American Journal of Civil Engineering
2026, Vol. 14, No. 5, pp. 324–330
https://doi.org/10.11648/j.ajce.20261405.13

Research Article

Benchmarking YOLOv8 Variants for Automated Infrastructure Assessment

Nnanna Ekedebe*

Department of Information Technology, Federal University of Technology, Owerri, Nigeria

Abstract

Structural cracks in buildings and bridges pose a serious safety concern in Nigeria, with manual inspection methods being slow, subjective and impractical at scale. Undetected structural cracks can lead to building collapses, which have become a recurring and deadly problem in Nigeria's construction sector. This study benchmarks three YOLOv8 variants – nano, small and medium under identical training and evaluation conditions using a 500-image Building Crack Detection dataset sourced from Roboflow Universe, with models evaluated on mAP@50, mAP@50-95, Precision, Recall and Inference Speed. All three models were trained for 50 epochs under the same conditions to ensure a fair comparison across variants. All three models achieved a mAP@50 of at least 99.4% and 100% recall, with YOLOv8s and YOLOv8m both achieving the highest mAP@50-95 at 99.50% and YOLOv8n achieving the fastest inference speed at 3.1 ms per image. These results show that all three models were highly accurate and rarely missed a crack, with the main difference between them being how fast each one could process an image. YOLOv8n was identified as the most optimal model for Nigerian infrastructure deployment, demonstrating the best balance of accuracy and speed for resource-constrained settings. Although the larger models were slightly more precise, the accuracy gap was small enough that speed became the deciding factor for practical deployment. This study demonstrates the viability of automated crack detection using YOLOv8 and provides a foundation for affordable and scalable structural monitoring in Nigeria.

Keywords

Crack Detection, YOLOv8, Structural Health Monitoring, Object Detection, Nigeria

1. Introduction

Over the years, automated structural crack detection using deep learning has gained significant attention, fuelled by growing concerns over aging infrastructure and the limitations posed by manual inspection. It is important to note that this field has progressed from traditional image processing techniques toward the use of Convolutional Neural Networks (CNNs) and YOLO based approaches capable of real-time detection [1]. This study sets out to benchmark three YOLOv8 variants to identify the most optimal model for crack detection deployment in the Nigerian infrastructural environment.

Structural cracks in buildings are most times early indicators that there might be a serious issue with the structure, which is why they cannot be left undetected. The manual detection method, which involves visual and sometimes physical examination by structural engineers, can be slow, subjective and dangerous, especially in the case of elevated structures [2].

Nigeria has faced her fair share of building collapses, and most of the contributing factors include aging infrastructure,

*Correspondence: Nnanna Ekedebe (nnanaekedebe@gmail.com)

Received: 13 July 2026; **Accepted:** 24 July 2026; **Published:** 18 September 2026

substandard construction materials and poor regulatory enforcement. This problem is further aggravated by Nigeria's limited structural engineering capacity, which makes inspection at scale impractical. As a result, there is a need for an automated, affordable and deployable structural monitoring approach [3].

Deep learning has enabled automated and highly accurate defect detection directly from images, and can at times be more accurate and faster in detection than human experts. The fact that models can be deployed on drone footage or mobile camera setups for field inspection makes them more suitable for resource-constrained deployment. YOLO models are particularly suitable because of their real-time speed, accuracy and lightweight design [4].

As most studies have rarely systematically benchmarked YOLOv8n, YOLOv8s and YOLOv8m under identical conditions for crack detection, the accuracy vs computational efficiency trade-off remains largely unexplored. Furthermore, most prior work has rarely framed crack detection within Nigerian infrastructure assessment and resource-constrained deployment settings, and this study aims to address this specific gap.

This study sets out to benchmark three YOLOv8 variants, nano, small and medium trained and evaluated under identical conditions, with a publicly available annotated crack detection dataset sourced from Roboflow Universe [5]. The metrics used for evaluation include mAP@50, Precision, Recall and Inference Speed, with the goal of identifying the optimal model that balances accuracy and efficiency for Nigerian infrastructure deployment.

The paper follows the following structure: Section 2 presents related work, Section 3 presents methodology, Section 4 presents results and discussion, and Section 5 presents the conclusion.

2. Related Work

This section presents a review of existing literature on structural crack detection, progressing from traditional inspection methods through image processing approaches to deep learning and YOLO based methods. This study fits within the YOLO based approach category.

