International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 11, November 2026
Received: 01-10-2026 Accepted: 05-11-2026 Published: 12-11-2026

REAL-TIME ROADWAY DEFECT LOCALIZATION WITH YOLO26: EFFICIENT VISION FOR SMART INFRASTRUCTURE

Mosapoina Rajesh [1], M. Sreedhar [2]
[1]PG Scholar, Department of ECE,
JNTUA College of Engineering, India
[2]Assistant Professor, Department of ECE,
JNTUA College of Engineering, India
Gmail : rajesheranna2@gmail.com, sreedhar470@gmail.com

Abstract

Despite the long and complicated history of the problem, pothole detection is still a significant challenge in intelligent transportation and road asset management due to the fact that irregular degradation of the road surface is directly related to safety, the integrity of vehicles, and maintenance prioritization. This paper uses the lightweight real-time object detector YOLO26n as a standard model for the pothole localization problem, and evaluates it on the two available datasets on Kaggle that include various road surface textures, lighting conditions, and pothole defect forms. The model is trained for 100 epochs using an image resolution of 640 images and a batch size of 16, on dual NVIDIA Tesla T4 GPUs, so it can be evaluated under real training conditions to understand how well the model performs in different scenarios.

Experimental results show a precision of 0.7567, recall of 0.6549, mAP50 of 0.7477, and mAP50-95 of 0.4786. The values show that YOLO26n is a reliable object detection method to define pothole regions at certain overlap values, whereas at smaller overlap values, it is challenging to achieve tighter localization because the geometry of the pothole is not regular, the contrast is low and the visibility of the pothole is unclear. Further validation results, precision-recall results, confidence curves and confusion matrices also validate the fact that the detector performs well on simple pothole cases and still fails on some cases due to missed detections and imprecise boundaries in complex cases.

The overall success of this work suggests that YOLO26n is a promising lightweight baseline model for pothole detection and can be applied to reproducible benchmarking on public datasets. The study also points to the need for further research in precise defect localization, which will be encouraged by segmentation, refinement strategies or improved post-processing for infrastructure monitoring applications.

Keywords

YOLO26, Pothole detection, Real-time object detection, Road infrastructure monitoring, Computer vision.

I. INTRODUCTION

The problem of road surface deterioration due to potholes is still a significant concern in road transportation safety, vehicle durability and maintenance planning, thus making the automatic assessment of the condition of the road surface an important research area in the monitoring of infrastructure. The key challenges with traditional inspection methods are high labor intensity, slow inspection speed, and inconsistent results. The vision-based approach can be a scalable solution to continuous and objective identification of potholes. The size, texture, shape, and visibility of the pothole are different in various scenarios, particularly in shadowed conditions, with blur and uneven pavement. Therefore, real-time detection of pothole is a challenging computer vision task.

This has led to significant improvements in road-defect analysis with the help of the family of object detectors 'You Only Look Once' (YOLO) in recent years. YOLO methods are attractive because they are single-stage object localization and classification systems that provide more rapid inference than many two-stage object detection systems, and can be used for traffic and infrastructure applications.

But the existing studies have still reported the shortcomings of the ability to detect small defects, to process low-contrast road images and to achieve both detection accuracy and model efficiency in reality[1],[2].

This paper is aimed at studying the YOLO26n model for the real time pothole detection in road images which is light weight in nature. Two public Kaggle datasets, pothole-detection-dataset-yolov11-optimized and potholes-detection-yolov8 are used to set up a reproducible evaluation setting for the benchmark. The training on Kaggle takes place over 100 epochs (cycles) with a size of 640 images and a batch size of 16 devices, which are dual NVIDIA Tesla T4 graphics processing units. This research aims to investigate if YOLO26n can overcome the limitation of the previous road-defect detection using YOLO

International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 11, November 2026
Received: 01-10-2026 Accepted: 05-11-2026 Published: 12-11-2026

models, and offer an effective and efficient solution for pothole detection.

The primary research question is whether YOLO26n can be designed to achieve an enhanced utility performance in terms of pothole detection for road infrastructure monitoring with lightweight features to be used in real-time. The goal of this work is to compare the effectiveness of YOLO26n when applied to the problem of pothole detection with some public annotated datasets and place it in the recent development of vision-based infrastructure monitoring systems. The study is significant due to the fact that most of the literature published in the recent years concerns the earlier YOLO generations and only a few studies have been published about the use of newer detector families for the purpose of pothole detection. The study's primary drawback is that the data it analyzes is public data, and that it uses controlled evaluations of models instead of field deployment, thus limiting its contribution to methodological and comparative.

