

Graph Neural Network-Based Structural Health Monitoring of Intelligent Laminated Damping Systems

Mirtins Parker

School of Computing, Clemson University, Clemson, SC, USA.
martins.parker842@clemson.edu

Elliot M. Schmidt

Department of Computer Science, University of Houston, Houston, TX, USA.
elliott646@uh.edu

Larry Salonen

Department of Computer Science, George Mason University, Fairfax, VA, USA.
larrysalonen@gmu.edu

Abstract

The integration of structural health monitoring (SHM) with intelligent laminated damping systems presents a significant opportunity to enhance the resilience, safety, and longevity of critical civil, aerospace, and mechanical infrastructures. This paper investigates the system-level design and deployment of graph neural network (GNN) architectures for the real-time assessment and prognostics of such smart composite structures. Intelligent laminated damping systems combine viscoelastic or constrained-layer damping treatments with embedded sensor arrays, generating spatially distributed, high-dimensional dynamic data that pose unique challenges for conventional signal processing. GNNs, by virtue of their relational inductive bias, are inherently suited to model the complex topological dependencies among sensor nodes, damage-sensitive features, and laminate layup configurations. The paper provides a comprehensive examination of the architectural choices, integration patterns, and trade-offs involved in building GNN-driven monitoring frameworks. Emphasis is placed on the fusion of physics-informed representations, the translation of laminate-scale phenomena into graph abstractions, and the orchestration of data pipelines from edge sensing to cloud analytics. Beyond technical design, the discussion extends to life-cycle governance, sustainability of monitoring infrastructures, adversarial robustness of learned models, algorithmic fairness in multi-asset management, and policy implications for the adoption of intelligent SHM systems. Case illustrations from aerospace skins, wind turbine blades, and bridge retrofits are used to ground the analysis. The work articulates a forward-looking perspective on how GNN-based monitoring can evolve from isolated prototypes to trusted, scalable, and equitable decision-support layers within the next generation of adaptive structural systems.

Keywords

structural health monitoring; graph neural networks; laminated damping; smart structures; sensor networks; infrastructure governance; robustness; fairness; edge intelligence; digital twin.

1. Introduction

Structural health monitoring has evolved from periodic visual inspections into a data-intensive discipline, driven by advances in sensor technology, wireless communication, and machine learning. The fundamental goal is to detect, localize, and quantify damage before it compromises structural integrity, thereby enabling condition-based maintenance and extending service life. In parallel, the deployment of intelligent laminated damping systems has attracted interest across aerospace, automotive, and civil engineering domains, where lightweight laminated composites with embedded viscoelastic layers can simultaneously provide load-bearing capacity and vibration attenuation. The marriage of these two domains creates a class of structures that are not only mechanically adaptive but also information-rich, offering the possibility of real-time, autonomous damage prognosis through dense sensor networks. However, realizing this vision requires analytical frameworks that can handle the complexity of distributed dynamic responses, spatially varying damping properties, and the intricate coupling between material degradation and sensor signals. Foundational reviews of SHM have established the taxonomy of vibration-based, impedance-based, and guided-wave methods [1, 2], while the evolution of smart structures has been surveyed extensively, emphasizing the integration of actuators, sensors, and control [3].

Concurrently, the treatment of damping in laminated plates and sandwich structures has received sustained theoretical and experimental attention, yielding models that capture frequency- and temperature-dependent behavior of viscoelastic cores [4]. These models provide the physical basis for interpreting sensor data, yet their use in large-scale monitoring remains limited due to computational cost and difficulty in updating parameters across numerous degrees of freedom. The rise of deep learning, particularly convolutional neural networks and recurrent architectures, has enabled end-to-end damage classification from raw sensor signals, bypassing traditional feature extraction pipelines. Nevertheless, these methods typically assume regular spatial grids or time-series inputs and do not natively exploit the relational structure among sensing points distributed across a laminate. Graph neural networks offer a compelling alternative by directly operating on graph-structured data, where sensing locations, ply interfaces, and material discontinuities are represented as nodes and edges. Early GCN formulations demonstrated powerful semi-supervised learning on graph data [5], opening pathways for modeling physical systems as interaction networks [7]. In the context of smart laminated damping systems, GNNs can capture both local damage-induced anomalies and global changes in energy dissipation pathways. The present work explores the system-level ramifications of this shift, spanning from architectural design and data integration to governance, sustainability, and fairness, with the aim of framing GNN-based SHM as a holistic socio-technical infrastructure rather than a mere algorithmic upgrade.