Manual inspection has long been the standard for structural assessment, with its limitations including subjectivity, high dependency on inspector experience, proneness to fatigue and inability to scale. It is also practically dangerous for elevated structures such as skyscrapers and bridges [6].

Early detection methods relied on edge detection, thresholding and morphological operations to identify cracks. These methods produced inconsistencies across varying lighting conditions and various surface types, and also struggled with thin cracks, background noise and texture variations, which motivated a shift toward approaches that combine convolutional feature extraction with image processing to characterise cracks more reliably [7].

CNNs brought significant improvement in crack detection over classical methods. It is also important to note that U-Net and similar segmentation models were applied for pixel-level crack mapping, with deep learning models proving more robust across varying conditions than traditional methods [8, 9].

YOLO models introduced real-time capabilities to object detection applicable to structural inspection. YOLOv5 and earlier variants showed strong results in crack and defect detection tasks. With YOLOv8, architectural improvements made it faster and more accurate than its predecessors. However, prior studies have rarely applied YOLOv8 specifically to crack detection and until recently, none had systematically benchmarked its variants [10, 11].

While deep learning and YOLO based approaches have advanced the field of crack detection significantly, existing studies have focused more on single model applications without systematic variant benchmarking. The accuracy versus efficiency trade-off and crack detection within Nigerian infrastructure assessment and resource-constrained deployment contexts have remained largely unexplored, and that is where this study comes in to address both gaps.

3. Materials and Methods

This study adopts an experimental research design, where three YOLOv8 variants are trained and evaluated under identical conditions to ensure a fair and unbiased comparison across all models. The entire pipeline was executed on Google Colab using a T4 GPU.

The dataset used is the Building Crack Detection dataset sourced from Roboflow Universe [5], comprising 500 total images across a single class, Cracks. The dataset was pre-split by Roboflow into training (360 images), validation (90 images) and test (50 images) sets. Images were preprocessed and resized to 640x640, with data augmentation automatically applied during export. The relatively small size of the dataset is acknowledged as a limitation of this study, as larger datasets may yield more generalisable results. Figure 1 presents sample images from the Building Crack Detection dataset, showing the variety of crack appearances and surface types that are present in the training data.