II. RELATED WORK

1 Vision-Based Pothole Detection

Pothole detection using vision has become a significant research topic in transportation engineering due to the fact that potholes have a direct impact on road safety, driving comfort, maintenance cost and reliability of the infrastructure. Previous reviews indicate that computer-based pothole detection is preferred over manual surveys due to the drawbacks of conventional surveys in terms of time, lack of consistency and the ability to cover large road networks[3],[4]. However, the detection of potholes is still a challenging visual task as the geometry of pothole defects is not fixed, the boundaries of these defects are not always clearly defined, high similarity in the surrounding asphalt, shadowing, and motion blur are some of the problems[5]. Under these conditions, pothole detection is more difficult than the general object recognition problem, and requires models to be able to find small, deformable and visually ambiguous objects in real-time.

2 Deep Learning and YOLO for Road Defects

In particular, the deep learning models have made considerable strides in the field of road-defect analysis, particularly those utilizing convolutional networks for detecting objects and segmenting them. The You Only Look Once (YOLO) family is popular because they do both localization and classification in a single stage pipeline, which is suitable to real-time applications in the infrastructure. Newer models such as YOLOv8 have shown better defect detection results in recent studies, but it has also been found that the models are limited by the same problems as previous models, such as: failure to detect small potholes, false signals from rough surface textures on roads and illumination

Table 1. Training configuration

the lack of robustness in difficult conditions[6],[7]. Other recent efforts have also explored pothole detection and estimation of depth, point cloud fusion, and RGB-D segmentation, increasing the physical characterization and complexity of the system, and relying on additional sensing systems[8],[9]. The current research focus reflects the fact that in most cases better performance is achieved by incorporating additional architectural components or additional input signals rather than by testing to see if newer light-weight detector families can achieve the same performance on their own.

IV. PROPOSED APPROXIMATION TECHNIQUE

1 Dataset Description

In this work we use two public datasets available on Kaggle: pothole-detection-dataset-yolov11-optimized and potholes-detection-yolov8, to create a pothole detection corpus for benchmark purposes. The first dataset is already split into training, validation and test sets and the second dataset is used to give the data a wider variety of road surface appearances and pothole shapes. Selection of public data enables repeatability and ensures consistency of the study with current work on pothole detection based on annotated road images to evaluate the model. The datasets are appropriate for object detection for they contain the pothole instances in the form of images with their respective labels, which is an input requirement of YOLO26n.

2 Experimental Setup

The model was trained on Kaggle using two NVIDIA Tesla T4 GPUs, 100 epochs, image size of 640 and batch size of 16. This setup has been optimized to allow the same detector to be visually examined in a controlled, repeatable, and practical training setting to evaluate its performance in a benchmark. The system design has been made as simple as possible to ensure study focus is on the system suitability in the pothole detection task and not on deployment of the system with different multimodal acquisition.

3 Workflow of the Study

According to the fact that the primary goal of this sector is to develop a lightweight detector, YOLO26n is selected as the detector as it is chosen to have excellent behavior for small targets and efficient inference behavior. The model is also equipped with features for optimizing for the small target and boosting training stability, both of which are theoretically beneficial for pothole detection as small potholes tend to small-sized appearance on the image and around the pothole, there is also noise from asphalt around the pothole. The result of this study is fine-tuning the Kaggle pothole datasets using the standard Ultralytics detection pipeline with the pretrained weights of YOLO26n.

2250-3676 ijesatj@gmail.com https://www.ijesat.com/

International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 11, November 2026
Received: 01-10-2026 Accepted: 05-11-2026 Published: 12-11-2026

4 YOLO26n Architecture and Technical Design

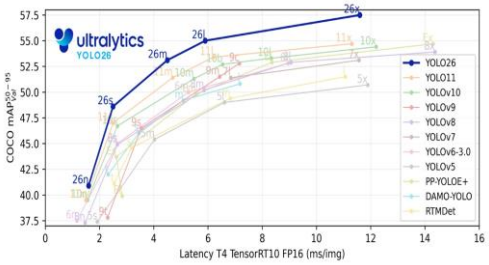
YOLO26n is selected in this study because it is designed as a lightweight end-to-end object detector optimized for efficient inference and improved small-object sensitivity. Unlike earlier YOLO variants that rely heavily on post-processing stages and computationally expensive detection refinement, YOLO26n incorporates an optimized lightweight detection pipeline suitable for real-time deployment scenarios. The model architecture combines compact feature extraction layers with improved feature aggregation strategies to preserve important spatial information associated with small and irregular pothole regions. The architecture integrates Small-Target-Aware Label Assignment (STAL), which improves the matching process between predicted regions and ground-truth annotations during training. This is particularly important for pothole detection because potholes often occupy limited image regions and exhibit irregular boundaries. YOLO26n also incorporates Progressive Loss Balancing (ProgLoss), which stabilizes optimization during training and reduces sensitivity to noisy surface textures, shadows, and low-contrast road conditions[16],[17],[18].

