

Physics-Informed Machine Learning for Damage Localization and Progression in Composite Structures

Kevin Steiner¹, and Olga Fink²

^{1,2} EPFL – Intelligent Maintenance and Operations Systems (IMOS)
Laboratory CH-1015 Lausanne
Switzerland

kevin.steiner@epfl.ch
olga.fink@epfl.ch

ABSTRACT

Composite materials are widely used in lightweight structures, but their heterogeneous architecture and interacting damage mechanisms make damage progression difficult to predict. As a result, composite components are often designed with conservative safety margins, increasing weight, cost, and inspection effort. Improved methods for structural health monitoring and damage-progression analysis are therefore needed to support earlier damage detection, residual-life assessment, and more efficient structural design. This research develops physics-informed machine-learning methods for monitoring and modeling damage evolution in composites. We propose a framework that addresses three complementary challenges: (1) damage detection and localization from sparse strain measurements using graph neural networks that generalize to unseen damage locations and sensor layouts, (2) microscale stress prediction in damaged composite microstructures using a deep-learning surrogate that can be integrated into micro-macro simulation workflows, and (3) macroscale damage modeling using a topologically informed neural network. Together, these contributions form a multiscale learning framework linking sparse experimental observations, microscale material response, and macroscale laminate damage evolution.

1. MOTIVATION AND PROBLEM STATEMENT

The high stiffness-to-mass ratio of composite materials makes them attractive for aerospace, wind energy, automotive, and civil engineering applications. However, their heterogeneous architecture and interacting damage mechanisms make damage progression difficult to predict, leading to conservative designs with large safety factors (Yang, Pei,

Yan, & Long, 2020). Because composite damage can remain hidden from visual inspection while still compromising structural integrity (Ogunleye, Rusnáková, Javořík, Žaludek, & Kotlánová, 2024), improved methods for damage localization and progression analysis are needed to support reliable structural health assessment and targeted inspections.

Current methods that address damage localization rely on dense or regularly spaced sensors, or controlled experimental conditions (Bergmayr, Höll, Kralovec, & Schagerl, 2023; Del Priore, Sbarra, Coppotelli, Lampani, & Pasquali, 2025), which limits their applicability to sparse, noisy, and irregular sensor layouts.

At the same time, also predicting damage evolution remains a major challenge. As damage initiation and growth occur across multiple spatial and temporal scales, macroscopic damage criteria often cannot capture the underlying microscale mechanisms (Hinton & Kaddour, 2012), while high-fidelity multiscale simulations remain computationally prohibitive for industrial-scale structures (Ghosh, Biswas, & Bhattacharyya, 2025). Although machine-learning surrogates can reduce this cost, existing approaches typically address either microscale responses in isolation (Yacouti & Shakiba, 2025) or structural-scale behavior without evolving microstructural damage (Ghosh et al., 2025).

Moreover, the application of current mesh-based machine-learning methods (Pfaff, Fortunato, Sanchez-Gonzalez, & Battaglia, 2021; Iparraguirre, Alfaro, Gonzalez, & Cueto, 2026) to damaged composite laminates remains relatively unexplored. The challenges of these methods to extrapolate to unseen geometries additionally limit their applicability within unified micro-macro damage-progression frameworks.

To address these challenges, this research develops a unified multiscale learning framework for composite damage monitoring and progressive damage analysis that spans sens-

Kevin Steiner et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ing, microstructural modeling, and structural-scale prediction. First, it proposes a graph neural network (GNN) for damage detection and localization from sparse strain measurements, enabling robust identification of damage under previously unseen damage locations and sensor layouts. Second, it introduces a physics-informed deep-learning surrogate for full-field stress prediction in damaged composite microstructures, providing a computationally efficient alternative to high-fidelity simulations while preserving the spatial resolution required for damage-progression analysis. Third, it develops a topology-aware graph representation of composite laminates at the macroscale, explicitly modeling the hierarchical relationships between plies and interfaces to capture damage-induced interactions and structural response. By bridging structural health monitoring, microscale damage mechanics, and macroscale structural modeling within a common learning paradigm, these contributions establish the foundations for scalable, data-driven prediction of damage initiation, localization, and evolution across the micro–macro scale hierarchy.

