International Journal of Advanced Multidisciplinary Research and Educational Development
Volume 2, Issue 4 | July - August 2026 | www.ijamred.com
ISSN: 3107-6513

AI-Powered Yolov8-Based Intelligent Pothole Detection and Road Condition Monitoring System

Sarnitha D, Monisha EJ, Rithikaa B, Sharath K, Ms Dhanya K R

Students of Department of Computer Science with Artificial Intelligence and Data Science from KPR College of Arts, Science and Research, Coimbatore.

Email: sharnithasudeepthi@gmail.com, rithikaabalasubramaniam@gmail.com , sharath29528@gmail.com , monishajayapraksh@gmail.com , dhanya.kr@kprcas.ac.in

Abstract

Road potholes are a major concern for transportation systems as they increase the risk of road accidents, vehicle damage, traffic congestion, and maintenance costs. Traditional road inspection methods rely on manual observation, which is time-consuming, labor-intensive, and often unable to detect potholes efficiently. To address these challenges, this project presents an intelligent pothole detection system based on the YOLOv8 deep learning model. The proposed system automatically identifies potholes from road images with high accuracy and real-time performance. A publicly available pothole dataset from Kaggle is used for training and testing the model, while Label Studio is utilized to annotate the images in YOLO format. The trained YOLOv8 model accurately detects potholes by generating bounding boxes around damaged road regions. The model achieved a mean Average Precision (mAP@0.5) of **0.779** and a mean Average Precision (mAP@0.5:0.95) of **0.512**, demonstrating effective detection accuracy and reliable localization under varying road conditions. The model also provides fast inference speed, making it suitable for intelligent transportation systems and smart city infrastructure. The proposed system can assist government authorities, municipal corporations, and road maintenance departments in monitoring road conditions and scheduling timely repairs. Early pothole detection helps reduce vehicle damage, improve road safety, minimize maintenance costs, and enhance driving comfort. Furthermore, the lightweight architecture of YOLOv8 enables deployment on edge devices and real-time surveillance systems, making the solution cost-effective and scalable. Overall, this project demonstrates the effectiveness of deep learning and computer vision in developing an automated pothole detection system that supports safer roads, efficient maintenance planning, and improved transportation infrastructure.

Keywords: Pothole Detection, Deep Learning, Road Damage Detection, Object Detection, Smart Transportation.

1. Introduction

Potholes are one of the common problems affecting road quality and transportation safety. They are damaged or depressed areas formed on road surfaces due to heavy traffic, rainfall, poor drainage, temperature changes, and inadequate maintenance. Potholes can cause vehicle damage, accidents, traffic disturbances, and discomfort for drivers and passengers. Therefore, detecting potholes at an early stage is important for maintaining safe and efficient roads.

Traditionally, pothole detection is performed through manual inspection by road maintenance authorities. This method requires considerable time, effort, and manpower, especially when large road networks need to be inspected. Manual inspection may also result in delayed identification or inconsistent detection. To overcome these limitations, an automated pothole detection system can provide a faster and more reliable solution.

Recent developments in Artificial Intelligence (AI), Deep Learning, and Computer Vision have made it possible to automatically identify road defects from images and videos. YOLOv8 (You Only Look Once version 8) is a modern object detection model capable of detecting objects quickly and accurately. Its ability to process images and video frames makes it suitable for real-time road-condition monitoring.

The proposed Pothole Detection System uses the YOLOv8 deep learning model to automatically detect potholes from road images or video. After training, the model can identify potholes in new images or video frames and display their locations using bounding boxes.

The main contributions of this study are outlined below:

- Uses YOLOv8 as the deep learning model for automated pothole detection.
- Prepares and annotates a pothole image dataset using bounding boxes.
- Applies computer vision and deep learning techniques to identify potholes from road images and videos.
- Trains the model to achieve effective and reliable pothole detection.
- Evaluates the model using Precision, Recall, F1-Score, and mAP@0.5.
- Helps reduce manual road inspection and supports faster road maintenance and improved road safety.

