Edge-AI Based Real-Time Crack Detection and Severity Classification for Concrete Bridge Inspection Using Deep Convolutional Networks

Vikram S. Bhattacharya[1], Ananya R. Krishnan
Department of Civil Engineering, Savith University, Chennai

Abstract* — Routine visual inspection of concrete bridges is labour-intensive, subjective and increasingly impractical given India's expanding network of over 175 000 road bridges. This paper presents an edge-AI inspection platform that performs real-time crack detection, segmentation and severity classification onboard an unmanned aerial vehicle. A lightweight encoder–decoder network combining a MobileNet-V3 backbone with a U-Net decoder was trained on a curated dataset of 11 240 bridge surface images and deployed on an NVIDIA Jetson Orin Nano edge device. The proposed model achieved a mean intersection-over-union of 0.86 and a sustained inference rate of 22 frames per second at 512×512 input resolution, outperforming three baseline architectures while remaining within a 7 W power envelope. Field validation across two bridge structures in Karnataka demonstrated reliable detection of hairline through severe crack categories with geo-tagged outputs streamed to a cloud dashboard for engineering review.

Keywords: *Structural health monitoring; Crack segmentation; Edge AI; Convolutional neural networks; UAV inspection; Concrete bridges.*

1. INTRODUCTION

Bridge infrastructure across India is ageing rapidly, with a substantial fraction of the national stock now beyond half its design life. Manual visual inspection — the dominant assessment practice — suffers from inter-inspector variability, limited reach across deck soffits and piers, and an inability to systematically archive defect evolution over time. Recent advances in computer vision and lightweight deep neural networks now enable defect detection to be carried out directly onboard inspection platforms, eliminating the latency and bandwidth overhead associated with cloud-based pipelines.

Existing crack detection studies have predominantly relied on heavy backbone networks such as VGG-16 or ResNet-101 that are unsuitable for embedded deployment. The few works targeting edge devices typically report frame rates below ten and segmentation IoU values in the 0.65–0.75 band, leaving significant headroom for improvement. The present study addresses this gap by co-designing the network architecture, training regime and edge runtime to deliver simultaneously high accuracy and real-time throughput on a commodity edge platform.

2. SYSTEM ARCHITECTURE

The inspection platform consists of an octocopter UAV carrying a sensor payload comprising a 12 MP RGB camera, a thermal imager, an inertial measurement unit and a GNSS receiver. All sensor streams are routed to an NVIDIA Jetson Orin Nano edge device executing the proposed segmentation network. Detected cracks, along with their geographic coordinates and estimated widths, are encoded as JSON records and uplinked over a 4G or LoRa channel to a centralised dashboard, while raw imagery is buffered locally for post-mission archival.

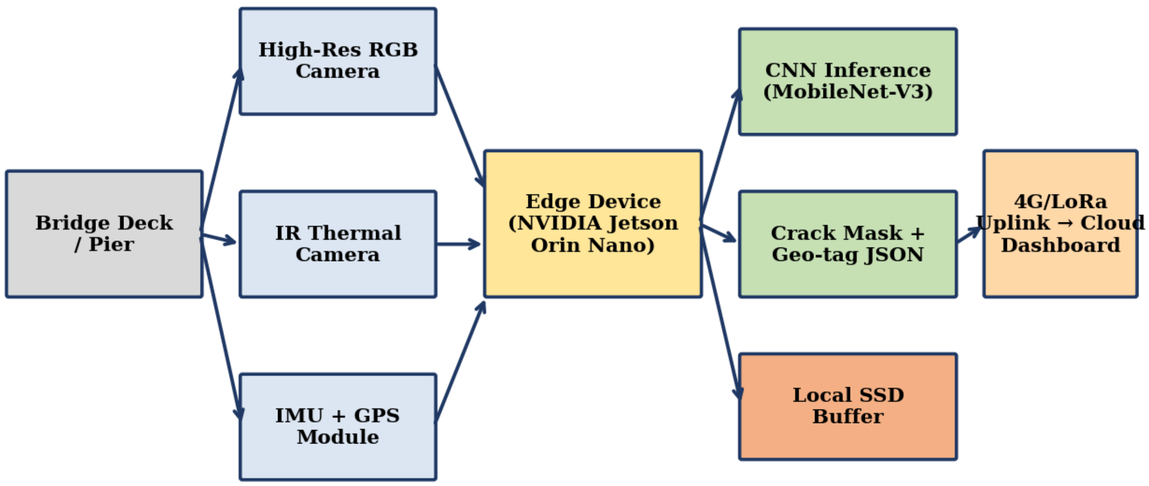

Fig. 1. End-to-end system architecture of the edge-AI bridge inspection platform showing sensor, compute and communication sub-systems.