Parameter	*Value*
Model	*YOLO26n*
Hardware	*Kaggle T4 x2 GPU*
Epochs	*100*
Image size	*640*
Batch size	*16*
Task	*Object detection*
Training mode	*Transfer learning / fine-tuning*

Another important characteristic of YOLO26n is the MuSGD optimization strategy, which improves convergence stability while maintaining lightweight computational behavior. The detector further adopts an end-to-end inference pipeline with reduced dependency on computationally expensive post-processing operations, enabling lower latency during real-time deployment.

The lightweight configuration of YOLO26n makes it suitable for intelligent transportation systems, embedded monitoring devices, edge computing platforms, and smart infrastructure applications where computational efficiency and inference speed are critical.

Figure 1. YOLO26 Models Benchmark Comparision.



5 Training Procedure

The standard workflow for object detection from Ultralytics is used to train, including image loading, optimizing, augmenting, and evaluating. The input resolution is set to 640 pixels, which provides a good compromise between preserving fine detail and smooth training performance on the available GPU configuration, with a batch size of 16. To avoid the effects of dataset imbalance and achieve the model convergence over the entire pothole corpus, a cyclic training procedure is applied over 100 epochs.To achieve the model convergence over the entire pothole corpus and minimize the effect of dataset imbalance, a cyclic training procedure is applied for 100 epochs.

6 Evaluation Metrics

The model is tested with the metrics of the standard detection problem, including the precision, recall, F1 score and mean Average Precision (mAP). The use of these metrics in pothole detection is popular as they encapsulate both class prediction quality and class localization quality. Because pothole detection is an on-line infrastructure problem, inference speed is also taken into consideration as an effective measure.

Precision, recall, and F1-score are defined as:

Precision=TPTP + FP *Precision=TP+FPTP* (1)

Recall=TPTP + FN *Recall=TP+FNTP* (2)

F1=2 × Precision × Recall Precision + Recall *F*1 = *Precision* + *Recall*2 × *Precision* × *Recall* (3)

Mean Average Precision is computed from the precision-recall curve as:

mAP = 1N ∑I = 1 NAP i*mAP* = *N*1*I* = 1∑*NAPi* (4)

where TP*TP* is true positive, FP*FP* is false positive, FN*FN* is false negative, and APi*APi* is the average precision for class *i*.

VI. RESULTS AND DISCUSSION

1 Quantitative Performance:

The YOLO26n model achieves a precision of 0.7567, recall of 0.6549, mAP50 of 0.7477, and mAP50-95 of 0.4786 on the pothole detection benchmark. The values prove that the model is well suited for pothole detection within normal detection conditions and that the localization of the irregular defects in the road is difficult. The precision score reveals that the majority of the predicted potholes are correct, and the recall score indicates that there was still a chance to miss some potholes, particularly in highly complex scenes. The disparity between mAP50 and mAP50-95 indicates that the detector has a higher performance on road pothole general recognition than on the accurate road boundary fitting, a typical problem for road defect detection.

2250-3676 ijesatj@gmail.com https://www.ijesat.com/

International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 11, November 2026
Received: 01-10-2026 Accepted: 05-11-2026 Published: 12-11-2026

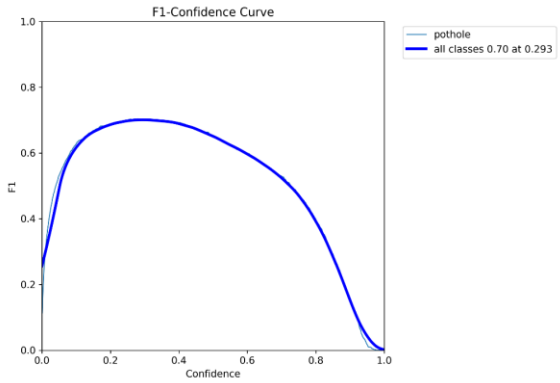


Figure 2. F1-Confidence Curve

This Precision-Confidence curve demonstrates the relationship between the two—high precision at high thresholds. The Recall-Confidence curve, however, suggests the opposite trend - the less confident they are, the better their recall will be. For practical pothole monitoring systems, the F1-Confidence curve indicates the optimum operating point is at a moderate level of confidence. For infrastructure monitoring, a slightly lower threshold is acceptable as those false detections may be more tolerable than failing to detect a dangerous pothole.

