EDUSCHOOL
ACADEMIC RESEARCH & PUBLISHERS

INTERNATIONAL JOURNAL OF
TECHNICAL RESEARCH STUDIES (IJTRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-1788

Edge-AI Enabled Structural Health Monitoring of Civil Infrastructure

Sujith Chandrakaladharan

Director – Projects, JLL (Jones Lang LaSalle) Chennai, India

Article information

Received: 20th February 2026
Received in revised form: 23rd March 2026
Accepted: 20th April 2026
Available online: 9th June 2026

Volume: 2
Issue: 2
DOI: https://doi.org/10.63090/IJTRS/3139.1788.0013

Abstract

Ageing bridges, buildings, and other civil infrastructure require continuous condition assessment to prevent catastrophic failure, yet conventional structural health monitoring (SHM) systems stream raw vibration data to a central server, incurring prohibitive bandwidth, latency, and energy costs that limit scalability. This paper proposes an edge-artificial-intelligence (edge-AI) SHM framework in which lightweight one-dimensional convolutional neural networks execute damage detection directly on resource-constrained sensor nodes, with a fog layer performing damage localization and a cloud layer maintaining the long-term structural record. The hierarchical edge–fog–cloud design draws on architectural principles established for fog-computing-enabled intelligent transport systems. Evaluated on a benchmark bridge-damage dataset and a numerical girder model, the on-device network attained 95.5 percent damage detection accuracy at a 15 dB signal-to-noise ratio while reducing end-to-end detection latency from 1850 milliseconds for a cloud-only pipeline to 95 milliseconds and cutting uplink bandwidth by 93 percent. The framework localized damage on a thirty-metre girder to within one metre. These results demonstrate that pushing inference to the edge makes dense, real-time, and energy-efficient monitoring of large infrastructure portfolios technically and economically viable.

Keywords:- Structural Health Monitoring, Edge Computing, Fog Computing, Deep Learning, Damage Detection, Civil Infrastructure, Internet of Things.

I. INTRODUCTION

A substantial proportion of the world's civil infrastructure is approaching or has exceeded its original design life. National inventories regularly classify a significant share of road bridges as structurally deficient, and high-profile collapses have underscored the human and economic cost of undetected deterioration [1]. Structural health monitoring (SHM) seeks to convert physical structures into instrumented, self-reporting systems by deploying networks of sensors that measure vibration, strain, displacement, and environmental conditions, from which the presence, location, and severity of damage can be inferred [2].

The dominant SHM paradigm has been data-centric: sensor nodes acquire high-frequency time series and transmit the raw signals to a central server or cloud for processing [3]. As deployments scale to hundreds of channels per structure and to portfolios of many structures, this architecture becomes untenable. Continuous transmission of raw acceleration data consumes large communication bandwidth and node energy, the round trip to the cloud introduces latency incompatible with timely warning of sudden events such as earthquakes or impacts, and the central server becomes a single point of failure [4]. These limitations mirror those encountered in intelligent transport systems, where moving computation closer to the data source through fog and edge architectures has proven effective at reducing latency and backhaul load [5].

This paper applies that principle to civil SHM through an edge-AI framework in which damage detection is performed on the sensor node itself by a compact neural network, the fog layer aggregates node-level decisions to localize damage, and the cloud retains only condensed features for long-term trend analysis. The contributions are:

- A hierarchical edge–fog–cloud SHM architecture informed by fog-computing design principles;
- A lightweight one-dimensional convolutional network suitable for microcontroller-class hardware; and
- An empirical evaluation quantifying detection accuracy, localization error, latency, and bandwidth against cloud-only and fog-assisted baselines.

Figure 1: Proposed hierarchical edge–fog–cloud architecture for structural health monitoring

