelopes from all actuator-sensor paths at six excitation frequencies (50 to 250 kHz). The resulting 3D envelope tensors are fed to a semi-supervised convolutional neural network (SSCNN) trained under the same proxy-label regime. A weighted averaging

ensemble (WAE-Fitness) selects the best ensemble configuration based on cohort-level Fitness computed on non-test units. The GW framework leverages the spatial richness of multi-path, multi-frequency interrogations to capture damage localization information unavailable from AE alone (Qing et al., 2007; Rizzo et al., 2007).

4. Data and Analytical Framework

The experimental dataset used in this study originates from the European H2020 ReMAP project, which conducted fatigue tests on CFRP T-stiffener composite panels representative of aeronautical structural components. Each specimen featured an IM7/8552 carbon fiber-reinforced epoxy skin with a single bonded T-shaped stiffener. Specimens were subjected to compression-compression (C-C) fatigue loading at 2 Hz with a load ratio $R = 10$, using a hydraulic MTS testing machine. To introduce realistic damage variability, controlled 10 J impacts were applied at different panel locations, and some specimens contained pre-inserted PTFE disbond defects of varying dimensions. The complete dataset properties are summarized in Table 1.

Table 1. Experimental specimen information and monitoring coverage (ReMAP dataset).

Parameter	Campaign 2019 (Units 1-5)	Campaign 2020 (Units 6-12)	AE Monitoring	GW Monitoring
Specimen material	CFRP IM7/8552	CFRP IM7/8552	12 units	5 units
Load range (kN)	-65 to -6.5	-65 to -6.5	Continuous	Every 5000 cycles
Loading frequency	2 Hz	2 Hz	—	—
Impact energy	10 J (various locations)	10 J (various locations)	—	—
Disbond size (mm)	None	15×20 to 30×30	—	—
Lifetime (cycles)	133K-438K	49K-756K	Full life	Full life

Figure 2 illustrates the composite T-stiffener panel geometry and the sensor network layout. Four Vallen Systeme VS900-M broadband AE sensors (frequency range 100–900 kHz) were bonded to each panel and operated continuously during fatigue loading. Eight piezoelectric transducers (PZTs) were distributed across the panel surface for GW acquisition, with each PZT acting sequentially as the actuator while the remaining seven served as receivers, yielding 56 unique actuator-sensor paths. GW signals were excited at six frequencies (50, 100, 125, 150, 200, and 250 kHz) at each acquisition step.

Composite T-Stiffener Panel with Sensor Network Layout




Figure 2. Composite T-stiffener CFRP panel with AE sensor (filled circles) and GW PZT transducer (squares) network layout.

A central challenge in the fusion of AE and GW is the inherent temporal asynchrony between the two modalities. AE data are windowed into time segments of 500 cycles (Frameworks 1 and 2) or 1000 cycles (Framework 3), producing dense HI time series with hundreds to thousands of observations per specimen. In contrast, GW acquisitions occur every 5000 cycles, yielding only 10–100 GW measurements per specimen depending on total lifetime. Figure 3 illustrates this rate mismatch and the zero-order-hold (ZOH) synchronization scheme employed to align the two timelines.




Figure 3. Temporal asynchrony between AE (dense) and GW (sparse) health indicator streams, and their alignment via zero-order-hold (ZOH) synchronization.

The ZOH synchronization assigns to each AE window at cycle c_k the GW HI value from the most recent preceding GW acquisition (the largest $c_g \leq c_k$). This nearest-preceding assignment preserves causality: no future GW information is introduced into the AE timeline at any window. The resulting synchronized time series provides aligned (AE-HI, GW-HI) pairs at the AE window grid, which are subsequently used as inputs to the inter-modality fusion regressors. It is important to note that ZOH introduces a systematic interpolation artifact near abrupt structural changes; future work may explore more sophisticated causal interpolation schemes such as Gaussian process-based temporal alignment (Rasmussen & Williams, 2006).

The data partitioning strategy follows a strict leave-one-out cross-validation (LOOCV) protocol at the unit level. For AE models, each fold holds out one unit for testing and one for validation, with the remaining 10 units used for training. For GW models, each fold holds out one unit for testing, one for validation, and trains on the remaining three. For inter-modality fusion, the three units instrumented with both AE and GW sensors define three folds; one unit is held out for testing in each fold, while the meta-learner is trained and validated on the other two. This conservative partition ensures that the test unit is fully excluded from all stages of model training, including unimodal pretraining, thereby providing an unbiased estimate of generalization performance.

