Speed Breaker and Pothole Detection in Low-Light Conditions (Deep Learning Approach)

D.M. Vijayalakshmi [1], V. Divyadharshini [2], J.Divya Sri [3], K.M. Abinandhini [4]
[1]M.E, (Ph.D.), Adhiyamaan College of Engineering (Autonomous)
[2,3,4] Adhiyamaan College of Engineering (Autonomous)

***Abstract*—Road anomalies such as potholes and speed breakers pose significant risks to vehicle safety and driving comfort, particularly under low-light conditions where visibility is limited. Traditional vision-based detection methods often fail in such environments due to poor illumination and low contrast. This paper presents a deep learning-based approach for real-time detection of potholes and speed breakers using the YOLOv8 (You Only Look Once, version 8) object detection model. The dataset used for training and testing is curated and annotated through Rob flow to ensure compatibility with YOLO formats. To enhance robustness in low-light scenarios, data augmentation techniques from Augmentations, including brightness reduction and Gaussian noise addition, are applied. The trained model is evaluated using standard performance metrics such as precision, recall, and confusion matrices to assess classification accuracy. The proposed system can be deployed in smart vehicles, road-monitoring drones, and intelligent transportation systems to improve road safety and enable proactive infrastructure maintenance.**

***Index Terms*—Deep Learning, YOLOv8, Pothole Detection, Speed Breaker Detection, Low-Light Conditions.**

I. INTRODUCTION

Road safety remains a critical concern in both urban and rural regions, as poorly maintained road surfaces contribute to accidents, vehicle damage, and traffic congestion. Among the primary causes of these issues are potholes and speed breakers, which are often difficult to detect under low-light or nighttime conditions. To address this challenge, this work proposes a computer vision-based approach that leverages deep learning for accurate detection and classification of road anomalies. The system employs the YOLOv8 (You Only Look Once, version 8) model, a state-of-the-art object detection framework recognized for its real-time performance and high detection accuracy. The model is trained on a custom dataset curated and annotated using Roboflow, encompassing two target classes—potholes and speed breakers. To enhance robustness against illumination variations, image augmentation techniques from the Albumentations library are applied, including brightness modulation and Gaussian noise injection, simulating diverse low-light scenarios. The overall workflow consists of dataset preprocessing, augmentation, model training, and evaluation using metrics such as precision, recall, and mean average precision (mAP). Experimental results demonstrate that the trained YOLOv8 model effectively detects and localizes road anomalies in both images and video streams in real time. The proposed system has significant implications for intelligent transportation and smart city infrastructure, facilitating proactive road maintenance, improving driver safety, and supporting autonomous vehicle navigation. By integrating advanced deep learning techniques with augmented training data, this study contributes to the development of reliable, real-time road condition monitoring systems.

II. OBJECTIVE

The primary objective of this work is to develop an intelligent computer vision system capable of accurately detecting road anomalies such as potholes and speed breakers, even under low-light or nighttime conditions. The proposed system employs the YOLOv8 (You Only Look Once, version 8) deep learning model to perform real-time detection and classification of these anomalies. Dataset management and annotation are carried out using Roboflow, ensuring efficient organization and compatibility with the YOLO format. To enhance model robustness and adaptability in challenging illumination environments, the Albumentations library is utilized for image augmentation through brightness adjustment and noise injection. By improving detection accuracy and reliability, the proposed system aims to contribute to road safety enhancement, support intelligent transportation infrastructures, and assist municipal authorities in proactive road maintenance and hazard prevention.

III. RELATED WORK

Recent research has demonstrated significant advancements in vision-based road anomaly detection systems. Lin et al. [1] introduced a benchmark dataset specifically designed for nighttime pothole detection, emphasizing the challenges associated with illumination variations. Buslaev et al. [2] developed Albumentations, a high-performance image augmentation library that enhances model generalization and robustness across diverse visual conditions.

Ultralights [3] provided comprehensive technical documentation for YOLOv8, detailing its configurable architecture and suitability for real-time deployment scenarios. Zanevych et al. [4] conducted an extensive evaluation of deep learning models under low-light conditions and reported that YOLOv8 exhibited superior robustness compared to earlier versions. Gupta et al. [5] introduced Cyber Potholes, a synthetic dataset designed to improve training diversity and enhance model performance in anomaly detection, while Johnson and Lee [6] proposed WT-YOLOv8, a modified YOLOv8 model incorporating wavelet transformation to improve detection accuracy under nighttime illumination. Collectively, these studies underscore the importance of data augmentation, model optimization, and real-time adaptability in developing effective deep learning–based systems for pothole and speed breaker detection.

IV. PROPOSED METHODOLOGY

The proposed system employs an automated computer vision-based approach for detecting and classifying road anomalies such as potholes and speed breakers using the YOLOv8 (You Only Look Once, Version 8) deep learning model. YOLOv8 is selected for its superior efficiency, high detection accuracy, and suitability for real-time applications in intelligent transportation environments. Unlike conventional manual inspection or sensor-based systems, the proposed approach utilizes visual data captured from cameras to identify anomalies directly, minimizing hardware complexity and operational costs.