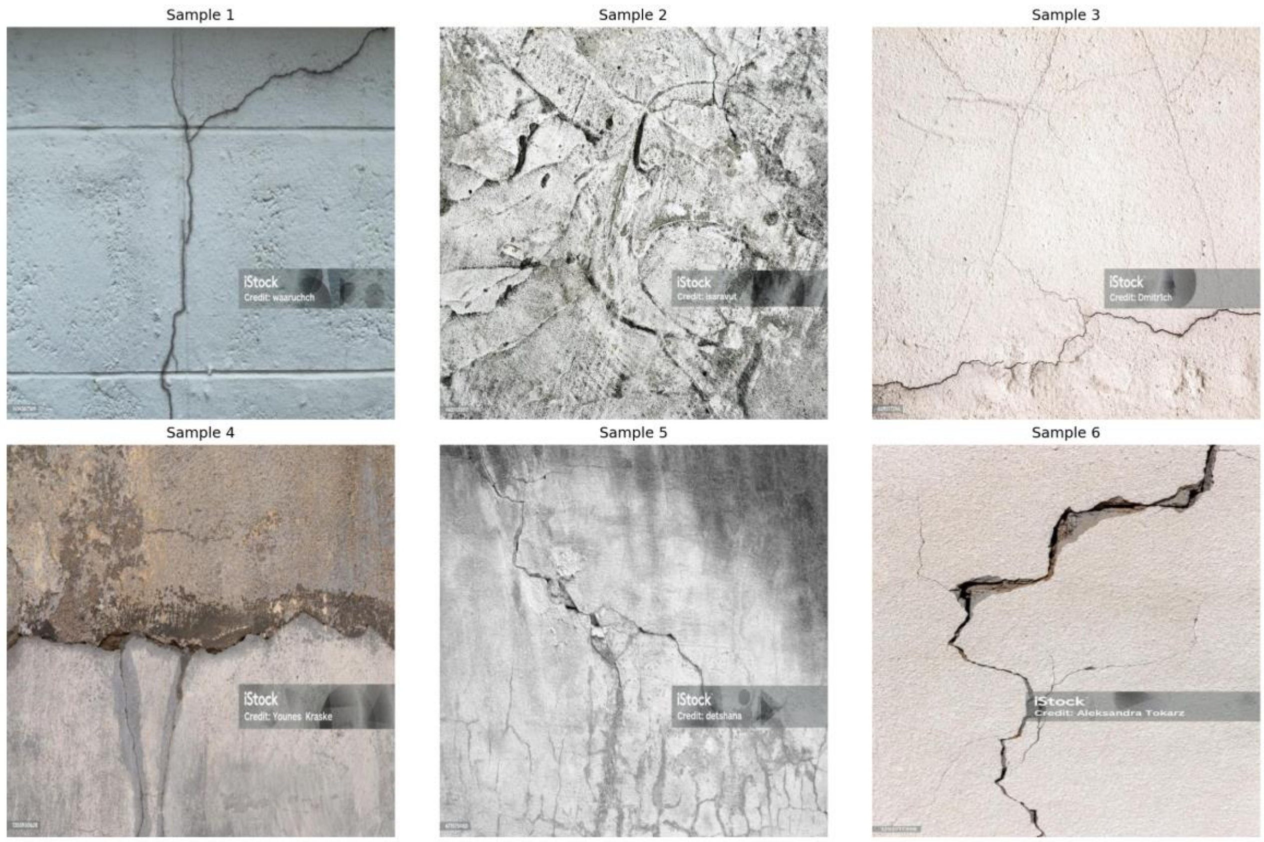

Figure 1. *Sample Training Images.*

Three YOLOv8 variants were selected for benchmarking. YOLOv8n (nano), with 3.2M parameters, was selected for its speed and lightweight design. YOLOv8s (small), with 11.2M parameters, was selected for its balance of speed and accuracy. YOLOv8m (medium), with 25.9M parameters, was selected for its higher accuracy. All three models were initialised with pretrained COCO weights [4].

All three models were trained under identical hyperparameters to ensure a fair comparison. These include 50 epochs, an image size of 640x640, a batch size of 16, and the AdamW optimizer, which is automatically selected by the Ultralytics framework. An early stopping patience of 15 epochs was also applied throughout training.

All models were evaluated on the same held-out test set using five metrics: mAP@50, mAP@50-95, Precision, Recall and Inference Speed. Inference Speed was included specifically to assess real-world deployment suitability in resource-constrained environments.

The results of all three variants were compiled into a single comparison table, with the best variant identified based on the highest mAP@50 combined with acceptable inference speed. Qualitative prediction visualizations were also included, with the discussion framed around Nigerian infrastructure deployment suitability.

4. Results and Discussion

All three YOLOv8 variants were evaluated on the same held-out test set of 50 images, with the results compiled into Table 1 for direct comparison.

Table 1. *YOLOv8 Variant Comparison for Crack Detection.*

Model	Parameters	mAP@50 (%)	mAP@50-95 (%)	Precision (%)	Recall (%)	Speed (ms)
YOLOv8n (Nano)	3.2M	99.41	98.70	97.83	100	3.1
YOLOv8s (Small)	11.2M	99.50	99.50	99.85	100	6.6
YOLOv8m (Medium)	25.9M	99.50	99.50	99.80	100	13.5

It is observed that all three model variants achieved a mAP@50 of at least 99.4%, with YOLOv8s and YOLOv8m reaching 99.50%.

While all three models achieved near-identical mAP@50 scores, YOLOv8s and YOLOv8m both achieved the highest mAP@50-95 at 99.50%, with YOLOv8n at 98.70%, demonstrating that the larger variants are more precise at tighter detection thresholds, as mAP@50-95 is a stricter evaluation metric. In terms of precision, YOLOv8s achieved 99.85%, YOLOv8m achieved 99.80%, and YOLOv8n achieved 97.83%. All three models also achieved a recall of 100%, indicating that no cracks were missed across the test set. Figure 2 presents a visual comparison of the three YOLOv8 variants across mAP@50, mAP@50-95, Precision and Recall.

