Automated Structural Health Monitoring of FRP-Confined Columns Using Deep Learning and Multimodal Sensor Data

Tobiasz Spychaj
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
spychaj287@uc.edu

Abstract

Fiber-reinforced polymer (FRP) confinement is an effective technique for strengthening reinforced concrete columns, but the resulting composite-concrete system can degrade through matrix cracking, fiber rupture, debonding, and concrete crushing that are often poorly observable from external inspection. Structural health monitoring (SHM) of such columns therefore requires more than periodic visual assessment or a single sensor type. This manuscript develops a methodological framework for automated SHM of FRP-confined columns that combines physics-informed feature construction, multimodal sensing, deep learning representation learning, and uncertainty-aware decision logic. The framework integrates hoop strain, acoustic emission, vibration response, and environmental measurements into a synchronized observation model; converts confinement mechanics into damage-sensitive features; and uses modality-specific encoders with cross-modal attention to produce a fused health representation. A training objective combines supervised damage-state estimation with physically motivated regularization, while Monte Carlo dropout provides an explicit predictive-distribution estimate for alert decisions. The central contribution is not a claim of demonstrated field performance but a coherent, testable architecture that connects confinement mechanics, sensor physics, deep learning fusion, and maintenance thresholds in one SHM pipeline. The framework clarifies how automated monitoring can distinguish progressive confinement damage from environmental and operational variability, and it identifies the calibration, transferability, and validation conditions that must be satisfied before deployment on instrumented columns.

Keywords

FRP-confined columns; structural health monitoring; multimodal data fusion; deep learning; uncertainty quantification; automated damage assessment

1. Introduction

FRP-confined concrete columns are widely used in retrofit and new construction because externally bonded or wrapped composites provide lateral confinement, increase axial load capacity, and improve ductility. Design-oriented stress-strain models describe the interaction between hoop strain in the FRP jacket and the lateral confinement pressure that suppresses concrete dilation [1], and refined formulations account for jacket stiffness, unconfined concrete strength, and ultimate strain limits [2]. Experimental and analytical studies of concrete columns confined by fiber composites established that confinement effectiveness depends on the composite jacket, the concrete core, and the bond condition between them [3]. Broader reviews of FRP applications in civil infrastructure emphasize that durability, inspection, and long-term performance remain central concerns in service [4]. These

characteristics make FRP-confined columns attractive structural elements but also create a monitoring problem: the critical damage state may develop beneath or within the jacket, while visible surface changes remain limited until late-stage rupture or debonding.

Structural health monitoring offers a route to continuous condition assessment, and vibration-based methods have long been used to detect changes in stiffness, boundary conditions, and mass distribution [5]. Machine learning can convert measured signatures into damage-sensitive decisions, particularly when damage-state labels are available or when unsupervised anomaly detection is appropriate [6]. Reviews of vibration-based condition monitoring document both the promise and the practical constraints of modal identification, environmental variability, and sensor placement [7]. Data science and engineering have expanded SHM toward heterogeneous sensing, data management, and automated inference [8]. Deep learning has been applied to vibration-based damage detection through one-dimensional convolutional networks that learn local waveform features directly from acceleration signals [9], and broader reviews trace the transition from traditional modal methods to machine learning and deep learning in civil structures [10]. Next-generation smart sensing further expands the available measurement space through dense, low-cost, and distributed sensing technologies [11].

Despite this progress, automated SHM for FRP-confined columns remains under-specified as an integrated methodology. Single-modality monitoring can be ambiguous: strain measurements reveal jacket expansion but may not localize cracking; acoustic emission detects active damage but is sensitive to threshold selection and background noise; vibration features capture global stiffness loss but can be confounded by temperature, humidity, and boundary conditions. Automated fabrication of in-situ confinement also introduces process and layup variability that monitoring must account for, because fabrication quality influences the as-built confinement condition [12]. A defensible monitoring framework therefore needs to fuse complementary measurements, embed confinement physics, quantify uncertainty, and convert probabilistic outputs into maintenance actions. The present paper develops such a framework. Its central contribution is a deep learning architecture and decision procedure for FRP-confined columns that uses physics-informed multimodal features, attention-based fusion, and uncertainty-aware alerting. The framework is presented as a methodological system rather than as an empirically validated field study; the analysis identifies the conditions under which the approach is physically meaningful, the failure modes that must be controlled, and the validation evidence required before operational use.

