

Incipient Anomaly Detection in Bridge Health Monitoring using Capacity-Constrained Classical and Hybrid-Quantum Autoencoders

Yesha Pandya^{*†1} , S N Pandya^{†2} , and Dr. N K Solanki² 

¹ Department of Computer Science and Engineering, School of Technology, GSFC University, Vadodara, India

(Email: pandya.yesha.n@gmail.com)

² Department of Applied Mechanics, Faculty of Technology, The Maharaja Sayajirao University of Baroda, Vadodara, India

(Email: sarika.pandya-polyappmech@msubaroda.ac.in, nksolanki-appmech@msubaroda.ac.in)

^{*} Corresponding author.

[†] These authors contributed equally to this work.

Abstract

This paper introduces a vibration-based machine learning framework for anomaly detection in bridges. Structural signals need to be isolated from stochastic noise for incipient damage detection in civil infrastructure. This study uses Continuous Wavelet Transform to map raw accelerometer telemetry from the openLAB reference bridge into 2D time-frequency domain to allow efficient damage localisation. The 2D matrices are passed through a convolutional frontend which employs pooling layers to filter environmental noise. An acoustic emission proxy is injected into a targeted sensor to simulate early-stage micro-cracking. Various classical and hybrid-quantum autoencoder topologies, including a physics-informed neural network, are designed and compared. All models share the same convolutional frontend and decoder; the hybrid quantum models replace the classical bottleneck layer with a variational quantum classifier. Note that all experiments were simulated on classical hardware using PennyLane.

Each sensor channel is constrained to a two-variable bottleneck to mirror current noisy intermediate-scale quantum hardware limitations. An isolated single-sensor architecture fails to differentiate damage transients from baseline noise and routine reconstruction errors. This limitation is overcome by expanding the architecture into a 6-sensor weight-tied Siamese convolutional network that maintains the strict bottleneck constraint while capturing geometric context of the bridge deck. Evaluation shows spatial fusion successfully recovers anomaly detection capacity for classical convolutional networks and unentangled quantum baselines. However, forcing inter-sensor quantum entanglement currently triggers gradient vanishing.

Keywords: Structural Health Monitoring, Quantum Machine Learning, Physics-Informed Neural Networks, Anomaly Detection, Multi-sensor Spatial Fusion

1 INTRODUCTION

Indian civil infrastructure is aging rapidly. Traditional maintenance such as schedule-based visual inspections are fundamentally reactive. These methods are incapable of identifying internal structural degradation before macroscopic yielding occurs. Continuous Structural Health Monitoring (SHM) using vibration telemetry provides a data-driven alternative. Machine learning algorithms can theoretically detect structural anomalies in real-time by processing the dynamic response of a bridge under operational loading.

Incipient anomalies are characterized by low-amplitude transient energy bursts that are easily buried beneath the macroscopic environmental and operational noise floor of the structure. Isolating these structural signals from stochastic noise is a computation challenge. Deep learning architectures such as autoencoders can theoretically isolate and identify these anomalies.

This paper evaluates and compares multi-sensor spatial fusion for classical and parameterized quantum topologies for early-stage damage detection. Both paradigms are benchmarked on real-world accelerometer telemetry recorded from a full-scale research bridge deck under operational loading [1]. The underlying codebase, network artifacts, and data processing pipeline are archived in an open-access repository [2].

1.1 Theoretical Rationale for using Hybrid Quantum Autoencoders

Physical structural failure events are extremely rare in operational infrastructure. Therefore, unsupervised learning has become the standard approach for anomaly detection in structural health monitoring (SHM). Classical convolutional autoencoders (CAEs) and variational autoencoders (VAEs) can learn baseline structural behavior without requiring labeled damage data.

However, when CAEs process multi-sensor arrays to capture spatial cross-correlations, their dense latent layers scale quadratically ($O(d^2)$) with feature dimension d . This increases the risk of overfitting since training data consists mainly of healthy operational measurements. Models become prone to memorize localized environmental noise and temperature-induced variance rather than learning generalizable structural dynamics.