Hyperparameter optimization employs Bayesian optimization (BO) for all model types (Snoek et al., 2012; Frazier, 2018). For AE-Framework 2, a physics-based BO (PBO) variant

combines an explicit prognostic criteria loss with an implicit regression loss, enabling the optimizer to explore architectures that satisfy degradation physics while fitting the proxy targets. For the inter-modality fusion models, BO is conducted under 5-fold cross-validation on the two training units, with mean squared error (MSE) as the objective. All ensemble strategies use weights computed exclusively on non-test units and applied unchanged to the held-out test unit.

5. Results

Figure 4 illustrates the health indicator evaluation criteria framework, showing representative HI trajectories demonstrating monotonicity, prognosability, and trendability properties. The visual clarity of these criteria guides the model development and selection process throughout the study.




Figure 4. Health indicator evaluation criteria visualization: (a) Monotonicity, (b) Prognosability, and (c) Trendability for representative multi-unit HI trajectories.

To comprehensively evaluate the proposed fusion framework, all three units instrumented with both AE and GW sensors were used in LOOCV, with each fold designating one unit as the test set. Table 2 presents the mean and standard deviation of HI evaluation criteria across the three folds for all fusion architectures and AE-GW framework combinations. The values are computed separately for cohort-level metrics (All) and test-unit-only metrics (Test).

Table 2. Performance metrics (mean $\pm$ SD over 3 folds) for inter-modality fusion models combined with each AE framework. Best Fitness-Test per AE framework shown in bold.

AE Framework	Fusion Model	Mo-All	Pr-All	Tr	Fitness-Test
FFT-SSLSTM (Framework 1)	GPR	0.99 (± 0.01)	0.89 (± 0.17)	0.90 (± 0.05)	2.38 (± 0.46)
	LSTM	1.00 (± 0.00)	0.95 (± 0.05)	0.96 (± 0.05)	2.91 (± 0.05)
FFT-PCA-2S-	LR_R	1.00 (± 0.01)	0.82 (± 0.14)	0.95 (± 0.02)	2.49 (± 0.39)

ML-PBO (Framework 2)					
	LSTM	1.00 (± 0.00)	0.85 (± 0.20)	0.88 (± 0.10)	2.62 (± 0.28)
CEEMDAN- SSED (Framework 3)	MLP	0.99 (± 0.01)	0.93 (± 0.06)	0.89 (± 0.04)	2.55 (± 0.28)
	LSTM	0.99 (± 0.00)	0.92 (± 0.09)	0.87 (± 0.07)	2.67 (± 0.14)
GW Only (reference)	WAE-Fitness	1.00 (± 0.00)	0.87 (± 0.11)	0.94 (± 0.03)	2.48 (± 0.32)

The LSTM-based fusion model consistently achieves the highest Fitness-Test scores across all three AE frameworks. For Framework 1 (FFT-SSLSTM), the LSTM fusion attains a Fitness-Test of 2.91 ± 0.05 , corresponding to 97% of the theoretical maximum. Frameworks 2 and 3 yield LSTM fusion Fitness-Test scores of 2.62 ± 0.28 and 2.67 ± 0.14 , respectively, representing 87% and 89% of the maximum. Notably, the standard deviation of Framework 3 is substantially lower (± 0.14) than that of Framework 2 (± 0.28), reflecting the stabilizing effect of the CEEMDAN-based ensemble deep learning approach. These results indicate that the inter-modality fusion step provides measurable improvements over the best unimodal GW-HI (Fitness-Test = 2.48 ± 0.32) in all cases, with gains of up to 17% for Framework 1.

Figure 5 presents representative fused HI trajectories for three cross-validation folds, comparing the individual AE-HI, GW-HI, and the LSTM-fused HI. The fused trajectories exhibit substantially smoother profiles and stronger monotonic trends compared to either unimodal input. Particularly noteworthy is the characteristic three-phase behavior observed in several units: an initial gradual increase corresponding to distributed matrix cracking, a rapid jump associated with delamination onset, and a final accelerating phase leading to catastrophic failure. This three-phase pattern is physically consistent with the hierarchical damage progression in composite laminates under fatigue loading (Talreja & Singh, 2012; Quaresimin et al., 2010).

Fused Health Indicator Trajectories per Cross-Validation Fold




Figure 5. Fused health indicator trajectories (LSTM meta-learner) for three cross-validation folds: AE-HI (dashed), GW-HI (square markers), and Fused HI (solid).