2. Physical and Monitoring Basis for FRP-Confined Columns

In a wrapped column, axial compression causes lateral expansion of the concrete core. The FRP jacket resists this expansion through hoop tension, producing a confining pressure that increases with jacket stiffness and hoop strain. A standard idealization relates the lateral confinement pressure to the jacket thickness, elastic modulus, hoop strain, and column diameter. For a circular column with a thin jacket, the first-order relation can be written as

$$f_l = 2t_f E_f \frac{\varepsilon_h}{D}$$

where f_l is the lateral confinement pressure, t_f is the total FRP thickness, E_f is the elastic modulus of the composite in the fiber direction, ε_h is the hoop strain, and D is the column diameter. The expression assumes a thin, linear-elastic jacket, negligible adhesive compliance, and uniform hoop strain around the circumference. These assumptions are useful for defining a physics-derived feature, but they are not exact for all wrapped or bonded

systems: debonding, lap-splice behavior, and nonuniform curvature can make local hoop strain differ from the circumferential average.

Damage in FRP-confined columns develops through interacting mechanisms. Matrix cracking and fiber-matrix debonding may occur before visible fiber rupture; concrete dilation and crushing alter the core stiffness and generate acoustic emission; and jacket debonding or rupture can cause abrupt local strain redistribution. Environmental exposure adds another layer of complexity because temperature and moisture can change resin properties, bond condition, and measured strain. Consequently, the monitoring state is not a single scalar. A useful representation includes at least the confinement utilization reflected by hoop strain, the active damage activity reflected by acoustic emission, the global dynamic response reflected by vibration features, and the environmental context that conditions all three. The physical meaning of each measurement is different, and the value of multimodal fusion lies in exploiting that difference rather than averaging heterogeneous signals into an opaque index.

The monitoring problem can be expressed as a state-estimation and decision problem. Let s(t) denote the latent structural state, including damage level, bond condition, and confinement effectiveness. Each sensor modality m produces a measurement y_m(t) that depends on s(t), sensor-specific parameters θ_m, environmental variables, and noise. The objective is to infer a damage state and a continuous damage index from a finite set of noisy, partially informative observations. This formulation clarifies why purely supervised classification is insufficient on its own: labels for intermediate damage states are difficult to obtain in real columns, and the mapping from sensor data to damage is conditioned by specimen geometry, jacket design, loading history, and exposure. A credible method must therefore combine data-driven learning with physical constraints and explicit uncertainty.

3. Methodological Framework

3.1 Multimodal Acquisition and Synchronization

The proposed framework assumes a column instrumented with three measurement groups. The first group measures strain at discrete locations using electrical resistance strain gauges or fiber optic sensors, with the objective of estimating hoop strain and its spatial variation. The second group measures acoustic emission through piezoelectric transducers, providing hit counts, energy, amplitude, rise time, and frequency content associated with cracking and debonding. The third group measures vibration response through accelerometers or dynamic strain sensors, from which modal frequencies, damping estimates, and time-frequency features are extracted. Environmental sensors measure temperature and relative humidity, and a data acquisition unit provides a common time base. The framework does not require a fixed number of sensors; instead, it requires that each channel be associated with a modality, a location, and a calibration record.

