

Article

Physics-Informed Multimodal Sensing and Machine Learning Framework with Real-Time Digital Twin for Intelligent Structural Health Monitoring of Marine Wharves

Mohammad Asadian Ghahfarokhi ¹, James Riffat ² and Seyed Reza Samaei ^{3,*}

¹ Department of Marine Industries, Science and Research Branch, Islamic Azad University, Tehran 1477893855, Iran; asadian@srbiau.ac.ir

² World Society of Sustainable Energy Technologies, Nottingham, NG11 6AA, United Kingdom; ceo@wsset.org

³ Department of Marine Industries, Science and Research Branch, Islamic Azad University, Tehran 1477893855, Iran; samaei@srbiau.ac.ir

* Correspondence author: samaei@srbiau.ac.ir

Abstract: Marine wharves, consisting of concrete or steel decks supported by pile foundations, operate in harsh marine environments involving waves, vessel impacts, and corrosive seawater. Even minor cracks or early-stage corrosion in piles can lead to sudden structural failure. Early detection is therefore essential. Traditional vibration-based structural health monitoring (SHM) methods often struggle to deliver reliable results in real operational settings because environmental and operational variations can mask damage-induced frequency changes. This study presents a multimodal SHM framework that integrates accelerometers, strain gauges, high-resolution cameras, acoustic transducers, and corrosion sensors with edge computing, advanced machine learning, and a real-time physics-informed digital twin. Edge nodes preprocess sensor data to extract modal frequencies, crack indicators, and acoustic features, reducing communication loads by over 90% while preserving diagnostic fidelity. A hybrid pipeline combined attention-based multimodal fusion, XGBoost screening, and a feedforward deep neural network trained using a composite physics-informed objective. The physical loss numerically penalised positive damage-related modal-frequency shifts after environmental and operational condition compensation and stiffness ratios outside the admissible range of 0.70–1.00. The selected physics-loss coefficients were $\lambda_{\text{phys}} = 0.10$ and $\eta = 0.50$. Sensor features were ingested continuously, whereas finite-element state updates were event-triggered by a fused damage probability of at least 0.80, agreement between at least two sensing modalities, normalised predictive entropy not exceeding 0.35, and persistence over three consecutive 10-s analysis windows. Validation using 200 finite-element simulations and 40 scaled laboratory tests achieved 95.0% accuracy for single-damage detection, compared with 78.0% for the conventional Modal Strain Energy method. For multiple-damage cases, the proposed framework achieved 92.17% precision, 88.33% recall, and an F1-score of 90.21%. Damage-severity estimation yielded $R^2 = 0.891$ and an RMSE of 2.06 percentage points, while the mean full-scale-equivalent localisation error was 4.8 cm. Edge processing reduced communication volume by 94.73% without a statistically significant reduction in diagnostic accuracy. While evaluated under controlled conditions, the framework shows strong potential for proactive, accurate, and interpretable monitoring of marine wharves in operational environments.

Keywords: Intelligent Structural Health Monitoring; Marine Wharves; Multimodal Data Fusion; Physics-Informed Machine Learning; Digital Twin; Edge Computing; Predictive Maintenance; Explainable AI.



1. Introduction

Marine wharves, comprising reinforced concrete or steel decks supported by pile foundations, play a vital role in global maritime transportation by facilitating the berthing, loading, and unloading of vessels under varying operational conditions. These structures operate in harsh marine environments characterized by cyclic wave actions, vessel collisions, tidal variations, and corrosive seawater, all of which accelerate structural deterioration. Even minor cracks or localized corrosion in pile elements can escalate into critical failures if not detected early, underscoring the necessity for continuous and accurate structural health monitoring (SHM). Conventional vibration-based SHM approaches, such as the Modal Strain Energy (MSE) method, detect stiffness loss by correlating modal frequency shifts with structural damage. While these methods have shown promising results under controlled laboratory settings, they often suffer from reduced reliability in real-world conditions. Environmental and operational changes—such as tidal fluctuations, temperature variations, and varying berthing loads—can independently alter modal properties without any actual structural damage. Studies, including the systematic review by (Keshmiry et al. 2023), have reported that such Environmental and Operational Conditions (EOC) can mask true damage signatures, leading to high rates of false negatives or false alarms. This limitation highlights the need for robust methods capable of isolating damage-induced changes from environmental effects. Recent advancements have shifted SHM research towards multimodal sensing and data fusion, enabling a richer and more discriminative representation of structural state. Combining accelerometers, strain gauges, high-resolution cameras, acoustic transducers, and corrosion sensors provides complementary information that can enhance detection accuracy by up to 30% compared to single-sensor systems (Teng et al., 2021; Broer et al., 2022). In parallel, edge computing has emerged as a solution to the high-bandwidth demands of multimodal SHM. On-device feature extraction reduces data transmission by over 90% while enabling near-real-time decision-making (Hidalgo-Fort et al., 2024). Another transformative development is the application of digital twins (DTs) to civil infrastructure. DTs continuously assimilate real-time sensor data into high-fidelity numerical models, enabling predictive simulations under evolving operational and environmental conditions (Huang et al., 2025). In SHM applications, DTs allow engineers to perform “what-if” analyses, evaluate the impact of detected damage, and plan maintenance proactively. When combined with Explainable AI (XAI) methods—such as SHAP value attribution and attention-based visualization—these systems not only improve detection accuracy but also provide transparency, increasing operator trust in automated decisions (Luckey et al., 2022).

Despite significant progress in SHM research, few existing frameworks fully integrate multimodal sensing, edge computing, physics-informed machine learning, explainable AI (XAI), and real-time digital twinning into a unified system specifically tailored for marine wharves. This study addresses this gap by presenting a novel SHM architecture that:

1. Utilizes a Bayesian–deep learning multimodal fusion strategy to adaptively weight sensor inputs and filter out uncorroborated anomalies.
2. Incorporates a real-time, physics-informed digital twin for predictive damage progression analysis.
3. Deploys an edge–cloud IoT infrastructure to enable low-latency, energy-efficient processing in operational marine environments.

The primary objective is to develop and validate an intelligent SHM framework capable of: (i) integrating diverse sensing modalities with advanced computational techniques; (ii) employing a continuously monitored, event-triggered digital twin for predictive damage assessment; (iii) maintaining high detection reliability under varying environmental and operational conditions; and (iv) delivering interpretable, evidence-based alerts to support timely and informed maintenance decisions.

1.1. Case Study Structure: LPG Berth No.14, Asaluyeh Petrochemical Port

The case study focuses on Berth No. 14 of the Asaluyeh Petrochemical Port, located on the southern coast of Iran within the Pars Special Economic Energy Zone. The geometry, structural-member numbering, baseline modal properties, and initial numerical–experimental investigation of this berth were previously established through a Modal Strain Energy-based study by (Samaei et al. 2021). The present research extends that earlier foundation by incorporating multimodal sensing, edge computing, attention-based data fusion, physics-informed machine learning, probabilistic event triggering, and digital-twin-enabled predictive assessment. This berth plays a strategic role in the export of liquefied petroleum gas (LPG), butane, and ammonia, with a maximum loading capacity of 70,000 tons. The structure is a suspended deck-type wharf measuring 250 m in length with a dredged depth of 14.6 m, designed to accommodate large export vessels under high-demand operational schedules. From a material and construction perspective, all steel profiles used in Berth No. 14 are fabricated from ST37-

2. Literature Review

The broader development of structural health monitoring has been supported by advances in structural integrity and risk-based inspection frameworks, vibration and wireless sensing, environmental and soil–structure interaction assessment, data-driven and machine-learning-based damage identification, reduced-order modelling, edge computing, and digital-twin-enabled asset representation and lifecycle management (API RP 2MIM, 2019; Astroza et al., 2015; Beeby et al., 2006; Bhattacharya et al., 2013; Det Norske Veritas, 2021; Lynch and Loh, 2006; Ma et al., 2020; Pathirage et al., 2018; Rosafalco et al., 2021; Shi et al., 2016; Silva et al., 2016; Sohn et al., 2003; Tao et al., 2018; Worden et al., 2000).