2. EXPECTED CONTRIBUTIONS

2.1. WP1: Damage Localization in Composite Panels

The first work package focuses on the development of a robust damage localization framework for composite panels using sparsely distributed strain gauges. A wide range of SHM approaches address damage localization in composite structures, including acoustic emission, active and passive guided wave, and strain-based methods (Singh & Sehgal, 2022). Strain sensing is particularly attractive for long-term monitoring (Ogunleye et al., 2024) because it can provide continuous measurements during normal operation without active interrogation of the structure as required by many guided-wave-based methods. Because damage-induced perturbations in the strain field are often highly localized, accurate localization typically relies on sufficient spatial coverage of the structure.

Recent advances in machine learning have enabled the direct inference of damage locations from the strain measurements (Bergmayr et al., 2023; Del Priore et al., 2025; Mylonas, Tsialiamanis, Worden, & Chatzi, 2023). However, current studies are still limited in several important aspects. First, even so-called sparse sensing configurations often maintain a high sensor density relative to the damage size, reducing the practical difficulty of the localization task. Second, evaluations are typically conducted under controlled laboratory conditions, with limited investigation of robustness to different magnitudes of measurement noise and sensor uncertainty. Third, and most importantly, existing methods rarely assess whether learned models can generalize to damage locations that are completely absent from the training data. This spatial generalization is critical in practice, where failures may occur in previously unseen regions of a structure.

To address these challenges, we propose a physics-informed GNN architecture for damage localization from sparse strain measurements. The approach exploits the inherently spatial nature of damage-induced strain perturbations through a graph representation of the sensor network. The model combines an inverse localization network with a pretrained forward plausibility model that predicts the strain response associated with the estimated damage state. Coupling these components during training constrains the predicted damage location to remain physically consistent with the measured strain field, improves localization robustness under sparse sensing conditions, and provides a mechanism to validate localization results through forward reconstruction of the sensor signal.

2.2. WP2: Stress Field Prediction in Damaged Composite Microstructures

The second work package focuses on developing machine-learning surrogate models for progressive damage simulation in composite microstructures, with the ultimate goal of enabling multiscale damage modeling at industrial scales. Composite failure emerges from interacting microscale mechanisms such as matrix cracking or fiber-matrix debonding, which may accumulate into delamination or structural failure. Capturing these processes is essential for reliable life prediction, structural health assessment, and the development of digital twins capable of forecasting structural degradation under operational loading conditions.

High-fidelity micro–macro simulations address this limitation by embedding representative volume elements (RVEs) into macroscale structural models and resolving local stress and strain fields during incremental loading. However, their widespread adoption is hindered by a fundamental computational bottleneck: the repeated solution of microscale boundary value problems at every load increment makes progressive multiscale damage simulations prohibitively expensive for industrial-scale components (Ghosh et al., 2025). To overcome this limitation, recent research has explored machine-learning surrogates for RVEs, including models that directly predict damage trajectories or high-resolution microscale fields (Yacouti & Shakiba, 2025).

However, most existing approaches assume undamaged microstructures or focus on predicting homogenized quantities and aggregate stress metrics. Furthermore, the presence of damage requires models to generalize across a potentially vast space of crack geometries, damage locations, and damage severities. To date, very few studies have addressed this challenge. Although initial efforts have integrated microscale surrogates into multiscale simulation frameworks (Zhou & Semnani, 2025), these approaches typically predict homogenized responses rather than the full stress tensor required for physics-based damage evolution. Consequently, the integra-

tion of machine-learning surrogates into progressive multi-scale damage workflows remains an open research challenge.

To address these limitations, we propose a GNN surrogate model for full-field stress prediction in damaged composite microstructures. The model will leverage local geometric features and equivariant message passing to improve generalization across previously unseen damage locations and geometries. In addition, we will investigate the integration of physics-informed inductive biases, including stress-equilibrium constraints and conservation laws, to improve robustness, generalization and physical consistency (Garban, Chemisky, Prulière, & Clément, 2025). Since no public benchmark currently exists for full-field stress prediction in damaged composite RVEs, we will create a new dataset spanning diverse microstructures, loading conditions, and damage scenarios. To demonstrate practical relevance, the proposed surrogate will be integrated into a micro-macro progressive damage simulation, where the predicted high-resolution stress fields are coupled with established damage criteria to assess whether expensive microscale solvers can be replaced while preserving the predictive fidelity required for progressive damage analysis. Finally, the workflow will be validated experimentally using pre-damaged composite coupons, providing a direct evaluation of machine-learning-driven multi-scale damage progression against physical observations.