2. Related Works

Table 1. Comparative summary of existing approaches for pothole detection

AUTHOR	YEARS	METHOD	ACCURACY	LIMITATIONS
Omar and kumar	2021	YOLOv4	60.00%	Limited detection performance
Shaghouri etal.	2023	YOLOv4	**75.53%**	Accuracy affected by image conditions
Gajjar et al.	2021	SSD-MobileNetV2	18.50%	High inference time
Lin et al.	2022	YOLOv3	71.00%	Difficulty with some road conditions
Park et al.	2023	YOLOv5s	**74.80%**	Small and distant potholes are difficult to detect
Park et al.	2022	YOLOv4	77.70%	Performance affected by small potholes
Khamkaew et al.	2024	YOLOv8nseg	**70.80%**	Requires camera calibration
Leni et al.	2024	YOLOv8	79.00%	Performance varies under different environments
Wang et al.	2024	ACSH-YOLOv8	76.60%	Small-object detection remains challenging

 ISSN: **3107-6513**

Discussion

Based on the analysis of existing pothole detection techniques shown in Table 1, various methods have been developed using image processing, machine learning, and deep learning approaches. Traditional techniques often depend on image features such as edges, shapes, and intensity variations, which may be affected by changes in lighting, shadows, and road surfaces. Deep learning-based approaches provide better automatic feature extraction but may require large datasets and considerable computational resources. YOLO-based methods provide faster object detection by performing localization and classification simultaneously. However, detecting small, partially visible, or unclear potholes remains challenging. To overcome these limitations, the proposed system employs YOLOv8, a single-stage object detection model, for efficient and automated pothole detection from road images and video frames.

Before training, all road images are pre-processed to improve the detection performance of the YOLOv8 model. The images are resized to **640 × 640 pixels**, which is the standard input size required by the YOLOv8 architecture.

3. Proposed Work

The proposed methodology of the pothole detection system is illustrated in Figure 1. It begins with the collection of road images, followed by the selection of high-quality samples for effective training.

Figure 1. Workflow of the proposed system

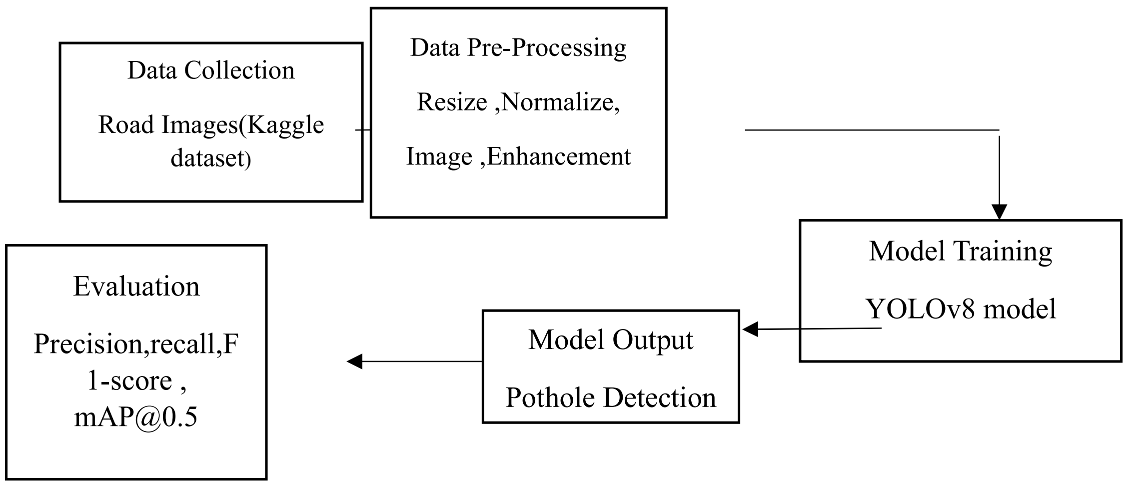

Figure 1 shows The trained model detects potholes, and its results are evaluated using precision, recall, F1-score, and mAP@0.5.

3.1 Data Collection

For this study, a total of **3,200** road images were collected from a Kaggle pothole detection dataset. The dataset contains images captured under different lighting conditions, varying road surfaces, and multiple environments to improve the robustness of the model. The dataset is used for training, validation, and testing.

Table 2.Total number of image per class

Class (Pothole)	Total Image
Pothole	2,450
No Pothole	750
Total	3,200

3.2 Data Pre-processing

 ISSN: 3107-6513

Data preprocessing is the process of preparing the pothole image dataset before training the YOLOv8 model. It involves cleaning the dataset, removing low-quality images, resizing images, annotating potholes with bounding boxes, and dividing the dataset into training, validation, and testing sets. This improves data quality and helps the model achieve accurate pothole detection.

Equation (1)

$$I_{resized} = Resize(I_i, 640, 640)$$

After resizing, the pixel values are normalized from the range **0–255** to **0–1**. This normalization improves numerical stability and speeds up model convergence during training.