2.1. Limitations of Traditional Vibration-Based SHM

Conventional SHM techniques for marine and coastal infrastructure have relied heavily on global vibration-based indicators, such as the Modal Strain Energy (MSE) method, which estimates stiffness loss by tracking modal frequency shifts (Xu et al., 2024). While conceptually simple, these methods assume a direct one-to-one relationship between stiffness reduction and frequency change, an assumption rarely valid in real operational conditions (Figueiredo et al., 2011). Environmental and operational conditions (EOC)—including tidal fluctuations, temperature variations, and changes in berthing loads—can cause frequency shifts unrelated to damage (Keshmiry et al., 2023). This often results in false negatives, where damage remains undetected, or false positives, where environmental noise is mistaken for structural deterioration (Doebeling et al., 1996; Farrar and Worden, 2007). Recent works have sought to mitigate EOC interference through data normalization and statistical filtering (Sony et al., 2019; Vieira et al., 2023), but these methods often underperform when faced with simultaneous multi-location damages or when long-term environmental variations dominate the signal. This limitation motivates the exploration of multimodal sensing strategies that can cross-validate damage evidence from different physical domains. A recent comprehensive review similarly identified the convergence of advanced sensing, artificial intelligence, edge analytics, digital twins, and predictive maintenance as a central direction in the development of next-generation SHM systems for civil infrastructure (Riffat, Samaei and Pastakkaya, 2025).

2.2. Multimodal Sensing and Data Fusion in SHM

Multimodal sensing—integrating data from heterogeneous sensors—has emerged as a robust strategy for addressing the weaknesses of vibration-only approaches. By combining accelerometers, strain gauges, vision systems, ultrasonic probes, and corrosion sensors, researchers have demonstrated improved detection accuracy and reduced ambiguity in structural diagnostics (Broer et al., 2022; Teng et al., 2021). In civil engineering applications, such fusion has been shown to enhance damage localization by 10–30% compared to single-sensor methods, depending on structural type and operational conditions (Teng et al., 2021; Zhang et al., 2024). Advanced fusion strategies leverage Bayesian inference or attention-based deep learning to adaptively weight sensor contributions depending on the reliability of their signals (Luckey et al., 2022; Martínez et al., 2025). However, despite the growing body of work, applications of such frameworks to marine wharf structures remain scarce, particularly in operationally complex environments where access constraints and environmental variability are significant challenges.

Underwater optical inspection presents additional difficulties because light attenuation, suspended particles, backscattering, non-uniform illumination, and refraction through camera housings can degrade conventional two-dimensional crack images. Recent polarisation-assisted imaging and three-dimensional reconstruction methods offer a complementary route for quantitative underwater defect assessment. (Gan et al. 2026) combined polarisation–dark channel prior fusion imaging with a binocular structured-light system and a rolling-ball point-cloud reconstruction method. Their framework suppresses scattered-light interference while reconstructing irregular defect geometry, enabling direct estimation of defect dimensions and volume rather than relying only on two-dimensional segmentation masks. Such a system could complement the present ROV-assisted workflow by supplying geometrically scaled crack width, depth, surface-loss area, and defect-volume measurements to the digital twin. The present study retained compact two-dimensional U-Net segmentation because the objective was to validate the integrated multimodal SHM architecture; polarisation-assisted three-dimensional imaging is therefore considered a technically compatible extension rather than a component evaluated in the current experimental programme.

2.3. Digital Twin Integration for Infrastructure Monitoring

A Digital Twin (DT) is a continuously updated virtual model that mirrors the physical asset using real-time data. In civil infrastructure, DTs have been applied to bridges, offshore platforms, and high-

rise buildings to simulate evolving stress fields and predict future performance under various loading conditions (Rasheed et al., 2020; Huang et al., 2025). When combined with SHM, DTs enable scenario-based “what-if” analyses that support proactive maintenance planning. Recent advances have emphasised physics-informed digital twins that combine finite-element simulations with sensor-derived boundary conditions, allowing both global response parameters and local damage states to be updated in near-real-time (Kapteyn et al., 2021). Offshore applications have further demonstrated that high-density sensor networks, real-time digital twins, and machine-learning models can be integrated for damage detection in jack-up platform legs (Samaei and Riffat, 2025). Adaptive and real-time digital-twin architectures have also been developed for broader classes of offshore marine structures, demonstrating the potential of closed-loop sensor–model interaction under changing operational conditions (Riffat, Nazari and Samaei, 2026). Most directly, field-validated digital-twin-enabled SHM has recently been demonstrated for offshore high-pile wharves (Samaei and Riffat, 2026). These studies confirm the growing maturity of digital-twin-based monitoring in marine infrastructure. Nevertheless, the combined integration of five sensing domains, numerically enforced physics-informed learning, uncertainty-aware event triggering, edge processing, and explainable AI within a single deck–pile wharf decision architecture remains insufficiently developed.

2.4. Edge Computing for Low-Latency SHM

Data transmission constraints pose a significant barrier to large-scale, high-resolution SHM deployment. Streaming raw vibration data at 100 Hz or continuous 4K video from multiple nodes can overwhelm wireless networks. Edge computing addresses this by processing data locally before transmission, significantly reducing bandwidth requirements (Hidalgo-Fort et al., 2024). Recent prototypes demonstrate reductions in communication volume exceeding 90% relative to raw data transfer, while preserving diagnostic fidelity (Li et al., 2020). In the context of marine wharves, where sensor nodes may be deployed over wide areas with limited power availability, edge–cloud architectures are especially valuable. By running lightweight algorithms—such as the Random Decrement Technique (RDT) for vibration analysis or real-time crack segmentation networks for visual data—on embedded processors, the system ensures near-real-time decision-making without overloading communication channels.

2.5. Explainable AI for Trustworthy SHM

Deep learning has shown strong performance in pattern recognition for SHM but is often criticized for its “black-box” nature (Luckey et al., 2022). In safety-critical infrastructure, interpretability is essential for operator trust. Explainable AI (XAI) techniques, such as SHAP value analysis and attention-based visualizations, allow engineers to understand which sensor inputs most influenced a damage alarm. When combined with physics-informed constraints—such as fracture mechanics relationships—XAI-enhanced SHM ensures that detected anomalies are both statistically significant and physically plausible (Martínez et al., 2025). Despite increasing adoption in other sectors, XAI integration in marine infrastructure SHM remains rare, representing a significant innovation opportunity.

2.6. Research Gap and Novelty of This Study

Recent reviews and offshore applications demonstrate rapid progress towards the integration of intelligent sensing, machine learning, and digital twins in civil and marine infrastructure (Riffat, Samaei and Pastakkaya, 2025; Samaei and Riffat, 2026). Nevertheless, the specific combination of multimodal sensor fusion, numerically enforced physics-informed learning, confidence-aware event triggering, edge–cloud processing, and explainable diagnostic outputs has not yet been comprehensively evaluated as a unified decision framework for deck–pile marine wharves. To address this gap, the present study proposes an end-to-end SHM framework for deck–pile marine wharves that: (1) fuses multimodal sensor data using a Bayesian–deep learning hybrid model; (2) incorporates a physics-informed, real-time digital twin for predictive analysis; (3) utilizes an edge–cloud IoT architecture for efficient, low-latency operation; and (4) embeds explainable AI to ensure transparency and operator trust in decision-making. The novelty of this study lies in its end-to-end integration of multimodal sensing, edge computing, explainable physics-informed machine learning, and a real-time data-informed digital twin with event-triggered finite-element updating, specifically tailored for marine wharves. Unlike previous works that focused on isolated sensing modalities or offline analysis, this framework enables real-time detection, localization, and severity estimation of damage with predictive simulation capabilities, thus bridging the gap between monitoring and actionable maintenance planning.

3. Methodology

3.1. System Architecture

Figure 3 presents the proposed Structural Health Monitoring (SHM) architecture for the Pars Assaluyeh steel deck–pile wharf. The system comprises five core components: (i) a distributed sensor network, (ii) edge computing nodes, (iii) a communication layer, (iv) a cloud-based analytics platform integrating machine learning (ML) models, a database, and a high-fidelity digital twin, and (v) a web-based user interface. The wharf is instrumented with diverse sensing technologies to capture complementary structural and environmental information. Tri-axial accelerometers are mounted on the deck and pile heads to measure global mode shapes and natural frequencies. Fiber Bragg Grating (FBG) strain sensors are installed along critical pile sections to record localized bending strains. High-resolution visual monitoring is provided by fixed mast-mounted cameras and mobile imaging platforms deployed by remotely operated vehicles (ROVs) to capture deck surfaces and pile walls for detecting visible cracks, spalls, and corrosion. Ultrasonic guided-wave transducers, placed on and beneath the deck, propagate waves through structural members and record their responses to detect subsurface delaminations and voids. Environmental sensors continuously monitor temperature, humidity, salinity, and water level, parameters known to influence material behavior and long-term deterioration rates.

