

Generative AI-Enhanced Surrogate Modeling for Fast Dynamic Analysis of Smart Laminated Structures with Viscoelastic Damping

Pankaj A. Banerjee

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
contactpankaj@ucf.edu

Sean A. Peters

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
hellosean@missouri.edu

Abstract

The increasing complexity of modern engineered systems demands rapid, high-fidelity structural analysis, particularly for smart laminated composites that incorporate viscoelastic damping materials and distributed sensing-actuation networks. Conventional full-order dynamic simulations, though accurate, are computationally prohibitive for real-time health monitoring, iterative design, and adaptive control. This paper presents a system-level investigation into generative artificial intelligence-enhanced surrogate modeling as a transformative approach to accelerate dynamic analysis of such multi-physics structures. Moving beyond purely computational mechanics, the work examines architectural trade-offs between data-driven generative models—generative adversarial networks and variational autoencoders—and physics-informed training strategies, emphasizing the joint role of accuracy, latency, and generalization. A comprehensive discussion is provided on integrating these surrogates within larger socio-technical infrastructures, including cloud-based digital twins, edge computing platforms, and automated decision pipelines. Governance challenges are examined through the lenses of data provenance, model drift, robustness against adversarial perturbations, and the ethical imperative of equitable access to advanced structural intelligence. Additionally, the paper addresses sustainability outcomes, regulatory policy considerations, and the certification hurdles that arise when AI-generated predictions inform safety-critical decisions. By synthesizing insights from smart materials, machine learning, systems engineering, and science and technology policy, this contribution delineates a pathway toward resilient, scalable, and responsible deployment of generative surrogate models for next-generation laminated structures.

Keywords

surrogate modeling; generative adversarial networks; smart laminated structures; viscoelastic damping; dynamic analysis; physics-informed machine learning; system-level governance.

1. Introduction

The design and operation of lightweight, high-performance structures increasingly rely on smart laminated composites that combine fiber-reinforced plies with integrated viscoelastic damping layers and piezoelectric transducers. These systems achieve vibration mitigation and energy dissipation through passive material rheology, while offering active control

capabilities via embedded sensor-actuator networks. Accurate prediction of their dynamic response under variable loading involves solving coupled partial differential equations that capture anisotropic elasticity, frequency-dependent viscoelasticity, and multi-physics electromechanical coupling. Full-scale finite element simulations often demand millions of degrees of freedom and extensive parameter sweeps, rendering them unsuitable for online monitoring, design space exploration, or digital twin synchronization. Addressing this computational bottleneck has motivated the development of surrogate models that approximate the input–output mapping with substantially reduced runtime [1]. However, conventional metamodels such as Gaussian processes or polynomial chaos expansions struggle to capture the high-dimensional, non-stationary, and strongly history-dependent behavior characteristic of viscoelastically damped smart laminates.

The advent of deep generative artificial intelligence, particularly generative adversarial networks (GANs) [2], provides a paradigm shift. Unlike discriminative surrogates that predict scalar or vector outputs, generative models learn the joint data distribution, enabling the synthesis of entire response fields and the quantification of prediction uncertainty through stochastic sampling. Physics-informed neural networks further augment this capability by embedding governing conservation laws as soft constraints during training, constraining the solution manifold even when labeled data are sparse [3]. These advances open the possibility of constructing surrogates that are fast enough for real-time inference, yet resilient to operating conditions not represented in the training set. The present study moves beyond algorithm development to interrogate the broader system-level implications of embedding generative AI surrogates within the lifecycle of smart laminated structures. We examine how such models transform system architecture, infrastructure requirements, governance frameworks, and sustainability outcomes. By treating the surrogate not as an isolated numerical tool but as a component of a larger socio-technical assembly, we identify critical trade-offs and propose directions for responsible implementation.

2. Background and Related Work

The analysis of laminated structures with viscoelastic damping has a rich history spanning continuum mechanics, finite element methods, and experimental identification. Early efforts established the modeling foundations for frequency-dependent viscoelastic constitutive laws and their incorporation into layered shell and beam theories, enabling the prediction of natural frequencies, loss factors, and time-domain relaxation behavior [4]. Concurrently, the integration of piezoelectric materials as distributed sensors and actuators was formalized through coupled electromechanical laminate theories, providing a basis for active vibration control [5]. Classical forced vibration analyses of three-layer damped sandwiches offered benchmark solutions that revealed the parametric sensitivity of damping efficacy to core thickness, material loss moduli, and boundary conditions [6]. To mitigate the large computational cost of fully coupled simulations, model order reduction methods based on component mode synthesis or Krylov subspace projection were introduced, preserving dominant dynamic features while alleviating runtime demands [7].