2. Foundations of Structural Health Monitoring and Smart Damping Systems

A comprehensive understanding of SHM requires situating it within the lifecycle of engineered structures. Traditional SHM systems are built around four sequential processes: operational evaluation, data acquisition, feature extraction, and statistical model development for damage diagnosis. Vibration-based methods have been extensively catalogued [2], and the modal parameters extracted from dynamic tests have long served as damage indicators. However, laminated structures with integrated damping layers exhibit non-proportional damping, which complicates mode shapes and natural frequencies, making conventional modal interpretation error-prone. Wireless sensor networks have been instrumental in reducing installation costs and enabling dense deployment [6], yet they introduce challenges related to synchronization, power management, and data loss. Smart structures research has

pushed toward active and semi-active damping control using piezoelectric materials and magnetorheological elastomers, but the emphasis has often been on control authority rather than on leveraging the same sensor suites for long-term monitoring [3].

Damping in laminated composites is typically introduced through constrained-layer treatments or co-cured viscoelastic interlayers. The resulting loss factors depend on temperature, frequency, and strain amplitude, and can evolve with material aging, debonding, and microcracking. Reviews on the modeling of viscoelastic damping in sandwich structures underscore the need for finite element formulations that incorporate complex moduli and frequency-domain analysis [4]. These models have been validated through laboratory experiments and used to optimize damping performance. Yet the translation of these high-fidelity models into field-deployed SHM remains an open challenge, because the boundary conditions, environmental variability, and aging phenomena in situ differ substantially from controlled laboratory settings. A crucial bridge is model updating, where numerical models are systematically adjusted to match measured data. Detailed numerical model updating procedures for smart laminated plates with embedded viscoelastic damping layers have been developed, linking finite element representations with experimental modal data through sensitivity-based optimization [8]. Such updating demonstrates how sensor-informed parameter tuning can refine the virtual representation of a damping system, a concept that GNN architectures can extend by learning the mapping between graph representations of sensor networks and latent damage variables without requiring explicit updating of full finite element models.

3. Graph Neural Networks for Structural State Inference

Graph neural networks emerged from the need to process data that does not reside on regular grids but is instead composed of entities and their pairwise relationships. The initial GNN model of Scarselli et al. [14] formalized the notion of recursive state propagation, allowing each node to aggregate information from its neighbors through message-passing layers. The resurgence of interest was fueled by graph convolutional networks, which simplified the convolutional filter to operate directly on the graph Laplacian or adjacency matrix, enabling efficient semi-supervised node classification [5]. Subsequent advances introduced attention mechanisms that dynamically weigh neighbor contributions, leading to graph attention networks [13], and inductive frameworks like GraphSAGE that generate embeddings by sampling and aggregating features from local neighborhoods, thus circumventing the need to retrain for unseen graphs [12]. A comprehensive review of GNN methods highlights their success across social networks, recommendation systems, biology, and physical simulation [15].

In structural health monitoring, the application of GNNs remains nascent but is gaining traction. The paradigm maps each sensor or patch of a laminated structure to a node, while edges encode geometric proximity, mechanical coupling, or sensor similarity. Damage-induced feature alterations propagate through the graph, and GNNs learn to detect these changes by modeling node- and edge-level anomalies. Sun et al. [10] demonstrated that a graph attention network could localize damage in beam-like structures using acceleration signals, outperforming traditional classifiers by leveraging the spatial dependencies between measurement points. The relational inductive bias inherent in GNNs makes them particularly well suited for composite laminates, where the anisotropic stiffness, layer-wise fiber orientations, and designated damping layers create complex interaction topologies that are not captured by grid-based geometries. By constructing graphs that reflect ply adjacency,

thickness, and material interfaces, a GNN can learn to encode the multi-scale mechanics without explicitly resolving governing equations. The attention mechanism [11, 13] allows the network to selectively focus on damage-sensitive regions, such as delamination boundaries or debonded damping layers, effectively prioritizing informative connections. This capability is crucial when the sensor density is non-uniform, a common scenario in retrofitted structures where legacy instrumentation coexists with new smart patches.

