

Foundation Model-Assisted Structural State Estimation for Smart Composite Laminates under Uncertain Environments

Jesse C. Young

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,
OR, USA.

jesseyoung95@oregonstate.edu

Siyuanran Xia

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.

xia704@uc.edu

Kabir D. Ghosh

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.

kabir1988@missouri.edu

Abstract

The increasing deployment of smart composite laminates in aerospace, energy, and civil infrastructure demands robust structural state estimation capable of coping with severe environmental and operational uncertainties. Traditional physics-based model updating and purely data-driven deep learning approaches each exhibit brittleness under distribution shift, motivating the search for more adaptive architectures. This paper investigates the integration of foundation models—large-scale pretrained time-series models—into a system-level framework for real-time strain, damage, and damping state inference in laminated composite structures instrumented with embedded sensor networks. We articulate a hybrid architecture that couples a foundation model backbone with physics-informed priors, domain-specific fine-tuning, and edge-cloud orchestration to reconcile representational generality with task-specific precision. The analysis foregrounds structural trade-offs including latency versus accuracy, data provenance and sensor heterogeneity, uncertainty quantification, and the tension between centralized model governance and on-device adaptation. The discussion extends to socio-technical dimensions, examining how foundation model-assisted estimation reshapes certification pathways, algorithmic fairness across heterogeneous structural populations, and the regulatory implications of automated decision-making in safety-critical systems. By centering the system-level perspective rather than algorithmic novelty, we provide a comprehensive roadmap for deploying foundation model capabilities in physically-grounded, uncertainty-aware infrastructure health management.

Keywords

structural health monitoring, smart composite laminates, foundation models, uncertainty quantification, digital twin, edge intelligence.

1. Introduction

Smart composite laminates are transforming the design of lightweight, high-performance structures by embedding active and passive sensing materials directly into the load-bearing plies. The ability to monitor internal strain distributions, interfacial delamination, and viscoelastic damping evolution in real time holds the promise of condition-based maintenance, extended service life, and reduced certification conservatism. At the core of this vision lies structural state estimation, the continuous inference of spatiotemporal mechanical fields from distributed sensor streams. Substantial progress has been made through physics-based finite element model updating strategies, which align numerical predictions with experimental modal data, and through data-driven methods that learn discriminative features directly from vibration or guided wave responses. Research has systematized the machine learning perspective on structural health monitoring, establishing that damage-sensitive features can be learned end-to-end [1]. Concurrently, the maturation of piezoelectric wafer active sensors has furnished low-profile, embeddable transduction elements capable of actuating and sensing Lamb waves across large areas, forming the hardware substrate for intelligent laminates [2]. Deep convolutional architectures have demonstrated impressive accuracy in vibration-based damage detection, yet their reliance on stationary training distributions renders them susceptible to performance degradation when material variability, thermal drift, or unmodeled loading scenarios change the sensor signatures [3].

This brittleness under distribution shift has motivated a broader rethinking of the modeling paradigm. Rather than training isolated task-specific models from scratch for each structural configuration and operational regime, a new wave of research explores the extraction of universal temporal representations through pretraining on massive, heterogeneous time-series corpora. Foundation models such as Lag-Llama have shown that a single pretrained probabilistic forecaster can achieve competitive zero-shot and few-shot performance across domains ranging from finance to energy to traffic [4]. The underlying insight—that large-scale autoregressive pretraining over diverse sequences can distill generic dynamics patterns—holds transformative potential for structural state estimation. By leveraging a pretrained backbone, downstream tasks on smart laminates could benefit from transferred knowledge of oscillatory modes, trend decomposition, and anomaly signatures learned from millions of unrelated time series, drastically reducing the data requirements and computational cost of fine-tuning per structure.