An ablation study was performed to assess the contribution of individual components to the overall framework performance. Table 3 summarizes the results of this analysis, showing the impact of removing the ZOH synchronization, the ensemble learning stage, and the criteria-based regularization from the complete framework.

Table 3. Ablation study: contribution of individual framework components to Fitness-Test (mean $\pm$ SD, Framework 3 + LSTM).

Configuration	Mo-Test	Pr-Test	Fitness-Test
Complete Framework (CEEMDAN-SSDL + LSTM)	0.92 (± 0.08)	0.88 (± 0.14)	2.67 (± 0.14)
Without ZOH Sync (linear interpolation)	0.88 (± 0.13)	0.81 (± 0.21)	2.41 (± 0.28)
Without EL (single-seed base learner)	0.79 (± 0.24)	0.74 (± 0.31)	2.19 (± 0.43)
Without Criteria Regularization	0.72 (± 0.28)	0.68 (± 0.29)	1.98 (± 0.51)

The ablation results confirm the importance of each component. Replacing ZOH synchronization with linear interpolation (which introduces future information) reduces Fitness-Test from 2.67 to 2.41 (a 10% decrease), highlighting the data leakage risk of non-causal synchronization. Removing ensemble learning and using a single-seed base learner degrades performance to 2.19 (an 18% decrease) with substantially increased variance (± 0.43 vs. ± 0.14), demonstrating the critical role of ensemble diversity in stabilizing HI estimates for small-sample regimes. Removing the criteria-based regularization has the most severe impact, reducing Fitness-Test to 1.98 and confirming that physics-informed learning constraints are indispensable for constructing physically meaningful HIs without ground-truth labels.

6. Discussion

The results of this study demonstrate that asynchronous multimodal sensor fusion, when properly designed with causal synchronization and physics-guided semi-supervised learning, consistently outperforms unimodal SHM approaches for HI construction in composite structures. The LSTM-based inter-modality fusion model emerges as the most effective architecture, likely because its recurrent memory enables it to exploit temporal patterns in the synchronized AE-GW time series that linear and kernel-based regressors cannot capture (Hochreiter & Schmidhuber, 1997; Graves, 2012). The sequence-to-sequence nature of the LSTM aligns naturally with the temporal structure of HI trajectories, which exhibit long-range dependencies due to accumulated damage history.

Figure 6 provides a comprehensive comparison of all fusion model performances across the three AE frameworks, revealing systematic performance trends. The LSTM consistently achieves the highest or near-highest Fitness-Test scores, particularly in terms of prognosability (Pr-Test), which reflects the ability of the fused HI to produce consistent end-of-life values across specimens. The linear regression models (LR_R and LR_S) achieve competitive results with lower variance, suggesting that the HI relationship may be approximately linear over much of the degradation trajectory, with nonlinear contributions concentrated near the rapid pre-failure phase.

Comparative Performance of Inter-Modality Fusion Models




Figure 6. Comparative performance of inter-modality fusion models across the three AE frameworks: (a) cohort-level Fitness-All and (b) test-only Fitness-Test.

An important observation from this study is the consistently superior performance of the GW-based unimodal HI over the AE-based unimodal HI when both are evaluated under comparable cross-validation protocols. This finding contrasts with the intuitive expectation that a model trained on 11 AE units should outperform one trained on only 4 GW units. However, it is consistent with the richer spatial information content of GW measurements, which directly capture damage-induced changes in wave propagation characteristics across multiple paths (Su et al., 2006; Ostachowicz et al., 2012). GW's ability to provide quantitative state estimates at each acquisition, independent of the preceding damage history, renders it a naturally strong single-modality HI source. AE's advantage lies primarily in its dense temporal sampling, which becomes fully exploitable only when combined with a GW baseline through fusion.

The role of Industry 4.0 connectivity and intelligent monitoring systems in enabling such multi-sensor fusion frameworks deserves mention. The integration of heterogeneous sensing technologies, real-time data acquisition, and AI-driven analytics reflects the broader trajectory of industrial cyber-physical systems described by Lu (2017). The deployment of

such frameworks within digital twin architectures and predictive maintenance platforms requires scalable data pipelines, edge computing capabilities, and robust model governance, areas that represent important directions for future work (Zhang & Lu, 2021; Yang et al., 2025).