Hybrid Classical-Quantum Autoencoders (HQAEs) circumvent this parameter explosion problem by replacing the dense bottleneck with a variational quantum circuit (VQC). Recent studies have reported QAEs can achieve competitive anomaly detection performance compared to deep classical autoencoders while working with substantially fewer trainable parameters [3-5]. Parameter complexity in a VQC scales linearly ($O(N \cdot L)$) with N qubits and L circuit layers [6]. This is because trainable weights are stored merely as single qubit rotation angles rather than large matrix grids.

Quantum bottlenecks also offer the geometric advantage of Hilbert space projection [7]. Encoding a latent vector into N qubits maps it to a Hilbert space of dimension 2^N .

$$\mathcal{H} \simeq \mathbb{C}^{2^N}$$

Where $\mathcal{H}$ represents the Hilbert space and $\mathbb{C}$ denotes the complex vector space.

This larger representation space may help separate subtle structural changes from environmental and operational variations that overlap in the compressed classical representation.

Directly encoding raw time-frequency representations, however, would require larger quantum circuits which are incompatible with current NISQ hardware [8]. Simulating these large circuits on classical hardware is computationally very expensive. We therefore use a classical 2D convolutional frontend to filter raw telemetry and compress spatial-spectral matrices down to low-dimensional latent vectors prior to quantum encoding.

2 RELATED WORK

Deep learning for structural dynamics has advanced considerably over the past decade. Early foundational work by Abdeljaber et al. [9] demonstrated the use of 1-D CNNs for real-time vibration-based damage detection. Subsequent studies transitioned to time-frequency representations because transient mechanical events are poorly localized in the raw time domain. Zhang et al. [10] showed that mapping acceleration signals into 2D frequency-domain images improved concrete deck crack detection by separating resonant frequencies from high-frequency noise.

Recent research has shifted towards use of unsupervised learning to handle the extreme data imbalance in SHM. Abraham et al. [11] and Kim et al. [12] deployed unsupervised deep learning to forecast normal structural dynamic responses and identify deviations from it. These classical models achieve high accuracy on macroscopic anomalies, but remain computationally expensive when scaled across dense multi-sensor data. Quantum machine learning (QML) has recently been explored to address classical capacity limitations. Alves et al. [13] extracted features from raw acceleration signals and encoded them into unsupervised quantum classifiers to detect anomalies.

However, a critical gap remains in the use of QML in SHM. Standard VQCs generally treat input parameters as independent mathematical features and apply rigid entanglement topologies without inherent geometric context. If deployed naively for multi-sensor telemetry, this mathematical formulation ignores physical reality: kinetic energy propagates through concrete and steel subject to strict spatial attenuation [14].

To address this gap, this study compares conventional quantum entanglement topologies with a physics-informed topology in which sensor-to-sensor distances are incorporated into the entanglement structure.

3 METHODOLOGY: DATA AND PROBLEM FORMULATION

The performance of machine learning models depends on the integrity of its input data. For SHM, algorithms need to map real-time kinetic energy of massive concrete and steel structures. This section describes the reference dataset, signal transformation, and the controlled anomaly proxy used to evaluate incipient damage detection.

3.1 The openLAB Reference Dataset

To establish a rigorous baseline, this study utilizes telemetry from the undamaged openLAB research bridge in Bautzen, Germany [1]. Data is acquired via piezoelectric accelerometers (PCB 393A03) routed through a Gantner Instruments Q-station.




Figure 1: (a) Spatial topology and (b) cross-sectional piezoelectric accelerometer array of the openLAB reference bridge [1].

The hardware employs a threshold trigger. Measurements initiate only when localized vibration exceeds $2 \times 10^{-4} \text{ m/s}^2$. Each triggered event captures a continuous 70-second burst sampled at 500 Hz. Upstream processing removes the DC offset and applies a 4th-order Butterworth filter to enforce a 0.5 to 100 Hz bandwidth.

For this study, node PE11 (located on the bottom edge of precast element span 1, girder 1) was designated as the isolated target for localized anomaly proxy injection.