2.3. WP3: Laminate-Scale Damage Propagation

The third work package focuses on developing a topology-aware GNN model for modeling damage evolution in ply-stacked composite structures at the laminate scale. It complements WP2 because many engineering decisions ultimately depend on how local damage processes propagate and interact at the laminate scale.

Classical progressive damage models typically require fine spatial discretizations, carefully calibrated constitutive laws, and repeated nonlinear simulations throughout the loading history (Ghosh et al., 2025). Consequently, their computational cost often limits their applicability in design-space exploration, uncertainty quantification, and large-scale optimization studies. Recent advances in machine learning have demonstrated the potential of graph-based and mesh-based surrogate models for approximating complex structural simulations on discretized domains (Pfaff et al., 2021). However, these approaches typically operate on fixed laminate configurations, and their ability to generalize across different stacking sequences, ply arrangements, and interface configurations remains largely unexplored.

This limitation is important because laminate topology is a primary driver of damage propagation. A surrogate model intended for engineering design should be able to transfer across laminate configurations rather than being restricted to a single layout. Achieving such generalization requires repre-

sentations that explicitly encode the hierarchical relationships between plies and interfaces. Recent advances in topological machine learning provide a promising foundation for developing such representations (Wei & Fink, 2025).

To address these challenges, we propose a topology-aware machine learning model that explicitly represents plies, interfaces, and structural features as interconnected entities. This representation aims to learn transferable damage-evolution models that generalize across previously unseen laminate designs. Furthermore, the topological representation provides a natural mechanism for incorporating physical priors, such as load transfer between adjacent plies and discontinuities of transverse strain across interfaces. Through the integration of topological and physical information, the proposed model aims to provide accurate, physically consistent, and computationally efficient predictions of laminate-scale damage evolution.

3. PROPOSED RESEARCH PLAN

The PhD research began in September 2025 and has made substantial progress during its first year. To date, efforts have focused on the development of the composite sandwich-panel dataset and the damage localization framework of WP1. Completion of the first WP is planned for September 2026, after which the focus will shift to the development of microscale surrogate models for progressive damage simulation in WP2. The third WP focuses on topology-aware modeling of laminate-scale damage evolution, and is scheduled to begin in September 2027. The overall research plan follows a structured three-year progression from damage localization, to microscale damage modeling, and finally to laminate-scale damage propagation.

3.1. Current Progress and Preliminary Results

To demonstrate the feasibility of the proposed research direction, we present the preliminary results from WP1. As a first step, we developed a new dataset for damage localization and severity assessment in composite aerospace structures. The dataset is based on a finite-element model of a sandwich-panel payload fairing measuring $800\text{ mm} \times 1020\text{ mm}$, and containing a circular access-door cutout. The structure was simulated under axial compression for one undamaged and more than one thousand damaged configurations. The considered damage types include scratches, wrinkles, and cracks of varying sizes in the facesheet. These defects were modeled as localized stiffness reductions through a 99 % reduction in the local facesheet thickness. To emulate a realistic monitoring setup, strain measurements were extracted at 32 sparsely distributed sensor locations, and to account for measurement uncertainty, independent Gaussian noise with a standard deviation of $15\text{ }\mu\text{S}$ was added to the strain signals.

To localize damage from sparse strain measurements, we de-




Figure 1. Damage localization errors without noise (left) and with sensor noise (right).

veloped a GNN architecture using message passing to capture localized strain distributions caused by damage. Rather than encoding absolute sensor coordinates as node features, we represent them through relative edge features, encouraging the model to learn physically meaningful spatial relationships. Damage localization is then performed through a spatially aware pooling operation that aggregates information across the sensing network. To evaluate the model's ability to generalize beyond the training distribution, an entire region in the lower right quadrant of the panel was excluded from training and used exclusively for testing. Despite this challenging setup, the proposed model achieved a mean localization error of 4.66 cm and consistently outperformed baseline architectures when predicting damage locations in previously unseen regions under noise-free conditions, as shown in Fig. 1. These results indicate that graph-based representations are a promising direction for achieving spatial generalization in structural health monitoring.