Beyond spatial topology, the temporal evolution of vibration signals can be encoded as dynamic graphs or as sequences of graph snapshots. Such spatiotemporal graph networks extend the framework to track damage progression over operational cycles. The interaction network model [7] provides a blueprint for reasoning about physical dynamics by inferring state changes from pairwise interactions, which is directly applicable to the recursive computation of strain transfer and energy dissipation across laminate plies. Integrating physics-informed priors into GNN layers, such as ensuring that edge messages respect approximate compliance relationships or damping trends, can improve sample efficiency and generalization. This design space requires balancing expressive power with computational tractability, a trade-off that becomes acute when moving from offline laboratory datasets to online, multi-channel sensing streams at the edge.

4. Intelligent Laminated Damping Systems: Sensorization and Model-Based Reasoning

An intelligent laminated damping system extends the traditional laminated composite by embedding or surface-mounting a network of sensors and, optionally, actuators. Typical choices include piezoelectric patches for guided wave excitation and reception, fiber Bragg gratings for strain and temperature monitoring, and MEMS accelerometers for low-frequency vibration. The sensor layout must be co-designed with the damping treatment to ensure that the monitored modes are sensitive to degradation in the viscoelastic layer. For instance, a constrained-layer damping scheme in a wind turbine blade skin might combine surface accelerometer arrays with embedded piezoelectric sensors near the root spar to capture both global blade bending and local skin delamination. The sensor signals inherently reflect the coupled electro-mechanical-damping behavior, and interpretation requires models that connect the measured voltages or displacements to the integrity of the damping interface.

Numerical simulations have been indispensable for understanding the dynamic response of such hybrids. Detailed finite element models of laminated plates with integrated viscoelastic layers and surface-bonded piezoelectric sensors provide the basis for virtual sensing and model validation [8]. These models incorporate frequency-dependent material properties and enable parametric studies on the influence of debonding size and location on modal quantities. However, their computational expense precludes real-time deployment. One approach is to pre-compute a library of damage scenarios and train a surrogate model, such as a neural network, to map sensor features to damage states. Deep convolutional networks have shown real-time damage detection capabilities using raw acceleration signals from simple beam and frame structures [9]. The challenge is that such 1-D and 2-D CNN architectures flatten the spatial relationships, potentially missing the directional dependencies that govern wave propagation in orthotropic laminates. GNNs, in contrast, preserve the native graph structure, allowing the algorithm to learn that a certain pair of sensors is more tightly coupled through the damping layer than another pair, and that damage at a point manifests as altered pairwise coherence or transfer entropy. This model-based reasoning is enriched by graph pooling and unpooling operations that can aggregate fine-scale sensor graphs into coarse representations suitable for classification or regression at the structural component level.

The signal processing chain also benefits from advances in graph signal processing, which extends Fourier analysis to graph domains. By defining the graph Laplacian from the sensor adjacency, one can compute graph Fourier transforms that separate smooth variations (global stiffness loss) from high-frequency graph modes (localized damage). GNNs can learn to combine such spectral features with attention-based aggregation, forming a multi-paradigm reasoning engine. The integration of physical knowledge, for example by enforcing that edge weights reflect mechanical impedance or damping loss factors, leads to physics-enhanced GNNs that require less labeled damage data, a critical advantage given the scarcity of in-service damage records for novel laminated systems.

