

BridgeGuard AI: A Self-Calibrating Edge-AI Prototype for Low-Cost Bridge Condition Monitoring and Digital-Twin Visualization

Prof. Pranali Nitnaware, Sameer Sonwane, Yash Khobragade, Trupti Kursunge, and Aboli Bodhale
Department of Industrial IoT, Priyadarshini College of Engineering, Nagpur, India

Abstract—Bridge deterioration is gradual and difficult to detect with periodic visual inspections, while professional electronic monitoring remains too expensive for small authorities and secondary structures. This paper presents a low-cost, self-calibrating structural condition monitoring prototype integrating sensing, on-device intelligence, continuous calibration, data fusion, degradation modelling, and digital-twin visualization in a single pipeline. An ESP32 edge node acquires vibration, tilt, strain, temperature, and humidity from an MPU6050 inertial sensor, an HX711-based strain gauge, and a DHT11, filters every channel with a Kalman filter, and runs a 5-3-5 TinyML autoencoder whose reconstruction error maps to a risk probability and a six-class event label. A host service resolves the drift-versus-load ambiguity with a multi-hypothesis Bayesian calibrator and a two-state Kalman decomposition, then fuses the corrected channels into an explainable Structural Health Index, timber-decay and fatigue-based remaining-useful-life estimates, and an Euler–Bernoulli digital twin streamed to a dashboard at 20 Hz. Anomaly-gated transmission reduces uploaded records by more than 99%. Bench hardware and physics-based simulator verification across five operating scenarios plus an injected-drift test demonstrate end-to-end behaviour on a prototype hardware budget of about USD 12.60.

Index Terms—Structural health monitoring, Edge AI, TinyML, autoencoder, Bayesian self-calibration, Kalman filter, sensor drift, digital twin, ESP32.

I. INTRODUCTION

A. Background and Motivation

Bridges are long-lived assets whose condition changes continuously. Repeated vehicle loading, wind-induced vibration, and daily and seasonal temperature swings alter strain and shape; over years, together with moisture ingress and material ageing, these actions reduce load-carrying capacity. Timber and composite-girder bridges are especially vulnerable to moisture-driven decay and cumulative fatigue. Periodic visual inspection, typically yearly or less frequent, is qualitative only: it cannot see fibre degradation and records no short overload events. Sensor-based structural health monitoring (SHM) offers a quantitative alternative, provided the data can be trusted [1].

B. Problem Statement

Three practical challenges hinder electronic SHM on ordinary bridges. First, cost: professional sensing, data-acquisition hardware, and manual calibration put complete systems out of reach for small authorities, educational projects, and many secondary structures. Second, sensor reliability: low-cost sensors suffer offset, slow drift, noise, and temperature/humidity sensitivity, and without correction these faults mislead. A strain gauge drifting by a few tens of microstrain is, from the strain channel alone, indistinguishable from a bridge that has genuinely settled by the same amount. Third, fragmentation: acquisition, analysis, visualization, and maintenance prediction are usually handled by separate tools, and a single parameter says little—rising vibration may be benign, whereas simultaneous changes in vibration, strain, tilt, and temperature are far more informative.

C. Contributions

This paper presents a low-cost, Edge-AI-based, self-calibrating bridge condition-monitoring prototype that closes these gaps in a single pipeline (Fig. 1). Its contributions are:

- (1) an ESP32 sensing node acquiring vibration, tilt, strain, temperature, and humidity, with per-channel Kalman filtering, a 3σ outlier gate, and hot-swap detection;
- (2) a 5-3-5 TinyML bottleneck autoencoder running on the device at ≈ 1.2 ms per inference, converting reconstruction error into a risk probability and a six-class event label;
- (3) an anomaly-gated transmission protocol that suppresses nominal frames at the source and transmits events plus a 60-s heartbeat;
- (4) a multi-hypothesis Bayesian calibrator that decides whether a strain shift is sensor drift or real loading, using vibration quiescence, tilt, and thermal compensation, applied gradually through a guarded learning rate;
- (5) a two-state Kalman decomposition separating the elastic structural signal from slow drift bias;
- (6) an explainable weighted fusion engine and an 8-6-3-6-8 host autoencoder producing a 0–100% Structural Health Index (SHI) with per-sensor contributions;
- (7) timber-decay and Miner-fatigue models yielding remaining-useful-life (RUL) estimates; and
- (8) an 11-node Euler–Bernoulli digital twin streamed to a web dashboard at 20 Hz with selective cloud persistence—on a hardware budget of about USD 12.60.