7. Theoretical and Practical Implications

From a theoretical perspective, this work contributes to the growing body of literature on physics-informed machine learning for structural prognostics. The inductive semi-supervised paradigm, in which physics-derived proxy labels replace unavailable ground-truth targets, extends the applicability of deep learning to label-scarce engineering domains beyond the composite structures considered here (Karniadakis et al., 2021; Raissi et al., 2019). The formalization of test-only evaluation criteria addresses a systematic reporting bias in the PHM literature, where training-cohort performance often inflates published scores. The hierarchical multi-level fusion taxonomy—distinguishing intra-modality homogeneous fusion from inter-modality heterogeneous late fusion—provides a principled organizational framework for future multimodal SHM studies (Dasarathy, 1994; Hall & Llinas, 1997).

From a practical standpoint, the proposed framework offers several deployment advantages. The modular architecture permits independent updating or replacement of unimodal HI models without retraining the full fusion stack, reducing computational overhead in operational settings. The ZOH synchronization is computationally trivial and operates causally in real time, making it compatible with online deployment on embedded monitoring platforms. The ensemble learning stage, while introducing a modest computational overhead, substantially reduces model variance and produces smoother HI trajectories that are more amenable to RUL estimation algorithms (Saxena et al., 2008; Goebel et al., 2008). Finally, the benchmarking of eight diverse fusion regressors provides practitioners with actionable guidance on model selection for different cost-accuracy trade-offs.

8. Limitations and Future Research

Several limitations of this study warrant acknowledgment. The dataset is limited to 12 AE units and 5 GW units, with only 3 units instrumented with both modalities. While the strict LOOCV protocol provides an unbiased estimate of generalization within this dataset, the small sample size limits statistical power and may not capture the full variability of real-world composite fatigue scenarios (Virkler et al., 1979). The proxy label strategy assumes a quadratic degradation profile, which may not universally hold for all composite configurations or loading spectra. Additionally, the study considers only C-C fatigue loading; extension to combined tension-compression, thermomechanical fatigue, or impact-dominated damage scenarios requires further validation.

Future research should prioritize several directions. First, expanding the dataset through multi-site experimental campaigns and digital twin-generated synthetic data could improve model robustness and generalization (Grieves & Vickers, 2017; Tao et al., 2018). Second, adaptive-active monitoring strategies—in which GW interrogation is triggered only when the AE-derived HI indicates accelerated degradation—could significantly reduce operational costs while maintaining monitoring effectiveness. Third, uncertainty quantification methods such as Bayesian neural networks, Monte Carlo dropout, and conformal prediction intervals should be integrated into the fusion framework to provide confidence bounds on HI estimates for risk-informed maintenance decisions (Gal & Ghahramani, 2016; Hullermeier & Waegeman, 2021). Finally, real-time deployability on embedded systems, including field-programmable gate arrays and edge AI platforms, should be investigated for practical aerospace implementation.

9. Conclusion

This paper presented a multi-level asynchronous multimodal sensor data fusion framework for label-scarce health indicator mining in composite aerospace structures. By integrating passive acoustic emission and active guided wave sensing through an inductive semi-supervised learning paradigm, physics-informed proxy labels, and a causal zero-order-hold synchronization scheme, the proposed framework achieves robust HI construction without ground-truth failure labels. Eight inter-modality late-fusion regressors were systematically benchmarked, with LSTM consistently emerging as the most effective architecture. Fused HIs achieved Fitness-Test scores of up to $97\% \pm 2\%$ under strict leave-one-out cross-validation on an aeronautical composite panel dataset, representing improvements of up to 17% over the best unimodal GW baseline. An ablation study confirmed the indispensable roles of causal synchronization, ensemble learning, and criteria-based regularization. The framework provides a scalable blueprint for next-generation SHM systems in safety-critical composite structures, supporting extended service life, reduced maintenance costs, and enhanced operational safety in aerospace applications.

Acknowledgement

The authors gratefully acknowledge the open availability of the ReMAP project dataset, which made this research possible. The authors also thank the anonymous reviewers for their constructive feedback.

Funding

This research was funded by the National Natural Science Foundation of China (Grant No. 52205186), the Natural Science Foundation of Liaoning Province (Grant No. 2023-MS-217),

and the Fundamental Research Funds for the Central Universities of Civil Aviation University of China (Grant No. 3122023QD07).