Despite this promise, directly porting foundation models into the context of uncertain composite structures introduces deep system-level challenges that go beyond predictive accuracy. The physics of laminated composites is governed by anisotropic constitutive laws, coupling phenomena between mechanical, thermal, and piezoelectric degrees of freedom, and highly localized damage mechanisms that generate subtle waveform changes. Environmental uncertainties—temperature gradients, moisture absorption, stochastic manufacturing imperfections—manifest as non-stationary, heteroscedastic noise in sensor data. The present paper proposes and analyzes a foundation model-assisted architecture for structural state estimation that explicitly addresses these challenges through a layered system design. Rather than advancing a single algorithm, we examine the structural trade-offs, governance frameworks, and deployment strategies necessary to make foundation models trustworthy and operationally viable for smart composite laminates. We build the argument by unfolding the sensing and uncertainty landscape of composite structures, articulating the foundation model paradigm for structural time-series, detailing a hybrid system architecture, and discussing robustness, fairness, and policy implications that arise when general-purpose pretraining intersects safety-critical physical systems.

2. Smart Composite Laminates and Sensing Infrastructure under Uncertainty

Smart composite laminates with integrated piezoelectric patches and viscoelastic layers have been modeled using detailed finite element discretizations combined with experimental modal analysis-based updating to capture damping properties and sensor-structure electro-mechanical coupling [5]. These physics-based digital replicas can predict how guided wave velocity, resonance frequencies, and impedance spectra shift with damage onset. However, the fidelity of such models is bounded by the extent to which boundary conditions, material parameters, and sensor degradation can be accurately characterized in the field. The challenge is compounded by the intrinsic variability of composite material systems, where fiber volume fraction, void content, and inter-laminar fracture toughness can exhibit coefficients of variation exceeding ten percent even within a single batch.

To enable continuous state awareness, engineered sensor networks such as the SMART Layer technology have been developed, embedding a dense array of piezoelectric elements between composite plies to provide in-situ ultrasonic and electromechanical impedance measurements [6]. These networks generate rich multi-modal data streams, but they are simultaneously sensitive to operational and environmental fluctuations that obscure the deterministic signatures sought by physics-based models. Uncertainty quantification in composites has therefore emerged as an essential discipline, addressing both aleatoric sources—inherent randomness in material properties and loading—and epistemic sources—model form errors and limited experimental data [7]. Classical stochastic finite element methods and perturbation techniques have been applied to propagate input uncertainties through laminated shell models, yet the computational expense limits their utility for real-time inference.

The central tension is that the very richness that makes smart laminates a compelling monitoring platform also creates a high-dimensional, non-stationary input space that exceeds the generalization capacity of conventional data-driven models trained on narrow laboratory conditions. When a structural component transitions from a climate-controlled hangar to a high-altitude flight environment with temperature swings exceeding eighty degrees Celsius, the relationship between sensor impedance and underlying stiffness can shift enough to trigger false alarms. This observation motivates a shift from isolated spectral feature extraction toward representations that absorb environmental context as an integral part of the learned latent space.

3. Foundation Model Paradigm for Structural Time-Series Analysis

Foundation models for time series depart from the tradition of bespoke model design by pretraining autoregressive or denoising transformations over vast collections of sequences, yielding reusable representations that encode generic dynamical motifs. Lag-Llama, for example, trains a transformer architecture on a corpus of diverse univariate and multivariate time series, learning to predict probability distributions over future tokens conditioned on lagged observations [4]. In parallel, models such as TimeGPT-1 have demonstrated that large-scale pretraining on public and proprietary time-series datasets can deliver zero-shot forecasting accuracy that rivals domain-specific statistical models [8]. These advances suggest that a pretrained foundation model could serve as a universal backbone for structural state estimation, capturing transient oscillation patterns, seasonal sensor drift, and slow damage accumulation without requiring an explicit finite element model for every monitored structure.

Transferring such models to the domain of smart composite laminates, however, demands careful consideration of the semantic gap between generic time-series objectives and the physics of structural diagnostics. Foundation models pretrained predominantly on economic indices, weather records, and web traffic logs learn representations that emphasize trend, seasonality, and abrupt spikes. Acoustic emission events, Lamb wave dispersion, and nonlinear modulation effects in damaged composites exhibit far richer spectral structure governed by partial differential equations. Physics-informed neural networks offer a viable template for bridging this gap by embedding conservation laws, compatibility conditions, and constitutive relations as soft constraints in the learning objective [9]. Integrating such inductive biases into the fine-tuning stage of a foundation model can steer the latent representations toward physically plausible subspaces, reducing the volume of task-specific data needed while preserving the generalization benefits conferred by the broad pretraining.