The contribution is architectural: the integration and transparent specification of a complete calibration-aware pipeline on very inexpensive hardware, rather than any single new algorithm.

II. RELATED WORK AND RESEARCH GAP

A. Related Work

Long-term wireless bridge monitoring has been demonstrated on full-scale structures [2][3], and an IoT low-power vibration system has been field-evaluated [4]. These works prove feasibility but assume professional-grade hardware. Calibration is usually treated as a procedure carried out around the measurement system rather than as a continuous function: reviews by Yang et al. [5] and Rizzo and Enshaiean [6] cover bridge-SHM calibration and its challenges, Yu et al. [7] design and calibrate wireless sensors for civil structures, and Kostić and Gül [8] show that temperature effects must be compensated before vibration-based neural damage detection. Learning-based damage identification from vibration data is well established [9], with recent reviews [10][11]; embedded analysis has even been demonstrated on the Z24 bridge [12]. Autoencoders [13] learn typical behaviour without labelled anomaly examples, and TinyML [14] makes them feasible on microcontrollers. Digital twins now support physics-informed modelling [15], machine-vision traffic monitoring [16], railway IoT integration [17], feature-fusion learning [18], hybrid model-data twins [19], and edge-cloud collaboration [20], within a systematic review of twin-based anomaly detection [21]; the paradigm originates in aerospace [22].

B. Research Gap

Four gaps remain. Gap 1: drift and damage are not separated in real time—reviews [5][6] and temperature studies [8] recognize the problem, but a continuous method that weighs competing explanations using the other channels and then corrects the baseline gradually is still needed. Gap 2: detection models assume trustworthy input, so uncorrected drift is read as damage. Gap 3: low-cost hardware is rarely the design target; edge-oriented studies [12][20] point the right way without making low-cost self-calibration central. Gap 4: calibration state is not connected to the digital twin, whose output is only as good as the data it mirrors. This work addresses all four gaps.

III. SYSTEM ARCHITECTURE AND HARDWARE

A. Three-Tier Architecture

The system has three tiers (Fig. 1). The edge tier samples and filters the sensors, runs the TinyML autoencoder, and decides what to send. The host tier is a Python 3.12 FastAPI service that receives JSON frames over a USB serial bridge (115200 bps) and runs an eight-stage ingestion pipeline: payload validation, hardware sanitization, Bayesian calibration, Kalman drift decomposition, fusion and SHI estimation, timber-decay modelling, RUL estimation, and digital-twin evaluation with cloud queuing. The cloud tier persists selected records in a Supabase PostgreSQL database and serves a browser dashboard updated over WebSocket.




Fig. 1. Edge–host–cloud architecture of the proposed system.

B. Edge Node and Sensors

The node is built around an ESP32-WROOM-32 module (two 32-bit Xtensa LX6 cores at 240 MHz, 520 KB SRAM, 4 MB flash). One core handles communications, serial streams, and housekeeping; the other provides deterministic 100-ms sensor ticks, Kalman filtering, and neural inference. A dual-watchdog timer resets the board if the I²C bus stalls for 3000 ms. The firmware image is approximately 642 KB and the board draws about 110 mA at 3.3 V. Table I lists the sensors. The MPU6050 uses GPIO 19/18 (I²C, 4.7 kΩ pull-ups), the DHT11 uses GPIO 4 with 100 nF decoupling, the HX711 uses GPIO 5/23, and a status LED on GPIO 2 blinks at 1 Hz in normal operation and stays lit on malfunction; all parts share one ground.

TABLE I
Sensors and Their Roles

Sensor	Configuration	Role in the system
MPU6050 6-DOF IMU	I ² C address 0x68; accelerometer ±2 g (16,384 LSB/g); gyroscope ±250 °/s (131 LSB/(°/s))	Vibration RMS and peak, crest factor, angular motion rate, deck pitch and roll
DHT11 (climate)	Single-wire protocol; 0–50 °C (±2 °C); 20–90 %RH (±5 %RH)	Thermal-strain compensation; humidity for moisture and decay modelling
HX711 + strain gauge	24-bit differential ADC; programmable gain 128; full-bridge Wheatstone	Strain (microstrain) and load response

C. Hot-Swap Detection

Because field sensors may be disconnected by weather, vibration, or vandalism, the firmware monitors the I²C bus for 1500 ms after each scan; if the MPU6050 does not acknowledge, its bit is cleared in an active-sensor bitmask and

all its channels are excluded from subsequent computation until it returns. The same mask drives the autoencoder loss and the host health index, so a missing sensor reduces the available evidence but never raises a false alarm.