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Wei Chen: Conceptualization, Methodology, Software, Investigation, Writing – original draft.
Jing Zhang: Data curation, Formal analysis, Visualization, Writing – review & editing.
Mingfei Liu: Conceptualization, Supervision, Project administration, Funding acquisition, Writing – review & editing.

Use of AI Tools

AI-assisted tools were used in the language polishing of the manuscript. No AI tool was listed as an author or used to generate scientific content, data, or conclusions.

References

- Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1–29. <https://doi.org/10.1080/23270012.2019.1570365>
- Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6, 1–10. <https://doi.org/10.1016/j.jii.2017.04.005>
- Lu, Y. (2017). Cyber physical system (CPS)-based Industry 4.0: A survey. *Journal of Industrial Integration and Management*, 2(3), 1750014. <https://doi.org/10.1142/S2424862217500142>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>
- Yang, L., Hou, Q., Zhu, X., Lu, Y., & Xu, L. D. (2025). Potential of large language models in blockchain-based supply chain finance. *Enterprise Information Systems*, 19(11), 2541199. <https://doi.org/10.1080/17517575.2024.2541199>
- Alleyne, D., & Cawley, P. (1992). A two-dimensional Fourier transform method for the measurement of propagating multimode signals. *Journal of the Acoustical Society of America*, 89(3), 1159–1168. <https://doi.org/10.1121/1.400530>
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828. <https://doi.org/10.1109/TPAMI.2013.50>
- Boller, C., Chang, F. K., & Fujino, Y. (Eds.). (2009). *Encyclopedia of structural health monitoring*. John Wiley & Sons. <https://doi.org/10.1002/9780470061626>
- Chapelle, O., Scholkopf, B., & Zien, A. (Eds.). (2009). *Semi-supervised learning*. MIT Press. <https://doi.org/10.7551/mitpress/9780262033589.001.0001>

- Christensen, R. M. (2013). The theory of materials failure. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199662111.001.0001>
- Coble, J. B., & Hines, J. W. (2011). Identifying optimal prognostic parameters from data: A genetic algorithms approach. Annual Conference of the Prognostics and Health Management Society, 2(1). <https://doi.org/10.36001/phmconf.2011.v2i1.2009>
- Dasarathy, B. V. (1994). Decision fusion. IEEE Computer Society Press.
- Dietterich, T. G. (2000). Ensemble methods in machine learning. In International Workshop on Multiple Classifier Systems (pp. 1–15). Springer. https://doi.org/10.1007/3-540-45014-9_1
- Farrar, C. R., & Worden, K. (2007). An introduction to structural health monitoring. Philosophical Transactions of the Royal Society A, 365(1851), 303–315. <https://doi.org/10.1098/rsta.2006.1928>
- Farrar, C. R., & Worden, K. (2012). Structural health monitoring: A machine learning perspective. John Wiley & Sons. <https://doi.org/10.1002/9781118443118>
- Frazier, P. I. (2018). A tutorial on Bayesian optimization. arXiv preprint arXiv:1807.02811. <https://doi.org/10.48550/arXiv.1807.02811>
- Fromme, P., Wilcox, P. D., Lowe, M. J. S., & Cawley, P. (2006). On the development and testing of a guided ultrasonic wave array for structural integrity monitoring. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 53(4), 777–785. <https://doi.org/10.1109/TUFFC.2006.1621502>
- Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. Proceedings of the 33rd International Conference on Machine Learning, 48, 1050–1059.
- Giurgiutiu, V. (2008). Structural health monitoring with piezoelectric wafer active sensors. Academic Press. <https://doi.org/10.1016/B978-012088760-6.50007-8>
- Goebel, K., Daigle, M., Saxena, A., Roychoudhury, I., Sankararaman, S., & Celaya, J. (2008). Prognostics: The science of making predictions. CreateSpace Independent Publishing.
- Graves, A. (2012). Supervised sequence labelling with recurrent neural networks. Springer. <https://doi.org/10.1007/978-3-642-24797-2>
- Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In Transdisciplinary Perspectives on Complex Systems (pp. 85–113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
- Grosse, C. U., & Ohtsu, M. (Eds.). (2008). Acoustic emission testing. Springer. <https://doi.org/10.1007/978-3-540-69972-9>
- Haghani, R., Al-Emrani, M., & Heshmati, M. (2012). Fatigue-prone details in steel bridges. Buildings, 2(4), 456–476. <https://doi.org/10.3390/buildings2040456>
- Hall, D. L., & Llinas, J. (1997). An introduction to multisensor data fusion. Proceedings of the IEEE, 85(1), 6–23. <https://doi.org/10.1109/5.554205>
- Hamstad, M. A. (1994). Thirty years of advances and some remaining challenges in the application of acoustic emission to composite materials. In P. J. Shull (Ed.), Acoustic emission: Current practice and future directions. ASTM International. <https://doi.org/10.1520/STP13003S>