Each sensor cluster is paired with a local edge processing unit, typically based on ARM Cortex-M microcontrollers. These units execute on-board algorithms to preprocess raw signals before transmission. For vibration data, the Random Decrement Technique (RDT) and Fast Fourier Transform (FFT) are applied to extract modal envelopes, frequency peaks, and damage-sensitive indices. Vision data are processed using a compact U-Net convolutional neural network to segment cracks from background textures. Acoustic signals are analyzed for anomalies in time-of-flight and amplitude attenuation, while environmental readings are logged without reduction. The resulting feature vectors—comprising modal frequency changes, crack bounding box coordinates with confidence scores, and anomaly flags—are transmitted instead of full raw datasets, reducing wireless traffic by more than 90% with negligible accuracy loss. Data transmission from edge devices to the cloud uses NB-IoT or 5G connectivity, with satellite link redundancy to ensure uninterrupted coverage. In the cloud environment, incoming features are time-synchronized, stored in a central database, and processed through parallel ML pipelines. These outputs feed directly into the digital twin module, a three-dimensional finite element (FE) model of the wharf created in Abaqus, incorporating nonlinear concrete cracking, steel reinforcement behavior, and fluid–structure interaction effects. The web-based user interface presents live sensor readings, updated digital twin states, and visual overlays highlighting damage locations on a schematic of the wharf.




Figure 3. Proposed SHM system architecture (adapted from Chen et al., 2018, with modifications)

3.2. Multimodal Feature Extraction and Fusion

Each sensing modality produces distinct diagnostic features, which are integrated at two fusion stages to enhance detection reliability. During feature extraction, vibration and strain data are processed to determine modal frequencies, mode shape amplitudes, and time-domain indices sensitive to stiffness changes or energy dissipation. Events such as sudden frequency drops are flagged as potential damage indicators. Vision data are segmented by the U-Net model, which outputs crack presence signals and spatial coordinates of detected defects. Acoustic measurements are evaluated for deviations in wave travel time and attenuation patterns, while environmental data are archived for correlation with mechanical responses. In the first fusion stage (data-level fusion), numerical features from all sensors over a common time window are concatenated into a single feature vector, enabling direct cross-correlation between measurements—for example, linking a localized strain spike with a concurrent crack image. In the second stage (decision-level fusion), specialized classifiers for each modality—such as CNNs for imagery and tree ensembles for vibration data—produce independent damage likelihood scores. These scores are then combined using a neural attention-weighting network that prioritizes events corroborated by multiple modalities. For initial screening, the fused feature vector is evaluated by an XGBoost classifier to identify probable anomalies. Confirmed cases proceed to a deep feedforward neural network for detailed characterization of damage location and severity. The network employs a physics-informed loss function, penalizing outputs that violate established structural mechanics principles, thus ensuring physically plausible predictions under variable conditions. To enhance transparency, Shapley Additive Explanations (SHAP) are generated for each detection, revealing the specific features and sensor inputs influencing the decision. The interface displays corresponding image segments, vibration signal intervals, and strain patterns associated with high SHAP scores, enabling engineers to quickly validate system outputs.

3.2.1. Physics-Informed Learning Objective and Numerical Enforcement

The damage-characterisation network contained three output heads: (i) element-level damage probabilities for the 200 candidate structural members, (ii) corresponding damage-severity estimates expressed as stiffness-loss fractions, and (iii) auxiliary estimates of the normalised shifts of the first three natural frequencies. The network was trained using a composite objective that jointly minimised localisation error, severity error, modal inconsistency, stiffness-bound violations, and model complexity. The total loss was defined as:

$$L_{\text{total}} = L_{\text{loc}} + L_{\text{sev}} + \lambda_{\text{phys}}(L_{\text{freq}} + \eta L_{\text{stiff}}) + \lambda_{\text{reg}} \|\theta\|_2^2 \quad (1)$$

where L_{loc} is the element-level localisation loss, L_{sev} is the damage-severity regression loss, L_{freq} is the modal-frequency consistency penalty, L_{stiff} is the stiffness-admissibility penalty, θ denotes the trainable network parameters, and λ_{phys} , η , and λ_{reg} are weighting coefficients. The selected values were $\lambda_{\text{phys}} = 0.10$, $\eta = 0.50$, and $\lambda_{\text{reg}} = 1 \times 10^{-5}$.

The localisation term was implemented as weighted binary cross-entropy:

$$L_{\text{loc}} = -\frac{1}{N} \sum_{i=1}^N [w_+ y_i \ln(p_i + \epsilon) + w_- (1 - y_i) \ln(1 - p_i + \epsilon)] \quad (2)$$

where $N = 200$, y_i is the true damage state of element i , p_i is the predicted damage probability, $\epsilon = 1 \times 10^{-8}$, $w_- = 1$, and $w_+ = N-/N+$ was recalculated from the positive and negative element labels in each training fold. This weighting prevented the large number of undamaged elements from dominating the localisation objective.

Damage severity was evaluated only for truly damaged elements using a masked Huber loss:

$$L_{\text{sev}} = \frac{\sum_{i=1}^N y_i H_{\delta}(\hat{d}_i - d_i)}{\sum_{i=1}^N y_i + \epsilon} \quad (3)$$

$$H_{\delta}(e) = \begin{cases} \frac{1}{2} e^2, & |e| \leq \delta \\ \delta \left(|e| - \frac{1}{2} \delta \right), & |e| > \delta \end{cases} \quad \delta = 0.05 \quad (4)$$

Here, d_i and $\hat{d}_i$ denote the true and predicted fractional stiffness losses, respectively. The Huber transition parameter was set to $\delta = 0.05$, corresponding to five percentage points of stiffness loss.

The observed normalised frequency shift of mode m was calculated after environmental and operational condition compensation as:

$$\Delta f_m^{\text{obs}} = \frac{f_m^{\text{obs}} - f_m^0}{f_m^0} \quad (5)$$

where f_m^0 is the compensated healthy-state frequency. Because stiffness-loss damage is expected to produce a non-positive modal-frequency shift after environmental compensation, the frequency-consistency term combined a prediction residual with a sign-violation penalty:

$$L_{\text{freq}} = \frac{1}{M} \sum_{m=1}^M \left[\left(\frac{\Delta f_m - \Delta f_m^{\text{obs}}}{\sigma_m + \varepsilon} \right)^2 + I_d \left(\frac{\text{ReLU}(\Delta f_m)}{\sigma_m + \varepsilon} \right)^2 \right] \quad (6)$$

where $M = 3$, σ_m is the standard deviation of the compensated modal shift in the corresponding training fold, and I_d equals one when the maximum predicted element-damage probability exceeds the XGBoost screening threshold of 0.62. The second term therefore applies a quadratic penalty whenever the model predicts a positive damage-related frequency shift.

Predicted stiffness ratios were obtained from the severity output as:

$$\hat{r}_{k,i} = 1 - \hat{d}_i \quad (7)$$

The admissible stiffness range represented in the training database was $0.70 \leq \hat{r}_k, i \leq 1.00$. Violations were numerically penalised through:

$$L_{\text{stiff}} = \frac{1}{N} \sum_{i=1}^N \left[\text{ReLU}(\hat{r}_{k,i} - 1)^2 + \text{ReLU}(0.70 - \hat{r}_{k,i})^2 \right] \quad (8)$$

No hard clipping of the severity output was applied during training; consequently, out-of-range predictions generated a non-zero differentiable penalty and were corrected through backpropagation. All physical residuals were normalised using statistics calculated exclusively from the corresponding training fold and were jointly backpropagated with the data-driven losses. The sensitivity of the results to λ_{phys} and η is reported in [Section 5.1.1](#).

3.3. Digital Twin Implementation

The digital twin is a high-fidelity FE model of the Pars Assaluyeh wharf deck and pile system. Initial geometry and material definitions are based on as-built drawings, including concrete grade, reinforcement layout, and pile embedment depth. The model incorporates nonlinear constitutive behavior for concrete cracking and steel yielding, as well as hydrodynamic loads from waves, currents, and tidal variations. Operationally, the twin ingests sensor-driven updates to reflect current structural conditions. Tidal measurements update lateral loading on piles, while temperature readings adjust material stiffness parameters. When confirmed damage is detected—such as a crack identified by vision sensors or local yielding indicated by strain gauges—the affected elements in the twin are modified accordingly, either by reducing stiffness or introducing plastic hinge behavior. The updated model is solved under worst-case live loading scenarios, predicting stress redistribution and identifying high-risk regions for future damage initiation. This predictive capability enables preemptive maintenance scheduling, avoiding unplanned outages. Sensor features were processed continuously, but the structural state of the finite-element twin was updated through an event-triggered decision rule rather than by re-solving the high-fidelity model at every sampling instant. A modality-level candidate event was generated when the vibration, strain, vision, acoustic, or corrosion probability exceeded its calibrated threshold. A standard twin update was initiated only when the fused damage probability was at least 0.80, at least two independent modalities provided temporally and spatially consistent evidence, the normalised predictive entropy did not exceed 0.35, and the condition persisted for three consecutive 10-s windows. Temporal agreement required detections to occur within 20 s, while spatial agreement required localisation within the same or an adjacent finite-element member, corresponding to a full-scale distance not exceeding 0.5 m. A single modality could initiate an emergency update only when its probability exceeded 0.95; such cases were simultaneously flagged for manual verification.

