



Research Article

Pothole Localization Using Fine-Tuned YOLOv8n for Road Maintenance Prioritization in Resource-Constrained Regions

Nnanna Ekedebe* 

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

Abstract

There is a significant safety and infrastructure challenge in resource-constrained regions like Nigeria and Africa posed by damaged roads and potholes. Over 70% of Nigeria's paved roads are damaged, and this contributes to road accidents and high vehicle repair costs. The current method of road inspection is manual, which is slow and reactive rather than proactive, making an automated detection system necessary. Deep learning models can scan large numbers of road images quickly and mark exactly where the damage is, which helps road agencies decide which roads to fix first. This study fine-tunes a pre-trained YOLOv8n object detection model on a publicly available pothole dataset comprising 665 images to automatically detect and localize road damage from images. The model was trained for 50 epochs using a free cloud GPU, with the dataset split into training, validation, and test sets. Data augmentation techniques such as flipping, rotating, and brightness adjustment were also applied to improve the model's ability to handle different road conditions. The model achieved an mAP@0.5 of 78.0%, precision of 81.4%, and recall of 67.5%, outperforming the dataset baseline of 74.0%. These results show that the model can correctly detect most potholes while keeping false detections low, though about one in three potholes was still missed, meaning the model works best as a support tool alongside human inspectors rather than a full replacement. These results also demonstrate that fine-tuning a lightweight pre-trained model on free cloud hardware represents a viable and accessible approach to automated road infrastructure monitoring in resource-constrained African settings.

Keywords

Pothole Detection, YOLOv8, Object Detection, Road Infrastructure Monitoring, Deep Learning, Nigeria

1. Introduction

The importance of road infrastructure to a functioning economy cannot be overemphasized; transportation underpins trade, commuting, emergency response, and daily life. Unfortunately, roads degrade constantly over time due to weather, traffic load, and even poor construction, which can lead to accidents, with serious road trauma now costing the world more

than \$2.2 trillion annually, or \$6.1 billion every day [1]. Potholes are not just an inconvenience but are a cause of vehicle damage, accidents, and even fatalities. Africa, and especially Nigeria, faces this problem acutely. The lack of repair of these faulty roads is not just an inconvenience alone; it is also tied to economic loss, that is repair of vehicles and transportation

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

Received: 13 July 2026; Accepted: 23 July 2026; Published: 18 September 2026



delay, and safety risks as well, especially accidents on motorcycles which are a very popular means of transportation in Nigeria. This is reflected in the scale of the problem nationally, with over 70% of Nigeria's paved road network reported as dilapidated [2].

The current approach for road surveys relies on inspection by road authorities, citizens' complaints, or periodic surveys. Manual inspections like these are slow and the measures are reactive rather than proactive, with actions taken mostly after complaints are made or an accident has happened. Covering thousands of kilometers of road manually is simply not feasible, especially with limited personnel and budget constraints.

Deep learning-based object detection has brought an answer to the scale and speed problem. It can process large amounts of videos and images consistently and automatically, shaving off the time and manpower it would take to manually carry out these inspections. Models like YOLO do not just classify into damaged or not damaged, they can locate exactly where the damage is in the image, which helps prioritize repairs more effectively. This paper sets out to fine-tune a pre-trained YOLOv8 object detection model for pothole and road damage identification, evaluate it using standard object detection metrics such as mAP, precision, and recall, and demonstrate its relevance to Infrastructure monitoring, especially in resource-constrained regions like Africa.

The paper follows the following structure: section two reviews the related work, section three describes the methodology, section four presents the results, section five the discussion, and section six the conclusion.

2. Related Work

There has been significant growth over the past decade in research into road damage detection using deep learning, with the field moving from manual and sensor based methods toward CNN and YOLO based approaches. This section reviews the most relevant work and identifies where this study fits in.

Ukhwah et al. [3] conducted early work that explored YOLO neural networks specifically for asphalt pavement pothole detection, establishing that deep learning could outperform traditional image processing methods for this task. However early work in this field struggled with real-time processing speed and localization accuracy, models could detect that damage existed but could not precisely pinpoint where in the image it was.

The next development in the field was transfer learning, where instead of training from scratch, pre-trained CNN models were adapted to pothole detection. This approach achieved 96% accuracy, outperforming models trained from scratch, which matters because transfer learning reduces data and computing requirements while improving accuracy. The limitation however was that the study focused on classification only and not localization. It could confirm a pothole existed but not identify where exactly it was in the image [4].

A comparative study evaluating YOLOv5, YOLOv6, and

YOLOv7 on road defect detection found that newer YOLO versions consistently improved on their predecessors in both detection accuracy and inference speed [5]. However no single study at that point had directly compared all versions up to YOLOv8 on pothole data specifically, leaving room to evaluate the most current and optimized version.