As the complexity of smart composite systems grew, the need for efficient surrogate models became pressing. Data-driven metamodeling initially centered on Gaussian processes and radial basis functions that interpolate simulation outputs across a design space. To handle high-dimensional uncertainty propagation, deep neural network surrogates were later developed, demonstrating tractable quantification of probabilities for expensive-to-evaluate models [8]. An important milestone for the present context was the development of a

numerical framework for model updating of smart laminated structures with viscoelastic damping that systematically calibrated finite element parameters against experimental modal data, enhancing predictive fidelity while acknowledging inherent material variability [9]. This line of inquiry underscored the necessity of combining physics-based models with measurement-driven adjustments, precisely the junction where generative surrogates can excel.

Generative adversarial networks were subsequently adapted to emulate complex physical systems by learning to sample from a distribution of displacement or pressure fields conditioned on input parameters, as demonstrated in fluid turbulence [10]. Alongside GANs, variational autoencoders provided a complementary probabilistic formulation, facilitating latent space interpolation and smooth traversals across structural configurations. Within the smart structures community, active constrained layer damping strategies were refined to blend passive viscoelastic dissipation with piezoelectric actuation, yielding designs that demand real-time control logic informed by rapid dynamic prediction [11]. Composite science contributed systematic damping characterization techniques using modal analysis, highlighting the interplay between lamination sequence, fiber orientation, and viscoelastic layer placement [12]. Stochastic model updating emerged as a rigorous method to incorporate parametric uncertainty into finite element models, moving beyond deterministic calibration toward statistically meaningful predictions [13]. Parallel developments in surrogate methodology showed that aggregation of multiple predictors through cross-validation-based selection could outperform any single metamodel [14], foreshadowing ensemble strategies of generative AI. Meanwhile, generative design frameworks that amalgamate topology optimization with deep generative models demonstrated how AI can autonomously propose novel structural configurations satisfying multiple performance criteria [15]. These diverse threads converge on the central proposition examined in this paper: generative AI-enhanced surrogates can fundamentally transform dynamic analysis of smart laminated structures by providing high-speed, uncertainty-aware response emulation that integrates with system-wide decision processes.

3. Generative AI-Enhanced Surrogate Modeling Architecture

The envisioned surrogate modeling architecture comprises three interconnected subsystems: a high-fidelity simulation engine for data generation, a generative model core that learns the mapping from design parameters and excitation conditions to evolution time-histories, and a deployment interface that exposes the surrogate to downstream consumers such as control algorithms, digital twins, and engineering dashboards. Training data are synthesized by executing a parametric sweep of a full-order finite element model that resolves ply-level kinematics, viscoelastic fractional-derivative or Prony-series representations, and linearized piezoelectric coupling. Each forward simulation delivers multi-channel temporal outputs—displacement, strain, and electric potential—sampled at thousands of nodes. This dataset serves as the empirical foundation on which the generative model is trained.

The generative core is architected as a conditional deep generative network, typically a conditional GAN or a conditional variational autoencoder [22]. In both paradigms, the generator network maps a latent random vector concatenated with a conditioning vector—comprising layup thickness, damping material loss factor, excitation frequency, or piezoelectric shunting circuit parameters—to a high-dimensional snapshot of the structural response. The discriminator or variational decoder evaluates fidelity, driving the generator to produce outputs indistinguishable from the finite element generated samples. Physics-

informed regularization is integrated by adding a penalty term that measures the residual of the governing elastodynamic and charge balance equations evaluated on the generated field. This hybrid loss ensures that the surrogate respects first principles even in sparsely sampled regions of the parameter space, dramatically reducing the amount of training data required relative to purely black-box models.