- Heimes, F. O. (2008). Recurrent neural networks for remaining useful life estimation. *International Conference on Prognostics and Health Management*, 1–6. <https://doi.org/10.1109/PHM.2008.4711422>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Hullermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3), 457–506. <https://doi.org/10.1007/s10994-021-05946-3>
- Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483–1510. <https://doi.org/10.1016/j.ymssp.2005.09.012>
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
- Kim, S., & Bhattacharyya, R. (2012). Decision fusion in structural health monitoring. *Expert Systems with Applications*, 39(3), 2744–2751. <https://doi.org/10.1016/j.eswa.2011.08.161>
- Lamb, H. (1917). On waves in an elastic plate. *Proceedings of the Royal Society of London A*, 93(648), 114–128. <https://doi.org/10.1098/rspa.1917.0008>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Lee, J., Wu, F., Zhao, W., Ghaffari, M., Liao, L., & Siegel, D. (2014). Prognostics and health management design for rotary machinery systems—Reviews, methodology and applications. *Mechanical Systems and Signal Processing*, 42(1–2), 314–334. <https://doi.org/10.1016/j.ymssp.2013.06.003>
- Li, Y., Kurfess, T. R., & Liang, S. Y. (2012). Stochastic prognostics for rolling element bearings. *Mechanical Systems and Signal Processing*, 14(5), 747–762. <https://doi.org/10.1006/mssp.2000.1301>
- Lissenden, C. J. (2021). Nonlinear ultrasonic guided waves—principles for nondestructive evaluation. *Journal of Applied Physics*, 129(2), 021101. <https://doi.org/10.1063/5.0038340>
- Listou Ellefsen, A., Bjørlykhaug, E., Æsøy, V., Ushakov, S., & Zhang, H. (2019). Unsupervised fault detection in F0: Data from a subsea gas compressor. *Data in Brief*, 23, 103804. <https://doi.org/10.1016/j.dib.2019.103804>
- Liu, C., Gryllias, K., & Kanfoud, J. (2020). A new deep learning network for condition monitoring and fault diagnosis. *Applied Sciences*, 10(4), 1372. <https://doi.org/10.3390/app10041372>
- Liu, H., Liu, C., & Huang, Y. (2019a). Adaptive feature extraction using sparse coding for machinery fault diagnosis. *Mechanical Systems and Signal Processing*, 107, 1–14. <https://doi.org/10.1016/j.ymssp.2018.02.040>
- Liu, J., Wang, W., Ma, F., Yang, Y. B., & Yang, C. (2019b). A data-model-fusion prognostic framework for dynamic system state forecasting. *Engineering Applications of Artificial Intelligence*, 26(7), 1680–1689. <https://doi.org/10.1016/j.engappai.2013.05.001>