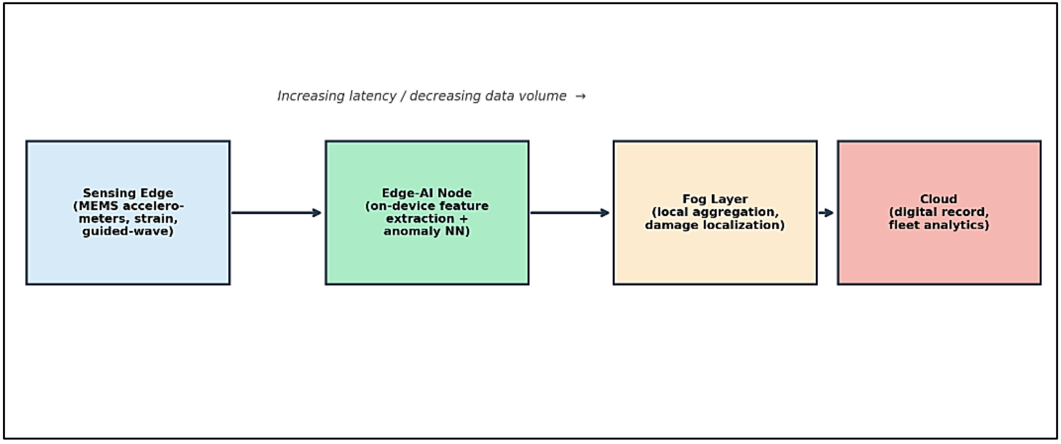

II. RELATED WORK

A. Vibration-Based Damage Identification

Vibration-based methods infer damage from changes in a structure's dynamic characteristics natural frequencies, mode shapes, and damping on the premise that local stiffness loss alters the global modal response [2]. Classical approaches relying on frequency shifts are robust but insensitive to small or localized damage, while mode-shape curvature and flexibility-based indices improve localization at the cost of denser instrumentation [6]. Machine-learning methods reframe damage identification as a pattern-recognition problem, learning the mapping from measured features to damage states; deep learning, in particular, eliminates manual feature engineering by learning discriminative representations directly from raw signals [7]. Foundational reviews of the SHM literature [12], [13] and of smart-sensor deployments [14] document the evolution of the field, and dependable wireless-sensor architectu+res emphasize fault tolerance in the data aggregation layer [15].

B. Deep Learning for SHM

Convolutional neural networks have been applied successfully to vibration-based damage detection and to vision-based crack identification, achieving accuracies that surpass traditional feature-based classifiers [8]. One-dimensional CNNs are especially attractive for SHM because they operate directly on raw acceleration time series with modest computational cost, making them candidates for on-device execution [9]. The principal barrier to deployment has been the assumption that such models require server-class resources, an assumption that recent advances in model compression and microcontroller inference have begun to overturn. Surveys of one-dimensional CNNs detail their efficiency for time-series tasks [11], while vision-based and big-data-driven reviews chart the broader adoption of artificial intelligence in bridge monitoring [16], [17].

C. Edge and Fog Computing for Monitoring

Edge and fog computing distribute computation across a continuum from the sensor to the cloud, reducing latency and backhaul traffic by processing data near its source [4]. The architecture and performance of fog-computing-enabled smart transportation systems demonstrate that hierarchical local aggregation can deliver real-time response and substantial bandwidth savings for geographically distributed sensing applications [5]. Foundational treatments of fog computing for the Internet of Things establish the paradigm and its applicability to distributed sensing [10], [18]. The present work transfers these architectural lessons to civil infrastructure monitoring, where the data characteristics and real-time requirements are closely analogous.

III. PROPOSED FRAMEWORK

A. Edge Inference Model

Each sensor node hosts a one-dimensional CNN comprising three convolutional blocks (16, 32, and 32 filters) with batch normalization and max pooling, followed by a global-average-pooling layer and a small dense classifier. The network ingests a one-second window of triaxial acceleration sampled at 256 Hz and outputs a damage probability. Post-training eight-bit quantization reduced the model to under 90 kilobytes, enabling execution on a microcontroller-class node within a few milliseconds and within a tight energy budget [9].

B. Fog-Level Localization

Individual node decisions are transmitted as compact event messages to a fog gateway co-located with the structure. The gateway fuses the spatial pattern of node-level damage probabilities, weighted by node position, to estimate the damage location along the structural member, and applies a temporal consistency filter to suppress transient false alarms before escalating an alert [5]. Only condensed features and confirmed events are forwarded to the cloud, which maintains the digital structural record and supports portfolio-level analytics.