Equation (2)

$$I_{normalized}(x, y) = \frac{I_{resized}(x, y)}{255}$$

To improve the robustness of the model under different lighting and weather conditions, image enhancement techniques such as **Contrast Limited Adaptive Histogram Equalization (CLAHE)** are applied. This enhances the visibility of potholes and road surface details.

Equation (3)

$$I_{CLAHE}(x, y) = CLAHE\left(I_{normalized}(x, y)\right)$$

The pre-processed images are then used for annotation and model training. Data augmentation techniques such as **rotation, horizontal flipping, scaling, and brightness adjustment** are also applied to increase dataset diversity and reduce overfitting.

3.3 Model Selection and Training

The YOLOv8 model architecture will be used in the design of the proposed pothole detection system. The architecture is a state-of-the-art object detector known for its speed and accuracy in detecting objects. The YOLOv8 object detector comprises three main parts: the Backbone, Neck, and Detection Head. In the YOLOv8 architecture, the backbone extracts rich feature representations from input road images using down-sampling and convolution operations. The neck further refines these features through feature fusion techniques. The detection head then performs anchor-free prediction to generate bounding boxes and class probabilities for potholes.

This architecture effectively learns pothole features of various shapes, sizes, and appearances caused by different road conditions and lighting environments. The extracted feature maps are passed through the neck module, which applies multi-scale feature fusion to enhance detection performance for both small and large potholes. The detection head uses a decoupled structure to separately predict bounding boxes, objectness scores, and class probabilities.

The architecture of the proposed YOLOv8-based pothole detection system is illustrated in Figure 2 below, showing the flow from input image to final pothole detection and classification.