3.2 Spatio Temporal Signal Processing

Standard Fast Fourier Transforms (FFT) isolate resonant frequencies but strip away temporal localization, rendering them insufficient for mapping transient acoustic emissions [15]. Therefore, a Continuous Wavelet Transform (CWT) using the Complex Morlet Wavelet was applied to map raw 1D signals into 2D time-frequency domain. The admissibility parameter is strictly set to $\omega_0 = 6.0$ to symmetrically balance time and frequency resolution according to the Heisenberg-Gabor limit.

The 70-second bursts were segmented into non-overlapping 2-second windows (1000 time steps at 500 Hz). The CWT maps these across 64 logarithmically spaced frequency scales. This converts 1D vibrations into 2D spatial matrices of shape (64, 1000). All kinetic energy values are then strictly bounded to $[0, 1]$ range using Min-Max normalization to optimize the tensors for stable gradient descent.




Figure 2: (a) 1D signal to CWT spectrogram mapping and (b) 6-sensor batch-channel stacking.

3.3 Physical Wave Mechanics and the Anomaly Proxy

The openLAB 2024 dataset provides measurements from a healthy structure and contains no damage events. So, synthetic structural degradation must be introduced to evaluate anomaly detection. Bridges do not typically fail instantaneously. Degradation develops progressively under repeated loading. When local stress exceeds the tensile capacity of a concrete or steel member, molecular bonds fracture. This localized failure releases an Acoustic Emission (AE), i.e., a burst of kinetic energy that travels through the girder as an elastic stress wave [16]. To replicate this, an AE proxy was injected directly into the normalized CWT matrices of the test data. This injection was applied only to the PE11 channel to simulate a micro-crack occurring immediately adjacent to that specific accelerometer.

The energy amplitude of this synthetic wave was restricted to 0.35 on the $[0, 1]$ scale to simulate incipient damage. Injecting a massive amplitude (e.g., 0.90) would simulate a catastrophic failure event which could be identified through a simple amplitude threshold. Capping the amplitude at 0.35 buries the anomaly deep within the bridge’s natural

environmental noise floor. The network cannot rely on raw signal clipping to detect it, and is forced to rely on complex geometric reconstruction.

3.4 Spectral Masking and Damage Severities

Amplitude dictates a fracture’s physical intensity and spectral width dictates its temporal footprint. A larger fracture generates a more sustained acoustic emission. To precisely quantify the detection limits of our architectures, a sliding spectral mask was applied to control the duration of the injected high-frequency burst. A localized high-frequency perturbation without a corresponding low-frequency global modal shift inherently signals early-stage damage in structural dynamics.

All masks evaluated represent incipient fatigue since the amplitude is capped at 0.35. They only differ in spectral width. Six different masks were tested:

1. Mask Width 2 (3.1% Spectral Coverage): Highly localized, isolated perturbation
2. Mask Width 3 (4.6% Coverage): Localized perturbation with increased spectral extent
3. Mask Width 4 (6.2% Coverage): Moderate spectral extent
4. Mask Width 5 (7.8% Coverage): Intermediate incipient degradation
5. Mask Width 6 (9.3% Coverage): Advancing anomaly that is still undetectable to visual inspection.
6. Mask Width 8 (12.5% Coverage): Upper physical boundary of the incipient phase, immediately preceding macroscopic failure.

Receiver Operating Characteristic (ROC) curves were plotted across all spectral mask widths to evaluate anomaly isolation boundaries across progressive damage severities.

4 ARCHITECTURAL FRAMEWORK AND TOPOLOGIES

4.1 Classical Bottleneck and Spatial Blindness

The architecture first processes single-channel CWT energy density matrix of shape $(1, 64, 1000)$ from node PE11. The deep classical frontend functions as a spatial-spectral low-pass filter to systematically strip stochastic operational noise. The encoder uses an initial 3×3 convolution to expand the input to 16 channels, followed by a 2×4 MaxPool. A second 3×3 convolution further expands the representation to 32 feature maps, and the subsequent 4×10 MaxPool produces a compressed $(32, 8, 25)$ tensor that flattens to 6400 classical features.