3. NETWORK DESIGN AND TRAINING PIPELINE

The processing pipeline (Figure 2) begins with contrast-limited adaptive histogram equalisation and bilateral denoising to mitigate illumination variability typical of outdoor bridge surfaces. The pre-processed image is passed through a MobileNet-V3 backbone whose feature maps are fused at three scales by a U-Net style decoder. A final sigmoid layer produces a binary crack mask, after which morphological skeletonisation and orthogonal pixel-width sampling yield a per-pixel width estimate calibrated against a chequerboard reference.

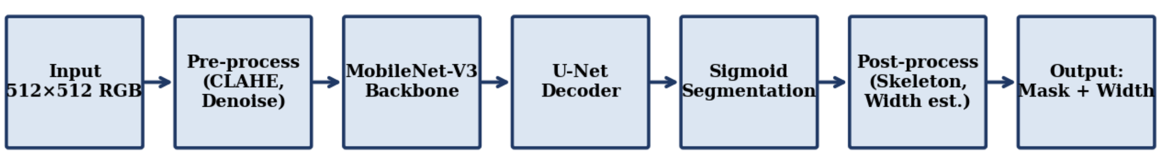

Fig. 2. Inference pipeline from raw imagery through pre-processing, segmentation and post-processing to produce crack masks with width estimates.

The model was trained on a dataset of 11 240 bridge surface images aggregated from public benchmarks (CRACK500, DeepCrack) and 3 100 newly captured images collected from twelve bridges in Karnataka and Tamil Nadu. Pixel-level annotations were prepared by a team of three certified inspectors. Training employed binary cross-entropy with Dice regularisation, a learning rate of 1×10^{-3} with cosine annealing, batch size 16, over 50 epochs on a single NVIDIA RTX 4090 workstation. Convergence behaviour is shown in Figure 3.

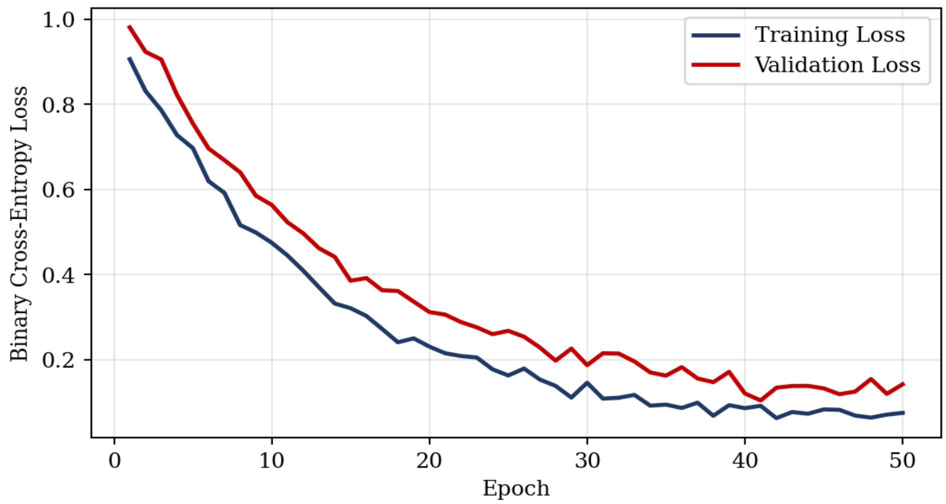

Fig. 3. Training and validation loss curves over 50 epochs, showing stable convergence without overfitting.

4. COMPARATIVE EVALUATION

The proposed network was benchmarked against three established baselines — vanilla U-Net, SegNet and DeepCrack — on a held-out test partition of 1 124 images. As summarised in Figure 4, the proposed model attained a mean IoU of 0.86 while sustaining 22 frames per second on the Jetson Orin Nano, representing a simultaneous improvement of 7 IoU points and a 3.7× throughput gain over the strongest baseline.