Sensor synchronization is treated as part of the observation model. If the structural state at time t is s(t), the measurement from modality m can be written as

$$y_m(t) = h_m(s(t), \theta) + \varepsilon_m(t) + b_m(t)$$

where h_m is the modality-specific measurement function, θ collects sensor and calibration parameters, ε_m(t) is random measurement noise, and b_m(t) is a slowly varying bias caused by temperature, aging, or installation effects. This separation matters because environmental drift can mimic damage in strain and vibration data. Preprocessing therefore includes timestamp alignment, resampling to a common decision interval, detrending, outlier

rejection, and calibration correction. The output is not a single concatenated vector but a set of synchronized modality windows with quality indicators for missing or saturated channels.

3.2 Physics-Informed Feature Construction

Deep networks can learn hierarchical representations from raw signals [13], but purely data-driven features are vulnerable when training data are limited or when operating conditions differ from those represented in the training set. For FRP-confined columns, physics-informed features reduce this vulnerability by encoding relationships that remain meaningful across specimens. For each time window i, a compact feature vector can include hoop strain, normalized confinement pressure, acoustic emission energy, and modal frequency shift:

$$\varphi_i = \left[\varepsilon_h, i, f_l, \frac{i}{f_c} o, E_A E, i, \Delta \frac{f_i}{f_0}, i\right]$$

Here ε_h,i is the window-averaged hoop strain, f_l,i is the confinement pressure computed from the first relation, f_co is the unconfined concrete compressive strength, E_AE,i is acoustic emission energy in the window, Δf_i is the change in a tracked modal frequency, and f_0,i is the corresponding baseline frequency. The ratio f_l,i/f_co expresses confinement utilization rather than absolute pressure, which is useful when comparing columns with different geometries. The acoustic energy term is not interpreted as a direct damage measure; it is a proxy for active cracking and debonding. The frequency shift term is sensitive to global stiffness change but also to temperature and boundary conditions, so it is paired with environmental variables in the network input rather than treated as a standalone indicator.

These features are not intended to replace learned representations. They are auxiliary channels that constrain the network toward physically plausible behavior. For example, a sequence in which hoop strain rises without acoustic activity or frequency change has a different physical interpretation from a sequence in which acoustic energy spikes while hoop strain remains stable. The network can learn such interactions, but the physics-informed features make the distinctions more explicit and allow the loss function to penalize predictions that contradict confinement mechanics.

3.3 Deep Learning Fusion Architecture

The fusion architecture uses modality-specific encoders followed by cross-modal attention. Each encoder maps a window of measurements from one modality into a latent vector. Temporal convolutional layers capture local waveform patterns, while recurrent or attention-based layers capture longer dependencies. Attention is appropriate because the relevance of a modality depends on the damage state: acoustic emission may dominate during active cracking, strain may dominate during progressive jacket expansion, and vibration may dominate when global stiffness changes. The attention mechanism assigns a normalized weight to each modality at each decision step. A compact form is

$$\alpha_m, i \; = \; softmax_m(u^{T} tanh(V \; h_m, i \; + \; c))$$

where h_m,i is the encoded representation of modality m at window i, V and u are learned projection matrices or vectors, c is a bias vector, and α_m,i is the resulting attention weight. The softmax is taken over modalities, so the weights sum to one at each window. The fused representation is then

$$z_i = \sum_m \alpha_m, iW_m h_m, i$$

where W_m is a modality-specific projection into the shared latent space and z_i is the fused health representation. This construction differs from simple concatenation because it allows the model to down-weight a noisy or uninformative modality without discarding it permanently. Multisensor fusion theory supports this general strategy: heterogeneous measurements should be combined according to their reliability, complementarity, and context rather than treated as interchangeable inputs [15]. The architecture can be implemented with temporal convolutions and self-attention blocks [14], and the attention weights provide a diagnostic signal for checking whether the model is relying on physically implausible channels.