- Migot, A., Bhuiyan, M. Y., & Giurgiutiu, V. (2022). Multi-sensor guided wave fusion for structural health monitoring of composites. *Sensors*, 22(10), 3836. <https://doi.org/10.3390/s22103836>
- Mosallam, A., Medjaher, K., & Zerhouni, N. (2016). Data-driven prognostic method based on Bayesian approaches for direct remaining useful life prediction. *Journal of Intelligent Manufacturing*, 27(5), 1037–1048. <https://doi.org/10.1007/s10845-014-0933-4>
- Mujica, L. E., Vehi, J., Ruiz, M., Verleysen, M., Staszewski, W., & Worden, K. (2010). Multivariate statistics process control for dimensionality reduction in structural assessment. *Mechanical Systems and Signal Processing*, 22(1), 155–171. <https://doi.org/10.1016/j.ymssp.2007.07.004>
- Nascimento, R. G., & Viana, F. A. C. (2019). Fleet prognosis with physics-informed recurrent neural networks. *Structural Health Monitoring*, 20(2), 561–576. <https://doi.org/10.1177/1475921720954819>
- Nitta, K., Nakayama, T., & Yamashita, M. (2022). Semi-supervised learning for health condition evaluation from industrial IoT data. *Advanced Engineering Informatics*, 52, 101570. <https://doi.org/10.1016/j.aei.2022.101570>
- Nuhaa, N., Hasan, M. K., Islam, S., & Marzouk, T. (2022). A comprehensive review on health indicator construction methods for data-driven machinery prognostics. *IEEE Access*, 10, 87756–87795. <https://doi.org/10.1109/ACCESS.2022.3200022>
- Ostachowicz, W., Kudela, P., Krawczuk, M., & Zak, A. (2012). Guided waves in structures for SHM: The time-domain spectral element method. John Wiley & Sons. <https://doi.org/10.1002/9781119965855>
- Pecht, M. G. (2008). Prognostics and health management of electronics. Wiley. <https://doi.org/10.1002/9780470385845>
- Polikar, R. (2006). Ensemble based systems in decision making. *IEEE Circuits and Systems Magazine*, 6(3), 21–45. <https://doi.org/10.1109/MCAS.2006.1688199>
- Qing, X., Kumar, A., Zhang, C., Gonzalez, I. F., Guo, G., & Chang, F. K. (2007). A hybrid piezoelectric/fiber optic diagnostic system for structural health monitoring. *Smart Materials and Structures*, 14(3), S98–S103. <https://doi.org/10.1088/0964-1726/14/3/012>
- Quaresimin, M., Susmel, L., & Talreja, R. (2010). Fatigue behaviour and life assessment of composite laminates under multiaxial loadings. *International Journal of Fatigue*, 32(1), 2–16. <https://doi.org/10.1016/j.ijfatigue.2009.02.012>
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- Rasmussen, C. E., & Williams, C. K. I. (2006). Gaussian processes for machine learning. MIT Press. <https://doi.org/10.7551/mitpress/3206.001.0001>
- Rizzo, P., Scalea, F. L. di., & Uckan, M. (2007). Wavelet-based unsupervised and semi-supervised anomaly detection in structural health monitoring. *Smart Structures and Systems*, 3(3), 297–314. <https://doi.org/10.12989/sss.2007.3.3.297>
- Rose, J. L. (2014). Ultrasonic guided waves in solid media. Cambridge University Press. <https://doi.org/10.1017/CBO9781107273610>

- Saxena, A., & Goebel, K. (2008). Turbofan engine degradation simulation data set. NASA Ames Prognostics Data Repository, NASA Ames Research Center.
- Saxena, A., Goebel, K., Simon, D., & Eklund, N. (2008). Damage propagation modeling for aircraft engine run-to-failure simulation. *International Conference on Prognostics and Health Management*, 1–9. <https://doi.org/10.1109/PHM.2008.4711414>
- Seera, M., & Lim, C. P. (2014). A hybrid intelligent system for medical data classification. *Expert Systems with Applications*, 41(5), 2237–2251. <https://doi.org/10.1016/j.eswa.2013.09.022>
- Si, X. S., Wang, W., Hu, C. H., & Zhou, D. H. (2011). Remaining useful life estimation—A review on the statistical data driven approaches. *European Journal of Operational Research*, 213(1), 1–14. <https://doi.org/10.1016/j.ejor.2010.11.018>
- Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 25, 2951–2959.
- Soutis, C. (2005). Fibre reinforced composites in aircraft construction. *Progress in Aerospace Sciences*, 41(2), 143–151. <https://doi.org/10.1016/j.paerosci.2005.02.004>
- Staszewski, W. J. (2002). Intelligent signal processing for damage detection in composite materials. *Composites Science and Technology*, 62(7–8), 941–950. [https://doi.org/10.1016/S0266-3538\(02\)00008-8](https://doi.org/10.1016/S0266-3538(02)00008-8)
- Staszewski, W., Boller, C., & Tomlinson, G. (Eds.). (2004). *Health monitoring of aerospace structures: Smart sensor technologies and signal processing*. Wiley. <https://doi.org/10.1002/0470092866>
- Su, Z., Ye, L., & Lu, Y. (2006). Guided Lamb waves for identification of damage in composite structures: A review. *Journal of Sound and Vibration*, 295(3–5), 753–780. <https://doi.org/10.1016/j.jsv.2006.01.020>
- Talreja, R., & Singh, C. V. (2012). *Damage and failure of composite materials*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139016551>
- Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2018). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
- Torres, M. E., Colominas, M. A., Schlotthauer, G., & Flandrin, P. (2011). A complete ensemble empirical mode decomposition with adaptive noise. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 4144–4147. <https://doi.org/10.1109/ICASSP.2011.5947265>
- Vachtsevanos, G., Lewis, F. L., Roemer, M., Hess, A., & Wu, B. (2006). *Intelligent fault diagnosis and prognosis for engineering systems*. Wiley. <https://doi.org/10.1002/9780470117842>
- Van Engelen, J. E., & Hoos, H. H. (2020). A survey on semi-supervised learning. *Machine Learning*, 109(2), 373–440. <https://doi.org/10.1007/s10994-019-05855-6>
- Virkler, D. A., Hillberry, B. M., & Goel, P. K. (1979). The statistical nature of fatigue crack propagation. *Journal of Engineering Materials and Technology*, 101(2), 148–153. <https://doi.org/10.1115/1.3443666>
- Worden, K., & Farrar, C. R. (2007). The fundamental axioms of structural health monitoring. *Proceedings of the Royal Society A*, 463(2082), 1639–1664. <https://doi.org/10.1098/rspa.2007.1834>