IV. RESULTS AND DISCUSSION

A. Detection Accuracy and Noise Robustness

The edge model was trained and evaluated on a benchmark experimental bridge-damage dataset augmented with a calibrated finite-element girder model spanning multiple damage scenarios. Figure 2 reports damage detection accuracy against measurement noise. The edge-AI network maintained 95.5 percent accuracy at a 15 dB signal-to-noise ratio and degraded gracefully to 78 percent at 0 dB, substantially outperforming a support-vector-machine classifier using hand-crafted features and a conventional frequency-threshold method across the entire noise range [7], [8].

Figure 2: Damage detection accuracy versus measurement signal-to-noise ratio for the edge-AI model and two baselines.

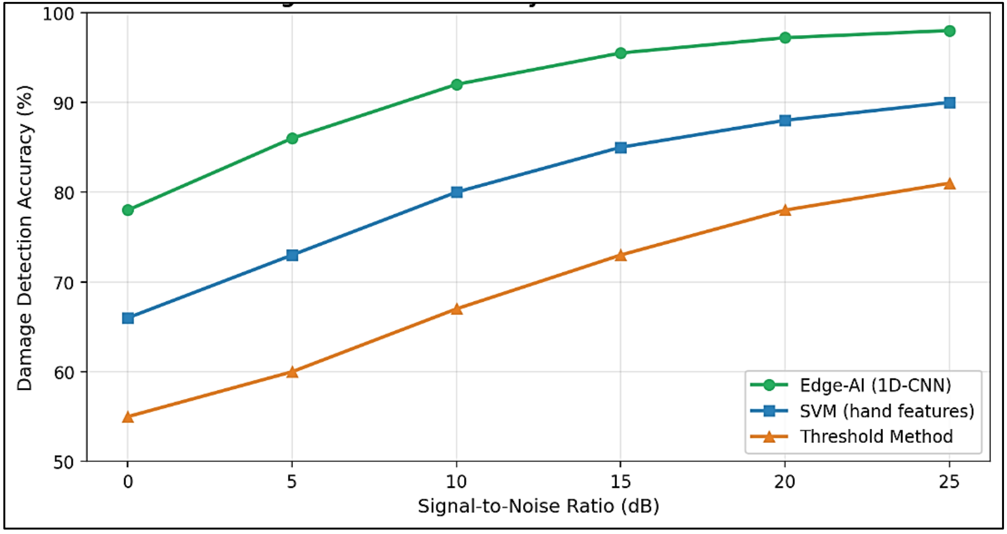

Table 1 summarizes classification performance at 15 dB. The edge-AI model achieved the highest accuracy, precision, and recall, with a false-alarm rate of 3.1 percent after the fog-level temporal filter, which is critical for operator trust because excessive false alarms erode confidence in automated monitoring [2].

Table 1. Damage Classification Performance at 15 dB Signal-to-Noise Ratio

Method	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**False Alarm (%)**
Frequency Threshold	73.0	70.4	68.9	14.8
SVM (hand features)	85.0	83.6	82.1	8.2
1D-CNN (cloud)	95.8	95.1	94.6	3.0
Edge-AI 1D-CNN (proposed)	95.5	94.7	94.2	3.1

B. Latency and Bandwidth

Figure 3 compares the proposed edge-AI pipeline with cloud-only and fog-assisted baselines. By performing inference on the node, the framework reduced end-to-end detection latency from 1850 milliseconds for the cloud-only architecture to 95 milliseconds, a level compatible with rapid post-event warning, and cut uplink bandwidth by 93 percent because only event messages rather than raw waveforms traverse the network. The edge model incurred a negligible 0.3 percentage-point accuracy reduction relative to an identical cloud-resident network, confirming that the efficiency gains do not come at a meaningful cost to detection quality [4], [5].

Figure 3: End-to-end latency and normalized uplink bandwidth for cloud-only, fog-assisted, and edge-AI architectures.