The most relevant prior work was a 2023 study that used YOLOv8 for automated real time detection and severity assessment of asphalt pavement distress, confirming YOLOv8 as the strongest performing YOLO variant for this task across diverse road condition images [6]. The limitation was that the study focused on severity classification and assessment rather than localization for repair prioritization, which is the gap this study addresses.

Building on this body of work, this study fine-tunes YOLOv8 on a publicly available pothole dataset and evaluates it using mAP, precision, and recall. While YOLOv8 has been applied to road damage detection in various contexts, its application to road infrastructure monitoring in resource-constrained African settings remains largely unexplored; that is the specific angle this study extends the existing literature toward.

3. Materials and Methods

This section shows the steps taken to fine-tune and evaluate a pre-trained YOLOv8 object detection model for pothole and road damage identification. It covers the following: dataset, data preprocessing, model selection, training configuration, and evaluation metrics.

3.1. Dataset

The dataset used in this study is the Roboflow Pothole Detection dataset, a publicly available annotated pothole dataset comprising 665 images collected across varied road conditions. The dataset is pre-annotated with bounding boxes in YOLO-compatible format, meaning no manual labeling was required. It contains a single damage class pothole and was split into 465 training images, 133 validation images, and 67 test images for the purposes of this study [7].

3.2. Data Training

Before training, the images were put through preprocessing steps to make them suitable for the model. Images were resized to 640×640 pixels, the standard input size expected by YOLOv8. Data augmentation was applied, that is, artificially expanding the dataset by flipping, rotating, and adjusting the brightness of existing images so the model learns to handle varied real-world conditions. As the dataset is officially divided into a train set and a test set, a validation set was carved out from the training portion, with the data split into 70% training, 20% validation, and 10% testing for the purposes of this study.

3.3. Model Selection and Architecture

YOLOv8 (You Only Look Once, version 8) was selected as the object detection model for this study, as it represents the most current and optimized version of the YOLO family at the time of this study, offering strong performance in both detection accuracy and inference speed [8]. The specific variant used is YOLOv8n (nano), a lightweight version well suited to fine-tuning on a single GPU without requiring large computational resources. It is important to note that the model is fine-tuned rather than trained from scratch; it begins with weights pre-trained on the COCO dataset and is then adapted specifically to the pothole and road damage detection task.

3.4. Training Configuration

The model was trained using the following configuration: a batch size of 16 images per step, an image input size of 640×640 pixels consistent with the preprocessing stage, and training was run for 50 epochs. The optimizer was automatically selected by YOLOv8 as AdamW with a learning rate of 0.002. All training was performed on Google Colab using a free-tier GPU instance, which provided sufficient computational resources for fine-tuning a lightweight YOLOv8n model on the Roboflow pothole dataset.

3.5. Evaluation Metrics

The model's performance was evaluated using standard object detection metrics. Mean Average Precision (mAP) measures overall detection accuracy across all damage classes

and confidence thresholds, with a perfect score of 1.0. Precision measures how many of the damage instances flagged by the model were actually correct, while Recall measures how many of the actual damage instances present in the images the model successfully identified. Intersection over Union (IoU) measures how well the predicted bounding box overlaps with the ground truth annotated box, with a threshold of 0.5 applied in this study, reported as mAP@0.5.

4. Results

This section presents the results obtained from fine-tuning YOLOv8n on the pothole dataset. The model was evaluated on a held-out test set of 67 images after training for 50 epochs, with results reported using mAP@0.5, precision, recall, and mAP@0.5-0.95.

4.1. Training Performance

It was noticed as the training went on that the model showed improvements across all 50 epochs, with the losses decreasing steadily throughout the training. By epoch 50 the box loss had dropped from 1.698 at epoch 1 down to 0.892, and classification loss dropped from 2.742 down to 0.622, this proved that the model was learning effectively. The training was completed in under 10 minutes on Google Colab GPU instance proving that fine tuning a lightweight YOLOv8n model can be done even when there are hardware constraints. The training loss curves are shown in Figure 1.




Figure 1. YOLOv8n Training Loss Curves over 50 Epochs.

4.2. Detection Performance

The following validation results were achieved during training: mAP@0.5: 78.0%, mAP@0.5-0.95: 50.8%, Precision: 81.4%, Recall: 67.5%. What these numbers mean is that during training, the model correctly identified potholes in 78% of the cases at a standard IOU threshold of 0.5, with a

precision score of 81.4%. This, therefore, means that the majority of the detections flagged by the model were genuine potholes rather than false alarms. It is also important to know that the recall of 67% meant that the model successfully found roughly two-thirds of the actual potholes present in the images. The Precision-Recall curve is presented in Figure 2.




Figure 2. Precision-Recall Curve for Pothole Detection.