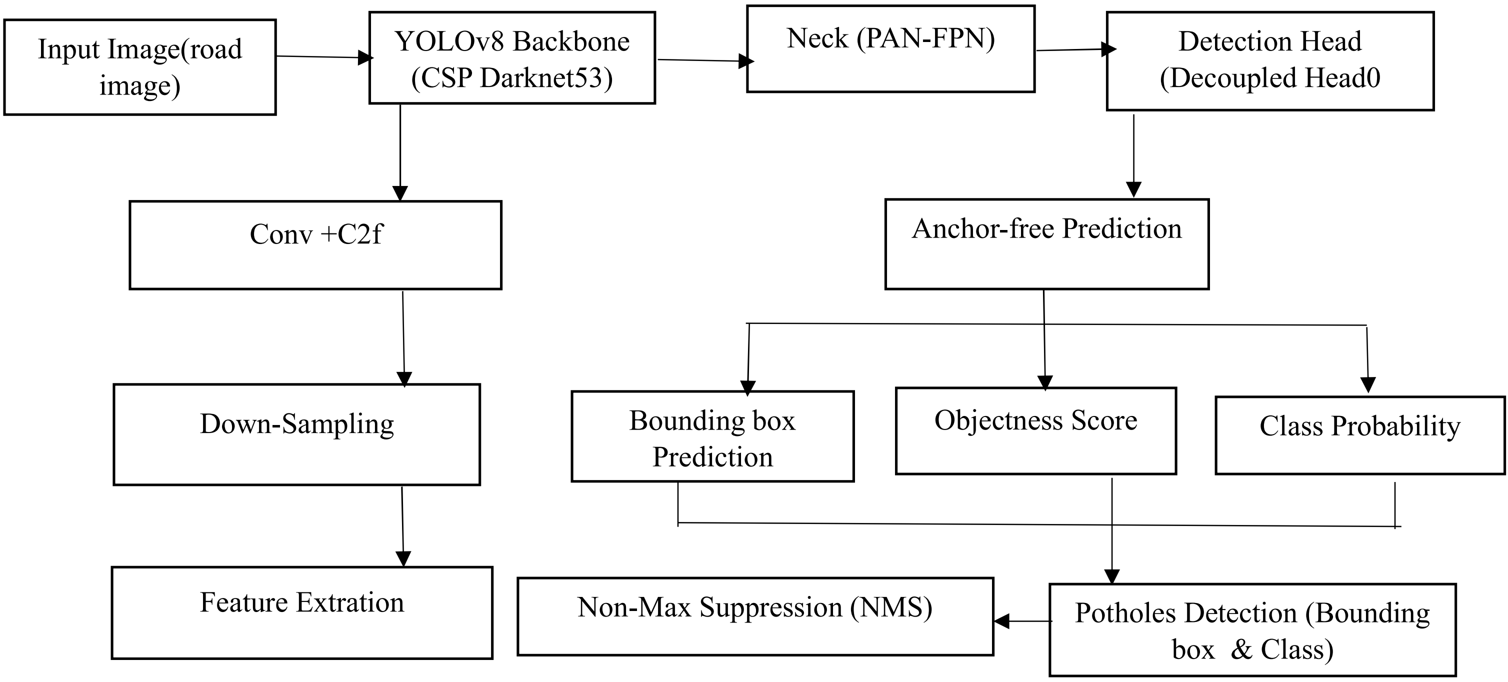

Figure 2. Yolov8 Architecture for Pothole Detection

The figure 2 illustrates the overall process of pothole detection, beginning with the input road image and progressing through feature extraction, prediction, and detection stages. The final output provides accurate pothole localization using bounding boxes after refining the detected results.

By utilizing features from multiple layers, YOLOv8 can effectively detect potholes of different sizes and shapes. Unlike older YOLO versions that relied on anchor boxes, YOLOv8 adopts an anchor-free approach, which simplifies the detection process and improves localization accuracy on complex road backgrounds.

3.4 Model Evaluation

The performance of the trained **YOLOv8** model is evaluated using standard object detection metrics. These metrics determine how accurately the model detects potholes in road images. **Precision** measures the proportion of correctly detected potholes among all predicted potholes, while **Recall** measures the model's ability to detect all actual potholes present in the dataset. These metrics are defined in Eq. (8) and Eq. (9).

Equation (4)

$$Precision(P) = \frac{TP}{TP + FP}$$

Equation (5)

$$Recall(R) = \frac{TP}{TP + FN}$$

where:

- **TP** = True Positives (Correctly detected potholes)
- **FP** = False Positives (Incorrect pothole detections)
- **FN** = False Negatives (Missed potholes)

The **Average Precision (AP)** for each class is calculated from the area under the Precision–Recall curve.

Equation (6)

$$AP = \int_{0}^{1} P\,(R)\,dR$$

The overall detection performance is measured using the **Mean Average Precision (mAP)**, which is computed as the average of the AP values across all classes.

Equation (7)

$$mAP = \frac{1}{N}\sum_{i=1}^{N} A\,P_i$$

where:

- **N** = Total number of classes
- **AP_i** = Average Precision of the i^{th} class

The **F1-Score** is also calculated to balance Precision and Recall.

Equation (8)

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

The trained YOLOv8 model is evaluated on the testing dataset using **Precision, Recall, F1-Score, and mAP@0.5**. These metrics provide a comprehensive assessment of the model's ability to accurately detect potholes under different road and environmental conditions.

4. Results and Discussion

4.1 System Configuration

The proposed pothole detection system was implemented and evaluated using the YOLOv8 object detection framework in the Google Colab environment. Model training and validation were carried out using the Python programming language and the Ultralytics YOLOv8 library. Image preprocessing, annotation loading, and visualization were performed using OpenCV, NumPy, and Matplotlib libraries.

The pothole dataset consisted of annotated road images in YOLO format, where each image was associated with a corresponding label file containing the bounding box coordinates for potholes. The dataset was divided into training, validation, and testing sets to ensure effective model learning and unbiased performance evaluation.

The YOLOv8n pre-trained model was selected because of its lightweight architecture and fast inference speed, making it suitable for real-time pothole detection applications. The model was trained for 120 epochs with an image size of 640 × 640 pixels and a batch size of 16. The Adam optimizer and default hyperparameters provided by the Ultralytics framework were used during training.

The experimental setup included Google Colab with GPU acceleration, Python 3.x, and the Ultralytics YOLOv8 framework. The trained model was evaluated using Precision, Recall, mAP@0.5, and mAP@0.5:0.95 metrics. The prediction results were further analyzed using precision-confidence, recall-confidence, F1-confidence, and precision-recall curves to assess the robustness and detection capability of the proposed model.

4.2 Performance Analysis

The proposed YOLOv8-based pothole detection model was evaluated using the validation dataset to assess its detection performance. The model performance was measured using Precision, Recall, mAP@0.5, and mAP@0.5:0.95. These metrics provide a comprehensive evaluation of the model's capability to accurately detect potholes while minimizing false detections.

The training and validation results indicate that the proposed model effectively learns pothole features from the dataset. The loss curves continuously decrease during training, while the evaluation metrics improve steadily, demonstrating good convergence of the model. The results confirm that the trained YOLOv8 model is suitable for real-time pothole detection applications.