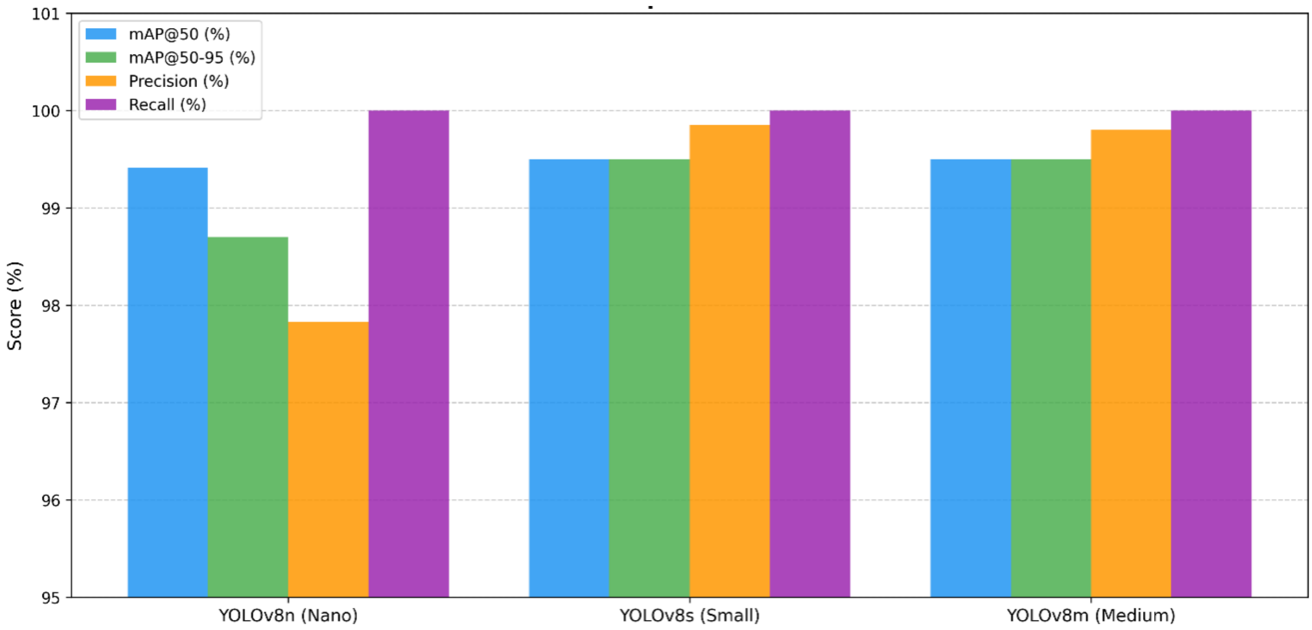

Figure 2. *Performance Bar Chart.*

In terms of inference speed, YOLOv8n was the fastest at 3.1ms per image, followed by YOLOv8s at 6.6ms per image, and YOLOv8m at 13.5ms per image. YOLOv8m is approximately 4.4 times slower than YOLOv8n; however, it is important to note that all three speeds remain suitable for real-world deployment. To put this in context, 3.1ms per image translates to approximately 323 frames per second, which is well above the threshold required for real-time crack detection during field inspection Regarding the accuracy versus speed trade-off, YOLOv8n offers the best balance, delivering near-identical accuracy to YOLOv8m while being 4.4 times faster. the difference in mAP@50-95 between YOLOv8n and YOLOv8m of 98.70% versus 99.50% represents a marginal gap that does not justify the speed cost especially in a resource constrained setting. For the Nigerian infrastructure context, YOLOv8n is recommended for drone or mobile camera deployment, as its low computational requirements make it well suited for edge devices such as Raspberry Pi, NVIDIA Jetson Nano and mobile phones, which are more readily available and affordable in resource-constrained environments like Nigeria. YOLOv8m, on the other hand, is recommended for scenarios where the highest possible accuracy is the priority and sufficient computational resources are available.

As seen in the prediction visualization figure, all three models successfully detected cracks across varying surface types, with consistent bounding box placement observed across all variants. Figure 3 presents the qualitative crack detection predictions generated by all three YOLOv8 variants on sample test images, demonstrating the consistent bounding box placement across models.