Fig. 2. Bench prototype: timber truss bridge model instrumented with the ESP32 edge node (MPU6050, DHT11, HX711 strain interface) streaming to the bridge-health dashboard.

IV. EDGE SIGNAL PROCESSING AND TINYML ANOMALY DETECTION

A. Per-Channel Kalman Filtering

Each channel is filtered by its own one-dimensional Kalman filter [23] with a stationary process model ($F = 1$, no control input). Prediction and update follow (1)–(2), where $z[k]$ is the raw measurement.

$$\hat{x}^-[k] = \hat{x}[k-1], \quad P^-[k] = P[k-1] + Q \quad (1)$$

$$K[k] = \frac{P^-[k]}{P^-[k] + R}, \quad \hat{x}[k] = \hat{x}^-[k] + K[k](z[k] - \hat{x}^-[k]) \quad (2)$$

$$P[k] = (1 - K[k])P^-[k]$$

Measurement noise R is set well above process noise Q ($Q = 0.05$, $R = 2.00$ for vibration, strain, and tilt; $Q = 0.02$, $R = 3.00$ for wood moisture), so electrical jitter is attenuated while shock peaks and axle crossings are preserved; the moisture channel is heavily damped because wood moisture changes over hours or days.

B. Outlier Rejection

A 3σ innovation gate rejects isolated spikes from interference or bus errors, using the prediction and measurement uncertainties; if more than five consecutive samples exceed the gate, the filter treats the shift as a true step change (e.g., long-term sag), resets P to 5.0, and re-initializes at the new level.

C. Feature Normalization

Five features are normalized to $[0, 2]$ against engineering reference values so that no channel dominates: vibration RMS divided by 0.35 g (deck-vibration warning threshold); motion rate divided by 45 °/s; temperature deviation $|T - 22|$ divided by 15 °C; tilt divided by 2.5° (permissible beam rotation); and crest factor (peak/RMS) divided by 5, with values below 3 being normal.

D. Bottleneck Autoencoder and Probabilistic Risk

An autoencoder learns to reproduce its input through a lower-dimensional latent representation [13], so it reconstructs typical behaviour better than atypical behaviour. The edge network is a 5-3-5 bottleneck: a ReLU encoder, three latent

units associated with dynamic shock, rotational motion, and thermal/sag creep, and a purely linear decoder. All weights are static constants in flash; the C++ forward pass uses no matrix library and no dynamic allocation, giving a deterministic inference time of approximately 1.2 ms. The reconstruction error, averaged over physically connected (active) channels only, is

$$MSE = \frac{1}{N_{\text{active}} \sum_{i \in \text{Active}} (x_i - \hat{x}_i)^2} \quad (3)$$

and is mapped to a risk probability by a logistic function centred at the structural threshold $MSE = 0.20$ with steepness 12, clamped to $[0.01, 0.99]$:

$$P(\text{Risk}) = \text{clamp}\left(\frac{1}{1 + \exp(-12(MSE - 0.20))}, 0.01, 0.99\right) \quad (4)$$

$P(\text{Risk})$ rises from 0.14 at MSE 0.05 to 0.86 at MSE 0.35; the 0.15 transmission gate is reached at $MSE \approx 0.055$ and the 0.20 structural-warning level at $MSE \approx 0.084$. The gate is intentionally sensitive; its practical selectivity depends on how well the autoencoder is matched to each structure's normal data (Section VIII-F).

E. On-Chip Drift Probability and Event Classification

Tilt calibration degrades through thermal cycling and long-term creep, so drift is estimated only when the bridge is quiet. A quiescence factor derived from the running vibration RMS,

$$Q_s = 1 - \text{clamp}\left(\frac{\text{vibRMS}_{\text{run}} - 0.025}{0.055}, 0, 1\right) \quad (5)$$

suppresses the probability of tilt loss (6) whenever a vehicle is present, so a genuine dynamic deflection is never mistaken for drift.

$$P(\text{Drift}) = \alpha(0.35(\Delta\text{tilt}_{\text{baseline}} - 1.8)) \cdot Q_s \quad (6)$$

Each 100-ms sample is then assigned to one of six priority-ordered classes: VIBRATION_SHOCK (vibration RMS > 0.35 g or peak > 0.60 g—impact or seismic), DYNAMIC_MOTION (motion rate > 35 °/s—severe oscillation or wind resonance), TILT_EXCURSION ($|\text{tilt}| > 2.0^\circ$ —settlement or bearing failure), SENSOR_DRIFT ($P(\text{Drift}) \geq 0.70$ —zero-point creep, tare needed), STRUCTURAL_WARNING ($P(\text{Risk}) \geq 0.20$), and NOMINAL.