3.4 Training Objectives and Uncertainty Quantification

The model outputs an ordinal damage-state probability and a continuous damage index. The training objective combines task loss, physical regularization, and parameter regularization:

$$L = L_task + \lambda_p\, L_phys + \lambda_r\, ||\theta||^2$$

Here L_task is the sum of a cross-entropy term for damage-state classification and a mean-squared-error term for damage-index regression, L_phys penalizes disagreement between predicted damage progression and physically admissible relations, λ_p and λ_r are nonnegative weighting coefficients, and θ denotes network parameters. The physical term can be defined as a soft penalty on monotonicity, on inconsistency between hoop strain and confinement pressure, or on damage-index changes that are unsupported by acoustic or vibration evidence. The exact form should be selected during calibration and should not be treated as universal. Its purpose is to keep the learned function within the region of behavior consistent with confinement mechanics.

Uncertainty is essential because an automated SHM system must distinguish a confident damage alert from a novel condition that the model has not learned. Predictive uncertainty can be separated into epistemic and aleatoric components [16], and Monte Carlo dropout provides a practical approximation to Bayesian model averaging [17]. With T stochastic forward passes, the predictive distribution is approximated by

$$p(y* | x*, D) \approx \left(\frac{1}{T}\right) \sum_{t=1}^{T} p(y* | x*, \theta_t)$$

where x is a new input window, D is the training data, θ_t is the parameter set obtained by dropout sampling, and y is the damage-state output. The mean of this distribution gives the estimated class probability, while its variance indicates model uncertainty. High predictive variance should suppress automated maintenance decisions and trigger additional inspection or data collection. This design is more defensible than a single deterministic softmax output, because it exposes the difference between a confident prediction and an extrapolation beyond the training distribution.

3.5 Automated Decision Logic

The final stage converts probabilistic outputs into a monitoring action. Let y_i take ordered values from 0 to 3, representing no observable damage, minor damage, moderate

damage, and severe damage. Let D_i be the continuous damage index and A_i be a binary alert. A conservative decision rule is

$$A_i = 1 if \mathrm{P}(y_i \geq 2 | x_i) > \tau and \Delta \frac{D_i}{\Delta} t > \kappa$$

where τ is a probability threshold, κ is a minimum damage-rate threshold, and ΔD_i/Δt is the estimated change in damage index over a specified interval. The rule requires both a high probability of moderate-or-severe damage and a positive trend, which reduces nuisance alarms caused by a single anomalous window. The thresholds τ and κ should be calibrated against inspection outcomes and design requirements, not chosen solely from model statistics. The output of the decision layer is therefore an alert with a timestamp, a fused confidence estimate, the contributing modalities, and the estimated damage trend.

4. Analytical Assessment and Technical Discussion

4.1 Damage Mechanisms and Sensor Complementarity

The framework is physically coherent only if each modality contributes information that is not fully redundant. Fiber optic sensors can provide distributed or quasi-distributed strain measurements with high resolution and immunity to electromagnetic interference, making them suitable for capturing hoop strain redistribution along a jacket [18]. Acoustic emission is sensitive to active microcracking and debonding, and acoustic emission testing has an established basis in concrete and composite structures [19]. Vibration features reflect global stiffness and boundary-condition changes, but they are less sensitive to local damage until the damage affects a significant portion of the structure. The combination is therefore complementary: strain indicates confinement demand; acoustic emission indicates active damage; vibration indicates global consequences; and environmental measurements condition all three. Cross-modal attention allows the model to weight these sources differently as the damage process evolves. For example, early debonding may produce acoustic activity with little global frequency change, while late-stage jacket rupture may produce abrupt strain redistribution and vibration change. A single modality would miss part of this progression.

4.2 Architecture and Data Efficiency

The proposed architecture is more complex than a single-input classifier, and complexity must be justified by data availability and decision requirements. Modality-specific encoders reduce the risk of negative transfer between heterogeneous signals, and attention provides an interpretable weighting mechanism. However, attention weights are not automatically causal explanations; they indicate where the network allocates representational capacity, not why damage occurred. The physics-informed features and regularization term provide a complementary constraint that is especially important when labeled damage states are scarce. In practice, the model should be trained with a hierarchy of objectives: self-supervised reconstruction of sensor windows, supervised classification where labels exist, and physically constrained regression of the damage index. Transfer learning from related structural systems can improve data efficiency, but it must be evaluated carefully because the relationship between acoustic emission and damage in a confined concrete column differs from that in mechanical systems. Reviews of machine learning for machine fault diagnosis show both the value and the limits of transferring diagnostic architectures across operating conditions [20]. The same caution applies here: architectural similarity does not guarantee physical transferability.