A critical architectural trade-off centers on the granularity of the surrogate. At one extreme, a single generative model can be trained across the entire operational envelope, requiring vast memory and extensive hyperparameter tuning but offering maximum generality. At the other extreme, an ensemble of specialized surrogates—each responsible for a limited sub-domain—can be aggregated, trading off increased orchestration complexity for improved local accuracy and reduced individual model size. Such aggregation strategies echo earlier findings from multiple-surrogate ensembles [14] and become particularly advantageous when deploying surrogates on edge devices with restricted computational resources. The choice of latent space dimensionality additionally governs the balance between expressiveness and overfitting; a low-dimensional latent space promotes smooth interpolations but may fail to capture multi-modal response distributions induced by bifurcations or mode switching.

Training a generative surrogate for viscoelastic systems presents unique difficulties because the response is both frequency and time dependent. Strategies such as concatenating the loading history into the conditioning vector or employing recurrent generator architectures with temporal attention are under active exploration. Furthermore, the inherent stochasticity of GAN training introduces convergence challenges that can be mitigated through careful application of Wasserstein loss formulations, spectral normalization, and gradient penalty mechanisms. Verification of the trained surrogate must include not only global error metrics such as the relative L2 norm of the displacement field but also the preservation of engineering quantities of interest: resonant frequencies, modal loss factors, and peak strain magnitudes. When these criteria are satisfied, the surrogate can accurately infer the full dynamic state in milliseconds, enabling responsive decision-making previously unattainable.

4. System-Level Integration and Infrastructure

Embedding generative AI surrogates within operational environments demands a holistic infrastructure that connects physical assets, simulation resources, data pipelines, and user-facing applications. A prevalent paradigm is the digital twin, which maintains a living computational replica of the physical structure, continuously updated through sensor streams [17]. The surrogate becomes the analytical engine inside the twin, producing on-demand predictions of how the structure will respond to anticipated loads or hypothetical damage scenarios. Cloud-based intelligent simulation platforms provide the elastic computational fabric required for initial model training and periodic retraining, as well as for orchestration of surrogate instances that serve multiple physical assets [18]. The architecture must address data provenance rigorously: all training runs, surrogate version artifacts, and inference results must be logged immutably to support auditability and regulatory compliance.

Edge computing introduces a further layer of deployment trade-offs. Within large civil infrastructure or aircraft platforms, network latency and bandwidth constraints may preclude round-trip communication to a cloud service for every control cycle. Lightweight surrogate models can be compressed via quantization, pruning, and knowledge distillation, then hosted on embedded systems co-located with piezoelectric sensor interrogators. This configuration enables real-time structural health monitoring where the surrogate rapidly evaluates the residual between predicted and measured response to detect anomalies, trigger alarms, or

adjust active damping parameters [21]. Governance of such distributed intelligence necessitates synchronization protocols so that edge-deployed models remain consistent with central infrastructure, and graceful degradation strategies are activated when connectivity is lost.

Interoperability is another system-level concern. The surrogate must consume design information from computer-aided engineering environments, sensor databases, and enterprise resource planning systems. Standardized interfaces, such as the Functional Mock-up Interface or open APIs built around representational state transfer principles, facilitate the plug-and-play integration of surrogate models into broader engineering workflows. From a life-cycle perspective, the surrogate evolves as the structure ages: changes in viscoelastic properties due to temperature, humidity, and material degradation can be tracked through model updating loops that periodically re-identify surrogate conditioning parameters using measured operational modal data. This closed-loop fusion of physics-based modeling, data-driven learning, and in-service monitoring elevates the structure to a cognitively aware state, capable of self-diagnosis and adaptive mitigation.

5. Governance, Robustness, and Fairness Considerations

The displacement of physics-based analysis with AI-generated predictions in safety-critical smart structures raises profound governance questions. Any surrogate model inherits the biases present in its training data, whether arising from incomplete sampling of the operational envelope, simulation idealizations, or systematic measurement errors. Bias can lead to overconfident predictions for underrepresented laminate configurations or loading scenarios, placing structural integrity at risk. A comprehensive fairness framework, as advanced in machine learning scholarship, must be operationalized in the engineering domain to ensure that no operational condition is disproportionately neglected [16]. This entails evaluating surrogate performance across well-defined subpopulations defined by excitation type, environmental severity, or geometric scale, and enforcing egalitarian accuracy benchmarks. The societal dimension of fairness extends to the equitable distribution of advanced structural intelligence, avoiding a scenario where only high-value assets in affluent regions benefit from real-time AI-driven monitoring, while critical infrastructure in less resourced settings remains reliant on outdated inspection regimes.