To first verify the reconstruction capability of proposed architectures, the CAE was evaluated under an unconstrained latent capacity of 8 continuous variables for PE11 node ($z_s \in \mathbb{R}^8$). The corresponding HQAE models we allocated $N_{\text{sensor}} = 8$ qubits for the same single-sensor case. Both models achieved a best validation mean squared error of ~ 0.0089 . This confirms their foundational mathematical soundness.

However, eight qubits per sensor would require 48 qubits for the six accelerometers. A state-vector simulation of N qubits requires 2^N complex amplitudes. This is computationally very expensive because hybrid quantum models are simulated on classical hardware. The latent allocation was therefore limited to two qubits per sensor, resulting in a 12-qubit global configuration.

Single-sensor models were also evaluated with the same two-qubit allocation ($z_s \in \mathbb{R}^2$) for direct one-to-one comparison with the multi-sensor architecture. This $N_{\text{sensor}} = 2$ budget compresses 64,000 input values to two latent variables; which severely starves the single-

sensor CAE and HQAE variants. In isolation, a single node mathematically fails to distinguish a low-amplitude acoustic emission from its own blurred baseline reconstruction error. This empirically validates single-sensor spatial blindness under extreme latent compression.

4.2 Siamese Multi-Sensor Spatial Fusion

To overcome single-sensor spatial blindness without increasing the two-qubit-per-sensor latent constraint, all six sensor representations are passed through a weight-tied Siamese frontend. Processing six independent (64, 1000) CWT matrices simultaneously increases parameter exposure and risks multi-channel overfitting. So, prior to feature extraction, the input tensor is dynamically folded across the batch and channel dimensions from $(B, 6, 64, 1000)$ into $(B \times 6, 1, 64, 1000)$. This forces all six accelerometer streams through identical classical convolutional filters simultaneously.




Figure 3: Siamese batch-channel folding of 6 spectrograms into the global latent graph

The six independent latent vectors are then unfolded and concatenated into a single 12-dimensional global latent tensor, which corresponds to the 12-qubit global register ($N = 12$) utilized in the quantum bottleneck. This global representation leverages spatial cross-correlation between sensors, where healthy telemetry from the 5 neighboring nodes can be used as reference.

4.3 HQAE Framework and Hardware Integration

The HQAE retains the same Siamese 2D-CNN encoder and deconvolutional decoder as the classical baseline; only the bottleneck is replaced by a 12-qubit VQC partitioned into six 2-qubit sensor registers. All quantum circuits are implemented and simulated on classical hardware using PennyLane; no physical quantum hardware is used. The CNN feature extraction runs on GPU (cuda), while the quantum state-vector simulation runs on the CPU (default.qubit). Custom autograd operators handle the GPU-CPU tensor transfer while preserving the PyTorch automatic-differentiation graph.

For sensor node index $s \in \{1, \dots, 6\}$ and latent component index $k \in \{1, 2\}$, the Siamese encoder outputs an unbounded 2-D classical feature vector

$$\mathbf{z}_s = [z_{s,1}, z_{s,2}]^T \in \mathbb{R}^2$$

Prior to quantum state preparation, each feature scalar $z_{s,k}$ is bounded to prevent Bloch sphere phase-wrapping via the transformation:

$$x_{s,k} = \pi \cdot \tanh(z_{s,k}), \quad \text{for } s \in \{1, \dots, 6\} \text{ and } k \in \{1, 2\}$$

This clamps rotation angles to $x_{s,k} \in [-\pi, \pi]$. Each bounded angle $x_{s,k}$ is then mapped onto the k -th qubit of sensor s 's dedicated sub-register via single-qubit Pauli- Y rotation gates (R_y angle embedding) applied to the initial computational ground state $|0\rangle$:

$$|\psi_s\rangle = R_y(x_{s,1}) |0\rangle \otimes R_y(x_{s,2}) |0\rangle \in \mathcal{H}_s \simeq \mathbb{C}^{2^2}$$

where $|\psi_s\rangle$ is the local 2-qubit state vector for sensor s and $\mathcal{H}_s \simeq \mathbb{C}^4$ represents the 4-dimensional local complex Hilbert state space.