The architectural choice between attaching a lightweight prediction head to a frozen foundation model backbone versus performing full parameter fine-tuning carries significant implications for deployment on resource-constrained edge devices embedded within composite structures. Full fine-tuning can absorb domain-specific dynamics more thoroughly but risks catastrophic forgetting of the general representations that confer robustness to unforeseen operational conditions. Parameter-efficient fine-tuning methods, such as low-rank adaptation, present a middle ground that preserves the bulk of the pretrained weights while allowing task-specific modulation, aligning well with the need for both adaptability and stability in long-lived structural assets.

4. System Architecture for Foundation Model-Assisted State Estimation

We propose a layered architecture that distributes inference across edge, fog, and cloud tiers to reconcile the computational demands of foundation models with the latency constraints of real-time structural monitoring. At the lowest tier, embedded sensor interfaces perform local signal conditioning, analog-to-digital conversion, and preliminary feature extraction using lightweight convolutional or recurrent encoders capable of running on microcontroller-class hardware. These edge encoders produce compressed latent vectors that are transmitted over low-power wireless links to a structural health monitoring gateway installed on or near the asset. The gateway hosts a distilled version of the foundation model backbone, obtained through knowledge distillation from the full pretrained model, enabling sub-second inference for anomaly detection and low-order state estimation. The full-scale foundation model resides in a cloud environment where periodic model updates, cross-structural federated aggregation, and computationally intensive physics-based simulations can be executed. This layered design embodies a deliberate trade-off between data fidelity and system responsiveness: safety-critical functions such as trip-wire alerts for abrupt damage must execute locally with bounded latency, while high-resolution spatial reconstruction and life-cycle prognosis can tolerate the round-trip delay of cloud computation [10].

Data governance within this architecture must contend with the heterogeneity of sensor modalities, sample rates, and metadata quality across a fleet of composite structures. The SMART Layer sensor networks produce guided wave signals sampled at megahertz rates, low-frequency strain gauge measurements, and occasional temperature readings from thermocouples, creating a multimodal data fusion challenge that few pretrained foundation models address natively. A middleware data fabric that normalizes these streams into a unified time-series format with associated contextual tags—including structural identifier, location, manufacturing batch, and maintenance history—is essential to enable effective fine-

tuning and to prevent data cascades where upstream inconsistencies propagate silently into downstream state estimates [13]. Moreover, the curation of representative fine-tuning datasets must explicitly sample across the operational envelope, from cold-soak ground idle to thermal load during climb, to mitigate distributional bias.

Edge intelligence further complicates the architecture because of the inherent tension between local model personalization and global model consistency. Federated learning frameworks can allow each structural asset to improve its own state estimator from local data while periodically sharing gradient updates with the cloud model, preserving data locality and reducing communication bandwidth [14]. However, federated aggregation must account for the non-i.i.d. nature of structural data, where two wings of the same aircraft may experience vastly different load histories, and a naive averaging of parameter updates could erase the specialized knowledge required for accurate per-component diagnostics.

5. Robustness, Trade-offs, and Uncertainty Management

The ultimate measure of a structural state estimation system is its robustness under distributional shifts that arise from material aging, repair, sensor failure, and environmental extremes. Foundation model-assisted architectures can enhance robustness through their capacity to treat environmental context as a conditioning variable rather than a perturbation. By attaching metadata embeddings that capture temperature, humidity, and load regime to the input sequence, the model can learn to disentangle reversible environmental effects from irreversible damage evolution. Transfer learning strategies that pre-adapt a foundation model on a corpus of physics simulation data—generated by solving stochastic finite element models across a sampled parameter space—provide a further safeguard by exposing the model to physically plausible failure modes before it encounters them in service [15].