4.3 Calibration, Transfer, and Deployment Constraints

Deployment requires calibration at multiple levels. Sensor calibration corrects gain, offset, and drift; model calibration aligns predicted probabilities with observed frequencies; and decision calibration sets τ and κ according to maintenance consequences. Environmental calibration is particularly important because temperature affects strain readings, resin modulus, and modal frequencies. A model trained in one season may learn a spurious association between temperature and damage if temperature is not explicitly represented. The observation model separates random noise from slowly varying bias, but the bias term must be estimated from baseline periods and updated when the structure undergoes legitimate changes such as repairs or load changes. Transfer across columns is also constrained by geometry, jacket thickness, concrete strength, and installation quality. A normalization scheme based on confinement utilization and structural slenderness can reduce, but not eliminate, these differences. The automated fabrication context reinforces this point: as confinement becomes more automated, process variability may decrease for a given system, but the monitoring model must still identify when as-built or in-service conditions depart from the assumed baseline [12].

4.4 Failure Modes and Validation Requirements

The framework has several identifiable failure modes. First, sensor faults can mimic damage: a debonded strain gauge may show an abrupt strain drop, and a loose accelerometer may shift modal estimates. Quality indicators and cross-modal consistency checks are therefore necessary. Second, distribution shift can produce overconfident predictions if the model is deployed on a column with a different jacket layout or loading history. Monte Carlo dropout provides an uncertainty signal, but it is not a substitute for out-of-distribution detection. Third, label scarcity can cause the model to learn environmental correlations rather than damage physics. Physics-informed features and regularization reduce this risk, but they do not remove it. Fourth, the decision thresholds may be poorly matched to inspection intervals or safety consequences. A high τ reduces false alerts but may delay detection; a low τ increases sensitivity but burdens inspectors. These trade-offs should be specified explicitly in the monitoring plan.

Validation should proceed in stages. Component-level validation should confirm that each sensor and preprocessing step produces stable features under controlled environmental variation. Model-level validation should test the architecture on held-out specimens and on temporally separated data, not only on random splits. Decision-level validation should compare alert timing and false-alarm rates against inspection records or known damage states. In the absence of field data, analytical validation can still examine whether the fusion logic is physically consistent, whether uncertainty increases under simulated sensor degradation, and whether the decision rule responds correctly to synthetic monotonic and abrupt damage scenarios. These checks do not replace empirical validation, but they establish the logical and physical conditions that a field study must satisfy.

5. Conclusion

This manuscript has developed a methodological framework for automated structural health monitoring of FRP-confined columns using deep learning and multimodal sensor data. The framework begins with confinement mechanics and defines a physics-derived confinement pressure that links hoop strain, jacket stiffness, and column geometry. It then specifies a synchronized multimodal observation model, physics-informed feature construction, modality-specific deep encoders with cross-modal attention, a composite training objective, and uncertainty-aware alert logic. The central argument is that FRP-

confined columns require fusion of complementary measurements because no single sensor modality captures the progression from matrix cracking and debonding to jacket expansion and global stiffness change. Deep learning provides the representational capacity to learn interactions across these modalities, while physics-informed features and regularization constrain the model toward behavior consistent with confinement mechanics. Uncertainty quantification and a two-condition decision rule reduce the risk of overconfident or transient alerts. The framework is not a substitute for empirical validation; its value lies in making the physical assumptions, data requirements, calibration steps, and failure modes explicit. Before deployment, the model must be tested on instrumented columns with known damage states, evaluated under environmental and operational variability, and calibrated against maintenance consequences. With those conditions satisfied, automated multimodal SHM can support more timely and defensible assessment of FRP-confined columns than periodic visual inspection or single-sensor monitoring alone.