4.3. Confusion Matrix Analysis

The confusion matrix confirms the model's strong true positive rate for pothole detection. The normalized confusion

matrix shows the model correctly classifying the majority of pothole instances, with a portion of background instances occasionally misclassified a common characteristic in single-class object detection tasks with varied backgrounds, as shown in Figure 3.




Figure 3. Normalized Confusion Matrix for Pothole Detection.

4.4. Sample Detection Results

The model was run on all 67 test images and successfully detected potholes across a range of road conditions and image sizes. Detection inference speed averaged 19.3ms per image, indicating that the model is capable of near real-time

processing, an important characteristic for practical deployment in road monitoring systems. Sample detection images showing bounding boxes accurately placed around pothole instances across varied road surfaces and lighting conditions are presented in Figure 4.



Figure 4. Sample Pothole Detection Results on Test Images.

Overall, the model achieved a strong performance on the pothole dataset with results demonstrating both accuracy and computational efficiency, and the fact that it was done in under 10 minutes proves that this approach is viable for real-world deployment especially in resource constrained settings like Nigeria and Africa as a whole.

5. Discussion

This section interprets the results presented above and discusses their broader implications for road infrastructure monitoring.

With a precision score of 81.4%, in a practical road monitoring system roughly 8 out of every 10 detections would be genuine potholes, greatly reducing the burden on road inspectors who would otherwise have to manually survey entire road networks. Although this is positive, a recall score of 67.5% means that roughly one in three potholes is still being missed, which suggests the model is best suited as a decision support tool rather than a fully autonomous system; human oversight is still required. Notwithstanding this, automated first-pass screening across thousands of kilometers of road remains far more efficient than purely manual inspection.

Looking back at the transfer learning work reviewed in Section 2 which achieved 96% classification accuracy, it is important to note that classification alone does not answer where the damage is. This study's model performed localization, placing bounding boxes around detected damage, which is a fundamentally harder task, making a direct accuracy comparison inappropriate. In terms of a more relevant benchmark, the

Roboflow baseline mAP@0.5 for this specific dataset stood at 74.0%, which this study's model outperformed at 78.0%.

The Federal Road Safety Corps reported 5,421 road deaths in Nigeria in 2024, a 7% increase from the previous year, with damaged roads playing a contributing role [9]. With agencies like FERMA responsible for tens of thousands of kilometers of road and lacking the resources for continuous manual monitoring, a detection system like this directly addresses a real gap in Nigeria and other African countries. This model gives road agencies a tool to prioritize where repairs are most urgently needed, rather than waiting for complaints or accidents to reveal the damage.

Beyond the Roboflow baseline, the precision-recall gap in this study becomes clearer when set against other recent YOLOv8-based pothole detection work. Khan et al. apply YOLOv8 within an autonomous-vehicle detection pipeline [10], and Chang et al.'s RAW-YOLOv8 adds deformable attention to improve multi-scale localization [11], while Bhavana et al.'s POT-YOLO, which layers edge segmentation pre-processing ahead of detection, reports a precision of 97.6% and recall of 93.52% [12] considerably higher than the 81.4% precision and 67.5% recall achieved here. What distinguishes these higher-performing models from this study's baseline fine-tune is less the underlying YOLOv8 architecture than the additional engineering layered on top of it, none of which was applied in this minimal fine-tuning approach.

The primary limitation of this study is the dataset, 665 images is relatively small for a production-level detection system, and the images come from countries like the United States, India, and Japan rather than African road conditions specifically. While some Africa-specific road datasets exist, such as

a Nigerian pothole dataset collected from Ekiti State, they remain limited in size and public availability, and the limited standardization of such datasets in annotation formats suitable for object detection training remains a barrier to Africa-specific model development further motivating future data collection efforts. Nigerian and African roads may have different visual characteristics in terms of surface textures, lighting conditions, and damage patterns that the current model has not been exposed to. Additionally, a missed detection rate of roughly one in three means the model is not yet suitable as a standalone safety-critical system without human verification.

Mirajkar et al. similarly applied YOLOv8 directly to road condition assessment without heavy architectural modification, achieving competitive results primarily by training on a substantially larger and more varied dataset [13], reinforcing that dataset scale and diversity, rather than model choice alone, remains a key lever for improving detection performance in future work.

Future work could focus on collecting or expanding Africa-specific road damage datasets to further fine-tune the model on locally relevant conditions, which would likely improve recall specifically. Architectural enhancements offer another promising direction: Yurdakul and Taşdemir show that adding lightweight attention and edge-aware convolution modules to a YOLOv8n-seg baseline can lift precision and recall by several points with minimal added complexity [14], while Karukayil et al. demonstrate that fusing vision with LiDAR depth data can extract pothole features that RGB imagery alone misses [15] — a particularly relevant direction for future deployment on Nigerian and African roads, where lighting and surface conditions vary widely. Deployment on mobile or edge devices mounted on moving vehicles for real-time road monitoring would also be a practical next step toward real-world use.