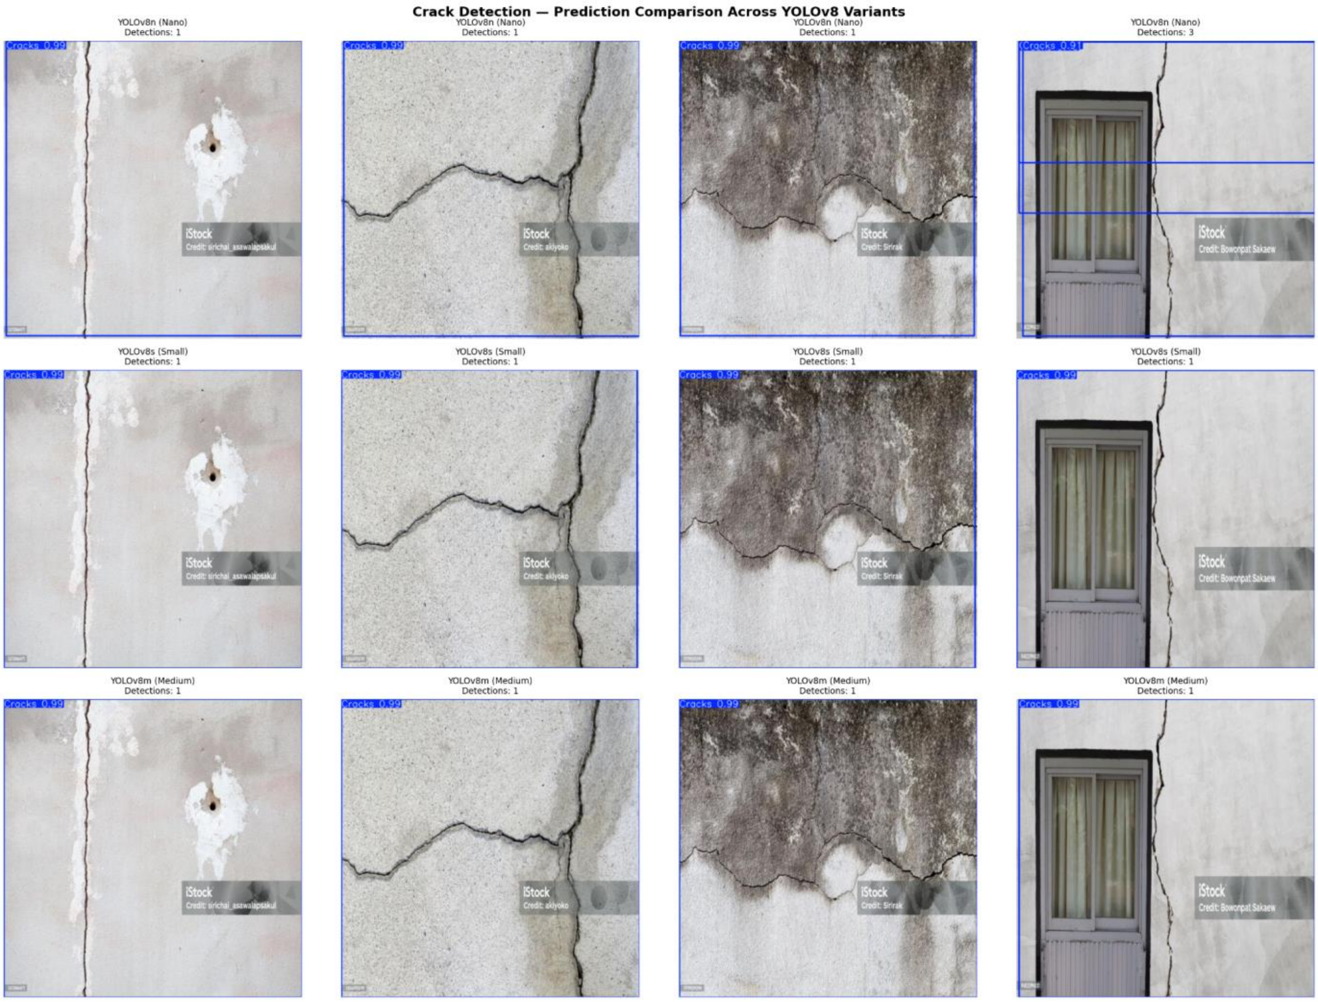

Figure 3. *Prediction Visualization.*

These findings can be situated within a small but growing body of work that has applied YOLOv8 to crack detection in civil infrastructure. Taffese et al. [10] benchmarked five YOLOv8 scales, from nano to extra-large, on a 1,170-image bridge crack dataset while also varying six optimizers, and found that YOLOv8 optimized with Stochastic Gradient Descent produced the strongest real-time performance; however, their evaluation is bridge-focused and does not consider building infrastructure or the specific deployment constraints of a developing economy such as Nigeria. Li et al. [11] took a complementary architectural approach, modifying YOLOv8 through the addition of attention mechanisms and deformable convolution to create a model they termed DCNA-YOLO, reporting a 2.0 percentage-point improvement in mAP over baseline YOLOv8 on bridge surface defects. Compared with these studies, the present work does not modify the YOLOv8 architecture; rather, it isolates model scale as the single controlled variable, training the nano, small and medium variants under identical hyperparameters on the same building-specific dataset, so that the accuracy-speed trade-off across scales can be attributed to scale alone rather than confounded by architectural changes or optimizer choice.

More recent work has continued to expand the YOLO-based crack detection space. Zhang et al. [12] proposed a YOLOv10-based model, YOLOv10-DECA, for real-time concrete crack detection, illustrating that newer YOLO releases are already being explored for this task; the present study nonetheless retains YOLOv8 because of its maturity, stable tooling and wide deployment support on constrained edge hardware, which are practical priorities for Nigerian field use. Owoeye et al. [13] integrated five attention modules into YOLOv8 for building crack detection and segmentation using a considerably larger dataset of 13,169 images with 19,386 annotations, achieving incremental gains of up to 3.5% in box precision over the unmodified baseline. The scale of their dataset relative to the 500 images used in the present study underlines a limitation discussed further below, but their results also indicate that attention-based refinements of YOLOv8 are a viable next step once a larger, Nigeria-specific crack dataset becomes available.