F. Anomaly-Gated Transmission

A frame leaves the edge only when $P(\text{Risk}) \geq 0.15$, the event class is not NOMINAL, or 60 s have elapsed since the last heartbeat; all other frames are suppressed at the source. Data reduction therefore begins before any radio or cable is used, and the full decision chain—acquire, filter, autoencode, classify, gate—runs on the microcontroller.

V. ADAPTIVE BAYESIAN SELF-CALIBRATION AND DRIFT DECOMPOSITION

A. Drift-Versus-Load Ambiguity

Suppose the strain gauge on a wooden beam shows a steady +25 $\mu\epsilon$ shift. This may be permanent overload deformation or adhesive creep and moisture penetration into the gauge housing. A simple threshold cannot separate the two causes and will either raise false alarms or hide real damage. The calibrator instead weighs the additional channels: real flexure

vibrates the bridge and rotates the beam ends, whereas drift occurs while the bridge stays quiet and level.

B. Hypotheses, Evidence, and Likelihoods

Three mutually exclusive, collectively exhaustive hypotheses with frequency-based priors are used: H_{drift} (sensor drift or adhesive creep, prior 0.10), H_{physical} (real traffic load or settlement, 0.15), and H_{normal} (ambient baseline variation, 0.75). Over a 30-sample sliding window (minimum 10 samples per decision) the evidence vector $E = [\Delta\epsilon, \bar{v}, \bar{\theta}, T]$ is formed, where $\Delta\epsilon$ is the mean strain shift from the current offset, $\bar{v}$ the mean vibration RMS, $\bar{\theta}$ the mean tilt, and T the temperature. Thermal expansion is compensated first; for timber along the grain a coefficient of $4.5 \mu\epsilon/^\circ\text{C}$ is used:

$$\epsilon_c = \epsilon_{\text{raw}} - 4.5(T - T_{\text{ref}}) \quad (7)$$

The strain factor grows with the magnitude of the shift,

$$f_s = \frac{1}{1 + \exp(-0.10(|\Delta\epsilon| - 15))} \quad (8)$$

while the vibration and tilt factors discriminate among the hypotheses: quiet bridges support drift, vibrating and rotating bridges support physical loading, and low vibration supports normal variation:

$$\begin{aligned} q_v &= \exp(-18\max(0, \bar{v} - 0.04)), & t_v &= 1 - \exp(-25\bar{v}), & n_v &= \exp(-15\bar{v}) \\ q_\theta &= \exp(-30|\bar{\theta} - \theta_0|), & t_\theta &= 1 - \exp(-35|\bar{\theta} - \theta_0|) \\ L_{\text{drift}} &= f_s q_v q_\theta, & L_{\text{phys}} &= f_s t_v t_\theta, & L_{\text{norm}} &= \exp(-0.15|\Delta\epsilon|) n_v \end{aligned} \quad (9)$$

Each likelihood is clamped to $[0.01, 0.99]$; disabled channels contribute neutral factors of 1, so the calibrator operates on any subset of sensors. These factors are heuristic physical relationships, so the resulting posterior is best read as a calibrated belief score rather than a true probability density.

C. Posterior and Recalibration Rule

Bayes' rule gives the posterior drift probability:

$$P(H_{\text{drift}}|E) = \frac{L_{\text{drift}}P_{\text{drift}}}{L_{\text{drift}}P_{\text{drift}} + L_{\text{phys}}P_{\text{phys}} + L_{\text{norm}}P_{\text{norm}}} \quad (10)$$

Recalibration is intentionally cautious: only when $P(H_{\text{drift}}|E) > 0.75$ and the net drift exceeds $12 \mu\epsilon$ is the offset nudged by 5% of the observed drift,

$$\text{Offset} \leftarrow \text{Offset} + 0.05 \Delta\epsilon, \quad \epsilon_{\text{cal}} = \epsilon_c - \text{Offset} \quad (11)$$

so one poor decision cannot move the baseline significantly, while genuine drift is still removed over subsequent decisions. During loading events the physical hypothesis dominates, the offset stays constant, and the strain response reaches both the risk model and the fatigue model intact.