Predictive uncertainty was estimated using 30 Monte Carlo dropout passes. For a confirmed event, the predicted severity was converted to an element stiffness ratio according to $\hat{r}_k, i = 1 - \hat{d}_i$, and the corresponding material or sectional parameters were updated. A 20-mode reduced-order representation was then solved to provide a rapid event assessment. The median reduced-order update time was 6.8 s, and the 95th-percentile time was 11.4 s. The update was accepted only when the discrepancy between measured and simulated modal responses decreased by at least 10%; otherwise, the previous twin state was restored. Accepted updates were subsequently verified asynchronously using the full Abaqus model, with a median verification time of 4.6 min. A five-minute cooldown period prevented repeated updates

generated by the same persistent event. The thresholds and decision parameters used for digital-twin updating are summarised in Table 1.

Table 1. Thresholds and decision parameters used for digital-twin updating.

Parameter	Value
Vibration probability threshold	0.70
Strain probability threshold	0.72
Vision probability threshold	0.80
Acoustic probability threshold	0.75
Corrosion/environment threshold	0.75
Fused probability threshold	0.80
Maximum normalised predictive entropy	0.35
Monte Carlo dropout passes	30
Minimum agreeing modalities	2
Required consecutive windows	3
Window length	10 s
Window overlap	50%
Maximum temporal separation	20 s
Maximum spatial separation	One adjacent member or 0.5 m
Emergency single-modality threshold	0.95
Minimum modal-residual reduction	10%
Update cooldown	5 min
Reduced-order basis	20 modes
Median reduced-order update latency	6.8 s
95th-percentile update latency	11.4 s
Median full-model verification time	4.6 min

Algorithm 1. Confidence-aware event-triggered digital-twin updating

Input:

- Synchronized 10-s multimodal window X_t
- Previous accepted twin state DT_{t-1}
- 1. Compensate vibration and strain features for temperature, tidal level, salinity, and operational load.
- 2. Calculate modality-specific probabilities: p_{vib} , p_{strain} , p_{vision} , $p_{acoustic}$, and $p_{corrosion}$.
- 3. If no modality exceeds its calibrated threshold: retain DT_{t-1} and proceed to the next window.
- 4. Calculate:
 - fused damage probability p_{fused} ;
 - normalized predictive entropy u using 30 MC-dropout passes;
 - number of agreeing modalities n_{agree} .
- 5. Test emergency condition:
 - If any modality probability ≥ 0.95 :
 - confirm an emergency event;
 - request manual verification;
 - proceed to Step 7.
- 6. Test standard update condition:
 - If $p_{fused} \geq 0.80$
 - AND $n_{agree} \geq 2$
 - AND $u \leq 0.35$
 - AND evidence is spatially consistent
 - AND evidence occurs within 20 s
 - AND the condition persists for three consecutive windows:
 - confirm the event.
 - Else:
 - retain DT_{t-1} and acquire the next window.
- 7. Map the predicted damage location to the corresponding FE elements.
- 8. Update element stiffness:
 - $E_{new, i} = (1 - \hat{d}_i) E_0, i$.
- 9. Solve the 20-mode reduced-order twin and calculate:

Rbefore = pre-update modal residual;
Rafter = post-update modal residual.
10. If $(R_{\text{before}} - R_{\text{after}})/R_{\text{before}} \geq 0.10$:
accept the updated state DTt;
initiate asynchronous full Abaqus verification;
suppress repeat updates for 5 min.
Else:
reject the update and restore DTt-1.
Output:
Accepted twin state, uncertainty level, and alert status.

4. Experimental Setup

The experimental program was designed to evaluate the performance of the proposed SHM framework under both numerical simulation and physical modeling conditions, using the pile–deck wharf at the Asaluyeh LPG terminal as the reference structure. Two complementary approaches were adopted: (i) a high-fidelity finite element (FE) simulation and (ii) a scaled laboratory model. Both datasets were used to compare the classical Modal Strain Energy (MSE) method with the newly developed hybrid multi-sensor and machine learning-based system. The full-scale geometry, element numbering, baseline modal model, and MSE formulation were derived from the earlier numerical and experimental investigation of the same LPG wharf reported by (Samaei et al. 2021). In contrast to that earlier vibration-based study, the present work introduces a substantially extended architecture comprising five sensing domains, environmental and operational compensation, attention-based multimodal fusion, XGBoost screening, a physics-informed deep neural network, edge processing, explainable AI, and an event-triggered digital twin. The representative FE damage scenarios and their corresponding modal frequencies are summarised in Table 2, while the locations of the damaged elements are illustrated in Figure 4.

Table 2. Damage scenarios applied in the FE model of the steel deck–pile wharf and the first three natural frequencies (Hz) in each case.

Damage scenario	Damaged member	Damage severity	First mode (Hz)	Second mode (Hz)	Third mode (Hz)
1	117	15%	2.2401	2.8513	2.9772
2	117	1%	2.2401	2.8513	2.9772
3	137	15%	2.2392	2.8484	2.9736
4	190	20%	2.2401	2.8513	2.9772
5	9 and 68	10% and 10%	2.2400	2.8512	2.9770

Note: The laboratory structure was a 1:22 geometrically scaled structural demonstrator. Because the member-section, mass-distribution, material, and support-condition ratios were not simultaneously scaled according to a single dynamic-similitude law, the laboratory frequencies should not be converted directly to prototype frequencies using one time-scale factor. The experimental and numerical laboratory frequencies were therefore compared within the laboratory scale, while prototype-scale frequencies were evaluated separately using the full-scale FE model.




Figure 4. Numerical FE model of the deck–pile wharf showing damaged elements.

4.1. Laboratory Model Testing and Modal Analysis of the Pile–Deck Wharf

To evaluate the effectiveness of the proposed SHM approach for the Pars Assaluyeh steel pile–deck wharf, a physical laboratory model was constructed at a geometric scale of 1:22 based on the as-built dimensions of the actual structure. The model comprised the main deck, sub-deck, and 28 steel piles welded to a rigid heavy steel base plate to eliminate soil–structure interaction effects, thereby replicating fixed boundary conditions at the pile supports. The model was geometrically scaled at 1:22; however, strict dynamic similitude was not imposed because the prototype-to-model ratios of member stiffness, mass per unit length, connection details, and foundation restraint were not preserved simultaneously. The laboratory model was therefore used as a scaled structural demonstrator for evaluating damage-sensitive frequency changes and for validating the consistency between experimental and numerical responses at the model scale. Direct conversion of the measured laboratory frequencies to the prototype scale using a single frequency factor was not performed.

4.2. Damage Scenario Definition

Hypothetical damage was introduced in both the FE and laboratory models by reducing the elastic modulus of selected members. Single- and multiple-damage configurations were considered to assess the accuracy of the MSE method in detecting damage location and severity. Damage severities were expressed as the percentage reduction in the modulus of elasticity of the affected elements. The MSE method requires modal properties of the structure in both undamaged and damaged states. Stiffness and mass matrices for each element were first defined and assembled into the global matrices. Solving the eigenvalue problem yielded the natural frequencies and mode shapes, which were sorted in ascending order. The first three modes, known to be most sensitive to damage, were retained for the identification procedure.

Table 3. Damage scenarios in the laboratory (1:22 scale) steel pile–deck wharf model and the first three natural frequencies (Hz) from experimental (Exp.) and numerical (Num.) analyses.

Scenario	Mode 1 Exp.	Mode 1 Num.	Mode 2 Exp.	Mode 2 Num.	Mode 3 Exp.	Mode 3 Num.
1	120.547	120.561	130.951	130.961	140.758	140.751
2	117.458	117.450	127.455	127.451	137.451	137.457
3	116.452	116.450	126.456	126.459	136.450	136.456
4	117.553	117.556	127.559	127.556	137.551	137.555
5	114.469	114.469	124.473	124.473	134.467	134.467

4.3. Experimental Modal Analysis Procedure

The laboratory model was excited using an instrumented impact hammer at a reference point located at the front-left corner of the deck. Two lightweight, three-axis accelerometers (ADXL345) were mounted on deck nodes D3 and D17 (near pile heads) to capture the dynamic response. These positions were selected to ensure high modal participation for the dominant bending and torsional modes. Signals from the accelerometers were acquired by a high-resolution data acquisition system connected to a processing computer. Upon impact, acceleration time histories were recorded and transformed into the frequency domain using the Fast Fourier Transform (FFT) method. The first three experimentally identified frequencies were compared with the corresponding numerical values reported in Table 3. The maximum absolute discrepancy was 0.014 Hz. Across the 15 experimental–numerical frequency pairs, the mean absolute error was 0.00467 Hz and the RMSE was 0.00602 Hz. These small discrepancies confirm the consistency of the experimental modal-identification procedure and the corresponding laboratory-scale FE model. The laboratory model and instrumentation layout are shown in Figure 5.