Prior to inter-sensor entanglement, the global 12-qubit state vector $|\Psi_0\rangle$ exists as a separable product state constructed via the multi-register tensor product ($\bigotimes$) across all six local sensor sub-registers:

$$|\Psi_0\rangle = \bigotimes_{s=1}^6 |\psi_s\rangle \in \mathcal{H}_{\text{global}} \simeq \mathbb{C}^{4096}$$

After applying parameterized unitary transformations $U(\theta)$, Pauli- Z expectation values are measured for individual qubits:

$$z'_{s,k} = \langle \Psi_0 | U^\dagger(\theta) (\sigma_z^{(s,k)}) U(\theta) | \Psi_0 \rangle \in [-1.0, 1.0]$$

where θ represents trainable parameters, $U^\dagger(\theta)$ is its conjugate transpose (adjoint), $\sigma_z^{(s,k)}$ is the Pauli- Z measurement operator for qubit k of sensor node s , and $\langle \Psi_0 | \dots | \Psi_0 \rangle$ denotes the quantum mechanical expectation value (inner product).

The resulting 12 expectation values are concatenated into a continuous vector ($z' \in \mathbb{R}^{12}$), transferred back to the GPU, separated into six 2-dimensional sensor representations, and passed to the classical decoder. Without inter-sensor entangling gates, the six local circuits constitute the HQAE[None] configuration.

4.4 Quantum Entanglement Topologies and Optimization Instability

The six physical sensor nodes are mapped to dedicated 2-qubit sub-registers, giving $C(6, 2) = 15$ potential spatial correlations between sensor pairs. PennyLane's in-built None and Strong entanglement were evaluated. Additionally, two domain-specific spatially informed topologies were designed to model spatial correlations across the bridge:

1. **bridge_hard** uses unparameterized Controlled-NOT (CNOT) gates between sensor sub-registers whose physical Euclidean separation satisfies $D_{ij} \leq 15.0m$
2. **bridge_soft** is a Physics-Informed Neural Network (PINN) that replaces rigid CNOT connections with parameterized Controlled- Y (CRY) gates [17]. with the applied rotation attenuated according to sensor separation. The rotation angle applied to each CRY gate is exponentially attenuated based on physical node separation:

$$\theta_{\text{applied}} = \theta_{\text{learned}} \cdot e^{-\gamma D_{ij}}$$

where learned θ_{learned} is the trainable parameter generated by the classical optimizer, D_{ij} is the physical Euclidean distance between bridge nodes i and j , and $\gamma = 0.1m^{-1}$ is the structural spatial attenuation coefficient. To enforce physical wave propagation constraints across the spatial graph, a wave mechanics regularization penalty is appended to the global loss function [20]:

$$L_{\text{Total}} = L_{\text{Recon}} + \lambda \left\| \frac{\partial^2 u}{\partial t^2} - c^2 \nabla^2 u \right\|^2$$

where L_{Recon} is the reconstruction Mean Squared Error (MSE), λ is a penalty scaling factor, u denotes the spatio-temporal elastic displacement field reconstructed by the decoder, c is the acoustic wave propagation velocity in concrete, and ∇^2 is the spatial Laplacian operator.

5 EXPERIMENTAL SETUP AND EVALUATION PROTOCOL

All classical and hybrid architectures are constrained to an identical optimization framework to ensure strict mathematical comparability.

5.1 Hyperparameter Constraints and Regularization

Models are trained using Adam optimizer [18] with a fixed learning rate of $\eta = 1 \times 10^{-3}$, batch size of 32, and a maximum of 15 epochs. Extended training was found to induce overfitting in the classical decoder. Reported metrics are obtained from three deterministic seeds (42,100 and 2026) to account for optimization variance. NISQ regularization measures applied to stabilize training across the classical-VQC interface are:

1. Classical encoder outputs $z_s \in \mathbb{R}^2$ were bounded via transformation:

$$x_{s,k} = \pi \cdot \tanh(z_{s,k})$$

. This clamps all rotation inputs strictly to $x_{\{(s,k)\}} \in [-\pi, \pi]$; which prevents Bloch sphere phase-wrapping and ensures 1-to-1 mapping onto single-qubit rotation gates (R_y).