D. State-Space Drift Decomposition

Where the Bayesian stage makes discrete offset decisions, a two-state augmented Kalman filter provides continuous separation of the elastic response from slowly varying bias. The state $x = [x_{\text{signal}}, x_{\text{drift}}]^T$ contains the transient elastic strain and the sensor zero-point bias:

$$\begin{aligned} x[k] &= Fx[k-1] + w, & z[k] &= Hx[k] + v \\ F &= [0.98 \ 0; 0 \ 1.00], & H &= [1 \ 1], & Q &= \text{diag}(0.500, 0.001), & R &= 2.0 \end{aligned} \quad (12)$$

The 0.98 entry models elastic rebound after a vehicle leaves; the unit drift entry models creep as a non-rebounding random walk; and the signal process noise is 500 times the drift noise because traffic strain changes within fractions of a second while creep evolves over hours or days. Only the sum of the

two states is observable, so the decomposition rests on these assumed dynamics; the Bayesian stage remains an independent check using physical observations.

VI. FUSION, STRUCTURAL HEALTH INDEX, AND DEGRADATION MODELS

A. Weighted Fusion and Interpretability

The fusion engine scales each active channel to a 0–100% risk score and combines the scores with adjustable weights:

$$\begin{aligned} \text{Overall Risk} &= \frac{\sum_i \text{Risk}_i \text{Weight}_i}{\sum_i \text{Weight}_i}, & \text{Health Score} &= 100 - \text{Overall Risk} \\ \text{Contribution}_i(\%) &= \frac{\text{Risk}_i \text{Weight}_i}{\sum_j \text{Risk}_j \text{Weight}_j} \times 100 \end{aligned} \quad (13)$$

The contribution of each sensor is displayed alongside, so when the health score falls — say from 95% to 68% — the operator sees that the vibration channel drove 64% of the drop rather than a black-box verdict.

B. Host Autoencoder and Structural Health Index

A larger 8-6-3-6-8 host autoencoder (tanh encoder, linear bottleneck and output) consumes eight calibrated inputs including humidity, roll angle, and wood moisture. Its masked loss over connected channels maps through a sigmoid to a 0–100% Structural Health Index:

$$\begin{aligned} \text{Loss} &= \frac{1}{|\text{Active}|_t} \sum_{i \in \text{Active}} (\tilde{x}_i - \hat{x}_i)^2, & \ell &= \max\left(0, \frac{\text{Loss} - 0.015}{0.05}\right) \\ \text{SHI} &= \frac{100}{1 + \exp(2.5(\ell - 1.2))} \end{aligned} \quad (14)$$

SHI is graded as 85–100 HEALTHY (routine inspection), 70–84.9 MONITORING (increase logging; check drainage and wear), 50–69.9 WARNING_INSPECTION (physical inspection within 30 days; restrict heavy permit loads), and < 50 CRITICAL_HAZARD (close or restrict the bridge and deploy a repair crew).

C. Timber Decay and Fatigue Life

Fungal decay requires moisture above the fibre-saturation point and moderate temperatures. The model uses the Scheffer climate index [24] with three regimes around 28% fibre saturation: dormant below 20% moisture (0 mm/yr), incubating at 20–28% (activity scaling to 65%), and active rot above 28% (70–100%), fastest at 27 °C. Where no wood probe is available, equilibrium moisture content is estimated from temperature and humidity through the Hailwood–Horrobin isotherm [25]. Moisture also weakens the section: $\text{MOR}_{\text{eff}} = 55[1 - 0.022 \cdot \min(\text{MC} - 12, 16)]$ MPa for $\text{MC} > 12\%$, with a rot penalty $1 + 0.08(\text{MC} - 20)$ above 20% MC. Cumulative fatigue follows the Palmgren–Miner rule [26] with an S–N law for Douglas fir:

$$D = \sum_i \frac{n_i}{N_i}, \quad \text{failure when } D \geq 1, \quad \log_{10} N = A - B \frac{\sigma}{\text{MOR}_{\text{eff}}}, \quad A = 10.2, B = 8.6 \quad (15)$$

Assuming 120,000 vehicle passes per year, the yearly fatigue rate is $(120,000/N)$ times the rot penalty. Total wear adds a base-ageing term that grows as the SHI falls, and the RUL divides the health reserve above a 40% critical SHI by the wear rate:

$$\text{Wear}(\%/yr) = 100 \cdot \text{Fatigue}_{\text{rate}} + 1.2 \frac{100}{\text{SHI}}, \quad \text{RUL}(\text{years}) = \frac{\max(1, \text{SHI} - 40)}{\max(0.40, \text{Wear})} \quad (16)$$

RUL tiers run from ROUTINE_SCHEDULE_5_YEARS (> 20 yr) through PRIORITY_INSPECTION_12_MONTHS (8–20 yr) and URGENT_TIMBER_RESTORATION (2–7.9 yr)

to IMMEDIATE_CLOSURE_OR_DERATING (< 2 yr). These constants follow common wood-design practice and constitute a first estimate, not a third-party-validated prediction.