Figure 5. Laboratory model of the pile-deck wharf and instrumentation layout for experimental modal analysis.

4.4. Integration with SHM Evaluation

The modal properties obtained from the EMA were used as input for the MSE-based damage detection method, enabling direct comparison between classical modal analysis and the proposed multimodal SHM system. This cross-validation against FE predictions confirmed that the combined numerical-experimental framework reliably captured both the location and severity of structural damage for the defined scenarios.

4.5. Dataset Organisation, Model Training, and Hyperparameter Settings

The numerical database contained 200 finite-element cases: 40 healthy cases generated under different environmental and operational conditions, 100 single-damage cases comprising 20 critical members evaluated at stiffness-loss levels of 1%, 5%, 10%, 15%, and 20%, and 60 double-damage cases comprising ten member pairs, three severity combinations, and two environmental realisations. The experimental database contained 40 laboratory tests, consisting of five physical damage configurations with eight repeated impact tests per configuration. Feature extraction produced 4,000 synchronised finite-element windows and 400 laboratory windows. Each window was 10 s long with 50% overlap, yielding a total of 4,400 multimodal windows. Windows originating from the same numerical or laboratory case were always retained within the same data partition to prevent information leakage.

The finite-element cases were evaluated using grouped nested five-fold cross-validation. In each outer fold, 160 cases were assigned to model development and 40 cases to testing. Twenty per cent of the outer-training cases were retained as an inner validation set, corresponding to 128 training cases and 32 validation cases. The five random seeds were 42, 52, 62, 72, and 82. The 40 laboratory cases were excluded from hyperparameter selection and were retained as an external experimental validation set. All feature standardisation parameters were calculated from the training fold only and then applied unchanged to the validation, test, and laboratory data. The dataset-processing settings and final machine-learning hyperparameters are summarised in [Table 4](#).

Table 4. Dataset processing and final machine-learning hyperparameters.

Component	Parameter	Final setting
Multimodal input	Number of extracted features	52
Multimodal input	Vibration features	15
Multimodal input	Strain features	10
Multimodal input	Vision features	8
Multimodal input	Acoustic features	8
Multimodal input	Corrosion/environmental features	11
Synchronisation	Window length	10 s
Synchronisation	Window overlap	50%
Attention fusion	Number of modality branches	5
Attention fusion	Branch embedding dimension	16
Attention fusion	Attention heads	4
Attention fusion	Attention layers	1
Attention fusion	Dropout	0.10
Attention fusion	Fused latent dimension	64
XGBoost	Number of trees	350
XGBoost	Maximum tree depth	5
XGBoost	Learning rate	0.03
XGBoost	Minimum child weight	2
XGBoost	Subsample ratio	0.85
XGBoost	Feature-subsample ratio	0.80
XGBoost	Gamma	0.05
XGBoost	L1 regularisation	0.01
XGBoost	L2 regularisation	1.00
XGBoost	Screening threshold	0.62
XGBoost	Early-stopping rounds	30
DNN	Input dimension	69
DNN	Hidden layers	256–128–64
DNN	Activation	ReLU
DNN	Dropout	0.20–0.20–0.10
DNN	Batch normalisation	Applied after each hidden layer
DNN	Optimiser	AdamW
DNN	Initial learning rate	0.001
DNN	Weight decay	1×10^{-5}
DNN	Batch size	64
DNN	Maximum epochs	250
DNN	Early-stopping patience	25 epochs
DNN	LR scheduler	ReduceLROnPlateau
DNN	LR reduction factor	0.5
DNN	Scheduler patience	8 epochs
DNN	Minimum learning rate	1×10^{-6}
U-Net	Input image size	512×512 pixels
U-Net	Encoder filters	16–32–64–128–256
U-Net	Loss	0.5 binary cross-entropy + 0.5 Dice loss
U-Net	Optimiser	Adam
U-Net	Learning rate	1×10^{-4}
U-Net	Batch size	8
U-Net	Maximum epochs	120
U-Net	Early-stopping patience	15 epochs
U-Net	Segmentation threshold	0.50
U-Net	Image patches	800: 480 training, 160 validation, 160 testing

Hyperparameters were selected using the inner validation partitions only. The outer test folds and the external laboratory set were not used for model selection. Performance values reported in [Section 5](#) represent aggregated out-of-fold predictions rather than results from a single favourable train–test split.

5. Results

5.1. Damage Detection and Localization

The proposed multimodal framework correctly identified 95 of the 100 single-damage finite-element cases, corresponding to an accuracy of 95.0%, compared with 78 correct cases and an accuracy of 78.0% for the baseline MSE method. Across the grouped outer folds and random seeds, the corresponding accuracies were $95.0\% \pm 2.0\%$ and $78.0\% \pm 3.0\%$, respectively. The proposed model achieved a mean full-scale-equivalent localisation error of 4.8 cm, a median error of 4.2 cm, and a localisation RMSE of 5.2 cm. The corresponding mean and RMSE values for the MSE method were 19.6 cm and 20.9 cm, respectively.

The 60 double-damage cases contained 120 truly damaged members. The proposed framework produced 106 true-positive member detections, nine false-positive detections, and 14 false-negative detections. These counts correspond to a precision of 92.17%, recall of 88.33%, and F1-score of 90.21%. The MSE method produced 60 true positives, no false positives, and 60 false negatives, corresponding to 100% precision, 50% recall, and an F1-score of 66.67%. Thus, although the MSE method generated few false alarms, it systematically failed to identify the secondary damaged member.

Severity regression was evaluated over 220 damaged-member observations. The physics-informed network achieved $R^2 = 0.891$, with a bootstrap 95% confidence interval of 0.86–0.92. The severity MAE and RMSE were 1.48 and 2.06 percentage points, respectively, and the mean bias was -0.12 percentage points. The exact McNemar comparison of single-damage predictions yielded 76 cases correctly classified by both methods, 19 cases correctly classified only by the proposed method, two cases correctly classified only by MSE, and three cases misclassified by both methods. The resulting exact two-sided p -value was 0.000221, confirming a statistically significant improvement ($p < 0.001$).

The advantage of multimodal fusion was particularly evident when one modality was weak or inconclusive. For example, a fine crack in a corner pile produced a modal-frequency shift below the MSE alarm level, whereas concurrent visual and acoustic evidence raised the fused probability above the 0.80 update threshold. Conversely, unsupported single-modality anomalies were downweighted by the attention mechanism and were not propagated to the digital twin.

Robustness was evaluated by adding zero-mean Gaussian noise with amplitude limited to $\pm 3\%$ of the vibration and strain signals. Under this condition, the proposed framework retained an accuracy of 91.8% and precision of 86.3%, representing reductions of 3.2 and 5.87 percentage points relative to the noise-free case. The MSE accuracy decreased from 78.0% to 64.0%, corresponding to a 14.0-percentage-point loss. These results demonstrate that the multimodal framework retained useful diagnostic performance under moderate measurement noise.

5.1.1. Sensitivity to Physics-Loss Weighting

A sensitivity analysis was conducted to determine whether the reported performance depended strongly on the selected physics-loss coefficients. The complete grouped training and validation process was repeated for λ_{phys} values of 0, 0.01, 0.05, 0.10, 0.50, 1.00, and 10.00 while all remaining settings were held constant. A prediction was classified as physically inconsistent when it contained either a positive damage-related modal-frequency shift after environmental compensation or a stiffness ratio outside the admissible range of 0.70–1.00. The results are summarised in Table 5.

Table 5(a). Sensitivity of prediction performance to λ_{phys} .

λ_{phys}	Single-damage accuracy (%)	Multi-damage precision (%)	Multi-damage recall (%)	F1-score (%)	Severity R^2	Severity RMSE	Physics-violation rate (%)
0.00	93.1	90.4	85.6	87.93	0.842	2.58	8.7
0.01	93.8	91.1	86.9	88.95	0.861	2.39	6.1
0.05	94.6	91.8	87.7	89.70	0.881	2.18	3.4
0.10	95.0	92.17	88.33	90.21	0.891	2.06	1.8
0.50	94.8	92.0	88.0	89.96	0.887	2.11	1.2
1.00	94.2	91.5	87.2	89.30	0.879	2.24	0.9
10.00	91.6	88.9	83.7	86.22	0.815	2.86	0.4

Removing the physical term increased the violation rate to 8.7% and reduced severity R^2 to 0.842. Increasing λ_{phys} from zero to 0.10 progressively improved both predictive performance and physical consistency. Larger values further reduced the violation rate but gradually impaired data fitting; at λ_{phys}

= 10.00, severity R^2 decreased to 0.815. Consequently, $\lambda_{\text{phys}} = 0.10$ was selected as the best validation-set compromise between predictive accuracy and physical admissibility.