Robustness to adversarial perturbations is an emergent challenge. In a generative surrogate, small, carefully crafted modifications to input parameters could in principle produce catastrophically erroneous response predictions, a vulnerability well-known in deep learning. Defending against such exploits requires adversarial training during surrogate development and runtime detection of out-of-distribution inputs using density estimation in the latent space or disagreement among ensemble members. Equally important is robustness to natural distribution shift, as structural aging and environmental exposure progressively modify the true dynamic behavior away from the training distribution. Continuous monitoring of surrogate prediction residuals via control charts with adaptive thresholds provides a statistical process control layer that triggers model re-identification well before prediction quality degrades below safety margins.

Uncertainty quantification emerges as the linchpin of trustworthy deployment. Surrogates must deliver not merely point estimates but well-calibrated prediction intervals. Generative models inherently offer sampling-based uncertainty, yet ensuring that these predictive distributions match empirical frequencies demands rigorous post-hoc calibration, for instance through conformal prediction or Bayesian meta-learning [19]. The system architecture must

propagate uncertainty downstream, so that a control system's confidence in a damping correction is commensurate with the surrogate's predictive confidence. Regulatory frameworks, still nascent, will likely evolve to require third-party validation of surrogate models using held-out physical experiments, akin to the certification process for aircraft software. Policy instruments such as AI impact assessments for infrastructure projects could mandate documentation of surrogate training data, bias audits, and robustness testing before deployment, embedding ethical reflection into engineering practice [20].

6. Sustainability and Policy Implications

From a sustainability standpoint, generative AI surrogates reduce the carbon footprint of structural analysis by compressing months of high-performance computing into minutes of inference on much lighter hardware. This energy efficiency is particularly consequential when scaled across fleets of wind turbine blades, aircraft components, or bridge girders that would otherwise each be subjected to repeated full-order simulations throughout their life cycle. Moreover, the ability to accurately predict damping performance and its degradation enables life-extension strategies that defer material-intensive retrofits, conserving resources and reducing waste. When integrated with model updating routines, the surrogate guides maintenance crews to apply targeted viscoelastic patch repairs only where loss factor deterioration crosses a critical threshold, minimizing unnecessary intervention.

Policy frameworks must be adapted to harness these benefits while mitigating risks. Standards bodies such as the International Organization for Standardization and national engineering societies should develop best-practice guidelines for validating and maintaining AI surrogates used in civil and mechanical engineering. These guidelines could prescribe minimum data diversity requirements, periodic re-certification intervals, and independent auditing protocols. International coordination is advisable because laminated structures often appear in globally traded products, and inconsistent regulatory approaches could fragment supply chains or create market asymmetries. Research funding agencies might prioritize open-access benchmark datasets for smart laminated structures that encompass various viscoelastic formulations and operating conditions, enabling robust comparison of surrogate methodologies across teams and preventing proprietary lock-in.

Looking ahead, the convergence of generative AI, smart materials, and structural health monitoring suggests a future where structures possess a native cognitive layer that continuously negotiates between energy dissipation, load-bearing, and communication needs. Policy discussions must broaden to encompass questions of accountability: when a surrogate-driven active damping system fails to prevent resonant fatigue damage, liability allocation between the model developer, system integrator, and asset operator becomes a challenging legal frontier. Proactive establishment of liability frameworks, possibly through insurance-linked model certification, will reduce uncertainty and encourage responsible innovation. Education and workforce development policies will equally need to expand engineering curricula to include data governance, algorithmic fairness, and AI system safety, ensuring that the next generation of structural engineers can critically interrogate rather than blindly trust black-box surrogates.

7. Conclusion

This paper has presented a system-level interdisciplinary examination of generative AI-enhanced surrogate modeling for the fast dynamic analysis of smart laminated structures with viscoelastic damping. By synthesizing insights from structural mechanics, deep generative

learning, infrastructure design, and technology governance, it mapped the architectural choices, integration strategies, and socio-technical implications that accompany the introduction of such surrogates into engineering practice. The proposed architecture leverages conditional generative models with physics-informed regularization, deployed across cloud and edge environments within digital twin frameworks, to deliver rapid, uncertainty-aware predictions that enable real-time control and adaptive maintenance. Critical trade-offs between model generality and specialization, latent space dimensionality and expressiveness, and centralized versus distributed inference were identified, each with direct impacts on scalability and robustness.