Explicit uncertainty quantification remains indispensable when decision thresholds cascade into maintenance actions with flight safety or structural integrity consequences. Bayesian neural networks that characterize both aleatoric and epistemic uncertainty have been combined with structural vibration features to provide prediction confidence intervals that can be used to gate automated alerts [16]. When a foundation model is fine-tuned in a Bayesian framework, the posterior distribution over its parameters captures the residual uncertainty arising from limited structural health data, enabling risk-aware decision making. This approach must be complemented by material-level uncertainty budgets that account for the documented variability in composite laminate properties [17]. The joint propagation of sensor noise, model uncertainty, and material variability through the estimation pipeline can be assessed via Monte Carlo dropout or deep ensembles at inference time, providing a multidimensional confidence map over the estimated strain and damage fields.

The trade-off between model generality and specialization emerges as a central design tension. A highly generic foundation model that has seen millions of time-series patterns may produce smooth, physically plausible predictions even under extreme sensor dropout, but it may miss the subtle precursor signatures of matrix micro-cracking that precede delamination in a specific laminate layup. Conversely, a heavily fine-tuned model can achieve outstanding accuracy on its target structure while losing the robustness to novel load cases that the pretraining was intended to confer. This trade-off must be managed through a model versioning and updating strategy that periodically re-anchors the fine-tuned model to a recently retrained foundation model, ensuring that advances in pretraining methodology and expansion of training corpora are harvested without destabilizing in-service performance.

6. Governance, Fairness, and Deployment Policy

Deploying foundation model-assisted state estimation in safety-critical infrastructures raises governance challenges that extend well beyond technical validation. The black-box nature of large transformer backbones conflicts with the certification requirements of airworthiness authorities and civil infrastructure codes, which demand transparency in the causal chain linking observable evidence to predicted structural states. Explainability techniques tailored to time-series foundation models—such as attention map analysis and input perturbation sensitivity—must be matured and standardized before regulatory bodies can accept model outputs as evidence in continued airworthiness decisions. The fusion of data-driven and physics-based approaches presents an opportunity for auditable hybrid reasoning, where a symbolic physics layer flags implausible model predictions and provides a fallback conservative estimate, thereby creating a certifiable safety envelope [20].

Fairness considerations enter the picture when foundation model-assisted systems are deployed across a fleet of structures with heterogeneous design, age, and maintenance histories. If the fine-tuning corpus over-represents newly manufactured composite specimens tested under ideal laboratory conditions, the resulting model may exhibit systematically degraded performance on aging airframes with repaired patches and sensor drift, leading to an inequitable distribution of safety margins across assets. Addressing this fairness dimension requires deliberate curation of training datasets that mirror the operational diversity of the installed base, as well as ongoing monitoring of per-asset prediction residuals to detect emergent bias [18].

The data supply chain for structural state estimation spans original equipment manufacturers, operators, maintenance providers, and third-party analytics platforms, creating a complex web of data ownership, licensing, and liability. Foundational models pretrained on openly available time-series data introduce additional layers of traceability, as the provenance of pretraining data may influence downstream model behavior in unforeseen ways. Emerging regulatory frameworks, notably the European Union’s risk-based approach to artificial intelligence, classify safety-critical infrastructure monitoring as a high-risk application subject to conformity assessments, transparency obligations, and human oversight requirements [19]. Navigating this landscape necessitates a governance framework that documents the entire model lifecycle, from pretraining data selection and fine-tuning procedures to deployment monitoring and decommissioning criteria, in a format auditable by both regulators and insurers.

The institutional dimension includes workforce transition and the evolution of engineering practice. As foundation models automate the low-level feature engineering that structural dynamics specialists traditionally performed, the role of the human expert shifts toward system-level oversight, exception handling, and continuous validation against first-principle models. This transformation demands updated curricula and professional certification pathways that equip engineers with the skills to interrogate, calibrate, and challenge the output of large-scale AI models, preventing automation complacency and preserving the critical human judgment that remains essential in safety-critical domains.