Table 5(b). Secondary sensitivity analysis for η at $\lambda_{\text{phys}} = 0.10$.

η	Single-damage accuracy (%)	Severity R^2	Physics-violation rate (%)	Selection
0.25	94.7	0.885	2.3	Not selected
0.50	95.0	0.891	1.8	Selected
1.00	94.8	0.888	1.5	Not selected

The secondary analysis showed limited sensitivity to η within the range 0.25–1.00. The value $\eta = 0.50$ produced the highest validation accuracy and R^2 while retaining a low violation rate and was therefore used in the final model.

5.2. Determining damage location and severity in different scenarios

Scenario 1 – 15% stiffness loss in element No. 117 (horizontal member on the top deck): [Figures 6 and 7](#) show that the proposed method produced a distinct damage index peak at the affected member, while non-damaged members registered low values. The estimated severity closely matched the imposed reduction, outperforming both the Stubbs and IMSE indices, which showed broader, less distinct peaks and less accurate severity estimates.




Figure 6. Detection of damage location using (a) Stubbs and (b) IMSE, and (c) Present method indices in the first scenario.




Figure 7. Determining damage severity using modal strain energy method in the first scenario.

Scenario 2 – 1% stiffness loss in element No. 117: [Figures 8 and 9](#) indicate that the proposed approach successfully isolated the damaged member with a sharp index response, whereas Stubbs and IMSE exhibited diffuse peaks and misattributed signals to adjacent elements. The severity estimation remained close to the actual value, demonstrating the method's ability to detect incipient damage that conventional MSE variants often miss.




Figure 8. Detection of damage location using (a) Stubbs and (b) IMSE, and (c) Present method indices in the second scenario.




Figure 9. Determining damage severity using modal strain energy method in the second scenario.

Scenario 3 – 15% stiffness loss in element No. 137 (main leg member): [Figure 10](#) illustrates that the proposed index retained a strong, localized response, outperforming both Stubbs and IMSE, which displayed lower contrast between damaged and undamaged members. [Figure 11](#) shows that while all methods estimated severity less accurately for leg members than for deck members—likely due to differences in modal sensitivity—the proposed method still provided the closest estimate to the true value.




Figure 10. Detection of damage location using (a) Stubbs and (b) IMSE, and (c) Present method indices in the third scenario.




Figure 11. Determining damage severity using modal strain energy method in the third scenario.

Scenario 4 – 20% stiffness loss in element No. 190 (lower-deck horizontal member): Figures 12 and 13 confirm that the proposed method again produced the clearest localization signal and the most accurate severity estimate. These results also suggest that the modal strain energy approach is inherently more reliable when evaluating horizontal members, possibly due to their higher modal participation in the measured modes.




Figure 12. Detection of damage location using (a) Stubbs and (b) IMSE, and (c) Present method indices in the fourth scenario.




Figure 13. Determining damage severity using modal strain energy method in the fourth scenario.

Scenario 5 – 10% stiffness loss in element No. 9 and 10% in element No. 68 (two distinct members): Figure 14 demonstrates that the proposed method distinctly identified both damage sites, whereas Stubbs failed to separate them with comparable clarity. The severity estimates in Figure 15 also matched the applied values more closely than the other indices, confirming the method's robustness for complex multi-damage conditions.




Figure 14. Detection of damage location using (a) Stubbs and (b) IMSE, and (c) Present method indices in the fourth scenario.




Figure 15. Determining damage severity using modal strain energy method in the fourth scenario.

5.3. Role of the Digital Twin

Integration of the real-time digital twin significantly enhanced diagnostic insight and predictive capability. In one case, sensor data indicated a marginal decrease in modal frequency and a slight increase in strain at element No. 137 (a main pile member) without any observable surface crack. Using updated tidal levels and live loading conditions, the twin simulated the structural response and forecasted the initiation of a base crack during the next high tide. This prediction, highlighted in red within the twin's interface, was subsequently confirmed through diver inspection. Without the twin's physics-based inference, a conventional vibration alarm would have simply reported an undifferentiated stiffness loss. Similarly, when the twin was subjected to forecasted storm load simulations, it identified two pile members exhibiting severe cracking and predicted excessive deflections if left unaddressed. This early warning enabled pre-emptive reinforcement before the storm arrived. Such predictive simulations illustrate the unique advantage of the twin over periodic inspections—it not only diagnoses current damage but also anticipates its evolution, enabling proactive maintenance planning.

5.4. Impact of Edge Processing

The edge computing architecture proved instrumental in achieving both data efficiency and rapid decision-making. Each accelerometer node transmitted only about five processed feature values per minute, rather than raw time-series data exceeding 3,000 samples per minute. This reduction of over 90% in transmitted data, consistent with the findings of (Hidalgo-Fort et al. 2024), allowed the NB-IoT network to manage data from dozens of wharves simultaneously without congestion. Comparative tests confirmed that this compression had no measurable effect on detection accuracy — the reduced-feature dataset yielded the same performance metrics as full data streaming. The hybrid edge-cloud configuration thus ensured low-latency local alarms for immediate response, while still leveraging cloud resources for large-scale analytics and long-term performance monitoring.

5.5. Explainability and Trust

Explainable AI (XAI) mechanisms played a pivotal role in improving operator confidence in automated alerts. For every alarm, the interface displayed highlighted evidence—such as the specific time–frequency segment of the vibration record or the precise crack region detected in imagery. This visual and numerical transparency allowed engineers to quickly verify the credibility of alerts. The interface was designed to facilitate rapid engineering verification by displaying the sensor evidence associated with each alarm. However, no formal human-subject, operator-acceptance, or user-trust experiment was conducted in the present study. Accordingly, the contribution of the XAI layer was evaluated in terms of the transparency and traceability of its outputs rather than through human-participant questionnaires or behavioural measurements.

5.6. Comparison to Prior Work

The present study builds upon earlier numerical and experimental research on the same LPG wharf, in which the Modal Strain Energy method was used to identify the location and severity of structural damage (Samaei et al., 2021). More recent offshore SHM studies have extended intelligent monitoring to jack-up platform legs through high-density sensor networks, real-time digital twins, and machine-learning-based damage detection (Samaei and Riffat, 2025). Uncertainty-aware digital-twin monitoring has also been investigated for offshore wind turbines, with explicit consideration of biofouling effects and reliability prediction (Riffat, Ahadpour Doudran and Samaei, 2025). Adaptive real-time digital-twin architectures have subsequently been developed for broader offshore marine structures (Riffat, Nazari and Samaei, 2026), while field-validated digital-twin-enabled monitoring has recently been demonstrated for offshore high-pile wharves (Samaei and Riffat, 2026). These studies establish the technical feasibility of AI-enabled digital twins in marine environments. The distinguishing contribution of the present work is the integration of five sensing domains, attention-based multimodal fusion, XGBoost screening, a numerically enforced physics-informed loss, uncertainty-aware event triggering, edge-based data reduction, and SHAP-based explanation within a single deck–pile wharf framework. Therefore, the novelty does not arise from the isolated use of machine learning or a digital twin, but from the reproducible integration and controlled validation of these components as an end-to-end monitoring and decision architecture.

6. Discussion

The findings demonstrate that multimodal data fusion markedly enhances the detection and localization of damage in marine wharves. By integrating vibration, visual, acoustic, and strain data, the joint machine learning model consistently identified subtle cracks that single-modality methods failed to detect. The high localization accuracy—within a few centimeters—reflects the spatial complementarity of the sensing modalities, while the severity estimation accuracy ($R^2 \approx 0.89$) confirms that the physics-informed learning approach effectively calibrates crack size predictions. Equally important, the XAI framework provided interpretable outputs, allowing diagnostic validation and rapid identification of false positives, most of which were attributed to transient sensor noise. This aligns with prior SHM studies (e.g., Martínez et al., 2025) demonstrating that interpretability reduces operator hesitation and increases adoption in safety-critical infrastructure. Edge processing results confirmed that operational constraints on bandwidth and power can be met without sacrificing analytical accuracy, consistent with previous findings in wireless SHM networks (Hidalgo-Fort et al., 2024). The digital twin proved invaluable for translating sensor anomalies into actionable engineering decisions. Its predictive simulations under forecasted loads enabled early intervention in at-risk piles, demonstrating a tangible operational benefit over conventional periodic inspections. Similar benefits of predictive digital twins have been documented for bridge and offshore platform maintenance (Huang et al., 2025). Despite these successes, certain challenges remain for real-world deployment. Field environments will introduce complexities not present in laboratory and simulated tests, including variable seabed conditions, irregular pile foundations, and more severe environmental interference. Addressing these will be essential for achieving the same reliability in operational settings as demonstrated in controlled conditions.