Metric	Value
Precision	0.7567
Recall	0.6549
mAP50	0.7477
mAP50-95	0.4786

Table 2 YOLO26n detection performance

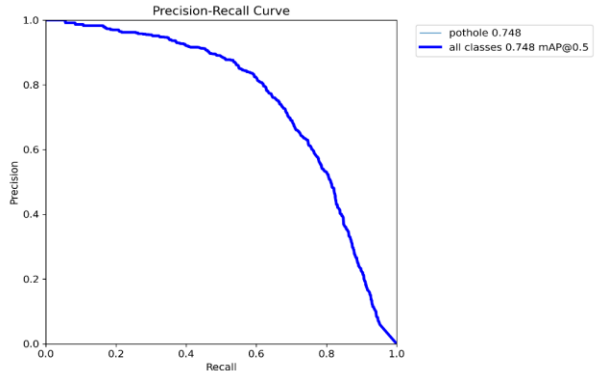


Figure 3. Precision-Recall Curve

The Precision-Recall curve exhibits a high precision at intermediate levels of recall, and the precision has a slight decline as the recall increases. This pattern shows that the detector is successful with those pothole images that are not as unclear, and less successful with the more complicated images that are not as easy to see. The shape of the curve is very common for datasets of road defects, like potholes, in a road that are collected from a variety of lighting, road surfaces, and weather conditions. The PR curve also demonstrates that the detection of potholes and avoidance of unnecessary false alarms is sufficiently well balanced by using YOLO26n.

Figure 9 (a)(b). Validation label samples

The prediction images show the validation results of YOLO26n which is able to detect pothole in various conditions of roads such as dry roads, wet roads, urban roads and damaged roads. In a few instances, the model produces confident detection of the large and small potholes. The model is further demonstrated to be robust in many of the scenes with the presence of partial occlusion and surface clutter, highlighting its potential in handling practical pothole screening.



Figure 4(a)(b). Validation label samples

However, in those validation images, there are a few images having duplicated detections, overlapping boxes and a lot of predictions in complicated scenes, this is why the mAP50-95 is lower.

The road condition and pothole shape diversity is confirmed by the images of the labels on the validation data set. This variety gives the benchmark a better scientific value, since it would be highly unlikely for a single road appearance to cause the model to over fit. Meanwhile, it makes localizing difficult, particularly in the case of broken edges or water in potholes.

VII. CONCLUSION

This study evaluated YOLO26n as a lightweight real-time detector for pothole identification using publicly available road-surface datasets. Experimental evaluation demonstrates that the detector achieves competitive pothole detection performance with a precision of 0.7567, recall of 0.6549, mAP50 of 0.7477, and mAP50-95 of 0.4786 while maintaining lightweight computational complexity suitable for real-time deployment.

2250-3676 ijesatj@gmail.com https://www.ijesat.com/

Received: 01-10-2026 Accepted: 05-11-2026 Published: 12-11-2026

Comparative benchmarking against previous lightweight YOLO nano variants further shows that YOLO26n achieves improved detection accuracy and lower computational overhead. The detector demonstrates strong capability for identifying pothole regions under varying road conditions, including irregular pavement textures, partial occlusion, and illumination variability.

Despite the positive results, accurate pothole boundary localization remains challenging in highly degraded road environments due to irregular defect geometry, shadow interference, water accumulation, and visually similar background structures. These limitations indicate that future research should investigate segmentation-assisted refinement, adaptive localization strategies, and multimodal sensing approaches to improve fine-grained defect characterization.

Overall, the study demonstrates that YOLO26n provides an effective balance between real-time inference efficiency and pothole detection accuracy, making it a practical candidate for intelligent transportation systems, smart-city infrastructure monitoring, and edge-based road inspection applications.

VIII. AUTHOR'S CONTRIBUTION

Conceptualization: Mosapoina Rajesh Methodology: Mosapoina Rajesh Investigation: Mosapoina Rajesh Writing – Original Draft: Mosapoina Rajesh. Supervision: M Sreedhar. Approval of the final text: Mosapoina Rajesh and M Sreedhar.