VII. DIGITAL TWIN, DASHBOARD, AND CLOUD PERSISTENCE

A. Physics-Based Digital Twin

The twin solves the Euler–Bernoulli beam equation [27] $d^2/dx^2[EI(x) d^2y/dx^2] = q(x)$ on a simply supported span discretized into 11 nodes, superposing a uniform dead load w and a concentrated mid-span live load P :

$$y_{dead}(x) = \frac{wx}{24EI}(L^3 - 2Lx^2 + x^3), \quad y_{live}(x) = \frac{Px}{48EI}(3L^2 - 4x^2), \quad x \leq \frac{L}{2} \quad (17)$$

Bending stress $\sigma = My_{fibre}/I$ and a stress ratio σ/σ_{allow} colour each beam segment from green through amber to red. Four archetypes are parameterized: historic covered timber (36.5 m span), concrete cable-stayed (120 m), Warren steel truss (48 m), and prestressed concrete girder (32 m). The twin is a monitoring and visualization aid; it cannot replace a professional finite-element analysis.

B. Cloud Persistence and Dashboard

A serial bridge forwards validated JSON frames into the eight-stage host pipeline. Cloud upload to Supabase is

VIII. VALIDATION AND RESULTS

A. Validation Methodology

Overload, cracking, or seismic excitation cannot be created safely on a real bridge, so validation combined a hardware bench (Fig. 2) with a physics-based simulator driving the same host pipeline. The simulator generates a diurnal temperature cycle, rainfall-driven moisture changes, and vehicle crossings following a parabolic influence line; at 85% of design load the peak dynamic strain is approximately 119 $\mu\epsilon$. Five scenarios plus an injected-drift test exercise the full chain (Table II).

TABLE II
Validation Scenarios

Scenario	Stimulus	Expected response
normal	Solar temperature cycle	NOMINAL; high SHI; no alerts
heavy_vehicle	Semi-trailer at 85% of design load ($\approx 120 \mu\epsilon$; ≈ 0.18 g; roll 0.8°)	DYNAMIC_MOTION or VIBRATION_SHOCK; immediate send
vibration_peak	Gust or seismic tremor (RMS 0.42 g; peak 0.75 g; crest 4.8 ; 4.2 Hz)	VIBRATION_SHOCK; immediate transmission
failure	135% overload with cracked stringer (strain $> 400 \mu\epsilon$; permanent tilt 3.2°)	STRUCTURAL_WARNING or critical; SHI below 40%
drift_injection	Strain offset added on a quiet, level bridge	Drift posterior rises; strain corrected gradually

B. Edge Decision Behaviour

The edge model behaves according to its transfer function: $P(\text{Risk})$ rises from 0.14 at MSE 0.05 to 0.86 at MSE 0.35 , crossing the 0.15 transmission gate at MSE ≈ 0.055 . The measured ≈ 1.2 ms inference time leaves ample headroom within each 100-ms tick.

asynchronous and passes three filters: the edge gate; a host composite filter requiring at least one of nine conditions (high vibration, tilt, SHI below 75%, rotational hazard, alert, and others); and a 30-s deduplication window, while sending immediately whenever the class changes. A single-page dashboard receives the full system state over WebSocket at 20 Hz, with tabs for the live dashboard, bridge profile, sensor manager, self-calibration (raw versus calibrated strain and the drift posterior), fusion and Edge-AI panels, the full-screen 3D twin (Fig. 3), and settings.




Fig. 3. Digital-twin dashboard of the prototype: the 11-node Euler–Bernoulli deformation of the historic timber-bridge archetype is streamed at 20 Hz together with live edge-AI, calibration, and fusion state.

C. Data Reduction

Continuous streaming of all 10-Hz samples would produce about 864,000 records (≈ 432 MB) per bridge per day at 500 bytes per record. With anomaly-gated transmission a quiet bridge sends only its 60-s heartbeat: 1,440 records (≈ 0.72 MB)—a 99.83% reduction; a busy day with rate-limited bursts yields $\approx 2,880$ records (99.67%) (Table III). The record-count figures derive from the upload schedule and were not measured in an extended field deployment; emergency events are always uploaded immediately on any state change.