5. System-Level Architecture and Multiscale Data Integration

Deploying GNN-based SHM in intelligent laminated damping systems necessitates a layered architecture that coordinates embedded sensing, edge computing, and cloud-based analytics. At the lowest layer, sensor nodes perform local signal conditioning and feature extraction. Intermediate gateways aggregate data and execute lightweight GNN inference for preliminary damage alarming. The cloud layer houses high-resolution models, digital twin synchronization, and fleet-wide learning. This architecture must be designed to handle heterogeneous data rates, from high-frequency guided wave acquisitions exceeding hundreds of kilohertz to low-frequency temperature and humidity records. The graph representation itself must therefore accommodate heterogeneous node types and dynamic edge attributes that reflect varying sampling intervals and sensor modalities.

Multiscale data integration is a principal system challenge. Laminated damping structures exhibit damage across length scales, from matrix microcracking and delamination at the fiber-matrix interface to large-area debonding of the damping layer. While a GNN can fuse signals from macro-scale accelerometers and micro-scale fiber optic strain sensors, the graph must encode hierarchical relationships: fine-scale nodes representing small material neighborhoods feed into coarse supernodes representing structural components. Hierarchical graph pooling architectures, such as DiffPool and its extensions, provide mechanisms to learn these coarsening operations, enabling the model to simultaneously classify local ply-level damage and global loss of damping performance. The training strategy must ensure that gradient signals propagate across scales, which demands careful initialization and possibly curriculum learning schemes where the network first learns coarse patterns before refining local sensitivity.

In practice, the creation of these graphs relies on a digital twin that stores the as-designed and as-built geometry, material properties, and sensor registrations. The digital twin for bridge infrastructure has been demonstrated as a platform for integrating inspection data, SHM inputs, and finite element models [20]. For an intelligent laminate, the twin should contain the stacking sequence, viscoelastic layer positions, sensor placements, and uncertainty bounds on material aging. The GNN then acts as a dynamic reducer that maps from the high-dimensional twin state and streaming sensor data to a low-dimensional damage index. Synchronization between the physical asset and its digital counterpart must be maintained through continuous data ingestion, and model drift must be detected and corrected via online learning or periodic recalibration using known input excitations, such as ambient wind or controlled actuation. The resulting closed-loop monitoring system can inform maintenance scheduling, load limitation advisories, and ultimately design feedback for future generations of intelligent laminates.

6. Deployment, Edge Intelligence, and Life-Cycle Governance

Realizing the vision of GNN-driven SHM in operational environments demands robust deployment strategies that acknowledge resource constraints, communication intermittency, and evolving cybersecurity threats. Edge intelligence frameworks are emerging as a means to decentralize deep learning inference, reducing the latency and bandwidth required to transmit raw vibration streams to a central server. Decentralized deep learning approaches for wireless SHM have proven effective for real-time damage detection using compact CNN models on embedded processors [19]. For GNNs, edge deployment is more challenging because the graph may span multiple physical sensor clusters, and neighbor aggregation requires cross-node communication. Distributed message-passing protocols can be implemented, where each sensor cluster holds a subgraph and exchanges aggregated embeddings with adjacent clusters over low-power wireless links. This design must balance the size of neighborhood aggregation hops against communication cost and power budgets, a non-trivial trade-off that influences detection latency and model accuracy.

The life-cycle governance of such systems extends from initial planning and sensor selection to long-term model maintenance and eventual decommissioning. Sustainability considerations favor sensor modalities with low environmental impact and energy-neutral operation, such as vibration energy harvesters integrated into the damping layer. The data pipeline itself must be governed by transparent protocols that define data ownership, access rights, and retention policies, particularly when monitoring involves public infrastructure or multi-owner assets like commercial aircraft. Edge and cloud data stores must comply with regional data sovereignty regulations, and the use of federated learning can allow model updates across a fleet of structures without centralizing sensitive operational data, as explored in general federated learning frameworks [22]. In a fleet of wind turbine blades with identical intelligent laminate designs, a global GNN model can be improved by aggregating locally trained weights while preserving the privacy of individual turbine operators. The governance structure must also address algorithm accountability, specifying how false alarms and missed detections are investigated and remediated.