2. Trainable quantum parameters are initialized from a near-zero interval:

$$\theta \approx U(-0.01, 0.01)$$

. This near-identity initialization allows the circuit to pass the state-prepared features with minimal initial distortion; and reduces the risk of initialization-induced barren plateaus. [19]

3. Backpropagation through the quantum circuit relies on the parameter-shift rule. So, gradient norm clipping ($|g|_2 \leq 1.0$) is applied before each optimization step to prevent high-variance quantum gradients from destabilizing classical feature extraction. [20]

5.2 Target Node Isolation Protocol

Because the evaluated anomaly is localized at PE11 while the remaining five nodes represent undamaged neighboring responses, averaging reconstruction error across all six channels would dilute the localized reconstruction error signal by a factor of $\frac{1}{6}$. This would mask subtle damage signatures and artificially inflate overall performance metrics. So, although all six channels contribute to spatial feature extraction and global loss computation during training, evaluation metrics are calculated strictly on the target accelerometer node (PE11):

$$\text{MSE}_{\text{target}} = \frac{1}{M} \sum_{m=1}^M \left\| \mathbf{X}_{\text{PE11}}^m - \hat{\mathbf{X}}_{\text{PE11}}^m \right\|_F^2$$

where $\mathbf{X}_{\text{PE11}}^m$ and $\hat{\mathbf{X}}_{\text{PE11}}^m \in \mathbb{R}^{64 \times 1000}$ denote the ground-truth and reconstructed energy density matrices respectively for target node PE11 on the m -th evaluation sample. $\|\cdot\|_F$ represents the Frobenius norm and M is the total number of evaluation samples. Discarding

the reconstruction error of undamaged neighboring nodes during evaluation isolates the model’s true localized anomaly sensitivity.

6 RESULTS AND DISCUSSION

AUC-ROC scores are evaluated across progressive spectral mask widths to quantify the detection boundaries for incipient anomalies. All AUC-ROC values in Tables 1 and 2 are reported as (mean $\pm$ standard deviation) across the three deterministic seeds. Table 1 presents the single-sensor results and Table 2 reports the corresponding multi-sensor spatial-fusion results.

To guarantee a strict equivalent ablation, every configuration across both tables is constrained to a 2-variable latent bottleneck per sensor node ($z_s \in \mathbb{R}^2$). This is matched to a NISQ simulation capacity limit of $N_{\text{sensor}} = 2$ qubits per node.

Table 1: Single-Sensor Evaluation (Target Node: PE11, $N_{\text{sensor}} = 2$ qubits)

Damage Severity	CAE (Baseline)	HQAE [None]	HQAE [Strong]
Mask 2 (3.1%)	0.620 ± 0.001	0.624 ± 0.010	0.634 ± 0.016
Mask 3 (4.6%)	0.673 ± 0.006	0.669 ± 0.006	0.678 ± 0.009
Mask 4 (6.2%)	0.720 ± 0.009	0.706 ± 0.009	0.720 ± 0.008
Mask 5 (7.8%)	0.752 ± 0.003	0.744 ± 0.007	0.747 ± 0.005
Mask 6 (9.3%)	0.781 ± 0.005	0.768 ± 0.010	0.777 ± 0.008
Mask 8 (12.5%)	0.819 ± 0.001	0.812 ± 0.011	0.813 ± 0.010

Table 2: Multi-Sensor Spatial Fusion (6-Sensor Array, $N_{\text{sensor}} = 2$ qubits per node, $N = 12$ total qubits)