The experiments also revealed an important limitation: performance degrades significantly in the presence of realistic sensor noise, particularly for very small defects. Further analysis showed that these damages often produce strain signatures that are comparable in magnitude to the measurement noise, making reliable localization difficult. This finding directly informed the next stage of the research. Specifically, we plan to introduce a damage-severity classification stage prior to localization. Since very small defects are often not operationally critical, grouping damage into severity categories can improve robustness, reduce false alarms, and better align the framework with maintenance decision-making requirements (Saxena et al., 2010).

Building on these encouraging initial results, the next phase of WP1 will focus on incorporating a physics-based plausibility model. This forward model will predict the expected strain measurements associated with a candidate damage location and severity. By coupling the localization network with this plausibility model during training, we aim to intro-

duce an additional physical consistency constraint that can reject implausible predictions and further improve robustness and generalization to unseen damage scenarios.

ACKNOWLEDGMENT

This research was funded by the Swiss State Secretariat for Education, Research and Innovation (SERI) under the project "Automated online monitoring of smart composite structures (ASSESS)". The authors would like to thank Beyond Gravity for providing the data for parts of this research.

REFERENCES

- Bergmayr, T., Höll, S., Kralovec, C., & Schagerl, M. (2023). Local residual random forest classifier for strain-based damage detection and localization in aerospace sandwich structures. *Composite Structures*, 304, 116331.
- Del Priore, E., Sbarra, R. G., Coppotelli, G., Lampani, L., & Pasquali, M. (2025). Damage Identification in Aeronautical Structures Using Strain Mode Shapes and Automatic Operational Modal Analysis. *Journal of Vibration Engineering & Technologies*, 13(8), 567.
- Garban, M. R. G., Chemisky, Y., Prulière, É., & Clément, M. (2025). *Physics-Informed Graph Neural Networks to Reconstruct Local Fields Considering Finite Strain Hyperelasticity*. arXiv.
- Ghosh, G., Biswas, D., & Bhattacharyya, R. (2025). Advancements in multiscale modeling of damage in composite materials: A comprehensive review. *Composites Part B: Engineering*, 307, 112819.
- Hinton, M., & Kaddour, A. (2012). The background to the Second World-Wide Failure Exercise. *Journal of Composite Materials*, 46(19-20), 2283–2294.
- Iparraguirre, M. M., Alfaro, I., Gonzalez, D., & Cueto, E. (2026). *MeshGraphNet-Transformer: Scalable Mesh-based Learned Simulation for Solid Mechanics*. arXiv.
- Mylonas, C., Tsialiamanis, G., Worden, K., & Chatzi, E. N. (2023). *Bayesian graph neural networks for strain-based crack localization*. arXiv.
- Ogunleye, R. O., Rusnáková, S., Javořík, J., Žaludek, M., & Kotlánová, B. (2024). Advanced Sensors and Sensing Systems for Structural Health Monitoring in Aerospace Composites. *Advanced Engineering Materials*, 26(22), 2401745.
- Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., & Battaglia, P. W. (2021). LEARNING MESH-BASED SIMULATION WITH GRAPH NETWORKS.
- Saxena, A., Roychoudhury, I., Celaya, J., Saha, S., Saha, B., & Goebel, K. (2010). Requirements Specification for Prognostics Performance - An Overview. In *AIAA Infotech@Aerospace 2010*.
- Singh, T., & Sehgal, S. (2022). Structural Health Monitoring of Composite Materials. *Archives of Computational*

Methods in Engineering, 29(4), 1997–2017.

- Wei, A., & Fink, O. (2025). Integrating physics and topology in neural networks for learning rigid body dynamics. *Nature Communications*, 16(1), 6867.
- Yacouti, M., & Shakiba, M. (2025). Integrated convolutional and graph neural networks for predicting mechanical fields in composite microstructures. *Composites Part A: Applied Science and Manufacturing*, 190, 108618.
- Yang, Z., Pei, C., Yan, H., & Long, L. (2020). Fatigue damage modeling of ceramic-matrix composites: A short review. *Material Design & Processing Communications*, 2(2), e129.
- Zhou, Y., & Semnani, S. J. (2025). A machine learning based multi-scale finite element framework for nonlinear composite materials. *Engineering with Computers*, 41(5), 2795–2831.