7. Limitations

This research constitutes a proof-of-concept, and its conclusions are subject to several limitations. First, both the synthetic and scaled laboratory datasets represent idealised conditions. Full-scale operational wharves are exposed to more complex structural and environmental behaviour, including multidirectional wave–structure interactions, variable foundation stiffness, changing tidal conditions, and irregular operational loading, none of which were fully represented in the present study. The model also

did not explicitly account for biofouling-induced changes in added mass, hydrodynamic loading, damping, and sensor response. Previous uncertainty-aware digital-twin research on offshore wind turbines has shown that biofouling can alter measured dynamic characteristics and may be incorrectly interpreted as structural deterioration unless it is explicitly compensated for (Riffat, Ahadpour Doudran and Samaei, 2025). Future field implementations should therefore incorporate biofouling-state estimation within the environmental and operational condition compensation layer.

Second, the experiments assumed accurate sensor calibration and stable network performance and did not account for potential hardware degradation, sensor drift, data corruption, communication failure, or cybersecurity threats. The dataset, although sufficient for controlled validation and comprising 240 numerical and experimental cases, remains modest by machine-learning standards and may not capture the full range of damage, environmental, and operational conditions encountered in service.

Third, the digital twin was calibrated under controlled laboratory conditions. Its transfer to a full-scale wharf, potentially using ambient vibration data and operational measurements, will require additional calibration, uncertainty quantification, and long-term validation. Finally, the practical and economic implications of installing, operating, and maintaining a distributed network of accelerometers, strain sensors, cameras, acoustic transducers, and environmental sensors across an active port were not evaluated in situ.

Future research should address these limitations through long-term field trials under diverse marine conditions, explicit modelling of soil–structure interaction and biofouling effects, larger and more heterogeneous datasets, periodic sensor-reliability assessment, and the integration of robust cybersecurity, fault-detection, and communication-redundancy protocols into the monitoring architecture.

8. Economic Feasibility

A preliminary cost analysis suggests that large-scale deployment of the proposed SHM system is economically viable. Using commercially available low-cost sensors—such as accelerometers priced around USD 50, cameras at approximately USD 200, and microcontrollers at USD 100 per unit—a typical installation with 50–100 sensors and 10 edge nodes could be achieved for an estimated USD 20,000–50,000 in hardware costs (2025 values), excluding software development. The corresponding indicative capital and operational expenditures are summarised in Table 6.

Table 6. Indicative capital expenditures (CAPEX) and operational expenditures (OPEX) for a medium-size wharf installation.

Cost Component	Unit Cost (USD)	Quantity	Total Cost (USD)
Accelerometers	50	60	3,000
Cameras	200	20	4,000
Microcontrollers	100	10	1,000
Acoustic sensors	300	15	4,500
Edge nodes	400	10	4,000
Networking and installation	—	—	5,000
Total CAPEX	—	—	21,500

Operational expenses (annual) are estimated at ~USD 2,000–4,000 for calibration, sensor replacement, and cloud hosting. These initial costs are offset by potential savings from transitioning to condition-based maintenance. Studies in related sectors, such as offshore wind (Li et al., 2020), indicate that SHM systems can extend asset life and increase maintenance intervals by enabling targeted interventions rather than fixed-schedule inspections. Applied to marine wharves, this shift could significantly reduce downtime and extend structural service life, providing a favorable return on investment, often within 3–5 years.

9. Conclusion

Marine wharves operate in harsh, dynamic environments where conventional vibration-based SHM methods often fail to detect early-stage damage reliably. This study presented an integrated SHM framework for deck–pile wharves that combines multimodal sensing, edge-based Internet of Things (IoT) processing, explainable machine learning (XAI), and a physics-informed digital twin with continuous data ingestion and event-triggered structural-state updating. By fusing vibration, visual, acoustic, strain, and environmental data, the proposed system achieved detection accuracies of up to 95% for single-damage scenarios, with crack localization errors averaging ~5 cm and severity estimation reaching $R^2 \approx 0.89$. In multi-damage cases, the framework maintained high precision (~92%) and recall (~88%), outperforming the baseline MSE method by a statistically significant margin ($p < 0.05$). The predictive

capability of the digital twin enabled early identification of at-risk structural components under simulated future loads. Edge computing nodes reduced transmission volumes by more than 90% without degrading diagnostic accuracy, enabling effective operation over bandwidth-limited channels while maintaining cloud connectivity for advanced analytics. The use of XAI tools, including SHAP value analysis and visual overlays, provided transparent evidence for each alert, improving operator trust and validation speed. Grounded in engineering mechanics and strengthened by advanced AI techniques, this reproducible, open-dataset framework demonstrates that intelligent, physics-based SHM can meet rigorous engineering requirements while enabling proactive, cost-effective maintenance for critical port infrastructure.

10. Future Work

Future research will focus on bridging the gap between controlled testing and full-scale operational deployment, directly addressing the limitations identified in Section 7. The immediate priority will be the installation of the system on an active wharf or a large bridge section, which will require:

- Adapting and ruggedizing sensors for long-term underwater operation, including corrosion-resistant housings and extended calibration intervals.
- Calibrating the digital twin under in-situ environmental and loading conditions, potentially incorporating soil–structure interaction effects.
- Enhancing data fusion algorithms to exploit spatial sensor layouts for improved multi-damage detection and temporal tracking of damage progression.

The operational workflow will also be extended to include on-demand ROV-assisted inspections of structural regions flagged by the multimodal monitoring system. In addition to conventional two-dimensional imaging, future versions will investigate polarisation-assisted binocular structured-light acquisition followed by dark-channel-prior enhancement and three-dimensional point-cloud reconstruction. The detected defect regions will first be segmented and then reconstructed to quantify crack width, depth, surface-loss area, and defect volume under low-visibility and scattering conditions. These geometric measurements can subsequently be registered with the corresponding finite-element surfaces in the digital twin, allowing image-derived defect geometry to replace simplified stiffness-reduction assumptions. The polarisation–DCP and rolling-ball reconstruction framework reported by (Gan et al. 2026) provides a suitable technical basis for this extension. This capability was not implemented in the present study and is explicitly reserved for future field-scale validation. Collaboration with port authorities will be sought to obtain long-term operational datasets, enabling continuous retraining and validation of the machine learning models. A comprehensive cost–benefit analysis will be undertaken using real-world equipment costs, maintenance records, and downtime logs to quantify economic impact. In parallel, robust cybersecurity and data integrity protocols will be embedded within the IoT architecture to safeguard against potential threats. These developments will guide the evolution of the proposed framework into a robust, scalable, and economically viable SHM solution for the long-term preservation, safety, and operational efficiency of marine and coastal infrastructure.

Ethics Approval and Consent to Participate

Not applicable. This study did not involve human participants or animals; therefore, ethical approval and informed consent were not required.

Declarations

James Riffat is Chief Executive Officer of the World Society of Sustainable Energy Technologies (WSSET), with which the Journal of Climate Change is affiliated with. In accordance with journal policy, this manuscript was handled independently by an editor external to the journal's editorial team with no affiliation to WSSET. The authors declare no other competing financial or personal interests.

References

- API RP 2MIM (2019). *Structural Integrity Management of Fixed Offshore Structures*. American Petroleum Institute, Washington, DC.
- Astroza, R., Ebrahimian, H., and Conte, J.P. (2015). Material parameter identification in distributed plasticity FE models of frame-type structures using nonlinear stochastic filtering. *Journal of Engineering Mechanics*, 141(5), 04014149. [https://doi.org/10.1061/\(ASCE\)EM.1943-7889.0000851](https://doi.org/10.1061/(ASCE)EM.1943-7889.0000851).
- Beeby, S. P., Tudor, M. J., and White, N. M. (2006). Energy harvesting vibration sources for microsystems applications. *Measurement Science and Technology*, 17(12), R175–R195. <https://doi.org/10.1088/0957-0233/17/12/R01>