References

1. Lam, L., & Teng, J. G. (2003). Design-oriented stress–strain model for FRP-confined concrete. Construction and Building Materials, 17(6–7), 471–489. https://doi.org/10.1016/S0950-0618(03)00045-X

2. Teng, J. G., Jiang, T., Lam, L., & Luo, Y. Z. (2009). Refinement of a design-oriented stress–strain model for FRP-confined concrete. Journal of Composites for Construction, 13(4), 269–278. https://doi.org/10.1061/(ASCE)CC.1943-5614.0000012

3. Mirmiran, A., & Shahawy, M. (1997). Behavior of concrete columns confined by fiber composites. Journal of Structural Engineering, 123(5), 583–590. https://doi.org/10.1061/(ASCE)0733-9445(1997)123:5(583)

4. Hollaway, L. C. (2010). A review of the present and future utilisation of FRP composites in the civil infrastructure with reference to their important in-service properties. Construction and Building Materials, 24(12), 2419–2445. https://doi.org/10.1016/j.conbuildmat.2010.04.062

5. Farrar, C. R., & Worden, K. (2007). An introduction to structural health monitoring. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 365(1851), 303–315. https://doi.org/10.1098/rsta.2006.1928

6. Worden, K., & Manson, G. (2007). The application of machine learning to structural health monitoring. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 365(1851), 515–537. https://doi.org/10.1098/rsta.2006.1933

7. Carden, E. P., & Fanning, P. (2004). Vibration based condition monitoring: A review. Structural Health Monitoring, 3(4), 355–377. https://doi.org/10.1177/1475921704047500

8. Bao, Y., Chen, Z., Wei, S., Xu, Y., Tang, Z., & Li, H. (2019). The state of the art of data science and engineering in structural health monitoring. Engineering, 5(2), 234–242. https://doi.org/10.1016/j.eng.2018.11.027

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. https://doi.org/10.1016/j.jsv.2016.10.043

10. Avci, O., Abdeljaber, O., Kiranyaz, S., Hussein, M., Gabbouj, M., & Inman, D. J. (2021). A review of vibration-based damage detection in civil structures: From traditional methods to Machine Learning and Deep Learning applications. Mechanical Systems and Signal Processing, 147, 107077. https://doi.org/10.1016/j.ymssp.2020.107077

11. Sony, S., Laventure, S., & Sadhu, A. (2019). A literature review of next-generation smart sensing technology in structural health monitoring. Structural Control and Health Monitoring, 26(3), e2321. https://doi.org/10.1002/stc.2321

12. Kumar, B. B. M., Hsueh, Y. C., Xin, Z., & Luo, D. (2021, July). Process and evaluation of automated robotic fabrication system for in-situ structure confinement. In The International Conference on Computational Design and Robotic Fabrication (pp. 368-379). Singapore: Springer Singapore.

13. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

14. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems 30 (pp. 5998–6008).

15. Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28–44. https://doi.org/10.1016/j.inffus.2011.08.001

16. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In Advances in Neural Information Processing Systems 30 (pp. 5574–5584).

17. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In Proceedings of the 33rd International Conference on Machine Learning (pp. 1050–1059).

18. Li, H. N., Li, D. S., & Song, G. B. (2004). Recent applications of fiber optic sensors to health monitoring in civil engineering. Engineering Structures, 26(11), 1647–1657. https://doi.org/10.1016/j.engstruct.2004.05.018

19. Grosse, C. U., & Ohtsu, M. (Eds.). (2008). Acoustic emission testing: Basics for research – applications in civil engineering. Springer.

20. Lei, Y., Yang, B., Jiang, X., Jia, F., Li, N., & Nandi, A. K. (2020). Applications of machine learning to machine fault diagnosis: A review and roadmap. Mechanical Systems and Signal Processing, 138, 106587. https://doi.org/10.1016/j.ymssp.2019.106587