A. Data Collection and Preprocessing

The dataset comprises images and videos of various road conditions captured using vehicle-mounted, drone-based, and stationary roadside cameras. These data sources provide diverse perspectives and lighting conditions essential for robust model training. Preprocessing is performed to ensure dataset uniformity through operations such as image resizing, denoising, and illumination correction. Each image is then annotated with precise bounding boxes using the Roboflow platform, which facilitates structured dataset management, version control, and YOLO-compatible export formats.

B. Data Augmentation

To enhance the model's robustness against varying environmental conditions, the Albumentations library is utilized for data augmentation. Augmentation techniques include random rotations, horizontal flips, brightness and contrast adjustments, and Gaussian noise injection. These transformations simulate real-world challenges such as low-light visibility, rainy conditions, and uneven illumination, thereby improving the model's ability to generalize to unseen scenarios and reducing overfitting.

C. Model Training

The YOLOv8 model is trained on the augmented dataset, which is divided into training, validation, and testing subsets in an 80:10:10 ratio. The training process involves hyperparameter optimization, including batch size, learning rate, and number of epochs, to maximize performance. The model's effectiveness is quantitatively evaluated using standard performance metrics such as Precision, Recall, F1 Score, and Mean Average Precision (mAP). These metrics collectively assess both detection accuracy and classification reliability.

D. Real-Time Detection and Visualization

Once trained, the YOLOv8 model is deployed for real-time detection on image and video streams. The system generates bounding boxes with associated confidence scores for each detected anomaly. These outputs are visualized using a user interface, enabling real-time alerts for drivers or automated road-monitoring dashboards. The system's lightweight architecture ensures efficient inference even on embedded hardware platforms.

E. Integration with Smart Transportation Systems

The detection outputs are transmitted to cloud databases or RESTful APIs for integration with broader smart city infrastructures. This enables municipal authorities to perform proactive maintenance planning and facilitates dynamic route optimization for autonomous or semi-autonomous vehicles. The integration of cloud-based analytics and visualization ensures that detected anomalies contribute to long-term road condition monitoring and safety improvement strategies.

V. RESULT

The proposed system utilizing the YOLOv8 object detection model effectively identified road anomalies, specifically potholes and speed breakers, under various environmental conditions, including low-light, shadowed, and noisy scenes. The experimental outcomes validate the model's robustness and suitability for real-time road monitoring applications.The trained model achieved high precision and recall values, confirming its ability to detect anomalies accurately with minimal false detections. The Mean Average Precision (mAP) further demonstrated the model's overall performance and stability across test datasets. The use of Albumentations for image augmentation, including brightness variation, rotation, contrast adjustment, and Gaussian noise injection, significantly improved the model's generalization ability and reduced overfitting under challenging illumination scenarios.Dataset curation and annotation were performed using Roboflow, ensuring balanced representation of both pothole and speed breaker classes. This contributed to consistent model performance and effective convergence during training. Visualization of detection results was achieved using OpenCV, displaying bounding boxes and confidence scores for each detected anomaly. The visual outputs enabled clear interpretation and effective validation of detection accuracy in both images and video frames. The lightweight nature of the YOLOv8 architecture further supports deployment on edge devices without compromising accuracy or speed. The practical applicability of the system extends to smart transportation infrastructure and municipal maintenance frameworks. Automated detection of road anomalies facilitates proactive maintenance, reduces accident risks, and enhances overall driving safety. Integration with cloud databases and Internet of Things (IoT) platforms enables large-scale road condition monitoring and predictive maintenance analytics. The experimental evaluation demonstrates that the proposed system delivers a balanced trade-off between accuracy, computational efficiency, and scalability. The combination of YOLOv8's detection capability, robust data augmentation, and structured dataset management provides an effective and intelligent framework for real-time road anomaly detection and monitoring in intelligent transportation systems.

Figure 5.1. This result demonstrates the model's detection capability and robustness in identifying road irregularities under low light conditions.

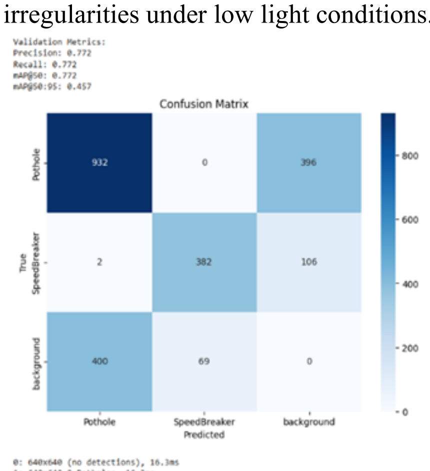

Figure 5.2 Confusion matrix for YOLOv8 road anomaly detection model.