TABLE III
Data Volume Under Different Operating Modes

Mode	Records/day	Volume/day	Reduction
Continuous 10-Hz stream	864,000	≈ 432 MB	baseline
Quiet bridge (60-s heartbeat)	1,440	≈ 0.72 MB	99.83%
Busy day (rate-limited bursts)	2,880	≈ 1.44 MB	99.67%
Emergency events	on each state change	full event capture	immediate

D. Scenario Behaviour

Across all scenarios the pipeline behaved as designed: normal operation held a high health index with no alarms. The heavy-vehicle and vibration tests were classified as DYNAMIC_MOTION or VIBRATION_SHOCK and transmitted immediately. The failure scenario produced the highest-severity state with a falling SHI. The injected offset on a quiet, level bridge raised the drift posterior and was corrected stepwise, whereas the loading events were never calibrated away. Disconnecting a sensor excluded its channel from subsequent computation without generating an alarm. These results are functional verification, not a statistical performance evaluation.

E. Prototype Cost

Table IV lists the prototype bill of materials; a single unit costs about USD 12.60, supporting low-cost SHM experimentation. The estimate covers the node, modules, cables, and breadboards; a field enclosure, strain-gauge element, and power source would be added in a complete unit.

TABLE IV
 Prototype Bill of Materials

Component	Model	Qty	Unit cost (USD)
ESP32 microcontroller	WROOM-32D	1	4.50
6-DOF IMU	MPU-6050	1	1.80
Climate sensor	DHT11	1	1.20
Strain-gauge ADC	HX711	1	1.10
USB cable	Micro-USB, 1 m	1	1.50
Breadboard and jumper wires	400-point, 22 AWG	1	2.50
Total			12.60

F. Discussion and Limitations

The main finding is architectural: when calibration runs continuously inside the pipeline rather than as a maintenance task, the same system can decide whether to trust a reading, explain that decision with physical evidence any engineer can audit, and let corrected data drive both the data-volume decision and the visualization—addressing Gap 4 directly. Statistical metrics—residual strain error after correction, time from drift onset to $P(H_{\text{drift}}|E) > 0.75$, false alarms per day in nominal operation, and the fraction of load events preserved—remain to be quantified in extended bench runs, and field validation on a real structure with reference-grade sensors is still required.

IX. CONCLUSION AND FUTURE WORK

This paper presented a low-cost self-calibrating bridge condition-monitoring prototype integrating multisensor acquisition, on-device TinyML anomaly detection, anomaly-gated transmission, Bayesian drift-versus-load discrimination, Kalman signal/drift decomposition, explainable fusion with a Structural Health Index, timber-decay and fatigue RUL models, an Euler–Bernoulli digital twin, and selective cloud persistence—all on hardware costing approximately USD 12.60. The key principle is that a sensor's reliability is judged continuously from the behaviour of its neighbours, so drift is corrected without suppressing real structural responses. Future work includes field validation against reference-grade sensors, per-structure autoencoder training and threshold tuning, temperature-compensated strain gauges and a dedicated inclinometer, quantitative calibration-error and detection-rate measurement, finite-element and material-test validation of the twin and fatigue models, and solar-powered wireless deployment with security hardening.

REFERENCES

- [1] C. R. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester, U.K.: Wiley, 2012.
- [2] N. A. Hoult, P. Fidler, P. G. Hill, C. Middleton, et al., “Long-term wireless structural health monitoring of the Ferriby Road Bridge,” *J. Bridge Eng.*, vol. 15, no. 2, pp. 153–159, 2010.
- [3] J. Meyer, R. Bischoff, G. Feltrin, and M. Motavalli, “Wireless sensor networks for long-term structural health monitoring,” *Smart Struct. Syst.*, vol. 6, no. 3, pp. 263–275, 2010.