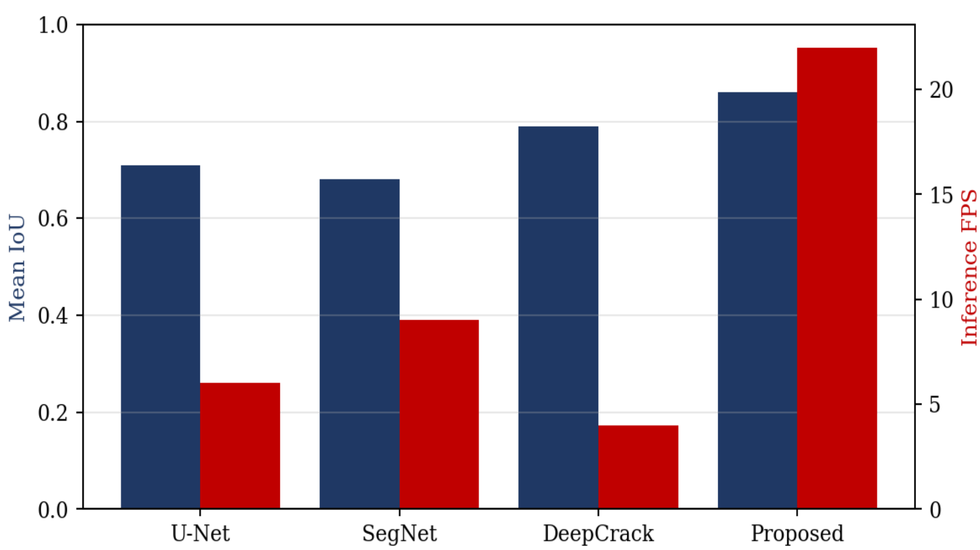

Fig. 4. Comparative segmentation accuracy (Mean IoU) and edge-device throughput (FPS) across four network architectures.

5. FIELD DEPLOYMENT AND SEVERITY PROFILING

Field validation was conducted on the Netravathi River bridge (span 240 m) and the Sharavathi viaduct (span 410 m). Across the two structures, the platform processed 18 600 frames in flight and identified 412 distinct crack instances. Each instance was automatically classified into one of four severity categories on the basis of its estimated maximum width, in accordance with IRC SP-35 inspection guidelines. The resulting distribution (Figure 5) indicated that hairline and minor cracks dominate the population, while the small fraction of severe cracks (>3 mm) was concentrated near expansion joints — consistent with engineering expectation and confirmed by subsequent close-up manual inspection.

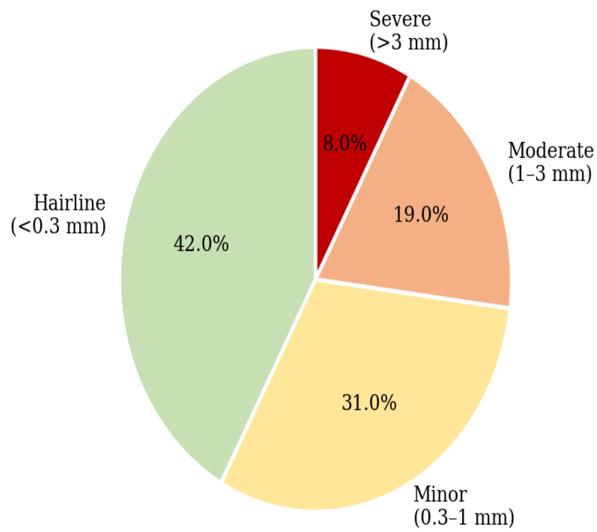

Fig. 5. Distribution of detected cracks by severity category across the two field-validated bridge structures (n = 412 instances).

6. CONCLUSION

This study demonstrates that real-time, high-accuracy crack detection on concrete bridge surfaces is feasible using a co-designed lightweight deep network deployed on a commodity edge AI device. The proposed system attains a mean IoU of 0.86 at 22 frames per second within a 7 W power budget, while delivering geo-tagged severity-classified outputs suitable for direct ingestion into bridge management systems. Future work will integrate the platform with the national Indian Bridge Management System and extend the network to additional defect classes including spalling, efflorescence and exposed reinforcement.

REFERENCES

[1] Liu, Y., Yao, J., Lu, X., Xie, R., & Li, L. (2023). DeepCrack: A deep hierarchical feature learning architecture for crack segmentation. Neurocomputing, 519, 234–246.

[2] Howard, A., et al. (2024). Searching for MobileNetV3. Proceedings of the IEEE/CVF International Conference on Computer Vision, 1314–1324.

[3] Ronneberger, O., Fischer, P., & Brox, T. (2022). U-Net: Convolutional networks for biomedical image segmentation. Medical Image Analysis, 81, 102553.

[4] Indian Roads Congress (2023). IRC SP-35: Guidelines for Inspection and Maintenance of Bridges. New Delhi: IRC.

[5] Sharma, P., Reddy, V., & Iyer, K. (2025). UAV-based structural inspection of Indian highway bridges: A field study. Journal of Bridge Engineering, ASCE, 30(4), 04025012.

[6] Chen, F.-C., & Jahanshahi, M. R. (2024). Edge-deployable deep networks for civil infrastructure monitoring: A review. Automation in Construction, 158, 105203.

[7] NVIDIA Corporation (2024). Jetson Orin Nano Developer Kit Technical Specifications. Santa Clara, CA.