Damage Severity	CAE (Baseline)	HQAE [None]	HQAE [Strong]	HQAE [Bridge Hard]	HQAE [Bridge Soft]
Mask 2	0.650 ± 0.015	0.629 ± 0.012	0.616 ± 0.010	0.617 ± 0.025	0.612 ± 0.048
Mask 3	0.705 ± 0.020	0.685 ± 0.012	0.664 ± 0.010	0.672 ± 0.035	0.647 ± 0.054
Mask 4	0.745 ± 0.016	0.729 ± 0.013	0.710 ± 0.010	0.709 ± 0.045	0.679 ± 0.064
Mask 5	0.788 ± 0.012	0.778 ± 0.017	0.749 ± 0.009	0.743 ± 0.054	0.712 ± 0.070
Mask 6	0.830 ± 0.013	0.819 ± 0.017	0.786 ± 0.010	0.770 ± 0.058	0.739 ± 0.078
Mask 8	0.888 ± 0.009	0.891 ± 0.018	0.839 ± 0.007	0.818 ± 0.063	0.801 ± 0.089

6.1 The Classical Spatial Regularization Advantage

The multi-sensor ablation shows improved performance for classical CAE compared to isolated-sensor baseline. At mask width 8, the spatial CAE scales to an AUC-ROC of 0.888, significantly outperforming the isolated baseline’s 0.819. This improvement shows that spatial

regularization inherently suppresses false positives associated with environmental noise and helps the architecture confidently isolate target node anomaly.

6.2 The Quantum Scaling Wall and Barren Plateaus

2 qubit HQAE[Strong] marginally outperforms the classical CAE at Mask Width 2 (AUC-ROC of 0.634 vs 0.620) in single-sensor ablations. But this advantage disappears in multi-sensor ablations, where it falls to AUC of 0.616 for the same mask width. In contrast, HQAE [None] remains comparatively stable and reaches AUC-ROC of 0.891 at Mask Width 8.

Strongly entangled configuration degrades despite the barren plateau mitigation strategies detailed in Section 5. Extending rigid CNOT topologies across 12 qubits inherently flattened the optimization landscape and rendered these regularizers largely ineffective. Consequently, the model failed to learn the multi-node spatial graph. Under the present NISQ simulation constraints, the entanglement structure appears to be a more significant limitation than qubit count alone.

6.3 The PINN Paradox

The physics-informed `bridge_soft` topology showed the highest optimization variance; its standard deviation increased to 0.089 at Mask Width 8. As defined in section 4.4, this topology incorporates distance-attenuated CRY gates and wave-mechanics-based loss penalty. However, this study injected a synthetic anomaly exclusively at node PE11, while the neighboring sensors remained unperturbed; which created a localized discontinuity that does not reflect physical wave propagation through the structure. This contrast between the imposed anomaly and propagation constraint gives a plausible explanation for the high variance and instability during optimization of the PINN.

7 CONCLUSION AND FUTURE WORK

This study showed that incorporating telemetry from spatially distributed sensors improves localized anomaly detection as compared to a single-sensor configuration. However, forcing global inter-sensor entanglement for larger circuits triggers barren plateaus despite mitigation measures and degrades model performance. Physics-informed topology showed high sensitivity to physical and spatial characteristics of anomaly.

A key limitation of this study was the use of synthetic anomaly injection which does not replicate the mechanical response of an actual structure. Future work would evaluate the proposed topologies using physically grounded data generated through finite-element simulations or a digital twin. Such models can accurately represent damage evolution and wave propagation across sensor networks. This provides a more realistic basis for testing the physics-informed networks.

REFERENCES


- [1] Jansen, A., Richter, B., Röder, R., Herbers, M., & Marx, S. (2024). Monitoring data of the openLAB research bridge (2024-02-01 to 2024-10-31) and building information [Data set]. OPARA Dataset Repository. <https://doi.org/10.25532/OPARA-660>
- [2] Pandya, Y. (2026). Spectra-cq [Source code]. GitHub. <https://github.com/yeshapan/spectra-cq>