IX. ACKNOWLEDGMENTS

The authors thank Jawaharlal Nehru Technological University Ananthapur College of Engineering (Autonomous) for infrastructure support, and M Sreedhar (Assistant Professor, Department of ECE, JNTUACEA) for guidance throughout this work.

X. REFERENCES

[1] Li, X., Zhang, Y., & Wang, L. (2026). YOLO-ERCD: An upgraded YOLO framework for efficient road crack detection. PMC Articles. https://pmc.ncbi.nlm.nih.gov/articles/PMC12846229/

[2] Zhang, C., Zhao, J., Cong, S., Yu, Y., Yue, X., Shen, Y., & Hui, Y. (2026). SDC-YOLOv8: An improved algorithm for road defect detection. PMC Articles. https://pmc.ncbi.nlm.nih.gov/articles/PMC12845961/

[3] Wang, L., Li, X., & Chen, H. (2024). Road defect detection based on improved YOLOv8s model. Scientific Reports, 14(1), Article 67953. https://doi.org/10.1038/s41598-024-67953-3

[4] Zhao, J., Zhang, C., Cong, S., Yu, Y., Yue, X., Shen, Y., & Hui, Y. (2025). Real time road defect detection using advanced YOLO models. Engineering Letters, 33(6), 1-10. https://www.engineeringletters.com/issues_v33/issue_6/EL_33_6_09.pdf

[5] Fan, R., Ozgenel, C. F., & Hosking, M. (2025). Pavement pothole detection system based on deep learning. Journal of Traffic and Transportation Engineering (English Edition). https://doi.org/10.1016/j.jtte.2025.01.112

[6] Shi, X., & Wang, J. (2024). Visible & thermal imaging and deep learning based approach for pothole detection. UGPTI Reports. https://www.ugpti.org/resources/reports/downloads/mpc24-559.pdf

[7] Mathavan, S., Ali, K., & Rahman, M. (2024). Pothole detection and dimension estimation using in-vehicle imagery. Developments in the Built Environment, 17, Article 100239. https://doi.org/10.1016/j.dibe.2024.100239

[8] Fan, R., Bocus, M. J., Zhu, Y., Jiao, J., Wang, L., Ma, F., Cheng, S., & Liu, M. (2024). Pothole detection for autonomous vehicles using deep learning. Frontiers in Built Environment, 9, Article 1323792. https://doi.org/10.3389/fbuil.2023.1323792

[9] Kumar, A., & Singh, R. (2024). A review of vision-based pothole detection methods using deep learning. Sensors, 24(17), Article 5594. https://doi.org/10.3390/s24175594

[10] Li, Y., Wang, H., & Zhang, X. (2023). Deep learning-based pothole detection for intelligent transportation: A YOLOv5 approach. International Journal of Advanced Computer Science and Applications, 14(12), 242-250. https://doi.org/10.14569/IJACSA.2023.0141242

[11] Varadharajan, S., Jose, S., Sharma, K., Wander, L., & Mertz, C. (2022). Pothole detection using deep learning: A real-time and AI-on-the-edge perspective. Journal of Field Robotics, 39(1), 102-133. https://doi.org/10.1155/2022/9221211

[12] Gupta, A., Sharma, R., & Kumar, P. (2021). Real-time pothole detection using deep learning. arXiv preprint arXiv:2107.06356. https://doi.org/10.48550/arXiv.2107.06356

[13] Maeda, H., Sekimoto, Y., Seto, T., Kashiyama, T., Omata, H., & Miyamoto, T. (2021). Application of various YOLO models for computer vision-based real-time pothole detection. Applied Sciences, 11(9), 4200. https://doi.org/10.3390/app11094200

[14] Nielsen, D., & Bonde, A. (2021). Real-time machine learning-based approach for pothole detection. Expert Systems with Applications, 184, Article 115562. https://doi.org/10.1016/j.eswa.2021.115562

[15] Zhang, L., Yang, F., Zhang, Y. D., & Zhu, J. (2019). Road crack detection using deep convolutional neural network. IEEE International Conference on Image Processing (ICIP), 3708-3712. https://doi.org/10.1109/ICIP.2016.7533052

[16] Ultralytics. (2026). YOLO26 Documentation. Ultralytics Docs. https://docs.ultralytics.com/models/yolo26/

[17] Ultralytics. (2026). Meet YOLO26: Better, Faster, Smaller YOLO Models. Ultralytics Blog. https://www.ultralytics.com/blog/meet-ultralytics-yolo26-a-better-faster-smaller-yolo-model

[18] Ultralytics. (2026). YOLO26 Official GitHub Repository. GitHub. https://github.com/ultralytics/ultralytics