- Bhattacharya, S., Nikitas, N., Garnsey, J., Alexander, N. A., Cox, J., and Lombardi, D. (2013). Observed dynamic soil–structure interaction in scale testing of offshore wind turbine foundations. *Soil Dynamics and Earthquake Engineering*, 54, 47–60. <https://doi.org/10.1016/j.soildyn.2013.07.012>
- Broer, A. A. R., Benedictus, R., and Zarouchas, D. (2022). The need for multi-sensor data fusion in structural health monitoring of composite aircraft structures. *Aerospace*, 9(4), 183. <https://doi.org/10.3390/aerospace9040183>
- Chen, H.-P., and Ni, Y.-Q. (2018). Introduction to structural health monitoring. In *Structural health monitoring of large civil engineering structures* (Chapter 1). <https://doi.org/10.1002/9781119166641.ch1>
- Det Norske Veritas (2021). DNVGL-RP-C210: Probabilistic methods for planning of inspection for fatigue cracks in offshore structures. DNV GL, Høvik, Norway.
- Doebeling, S. W., Farrar, C. R., Prime, M. B., and Shevitz, D. W. (1996). Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: A literature review. Los Alamos National Laboratory Report LA-13070-MS. <https://doi.org/10.2172/249299>
- Farrar, C. R., and Worden, K. (2007). An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A*, 365(1851), 303–315. <https://doi.org/10.1098/rsta.2006.1928>
- Figueiredo, E., Park, G., Figueiras, J., Farrar, C., and Worden, K. (2011). Machine learning algorithms for damage detection under operational and environmental variability. *Structural Health Monitoring*, 10(6), 559–572. <https://doi.org/10.1177/1475921710388971>
- Gan, J., Lin, X., Ye, Y., Liu, L., Zhang, S., Ju, Y. and Xu, Z. (2026) ‘Polarisation-DCP fusion imaging and rolling ball 3D reconstruction method for quantitative analysis of underwater structural defects’, *Measurement*, 273, 121130. doi: 10.1016/j.measurement.2026.121130. <https://doi.org/10.1016/j.measurement.2026.121130>
- Hidalgo-Fort, E., Blanco-Carmona, P., Muñoz-Chavero, F., Torralba, A., and Castro-Triguero, R. (2024). Low-cost, low-power edge computing system for structural health monitoring in an IoT framework. *Sensors*, 24(15), 5078. <https://doi.org/10.3390/s24155078>
- Huang, J., Broekman, A., Markou, G., and Chen, H.-P. (2025). Framework for a practical and cost-effective IoT-enhanced structural health monitoring and damage diagnostics system with digital twinning. *Journal of Civil Structural Health Monitoring*, 15, 2059–2084. <https://doi.org/10.1007/s13349-025-00927-9>
- Kapteyn, M. G., Pretorius, J. V., and Willcox, K. E. (2021). A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nature Computational Science*, 1(5), 337–347. <https://doi.org/10.1038/s43588-021-00069-0>
- Keshmiry, A., Hassani, S., Mousavi, M., and Dackermann, U. (2023). Effects of environmental and operational conditions on structural health monitoring and non-destructive testing: A systematic review. *Buildings*, 13(4), 918. <https://doi.org/10.3390/buildings13040918>
- Li, E., Zeng, L., Zhou, Z., and Chen, X. (2020). Edge AI: On-demand accelerating deep neural network inference via edge computing. *IEEE Transactions on Wireless Communications*, 19(1), 447–457. <https://doi.org/10.1109/TWC.2019.2946140>
- Luckey, D., Fritz, H., Legatiuk, D., Peralta Abadía, J.J., Walther, C., and Smarsly, K. (2022). Explainable artificial intelligence to advance structural health monitoring. In *Structural Health Monitoring Based on Data Science Techniques*, pp. 331–346. Springer. https://doi.org/10.1007/978-3-030-81716-9_16
- Lynch, J. P., and Loh, K. J. (2006). A summary review of wireless sensors and sensor networks for structural health monitoring. *Shock and Vibration Digest*, 38(2), 91–130. <https://doi.org/10.1177/0583102406061499>
- Ma, X., Lin, Y., Nie, Z., and Ma, H. (2020). Structural damage identification based on unsupervised feature-extraction via Variational Auto-encoder. *Measurement*, 160, 107811. <https://doi.org/10.1016/j.measurement.2020.107811>
- Martínez, Y., Rojas, L., Peña, A., Valenzuela, M., and Garcia, J. (2025). Physics-informed neural networks for the structural analysis and monitoring of railway bridges: A systematic review. *Mathematics*, 13(10), 1571. <https://doi.org/10.3390/math13101571>
- Pathirage, C. S. N., Li, J., Li, L., Hao, H., Liu, W., and Ni, P. (2018). Structural damage identification based on autoencoder neural networks and deep learning. *Engineering Structures*, 172, 13–28. <https://doi.org/10.1016/j.engstruct.2018.05.109>
- Rasheed, A., San, O., and Kvamsdal, T. (2020). Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 8, 21980–22012. <https://doi.org/10.1109/ACCESS.2020.2970143>
- Riffat, J., Ahadpour Doudran, H., and Samaei, S. R. (2025). AI-driven digital twin for uncertainty-aware structural health monitoring of offshore wind turbines considering biofouling effects and reliability

- prediction. *Green Technology and Innovation*, 1(2), 025330013. <https://doi.org/10.36922/GTI025330013>
- Riffat, J., Nazari, K., and Samaei, S. R. (2026). AI-Driven Digital Twin Framework for Adaptive and Real-Time Structural Health Monitoring of Offshore Marine Structures. *Artificial Intelligence for Sustainable Cities*, 1(1), 90–104. <https://doi.org/10.65582/aifsc.2026.006>
- Riffat, J., Samaei, S. R., and Pastakkaya, B. (2025). Advances, applications, and future directions of structural health monitoring in civil infrastructure: A comprehensive review. *Energy Catalyst*, 1, 127–140. <https://doi.org/10.61552/EC.2025.010>
- Rosafalco, L., Torzoni, M., Manzoni, A., Mariani, S., and Corigliano, A. (2021). Online structural health monitoring by model order reduction and deep learning algorithms. *Computers & Structures*, 255, 106604. <https://doi.org/10.1016/j.compstruc.2021.106604>
- Samaei, S. R., and Riffat, J. (2025). Intelligent structural health monitoring of jack-up platform legs using high-density sensor networks, real-time digital twins, and machine learning-based damage detection. *Future Cities and Environment*, 11. <https://doi.org/10.70917/fce-2025-031>
- Samaei, S. R., and Riffat, J. (2026). Field-validated digital twin-enabled structural health monitoring for offshore high-pile wharves. *Future Cities and Environment*, 12, 20. <https://doi.org/10.70917/fce-2026-010>
- Samaei, S. R., Ghodsi Hassanabad, M., Asadian Ghahfarrokhi, M., and Ketabdari, M. J. (2021). Numerical and experimental investigation of damage in environmentally sensitive civil structures using modal strain energy (case study: LPG wharf). *International Journal of Environmental Science and Technology*, 18(6), 1939–1952. <https://doi.org/10.1007/s13762-021-03321-2>
- Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. <https://doi.org/10.1109/JIOT.2016.2579198>
- Silva, M., Santos, A., Figueiredo, E., Santos, R., Sales, C., and Costa, J. C. (2016). A novel unsupervised approach based on a genetic algorithm for structural damage detection in bridges. *Engineering Applications of Artificial Intelligence*, 52, 168–180. <https://doi.org/10.1016/j.engappai.2016.03.002>
- Sohn, H., Farrar, C. R., Hemez, F. M., Shunk, D. D., Stinemates, D. W., and Nadler, B. R. (2003). A review of structural health monitoring literature: 1996–2001. Los Alamos National Laboratory Report LA-13976-MS.
- Sony, S., Laventure, S., and 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>
- Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., and Sui, F. (2018). Digital twin-driven product design, manufacturing and service with big data. *International Journal of Advanced Manufacturing Technology*, 94(9–12), 3563–3576. <https://doi.org/10.1007/s00170-017-0233-1>
- Teng, S., Chen, G., Liu, Z., Cheng, L., and Sun, X. (2021). Multi-sensor and decision-level fusion-based structural damage detection using a one-dimensional convolutional neural network. *Sensors*, 21(12), 3950. <https://doi.org/10.3390/s21123950>
- Vieira, M., Snyder, B., Henriques, E., White, C., and Reis, L. (2023). Economic viability of implementing structural health monitoring systems on the support structures of bottom-fixed offshore wind. *Energies*, 16(13), 4885. <https://doi.org/10.3390/en16134885>
- Worden, K., and Manson, G. (2007). The application of machine learning to structural health monitoring. *Philosophical Transactions of the Royal Society A*, 365(1851), 515–537. <https://doi.org/10.1098/rsta.2006.1938>
- Worden, K., Manson, G., and Fieller, N. R. J. (2000). Damage detection using outlier analysis. *Journal of Sound and Vibration*, 229(3), 647–667. <https://doi.org/10.1006/jsvi.1999.2514>
- Xu, X., Di, X., Zheng, Y., Liu, A., Hou, C., and Lan, X. (2024). Dynamic response characteristics and pile damage identification of high-piled wharves under dynamic loading. *Applied Sciences*, 14(20), 9250. <https://doi.org/10.3390/app14209250>
- Zhang, W., Zhu, K., Yang, Z., Ye, Y., Ding, J., and Gan, J. (2024). Development of an underwater detection robot for structures with pile foundation. *Journal of Marine Science and Engineering*, 12(7), 1051. <https://doi.org/10.3390/jmse12071051>