- [3] Sakhnenko, A., O'Meara, C., Neukart, F., & Bernabé-Moreno, J. (2021). Hybrid classical-quantum autoencoder for anomaly detection. *Quantum Machine Intelligence*, 3(2), 27. <https://doi.org/10.1007/s42484-021-00052-8>
- [4] Woźniak, S., Belis, V., Puljak, E., Grossi, M., Pierini, M., & Sofia, A. (2023). Anomaly detection in high-energy physics using a quantum autoencoder. *Physical Review D*, 108(3), 036006. <https://doi.org/10.1103/PhysRevD.108.036006>
- [5] Tschärke, K., Wendlinger, M., Ahouzi, A., Bhardwaj, P., Amoi-Taleghani, K., Schrödl-Baumann, M., & Debus, P. (2025). Quantum autoencoder for multivariate time series anomaly detection. *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2470–2481. <https://doi.org/10.1109/QCE65121.2025.00268>
- [6] Benedetti, M., Lloyd, S., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4), 043001. <https://doi.org/10.1088/2058-9565/ab4eb5>
- [7] Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4), 040504. <https://doi.org/10.1103/PhysRevLett.122.040504>
- [8] Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- [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] Zhang, Q., Barri, K., Babanajad, S. K., & Alavi, A. H. (2021). Real-time detection of cracks on concrete bridge decks using deep learning in the frequency domain. *Engineering*, 7(12), 1786–1796. <https://doi.org/10.1016/j.eng.2020.07.026>
- [11] Abraham, J. N., Tran, M. Q., Jayaraj, J. S., Matos, J. C., Valluzzi, M. R., & Dang, S. N. (2026). Unsupervised learning-based anomaly detection for bridge structural health monitoring: Identifying deviations from normal structural behaviour. *Sensors*, 26(2), 561. <https://doi.org/10.3390/s26020561>
- [12] Kim, S., Kim, H., & Hwang, Y. (2025). Data-driven dynamic response forecasting and anomaly detection in long-span bridges. *Journal of Civil Structural Health Monitoring*, 15, 3045–3062. <https://doi.org/10.1007/s13349-025-00952-8>
- [13] Alves, V. H. M., Gomes, R. F. I., & Cury, A. (2025). New perspectives on structural health monitoring using unsupervised quantum machine learning. *Mechanical Systems and Signal Processing*, 229, 112489. <https://doi.org/10.1016/j.ymssp.2025.112489>
- [14] Grosse, C. U., & Ohtsu, M. (2008). *Acoustic emission testing*. Springer Science & Business Media. <https://doi.org/10.1007/978-3-540-69972-9>
- [15] Oppenheim, A. V., & Schaffer, R. W. (2010). *Discrete-time signal processing* (3rd ed.). Prentice Hall. <https://www.pearson.com/en-us/subject-catalog/p/discrete-time-signal-processing/P200000000305/9780134681171>
- [16] Abdul Kudus, S., Misnan, M. F., Muhammad, H., Sawahli, A. M. S., Ahmad, Z., & Shahid, K. A. (2026). Acoustic emission-based assessment of damage evolution in bridge systems:

- Methods, indicators, and advances in structural mechanics. *Journal of Nondestructive Evaluation*, 45, 100. <https://doi.org/10.1007/s10921-026-01388-w>
- [17] 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. <https://doi.org/10.1016/j.jcp.2018.10.045>
 - [18] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations (ICLR)*, San Diego, CA. <https://arxiv.org/abs/1412.6980>
 - [19] Grant, E., Wossnig, L., Ostaszewski, M., & Severini, S. (2019). An initialization strategy for addressing barren plateaus in parametrized quantum circuits. *Quantum*, 3, 214. <https://doi.org/10.22331/q-2019-12-09-214>
 - [20] Schuld, M., Bergholm, V., Gogolin, C., Izaac, J., & Killoran, N. (2019). Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3), 032331. <https://doi.org/10.1103/PhysRevA.99.032331>

Appendix A: Extended ROC-AUC Sensitivity Curves

The following sensitivity curves detail the ROC-AUC performance across all evaluated damage severity levels (Mask 2 through Mask 8). The left column illustrates the single-sensor baseline ablations, while the right column demonstrates the enhanced stabilization achieved via multi-sensor spatial fusion.



Incipient Anomaly Detection in Bridge Health Monitoring