Within the Nigerian context specifically, prior work has largely addressed the causes and governance of building collapse [3] or the broader adoption of artificial intelligence tools for construction-site safety monitoring [14], rather than providing an empirically benchmarked, deployment-ready detection model. Bala Muhammad et al. [15] similarly call for greater integration of structural health monitoring and building information modelling to prevent building collapse in Nigeria, but stop short of evaluating a specific detection algorithm. Likewise, broader structural health monitoring reviews [6] highlight the promise of computer vision for building inspection but do not benchmark object detection models under the resource-constrained conditions typical of the Nigerian construction sector. By contrast, this study contributes a direct, like-for-like comparison of three deployable YOLOv8 scales on building-specific crack imagery, explicitly framed around the computational constraints of Nigerian field deployment, and arrives at a concrete recommendation, YOLOv8n, rather than a general discussion of feasibility. This constitutes the specific novelty of the present work relative to the existing literature.

It is important to acknowledge the 500-image dataset as a limitation of this study, as the results may not fully generalise to all crack types and surface conditions found across Nigerian infrastructure. Future work should incorporate larger and more diverse datasets to improve the generalisability of the findings.

5. Conclusion

With the goal of identifying the most suitable model for structural crack detection in Nigerian infrastructure, three YOLOv8 variants which are nano, small and medium were benchmarked under identical training and evaluation conditions.

It was observed that all three models achieved a mAP@50 of at least 99.4% and 100% recall on the test set, with YOLOv8s and YOLOv8m both achieving the highest mAP@50-95 at 99.50% and YOLOv8n achieving the fastest inference speed at 3.1ms per image. With YOLOv8n balancing both accuracy and speed for resource-constrained and edge device settings, it is recommended as the optimal model for Nigerian infrastructure deployment.

This study provides a foundation for affordable and scalable structural monitoring in Nigeria, by demonstrating that automated crack detection using YOLOv8 is both viable and accurate.

With the 500-image dataset acknowledged as a limitation, future work should explore larger Nigeria-specific datasets, extend benchmarking to YOLOv8l and YOLOv8x variants, and pursue real-world field testing on Nigerian buildings and bridges.

Abbreviations

mAP	Mean Average Precision
GPU	Graphics Processing Unit
CNN	Convolutional Neural Network
YOLO	You Only Look Once
COCO	Common Objects in Context

Acknowledgments

I acknowledge the support provided by Federal University of Technology Owerri (FUTO), through the Information Communication Technology (ICT) research center.

Author Contributions

Nnanna Ekedebe: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing

Conflicts of Interest

The author declares no conflicts of interest.