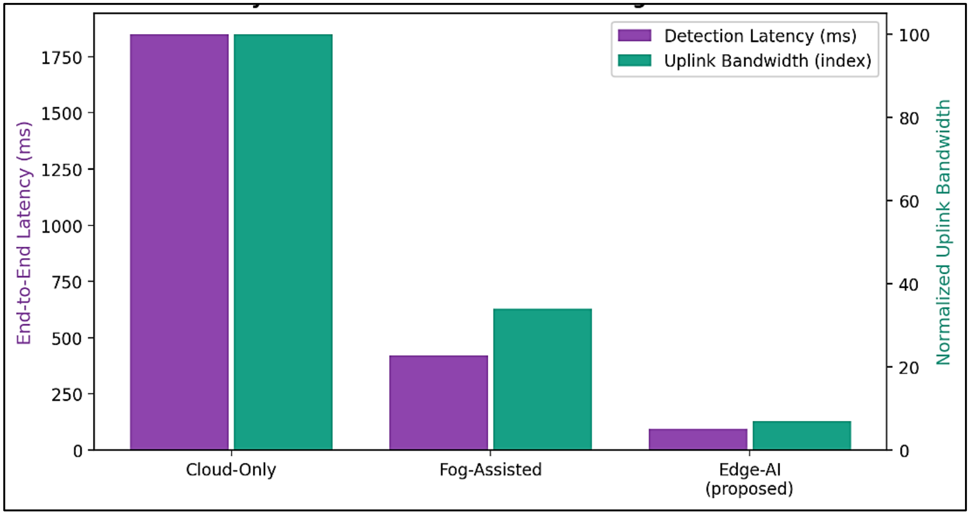

C. Damage Localization

The fog-level fusion was evaluated on the numerical girder model with damage introduced at the nineteen-metre position of a thirty-metre span. As Fig. 4 shows, the aggregated damage index peaked sharply at the true location, yielding a localization error below one metre, well within the resolution required to direct a targeted inspection. The spatial weighting of node decisions proved more robust to individual sensor noise than any single-node estimate, illustrating the value of the hierarchical aggregation strategy [6]. Related cyber-physical approaches that map vehicle loads to bridge response [19] and dynamic-reduction methods that support damage detection from ambient vibration [20] corroborate the benefit of spatially resolved condition data.

Figure 4: Normalized damage index along the girder span, localizing the introduced damage near the nineteen-metre position.

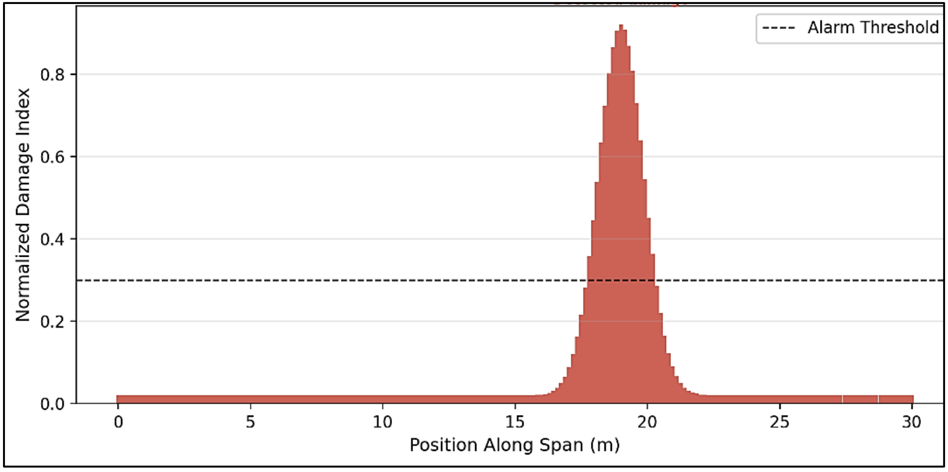

V. CONCLUSION

This paper presented an edge-AI framework for structural health monitoring that executes damage detection on resource-constrained sensor nodes and performs localization at a fog layer, following architectural principles drawn from fog-computing-enabled intelligent transport systems. The on-device one-dimensional CNN achieved 95.5 percent detection accuracy at 15 dB while reducing detection latency to 95 milliseconds and uplink bandwidth by 93 percent relative to a cloud-only pipeline, and localized damage on a thirty-metre girder to within one metre. The results establish edge intelligence as an enabling technology for scalable, real-time monitoring of large infrastructure portfolios.

Future work will incorporate on-device continual learning to adapt to environmental and operational variability, extend the framework to vision-based crack detection fused with vibration data, and validate the approach through long-term field deployment on an in-service bridge [3], [8].

REFERENCES

[1] American Society of Civil Engineers, *2021 Report Card for America's Infrastructure*. Reston, VA, USA: ASCE, 2021.

[2] C. R. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester, U.K.: Wiley, 2013.