Furthermore, the integration of GNN-based SHM into existing asset management workflows demands interoperability with enterprise systems such as computerized maintenance management software and structural reliability databases. This requires standardized APIs and semantic data models, potentially aligned with openBIM standards for civil infrastructure or emerging industrial internet of things frameworks. The economic viability of widespread adoption hinges on demonstrating a clear return on investment through reduced inspection costs, avoidance of catastrophic failures, and extended service intervals, metrics that must be empirically validated across pilot deployments.

7. Robustness, Fairness, and Policy Implications for Advanced SHM Infrastructures

As GNN models become embedded in critical decision-making pipelines, their robustness to adversarial perturbations and distributional shifts must be thoroughly examined. Adversarial attacks on graph neural networks have been shown to degrade node classification performance through subtle modifications to node features or graph structure [21]. In an SHM context, an attacker might spoof sensor readings or inject fake data packets that create ghost damage indications or mask real cracks. Defensive strategies include adversarial training, randomized smoothing of graph representations, and anomaly detection on the input manifold to reject out-of-distribution samples before they reach the GNN. Robustness provisions should be incorporated at the system design stage, not retrofitted, and should be verified through red-teaming exercises that simulate plausible cyber-physical attack scenarios.

Beyond security, fairness considerations arise when intelligent SHM systems are deployed across geographically and socio-economically diverse infrastructure portfolios. If a municipal bridge network equips only newly built or high-traffic bridges with advanced GNN monitoring while leaving older, underserved neighborhood structures with periodic visual inspections, the resulting safety outcomes become skewed. The design of monitoring deployment plans must address equitable distribution of sensing resources, influenced by risk profiles, traffic volumes, and community vulnerability indices. Fairness in machine learning has been extensively surveyed, identifying sources of bias from data collection, model selection, and feedback loops [18]. For SHM, a fair policy would ensure that the benefits of early damage detection and predictive maintenance are not reserved exclusively for assets in affluent areas. Algorithmic fairness constraints could be integrated into the GNN training objective to penalize performance disparities across asset groups, although this must be balanced against overall detection accuracy. The interpretability of GNN decisions is essential to justify such trade-offs to stakeholders. Techniques for interpreting deep neural networks, including layer-wise relevance propagation and Shapley value approximations for graphs, can provide explanations of which sensor nodes and ply interfaces contributed to a damage alert [25], enabling accountable maintenance decisions.

Policy frameworks are currently nascent for AI-augmented SHM. Regulatory bodies such as the Federal Aviation Administration for aircraft and national transportation agencies for bridges require clear certification pathways for digital monitoring systems that influence load ratings and maintenance intervals. A performance-based certification could be based on extensive simulation and field validation, possibly using digital twin environments to stress-test GNN models against historical damage cases. Standards organizations need to develop benchmarks and reference datasets for laminated composite damage detection, covering various layups, damping treatments, and environmental conditions, which are currently absent. The integration of intelligent fault detection concepts with the broader reliability and risk assessment frameworks [23, 24] will be key to establishing trust. Furthermore, the environmental footprint of the monitoring infrastructure itself must be considered; the production of sensors, edge devices, and cloud computing capacity has embodied energy and material costs that should be justified by the safety and longevity gains. A holistic sustainability assessment quantifies the net reduction in material consumption achieved by avoiding premature structural replacements and repairs enabled by early damage detection.

8. Conclusion

Graph neural network-based structural health monitoring of intelligent laminated damping systems represents a convergence of materials engineering, sensor networks, and relational machine learning that transcends conventional vibration-based diagnostics. This paper has articulated a system-level perspective that frames GNNs not merely as replacement classifiers but as integrating agents that fuse physical knowledge, sensor topology, and multi-scale damage mechanisms into a coherent inference architecture. The intelligent laminate, with its viscoelastic damping layers and embedded sensing, provides a natural graph structure where nodal interactions mirror mechanical coupling and energy dissipation pathways. By leveraging the inductive biases of graph convolutional and attention-based propagation, monitoring systems can achieve sensitivity to subtle delamination and debonding that eludes grid-oriented deep learning.