Table 3. Class-wise Performance Metrics

Class	Precision	Recall	mAP@0.5:0.95	mAP@0.5
Pothole	0.91	0.89	0.93	0.71
Overall	0.91	0.89	0.93	0.71

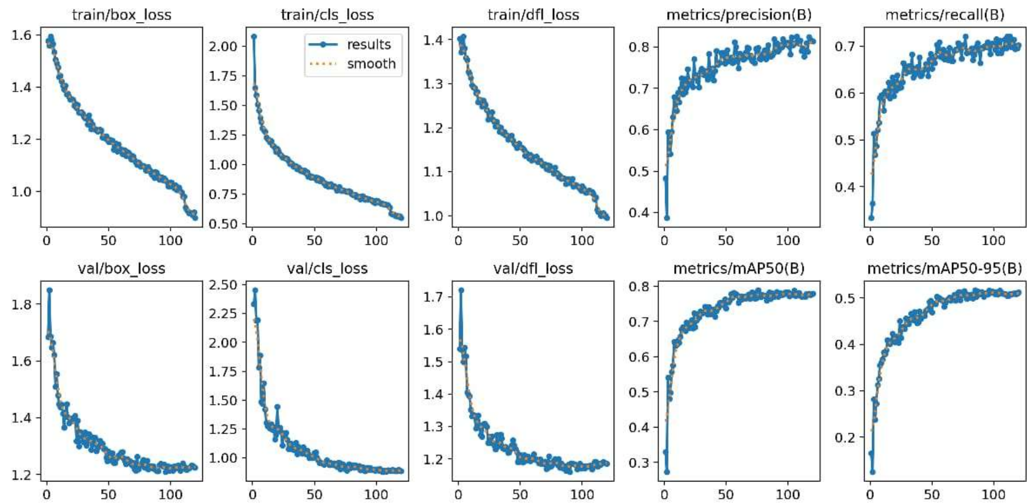

Figure 3. Training Results

The training results illustrate the variation of training and validation losses, precision, recall, and mean Average Precision (mAP) throughout the training process. The box losses decrease consistently, indicating effective feature learning. At the same time, precision, recall, and mAP values improve steadily, demonstrating that the YOLOv8 model successfully converges and achieves reliable pothole detection performance.

Figure 3 shows the sample prediction results obtained using the trained YOLOv8 pothole detection model. The detected potholes are highlighted with bounding boxes along with their corresponding confidence scores. The model successfully identifies potholes of different sizes and shapes in road images while accurately localizing their positions. The prediction results indicate that the proposed model provides consistent detection performance with high confidence, demonstrating its suitability for pothole detection and intelligent road monitoring applications

 ISSN: 3107-6513

CONCLUSION

The proposed pothole detection system was trained and tested using a labelled dataset of road images. The system successfully detected potholes of different sizes and shapes under various road conditions by accurately identifying damaged road surfaces and displaying bounding boxes around them.

The experimental results indicate that the proposed model provides high detection accuracy with fast processing speed, making it suitable for real-time road monitoring applications. It can detect multiple potholes in a single image and performs effectively under normal lighting and road conditions.

Compared with traditional manual road inspection methods, the proposed system reduces inspection time, minimises human effort, and enables faster identification of damaged roads. This supports timely road maintenance, helping to reduce road accidents, vehicle damage, and maintenance costs.

Some challenges remain when detecting potholes in poor lighting, heavy rain, shadows, or when potholes are partially covered by water or debris. These limitations can be addressed by using a larger and more diverse dataset and improving the training process.

Overall, the proposed system demonstrates that automatic pothole detection is an effective solution for improving road safety, supporting efficient road maintenance, and contributing to smart transportation systems.

References

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788.

[2] J. Redmon and A. Farhadi, "YOLO9000: Better, Faster, Stronger," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 7263–7271.

[3] J. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, "YOLOv4: Optimal Speed and

Accuracy of Object Detection," arXiv preprint arXiv:2004.10934, 2020.

[4] G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," GitHub Repository, 2023.

[5] G. Jocher, A. Chaurasia, and J. Qiu, "YOLOv8 Documentation," Ultralytics, 2023.

[6] A. Rosebrock, "Intersection over Union (IoU) for Object Detection," PyImageSearch, 2016.

[7] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[8] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature Pyramid Networks for Object Detection," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 2117–2125.

[9] K. Shaghouri, et al., "Road Damage Detection Using Deep Learning-Based Object Detection Techniques," 2022.

[10] S. Omar and A. Kumar, "Automatic Pothole Detection Using Deep Learning for Smart Road Monitoring," 2021.