7. Conclusion

This paper has examined the integration of foundation models into the structural state estimation pipeline for smart composite laminates, framing the problem as a multidimensional systems engineering challenge rather than a narrowly defined algorithmic task. We have

argued that the confluence of embedded sensor networks, pretrained time-series models, and physics-informed machine learning opens a viable path toward robust, uncertainty-aware monitoring of composite structures in environments where traditional model updating and conventional deep learning falter. By architecting a tiered deployment that balances edge responsiveness with cloud-scale modeling capacity, and by embedding explicit uncertainty quantification and physics constraints, the proposed framework addresses the most pressing technical risks.

At the same time, the analysis has illuminated deeper structural trade-offs between generalization and specialization, between autonomy and auditability, and between data efficiency and equitable performance across diverse structural populations. The certification of such systems will require a governance apparatus that is currently nascent, demanding concerted efforts to develop explainability standards, fairness auditing protocols, and lifecycle documentation practices for foundation model-assisted physical systems. Future work should pursue large-scale experimental validations on instrumented composite testbeds subjected to realistically varying environmental chambers, as well as the development of open benchmark suites that span diverse layups, damage modes, and operational profiles to drive rigorous comparative evaluation. Ultimately, the value of foundation models in structural health monitoring will be realized not by replacing physics but by enveloping it in a data-driven shell that adapts gracefully to the uncertainties that physical systems inevitably encounter.

References

1. Farrar, C. R., & Worden, K. (2012). *Structural health monitoring: A machine learning perspective*. John Wiley & Sons.
2. Giurgiutiu, V. (2014). *Structural health monitoring with piezoelectric wafer active sensors* (2nd ed.). Elsevier.
3. 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.
4. Rasul, K., Sheikh, A.-S., Schuster, I., Bergmann, U., & Vollgraf, R. (2023). Lag-Llama: Towards foundation models for probabilistic time series forecasting. In *Advances in Neural Information Processing Systems*, 36.
5. 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.
6. Lin, M., Qing, X., Kumar, A., & Beard, S. J. (2001). SMART Layer and SMART Suitcase for structural health monitoring applications. In *Proceedings of SPIE*, 4332, 98–106.
7. Sriramula, S., & Chryssanthopoulos, M. K. (2009). Quantification of uncertainty modelling in stochastic analysis of FRP composites. *Composites Part A: Applied Science and Manufacturing*, 40(11), 1673–1684.
8. Garza, A., & Mergenthaler-Canseco, M. (2023). TimeGPT-1. arXiv preprint, arXiv:2310.03589.
9. 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.

10. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
11. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
12. Worden, K., & Manson, G. (2007). The application of machine learning to structural health monitoring. *Philosophical Transactions of the Royal Society A*, 365(1851), 515–537.
13. Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. M. (2021). “Everyone wants to do the model work, not the data work”: Data cascades in high-stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
14. Bhuiyan, M. Z. A., Wu, J., Weiss, G. M., & Hayajneh, T. (2020). Edge intelligence for the Internet of Things. *IEEE Access*, 8, 69007–69022.
15. Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2018). Transfer learning for time series classification. In *Proceedings of IEEE International Conference on Big Data*.
16. Ye, X., Jin, T., & Chen, Z. (2021). Bayesian neural network-based structural damage detection under uncertainty. *Mechanical Systems and Signal Processing*, 146, 107037.
17. Lekou, D. J., & Philippidis, T. P. (2009). Mechanical property variability in FRP laminates and its effect on failure prediction. *Composites Part B: Engineering*, 40(7), 557–564.
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. Buiten, M. C. (2019). Towards intelligent regulation of artificial intelligence. *European Journal of Risk Regulation*, 10(1), 41–59.
20. Barthorpe, R. J., Worden, K., & Cross, E. J. (2013). On the fusion of physics-based and data-driven approaches within structural health monitoring. In *Topics in Modal Analysis I*, Volume 7 (pp. 183–191). Springer.