The full realization of such systems requires careful attention to architectural design, covering hierarchical graph construction, edge-cloud partitioning, and real-time message-passing

protocols. Sustainability and life-cycle governance demand that data pipelines, model updates, and hardware provisioning are treated as integral parts of the monitoring ecosystem, not as afterthoughts. Robustness against adversarial manipulation and fairness across diverse asset fleets must become first-class design requirements, supported by interpretability and auditing capabilities. The policy landscape must evolve in tandem, producing standards and certification frameworks that accommodate continuously learning models while maintaining safety assurance. As intelligent laminated structures propagate from aerospace skins and wind turbine blades to civil infrastructure retrofits, GNN-based SHM will play a central role in transforming these passive composites into self-aware, self-diagnosing members of a resilient built environment.

References

1. Farrar, C. R., & Worden, K. (2007). An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A*, 365(1851), 303–315.
2. Doebling, S. W., Farrar, C. R., & Prime, M. B. (1998). A summary review of vibration-based damage identification methods. *Shock and Vibration Digest*, 30(2), 91–105.
3. Chopra, I. (2002). Review of state of the art of smart structures and integrated systems. *AIAA Journal*, 40(11), 2145–2187.
4. Moreira, R. A. S., & Dias Rodrigues, J. (2010). A review on the modelling of viscoelastic damping in sandwich structures and laminated plates. *Mechanics of Composite Materials*, 46(3), 245–262.
5. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *Proceedings of the International Conference on Learning Representations*.
6. Lynch, J. P., & Loh, K. J. (2006). A summary review of wireless sensors and sensor networks for structural health monitoring. *Shock and Vibration Digest*, 38(2), 91–128.
7. Battaglia, P. W., Pascanu, R., Lai, M., Rezende, D. J., & Kavukcuoglu, K. (2016). Interaction networks for learning about objects, relations and physics. *Advances in Neural Information Processing Systems*, 29, 4502–4510.
8. Lu, J., Zhan, Z., Liu, X., & Wang, P. (2018). Numerical modeling and model updating for smart laminated structures with viscoelastic damping. *Smart Materials and Structures*, 27(7), 075038.
9. Abdeljaber, O., Avci, O., Kiranyaz, S., Gabbouj, M., & Inman, D. J. (2017). Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks. *Journal of Sound and Vibration*, 388, 154–170.
10. Sun, D., Wang, Q., & Li, H. (2021). Graph attention network-based approach for structural damage detection. *Mechanical Systems and Signal Processing*, 148, 107157.
11. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
12. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 1024–1034.

13. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. *Proceedings of the International Conference on Learning Representations*.
14. Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., & Monfardini, G. (2009). The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1), 61–80.
15. Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., & Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI Open*, 1, 57–81.
16. Kosmatka, J. B., & Liguore, S. L. (1993). Review of methods for analyzing constrained-layer damped structures. *Journal of Aerospace Engineering*, 6(3), 197–225.
17. Frangopol, D. M., & Soliman, M. (2016). Life-cycle of structural systems: recent advances and future directions. *Structure and Infrastructure Engineering*, 12(1), 1–20.
18. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
19. Avci, O., Abdeljaber, O., Kiranyaz, S., & Inman, D. J. (2021). Wireless and real-time structural damage detection: A decentralized deep learning approach. *Structural Control and Health Monitoring*, 28(4), e2687.
20. Zhang, C., Hammad, A., & Motamedi, A. (2020). Digital twin for the structural health monitoring of bridge infrastructure. *Journal of Civil Structural Health Monitoring*, 10, 115–127.
21. Zügner, D., Akbarnejad, A., & Günnemann, S. (2018). Adversarial attacks on neural networks for graph data. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2847–2856.
22. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
23. Faber, M. H., & Stewart, M. G. (2003). Risk assessment for civil engineering facilities: critical overview and discussion. *Reliability Engineering and System Safety*, 80(2), 173–184.
24. Worden, K., & Dulieu-Barton, J. M. (2004). An overview of intelligent fault detection in systems and structures. *Structural Health Monitoring*, 3(1), 85–98.
25. Montavon, G., Samek, W., & Müller, K. R. (2018). Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73, 1–15.