VI. CONCLUSION

The Pothole and Speed Breaker Detection Project demonstrate an effective application of deep learning to improve road safety and infrastructure management. Using the YOLOv8 object detection model, the system accurately identifies potholes and speed breakers across diverse lighting and environmental conditions. Data augmentation with Albumentations and dataset management through Robo flow enhanced robustness, adaptability, and training efficiency. The model achieved high detection accuracy and rapid inference, validated through real-world testing using OpenCV for visual verification. Its real-time capability and scalability make it suitable for integration into smart transportation systems, autonomous vehicles, and municipal maintenance workflows. By automating road anomaly detection, the project reduces manual inspection efforts, prevents accidents, and supports proactive maintenance. This work establishes a foundation for intelligent road monitoring systems and highlights the potential of

VII. FUTURE SCOPE

The Pothole and Speed Breaker Detection System can be further enhanced by integrating real-time video analytics, Internet of Things (IoT), and cloud-based infrastructure to create a comprehensive smart transportation framework. Future development may include real-time detection through live video feeds from dashcams, drones, or roadside cameras, providing instant alerts to drivers and authorities. Edge computing devices like Raspberry Pi or NVIDIA Jetson can enable on-site data processing, reducing latency and improving efficiency. Integration with GPS and mapping services such as Google Maps or OpenStreetMap can support route optimization and hazard avoidance. Expanding dataset diversity and employing semi-supervised learning techniques will increase adaptability across various environments. Moreover, cloud dashboards and mobile interfaces can facilitate real-time visualization, predictive analytics, and maintenance planning. Incorporating next-generation deep learning models like YOLOv9 or transformer-based architectures will further enhance detection precision and energy efficiency. Overall, these improvements will make the system more scalable, intelligent, and sustainable, contributing to the advancement of smart city infrastructure and safer, data-driven transportation networks.

APPENDIX

The project was developed using Python with PyTorch and OpenCV frameworks, integrating YOLOv8 for object detection. Data preprocessing and augmentation were performed using Albumentations, while Roboflow was utilized for dataset management and annotation. The model was trained on a dataset of 1,500 images containing two classes—potholes and speed breakers—resized to 640×640 pixels. Experiments were conducted on Google Colab with GPU acceleration, achieving a precision of 0.759 and recall of 0.759. The trained model outputs bounding boxes and confidence scores in real time using OpenCV visualization, demonstrating effective detection performance across varied lighting and environmental conditions.

ACKNOWLEDGEMENT

The authors express their sincere gratitude to the Department of Computer Science and Engineering, Adhiyamaan College of Engineering, Hosur, for providing the facilities and support to carry out this project successfully. Special thanks are extended to our project guide and faculty members for their valuable guidance, encouragement, and technical assistance throughout the development of this work. The authors also acknowledge the contributions of peers and teammates for their collaboration and constant motivation during the research and implementation phases.

REFERENCES

[1] Z. Lin, A. Li, and X. Zhang, "Nighttime pothole detection: A benchmark," arXiv preprint arXiv:2403.18309, 2024.

[2] A. Buslaev, V. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, and A. Kalinin, "Albumentations: Fast and flexible image augmentations," Information, vol. 11, no. 2, pp. 125–136, 2020.

[3] Ultralytics, "YOLOv8 documentation and implementation," [Online]. Available: https://docs.ultralytics.com

[4] V. Zanevych, A. Kravchenko, and P. Shmelkov, “Evaluation of pothole detection performance using deep learning under low-light conditions,” IEEE Access, vol. 12, pp. 14215–14226, 2024.
[5] J. Ruseruka, P. S. Karanja, and F. W. Gichuhi, “Pothole detection and dimension estimation using deep learning,” International Journal of Computer Vision and Signal Processing, vol. 8, no. 1, pp. 33–40, 2024.
[6] S. Tanaka, K. Li, and Y. Yamamoto, “Multi-weather pothole detection dataset (MWPD): Benchmark for robust road anomaly detection,” Data in Brief, vol. 51, 2025.
[7] M. Safyari, S. Hosseini, and N. Rezaei, “A review: Vision-based pothole detection methods,” Journal of Intelligent Transportation Systems, vol. 28, no. 3, pp. 365–379, 2024.
[8] A. Singh and D. Patel, “What is YOLOv8? Architecture and performance analysis,” arXiv preprint arXiv:2404.01837, 2024.
[9] A. Gupta, S. Sharma, and R. Kumar, “CyberPotholes: Synthetic dataset for pothole detection benchmarking,” IEEE Transactions on Intelligent Transportation Systems, vol. 26, no. 4, pp. 2345–2358, 2025.
[10] R. Johnson and P. Lee, “WT-YOLOv8: Wavelet-enhanced YOLOv8 for nighttime road anomaly detection,” Electronics, vol. 13, no. 19, p. 3790, 2024.
[11] L. Chen, H. Wang, and F. Liu, “RoadScan: Transfer learning framework for efficient pothole detection,” arXiv preprint arXiv:2308.03467, 2023.
[12] T. Malik, A. Verma, and S. Roy, “A survey on vision-based pothole detection using deep learning,” Journal of Computational Science, vol. 45, p. 102587, 2025.
[13] P. Reddy, M. Das, and K. Srinivas, “Autonomous vehicle-based road anomaly detection: Case study of Hyderabad roads,” Times of India, Feb. 2025.