References

[1] Zhang X, Wang H, Hsieh YA, Yang Z, Yezzi A, Tsai YC. Deep Learning for Crack Detection: A Review of Learning Paradigms, Generalizability, and Datasets. arXiv preprint arXiv: 508.10256. 2025. https://arxiv.org/abs/2508.10256

[2] Farahzadi L, Odeh I, Kioumarsi M, Shafei B. Automated Image-Based Condition Assessment of the Built Environment: A State-of-the-Art Investigation of Damage Characteristics and Detection Requirements. Results in Engineering. 2025, 26, 104978. https://doi.org/10.1016/j.rineng.2025.104978

[3] Okunola OH. Quantifying Frequent Building Collapse and Disaster Risk Reduction in Nigeria. Africa in Focus, Brookings Institution. 2022. https://www.brookings.edu/articles/quantifying-frequent-building-collapse-and-disaster-risk-reduction-in-nigeria

[4] Jocher G, Chaurasia A, Qiu J. Ultralytics YOLOv8 (Version 8.0.0) [Computer software]. GitHub. 2023. https://github.com/ultralytics/ultralytics

[5] Gobinath G. Building Crack Detection Dataset. Roboflow Universe. 2024. https://universe.roboflow.com/gobinath-g-lncll/building-crack-detection-dmhfv

[6] Mociran HA, Lăpuște AV. Structural Health Monitoring of Buildings Using Computer Vision: A State-of-the-Art Review. Engineering and Technology Quarterly Reviews. 2025, 8(2), 447. https://doi.org/10.5281/zenodo.17625389

[7] Qayyum W, Ehtisham R, Bahrami A, Mir J, Khan QUZ, Ahmad A, Özkılıç YO. Predicting Characteristics of Cracks in Concrete Structure Using Convolutional Neural Network and Image Processing. Frontiers in Materials. 2023, 10, 1210543. https://doi.org/10.3389/fmats.2023.1210543

[8] Ali R, Chuah JH, Abu Talip MS, Mokhtar N, Shoaib MA. Structural Crack Detection Using Deep Convolutional Neural Networks. Automation in Construction. 2022, 133, 103989. https://www.sciencedirect.com/science/article/abs/pii/S0926580521004404

[9] Nyathi MA, Bai J, Wilson ID. Deep Learning for Concrete Crack Detection and Measurement. Metrology. 2024, 4(1), 66-81. https://doi.org/10.3390/metrology4010005

[10] Taffese WZ, Sharma R, Afsharmovahed MH, Manogaran G, Chen G. Benchmarking YOLOv8 for Optimal Crack Detection in Civil Infrastructure. arXiv preprint arXiv:2501.06922. 2025. https://arxiv.org/abs/2501.06922

[11] Li T, Liu G, Tan S. Superficial Defect Detection for Concrete Bridges Using YOLOv8 with Attention Mechanism and Deformation Convolution. Applied Sciences. 2024, 14(13), 5497. https://doi.org/10.3390/app14135497

[12] Zhang C, Peng N, Yan J, Wang L, Chen Y, Zhou Z, Zhu Y. A Novel YOLOv10-DECA Model for Real-Time Detection of Concrete Cracks. Buildings. 2024, 14(10), 3230. https://doi.org/10.3390/buildings14103230

[13] Owoeye S, Durodola F, Abdulkareem S, Omotainse O. Integrating Attention Modules with YOLOv8 for Enhanced Crack Detection and Segmentation. Kurdistan Journal of Applied Research. 2026, 11(1), 121-142. https://doi.org/10.24017/science.2026.1.9

[14] Adamu II, Okanlawon TT, Oyewobi LO, Shittu AA, Jimoh RA. Revolutionising Construction Safety: Benefits of Harnessing Artificial Intelligence Tools for Dynamic Monitoring of Safety Compliance on Construction Projects in Nigeria. International Journal of Building Pathology and Adaptation. 2024. https://doi.org/10.1108/IJBPA-02-2024-0050

[15] Bala Muhammad I, Omoebamije O, Bala Muhammad M, Adam Nabage Y. Smart Technology Integration: SHM and BIM for Preventing Building Collapses. African Journal of Advances in Science and Technology Research. 2024, 17(1), 69-88. https://doi.org/10.62154/ajastr.2024.017.010427