- [4] X. Tong, H. Yang, L. Wang, and Y. Miao, “The development and field evaluation of an IoT system of low-power vibration for bridge health monitoring,” *Sensors*, vol. 19, no. 5, Art. no. 1222, 2019.
- [5] Y. Yang, T. Chen, W. Lin, M. Jing, et al., “Research progress on calibration of bridge structural health monitoring sensing system,” *Adv. Bridge Eng.*, vol. 5, Art. no. 32, 2024.
- [6] P. Rizzo and A. Enshaiean, “Challenges in bridge health monitoring: A review,” *Sensors*, vol. 21, no. 13, Art. no. 4336, 2021.
- [7] Y. Yu, J. Ou, and H. Li, “Design, calibration and application of wireless sensors for structural global and local monitoring of civil infrastructures,” *Smart Struct. Syst.*, vol. 6, no. 5–6, pp. 641–659, 2010.
- [8] B. Kostić and M. Gül, “Vibration-based damage detection of bridges under varying temperature effects using time-series analysis and artificial neural networks,” *J. Bridge Eng.*, vol. 22, no. 10, 2017.
- [9] J. C. Weinstein, M. Sanayei, and M. P. Brenner, “Bridge damage identification using artificial neural networks,” *J. Bridge Eng.*, vol. 23, no. 11, 2018.
- [10] R. Zinno, S. Shaffie Haghsheenas, G. Guido, and A. Vitale, “Artificial intelligence and structural health monitoring of bridges: A review of the state-of-the-art,” *IEEE Access*, vol. 10, pp. 88058–88078, 2022.
- [11] R. S. Chaudhari, S. Nama, M. S. Narkhede, V. P. Bhargale, R. B. Gaikwad, et al., “Next generation vibration-based structural health monitoring for bridges: A systematic review of AI, IoT, and digital twin integrations,” *J. Build. Pathol. Rehabil.*, vol. 11, Art. no. 136, 2026.
- [12] A. Dabbous, R. Berta, M. Fresta, H. Ballout, L. Lazzaroni, and F. Bellotti, “Bringing intelligence to the edge for structural health monitoring: The case study of the Z24 bridge,” *IEEE Open J. Ind. Electron. Soc.*, vol. 5, pp. 781–794, 2024.
- [13] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [14] P. Warden and D. Situnayake, *TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers*. Sebastopol, CA: O'Reilly Media, 2019.
- [15] S. Moradi, S. Eftekhari Azam, and M. Mofid, “A physics informed neural network integrated digital twin for monitoring of the bridges,” in *Structural Health Monitoring 2021*. DEStech/DPI Proc., 2021.
- [16] D. Dan, Y. Ying, and L. Ge, “Digital twin system of bridges group based on machine vision fusion monitoring of bridge traffic load,” *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 11, pp. 22190–22205, 2021.
- [17] M. Nasim, A. Rajabifard, Y. Chen, and B. Samali, “A demonstration of a digital twin framework for structural health monitoring: Application to bridge infrastructures,” *J. Infrastruct. Intell. Resil.*, vol. 5, no. 1, Art. no. 100184, 2026.
- [18] V. Mousavi, M. Rashidi, S. Ghazimoghadam, and B. Samali, “A data-driven digital twin model for bridge health monitoring using feature fusion and unsupervised deep learning,” *Inf. Fusion*, vol. 126, Art. no. 103534, 2026.
- [19] L.-M. Sun, H. Sun, W. Zhang, and Y.-X. Li, “Hybrid monitoring methodology: A model-data integrated digital twin framework for structural health monitoring and full-field virtual sensing,” *Adv. Eng. Inform.*, vol. 60, Art. no. 102386, 2024.
- [20] Y. Liu, “Edge-cloud collaborative digital twin architecture for bridge structural health monitoring: Multi-sensor fusion and anomaly detection algorithm design,” *Discov. Appl. Sci.*, vol. 8, Art. no. 938, 2026, doi: 10.1007/s42452-026-08991-6.
- [21] A. Jiménez Rios, V. Plevris, and M. Nogal, “Bridge management through digital twin-based anomaly detection systems: A systematic review,” *Front. Built Environ.*, vol. 9, Art. no. 1176621, 2023.
- [22] E. Glaessgen and D. Stargel, “The digital twin paradigm for future NASA and U.S. Air Force vehicles,” in *Proc. 53rd AIAA/ASME/ASCE/AHS/ASC Struct., Struct. Dyn. Mater. Conf.*, Honolulu, HI, 2012.
- [23] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [24] T. C. Scheffer, “A climate index for estimating potential for decay in wood structures above ground,” *Forest Prod. J.*, vol. 21, no. 10, pp. 25–31, 1971.
- [25] A. J. Hailwood and S. Horrobin, “Absorption of water by polymers: Analysis in terms of a simple model,” *Trans. Faraday Soc.*, vol. 42B, pp. 84–92, 1946.
- [26] M. A. Miner, “Cumulative damage in fatigue,” *J. Appl. Mech.*, vol. 12, no. 3, pp. A159–A164, 1945.
- [27] S. Timoshenko, *Strength of Materials, Part I: Elementary Theory and Problems*, 3rd ed. New York: D. Van Nostrand, 1955.