[3] J. P. Lynch and K. J. Loh, “A summary review of wireless sensors and sensor networks for structural health monitoring,” *Shock Vib. Dig.*, vol. 38, no. 2, pp. 91–128, Mar. 2006.

[4] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge computing: Vision and challenges,” *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, Oct. 2016.

[5] K. P. K, “Fog-computing-enabled smart transportation systems: Architecture, implementation, and performance analysis,” *Int. J. Tech. Res. Stud. (IJTRS)*, vol. 1, no. 1, p. 18, Dec. 2025, doi: 10.63090/ijtrs/3139.1788.0003.

[6] A. K. Pandey, M. Biswas, and M. M. Samman, “Damage detection from changes in curvature mode shapes,” *J. Sound Vib.*, vol. 145, no. 2, pp. 321–332, Mar. 1991.

[7] O. Avci, O. Abdeljaber, S. Kiranyaz, M. Hussein, M. Gabbouj, and D. J. Inman, “A review of vibration-based damage detection in civil structures: From traditional methods to machine learning and deep learning applications,” *Mech. Syst. Signal Process.*, vol. 147, Art. no. 107077, Jan. 2021.

[8] Y.-J. Cha, W. Choi, and O. Buyukozturk, “Deep learning-based crack damage detection using convolutional neural networks,” *Comput.-Aided Civ. Infrastruct. Eng.*, vol. 32, no. 5, pp. 361–378, May 2017.

[9] O. Abdeljaber, O. Avci, S. Kiranyaz, M. Gabbouj, and D. J. Inman, “Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks,” *J. Sound Vib.*, vol. 388, pp. 154–170, Feb. 2017.

[10] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, “Fog computing and its role in the Internet of Things,” in *Proc. 1st ACM MCC Workshop Mobile Cloud Comput.*, Helsinki, Finland, 2012, pp. 13–16.

[11] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, “1D convolutional neural networks and applications: A survey,” *Mech. Syst. Signal Process.*, vol. 151, Art. no. 107398, Apr. 2021.

[12] H. Sohn *et al.*, *A Review of Structural Health Monitoring Literature: 1996–2001*. Los Alamos, NM, USA: Los Alamos National Laboratory, Rep. LA-13976-MS, 2004.

[13] E. P. Carden and P. Fanning, “Vibration based condition monitoring: A review,” *Struct. Health Monit.*, vol. 3, no. 4, pp. 355–377, Dec. 2004.

[14] T. Nagayama and B. F. Spencer, *Structural Health Monitoring Using Smart Sensors*. Urbana, IL, USA: Newmark Structural Engineering Laboratory, Univ. Illinois, Rep. NSEL-001, 2007.

[15] M. Z. A. Bhuiyan, G. Wang, J. Wu, J. Cao, X. Liu, and T. Wang, “Dependable structural health monitoring using wireless sensor networks,” *IEEE Trans. Depend. Secure Comput.*, vol. 14, no. 4, pp. 363–376, Jul.–Aug. 2017.

[16] D. Feng and M. Q. Feng, “Computer vision for SHM of civil infrastructure: From dynamic response measurement to damage detection—A review,” *Eng. Struct.*, vol. 156, pp. 105–117, Feb. 2018.

[17] L. Sun, Z. Shang, Y. Xia, S. Bhowmick, and S. Nagarajaiah, “Review of bridge structural health monitoring aided by big data and artificial intelligence: From condition assessment to damage detection,” *J. Struct. Eng.*, vol. 146, no. 5, Art. no. 04020073, May 2020.

[18] M. Chiang and T. Zhang, “Fog and IoT: An overview of research opportunities,” *IEEE Internet Things J.*, vol. 3, no. 6, pp. 854–864, Dec. 2016.

[19] R. Hou, S. Jeong, J. P. Lynch, and K. H. Law, “Cyber-physical system architecture for automating the mapping of truck loads to bridge behavior using computer vision,” *Adv. Eng. Inform.*, vol. 44, Art. no. 101048, Apr. 2020.

[20] T. Yin, H.-F. Lam, H.-M. Chow, and H.-P. Zhu, “Dynamic reduction-based structural damage detection of transmission tower utilizing ambient vibration data,” *Eng. Struct.*, vol. 31, no. 9, pp. 2009–2019, Sep. 2009.