6. Conclusion

This study fine-tuned YOLOv8n for pothole detection and achieved an mAP@0.5 of 78.0%, demonstrating that lightweight pre-trained models can be effectively adapted for road damage detection. The results confirm that this approach is viable for road infrastructure monitoring in resource-constrained areas such as Nigeria and Africa in general. Future work should focus more on collecting Africa-specific road damage data to improve model generalization, and on deploying the model on edge devices for real-time road monitoring.

Abbreviations

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

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, Writing – review & editing

Conflicts of Interest

The author declares no conflicts of interest.

References

- [1] International Road Assessment Programme (iRAP). Safety Insights Explorer. Available from: <https://irap.org/safety-insights/global-toll/> (accessed 13 July 2026).
- [2] Autorods. Over 70% of Nigeria's Paved Roads Dilapidated. Available from: <https://autorods.com.ng> (accessed 13 July 2026).
- [3] Ukhwah, E. N., Yuniarno, E. M., Suprpto, Y. K. Asphalt Pavement Pothole Detection Using Deep Learning Method Based on YOLO Neural Network. 2019 International Seminar on Intelligent Technology and Its Applications (ISITIA). 2019, 35-40. <https://doi.org/10.1109/ISITIA.2019.8937176>
- [4] Vinodhini, K. A., Sidhaarth, K. R. A. Pothole Detection in Bituminous Road Using CNN with Transfer Learning. Measurement: Sensors. 2024, 31, 100940. <https://doi.org/10.1016/j.measen.2023.100940>
- [5] Yusof, N. I. M., Sophian, A., Mohd Zaki, H. F., Bawono, A. A., Embong, A. H., Ashraf, A. Assessing the Performance of YOLOv5, YOLOv6, and YOLOv7 in Road Defect Detection and Classification: A Comparative Study. Bulletin of Electrical Engineering and Informatics. 2024, 13(1), 350-360. <https://doi.org/10.11591/eei.v13i1.6317>
- [6] Addanki, A., Lin, J. Pothole Detection with YOLOv8. Available from: <https://www.researchgate.net/publication/376685010> (accessed 13 July 2026).
- [7] Dwyer, B. Pothole Detection Dataset. Roboflow Universe. Available from: <https://universe.roboflow.com/brad-dwyer/pothole-voxrl> (accessed 13 July 2026).
- [8] Ultralytics. YOLOv8 Documentation. Available from: <https://docs.ultralytics.com> (accessed 13 July 2026).
- [9] Federal Road Safety Corps (FRSC). 2024 Annual Road Safety Report. Available from: <https://punchng.com/road-crashes-claimed-5421-lives-in-2024-frsc/> (accessed 13 July 2026).

- [10] Khan, M., Raza, M. A., Abbas, G., Othmen, S., Yousef, A., Jumani, T. A. Pothole Detection for Autonomous Vehicles Using Deep Learning: A Robust and Efficient Solution. *Frontiers in Built Environment*. 2024, 9, 1323792. <https://doi.org/10.3389/fbuil.2023.1323792>
- [11] Chang, J., Chen, Z., Xia, E. Improved YOLOv8 Method for Multi-scale Pothole Detection. *Advanced Intelligent Computing Technology and Applications (ICIC 2024)*, Lecture Notes in Computer Science, vol. 14872. 2024. https://doi.org/10.1007/978-981-97-5612-4_33
- [12] Bhavana, N., Kodabagi, M. M., Muthu Kumar, B., Ajay, P., Muthukumaran, N., Ahilan, A. POT-YOLO: Real-Time Road Potholes Detection Using Edge Segmentation-Based YOLOv8 Network. *IEEE Sensors Journal*. 2024, 24(15), 24802-24809. <https://doi.org/10.1109/JSEN.2024.3399008>
- [13] Mirajkar, R., Yenikar, A., Nawalkar, S. Enhanced Pothole Detection in Road Condition Assessment Using YOLOv8. *International Conference on Wireless Innovation and Technology (ICWITE)*. 2024, 429-433. <https://doi.org/10.1109/ICWITE59797.2024.10502437>
- [14] Yurdakul, M., Taşdemir, Ş. An Enhanced YOLOv8 Model for Real-Time and Accurate Pothole Detection and Measurement. Available from: <https://arxiv.org/abs/2505.04207> (accessed 13 July 2026).
- [15] Karukayil, A., Quail, C., Auat Cheein, F. Deep Learning Enhanced Feature Extraction of Potholes Using Vision and LiDAR Data for Road Maintenance. *IEEE Access*. 2024, 12, 184541-184549. <https://doi.org/10.1109/ACCESS.2024.3512783>