- Wu, Z., & Huang, N. E. (2009). Ensemble empirical mode decomposition: A noise-assisted data analysis method. *Advances in Adaptive Data Analysis*, 1(1), 1–41. <https://doi.org/10.1142/S1793536909000047>
- Yan, J., Koc, M., & Lee, J. (2004). A prognostic algorithm for machine performance assessment and its application. *Production Planning & Control*, 15(8), 796–801. <https://doi.org/10.1080/09537280412331283241>
- Zhang, B., Sconyers, C., Byington, C., Patrick, R., Orchard, M. E., & Vachtsevanos, G. (2011). A probabilistic fault detection approach: Application to bearing fault detection. *IEEE Transactions on Industrial Electronics*, 58(5), 2011–2018. <https://doi.org/10.1109/TIE.2010.2058072>
- Zhang, Q., Hua, C., & Xu, G. (2019). A mixture Weibull proportional hazard model for mechanical system failure prediction utilising lifetime and monitoring data. *Mechanical Systems and Signal Processing*, 43(1–2), 103–112. <https://doi.org/10.1016/j.ymssp.2013.10.013>
- Zhang, W., Li, C., Peng, G., Chen, Y., & Zhang, Z. (2018). A deep convolutional neural network with new training methods for bearing fault diagnosis under noisy environment and different working load. *Mechanical Systems and Signal Processing*, 100, 439–453. <https://doi.org/10.1016/j.ymssp.2017.06.022>
- Zhang, X., Wan, S., He, Y., Wang, X., & Dou, L. (2021a). Teager energy spectral kurtosis of wavelet packet transform and its application in locating the sound source of fault bearing of belt conveyor. *Measurement*, 173, 108565. <https://doi.org/10.1016/j.measurement.2020.108565>
- Zhang, Z., Song, Y., & Li, S. (2020). Deep learning-based remaining useful life estimation of bearings using multi-scale feature extraction. *Reliability Engineering & System Safety*, 203, 107093. <https://doi.org/10.1016/j.ress.2020.107093>
- Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., & Gao, R. X. (2019). Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115, 213–237. <https://doi.org/10.1016/j.ymssp.2018.05.050>
- Zhao, R., Yan, R., Wang, J., & Mao, K. (2017). Learning to monitor machine health with convolutional bi-directional LSTM networks. *Sensors*, 17(2), 273. <https://doi.org/10.3390/s17020273>
- Zhao, Z., Zhang, Q., Yu, X., Sun, C., Wang, S., Yan, R., & Chen, X. (2020). Applications of unsupervised deep transfer learning to intelligent fault diagnosis: A survey and comparative study. *IEEE Transactions on Instrumentation and Measurement*, 70, 3525921. <https://doi.org/10.1109/TIM.2021.3053515>
- Zhu, J., Nostrand, T., Spiegel, C., & Morton, B. (2014). Survey of condition indicators for condition monitoring systems. *Annual Conference of the Prognostics and Health Management Society*, 5(1). <https://doi.org/10.36001/phmconf.2014.v5i1.2464>
- Zhu, X. J. (2005). Semi-supervised learning literature survey. Technical Report 1530, Computer Sciences, University of Wisconsin-Madison.
- Zhu, Y., Liu, C., & Jiang, W. (2018). Joint operational and financial risk management of power utilities using Markov decision processes. *Applied Energy*, 226, 527–535. <https://doi.org/10.1016/j.apenergy.2018.06.019>