Moving beyond algorithmic performance, the analysis foregrounded governance imperatives encompassing fairness, adversarial robustness, uncertainty calibration, and certification. These imperatives are not secondary concerns but foundational prerequisites for safe and socially equitable adoption of AI in critical structural systems. The sustainability gains—both in reduced computational energy consumption and extended asset service life—are significant but require supportive policy mechanisms, including open benchmarking and standardized validation protocols. As the technologies mature, interdisciplinary collaboration among engineers, computer scientists, ethicists, and policymakers will be essential to steward generative surrogates toward outcomes that are technically sound, ethically defensible, and broadly beneficial. The convergence of smart materials and artificial intelligence thus heralds not merely an acceleration of numerical analysis, but a paradigm shift in how societies design, maintain, and govern the built environment.

References

1. Simpson, T. W., Poplinski, J. D., Koch, P. N., & Allen, J. K. (2001). Metamodels for computer-based engineering design: survey and recommendations. *Engineering with Computers*, 17(2), 129–150.
2. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
3. 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.
4. Trindade, M. A., Benjeddou, A., & Ohayon, R. (2000). Modeling of frequency-dependent viscoelastic materials for active-passive vibration damping. *Journal of Vibration and Acoustics*, 122(2), 169–174.
5. Lam, K. Y., Peng, X. Q., Liu, G. R., & Reddy, J. N. (1997). A finite-element model for piezoelectric composite laminates. *Smart Materials and Structures*, 6(5), 583–591.
6. Mead, D. J., & Markus, S. (1969). The forced vibration of a three-layer, damped sandwich beam with arbitrary boundary conditions. *Journal of Sound and Vibration*, 10(2), 163–175.
7. Balmès, E. (1996). Parametric families of reduced finite element models: theory and applications. *Mechanical Systems and Signal Processing*, 10(4), 381–394.

8. Tripathy, R. K., & Bilonis, I. (2018). Deep UQ: Learning deep neural network surrogate models for high dimensional uncertainty quantification. *Journal of Computational Physics*, 375, 565–588.
9. 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.
10. Ghorbanian, P., & Bagheri, S. (2022). Generative adversarial networks for surrogate modeling of complex physics: Applications to turbulence. *Journal of Computational Physics*, 448, 110740.
11. Baz, A., & Ro, J. (1995). Optimum design and control of active constrained layer damping. *Journal of Vibration and Acoustics*, 117(1), 25–32.
12. Berthelot, J. M., & Sefrani, Y. (2004). Damping analysis of composite materials and structures. *Composite Structures*, 65(3–4), 421–432.
13. Mares, C., Mottershead, J. E., & Friswell, M. I. (2006). Stochastic model updating: Part 1—Theory and simulated example. *Mechanical Systems and Signal Processing*, 20(7), 1674–1695.
14. Viana, F. A. C., Haftka, R. T., & Steffen, V. (2009). Multiple surrogates: how cross-validation errors can help us to obtain the best predictor. *Structural and Multidisciplinary Optimization*, 39(4), 439–457.
15. Oh, S., Jung, Y., Kim, S., Lee, I., & Kang, N. (2019). Deep generative design: Integration of topology optimization and generative models. *Journal of Mechanical Design*, 141(11), 111405.
16. 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), Article 115, 1–35.
17. 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.
18. Guo, L., Wang, Y., Yang, Z., Liu, Q., & Gao, J. (2021). A cloud-based intelligent simulation platform for multi-disciplinary collaborative design. *Advanced Engineering Informatics*, 50, 101399.
19. Marrel, A., Iooss, B., Laurent, B., & Roustant, O. (2009). Calculations of Sobol indices for the Gaussian process metamodel. *Reliability Engineering & System Safety*, 94(3), 742–751.
20. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
21. Giurgiutiu, V. (2005). Tuned lamb wave excitation and detection with piezoelectric wafer active sensors for structural health monitoring. *Journal of Intelligent Material Systems and Structures*, 16(4), 291–305.
22. Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.