International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 2, Mar-Apr 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

Review on Deep Learning-Based Road Hazard Detection Using YOLOV8

Khadeeja[1], G Mohammed Kaif[1], Janakiram R S[1], and B Raghuram[1]

[1]*Ballari Institute of Technology and Management, Ballari, Karnataka, India*

khadeejahameed.1705@gmail.com, mohammedkaifrehman24@gmail.com, janakiram1715@gmail.com, raghuram2949@gmail.com

Abstract— Road hazards such as potholes and speed humps pose significant challenges to transportation safety, vehicle integrity, and infrastructure maintenance. Traditional manual inspection methods are inefficient and spatially inconsistent, necessitating automated solutions. This literature survey examines recent advances in deep learning-based road hazard detection systems, with particular emphasis on YOLOv8-based real-time detection frameworks, severity assessment methodologies, and practical deployment architectures. The review encompasses classification approaches using convolutional neural networks, real-time localization using YOLO family detectors, segmentation-based characterization methods, multimodal RGB-depth fusion techniques, and edge-cloud deployment strategies. Key findings indicate that YOLOv8 has emerged as the dominant architecture for real-time pothole detection, achieving precision rates exceeding 95% while maintaining inference speeds suitable for in-vehicle deployment. The survey identifies critical research gaps in standardized severity assessment, dataset availability, and deployment optimization for resource-constrained environments.

Keywords— YOLOv8, Road Hazard Detection, Deep Learning, Pothole Detection

I. INTRODUCTION

Road surface defects, particularly potholes and speed humps, represent persistent challenges to transportation infrastructure worldwide. These hazards contribute to vehicle damage, traffic disruption, maintenance inefficiency, and road accidents. Conventional inspection practices rely on manual surveys or complaint-driven workflows, which are laborious, delayed, and spatially inconsistent. Consequently, recent research has increasingly approached road hazard detection as a computer vision problem, with deep learning emerging as the dominant methodological paradigm. The evolution of road hazard detection research spans four interconnected objectives: presence detection, localization, severity characterization, and operational deployment. Early work focused primarily on binary classification—determining whether a pothole exists in an image. Recent advances address more complex questions: precise localization within frames, quantitative severity assessment, real-time processing capabilities, and integration with intelligent transportation systems. The literature reveals a clear progression from handcrafted feature engineering to end-to-end deep learning, from offline batch processing to real-time video analysis, and from pure detection to comprehensive hazard characterization. This survey examines the current state of deep learning-based road hazard detection with specific focus on methods and architectures relevant to real-time, deployment-ready systems. We analyze classification-based approaches, YOLOv8 detection frameworks, segmentation methods for precise characterization, multimodal sensing techniques, and lightweight architectures for edge deployment. The review is structured to support researchers and practitioners developing practical road safety applications

II. FROM CLASSICAL VISION TO DEEP LEARNING

Traditional computer vision approaches to pothole detection relied on edge detection, texture descriptors, grayscale transformations, Hough transforms, and handcrafted classifiers such as Support Vector Machines and decision trees. Chandana B. D. [1] notes that while these methods offered computational simplicity, they exhibited high sensitivity to illumination changes and background noise, reducing robustness in realistic road environments. The broader survey by Ahmad et al. [4] demonstrates that classical image processing and shallow machine learning approaches produced acceptable but unstable performance across varied surfaces, shadows, and pavement textures. Alternative sensing paradigms using vibration-based and accelerometer systems offered low-cost deployment but could only infer road irregularity indirectly. These methods often confused potholes with bumps, uneven surfaces, or other transient disturbances. Baroudi et al. [3] observe that non-visual sensing systems may detect anomalies but do not provide rich characterization such as precise area, boundary geometry, or depth.

This limitation becomes critical when municipal authorities require actionable information for repair prioritization. The transition to deep learning, particularly Convolutional Neural Networks (CNNs) and the YOLO family of detectors, addressed many deficiencies of classical methods by enabling automated feature extraction from raw visual data. CNNs reduced dependence on handcrafted descriptors, while one-stage detectors such as YOLO dramatically improved inference speed, making real-time roadside or in-vehicle deployment feasible.

III. CNN-BASED IMAGE CLASSIFICATION APPROACHES

Several studies treat pothole detection as an image classification problem. Chandana B. D. [1] presents a CNNbased system trained on curated road images labeled as pothole and non-pothole. The model produces binary predictions and maps confidence scores to qualitative severity labels (Minor, Moderate, Major). This approach demonstrates an accessible end-to-end pipeline extending beyond pure detection to include web deployment, report generation, and demonstration-level location mapping. However, confidence-based severity mapping serves only as a proxy for actual physical severity and cannot reliably estimate dimensions. Ahmad et al. [4] expand the classification perspective by investigating size-based pothole categorization using transfer learning with ResNet18, ResNet50, and MobileNetV2. Their study classifies pavement images into Normal, Small Pothole, and Large Pothole categories, with size determined using surface area relative to tire contact area. A valuable contribution of this work is analysis of image acquisition height. Images captured from approximately 3.5 feet yielded markedly higher classification accuracy than those captured from 2 feet, suggesting that viewpoint geometry, field of view, and scene context significantly influence learned representations. MobileNetV2 emerged as the strongest model, reaching 98in three-class categorization under favorable capture conditions. While classification-based studies demonstrate that CNNs can perform well when tasks are constrained and datasets well-curated, they expose a structural limitation: classification does not inherently localize potholes within images and is therefore less suited for multi-pothole scenes, road-wide scanning, or quantitative maintenance planning.

IV. YOLOV8 AND REAL-TIME LOCALIZATION

The most pronounced trend in recent literature is widespread adoption of YOLOv8 for pothole localization. Both Reddy et al. [7] and Vinoth et al. [8] describe YOLOv8-based real-time systems processing images and video streams to detect potholes with bounding boxes and confidence scores. These studies emphasize YOLOv8's anchor-free detection, lightweight design, enhanced backbone, and low-latency inference, making it particularly appropriate for dynamic road environments. Reddy et al. [7] report strong metrics for a custom pothole dataset, including precision of 95.2and mAP@0.5 of 96.1performance in the range of 15–20 FPS for webcam input. Similarly, Vinoth et al. [8] report average detection accuracy of 94.2While these works are not as methodologically elaborate as some segmentation or multimodal studies, they confirm that YOLOv8 has become a practical baseline for real-time pothole detection. Bhavana et al. [2] advance this work through POT-YOLO, an edge segmentation-based YOLOv8 framework enhanced by preprocessing with Contrast Stretching Adaptive Gaussian Star Filtering. Their architecture incorporates MBConv and E-SPPF modules to improve feature extraction and multiscale fusion. The reported classification accuracy of 99.10refinement can substantially improve performance, particularly under noisy and varied pavement conditions. This study broadens the anomaly concept beyond potholes to include cracks, oil stains, patches, and pebbles, reflecting the complexity of real roads.

V. SEGMENTATION FOR FINE-GRAINED CHARACTERIZATION

While bounding-box detectors identify candidate pothole regions effectively, segmentation-based methods provide more precise delineation of defect boundaries. This distinction becomes crucial when objectives shift from detection alone to quantification of area, damage percentage, or maintenance severity. Baroudi et al. [3] make one of the most substantial contributions by using a pre-trained YOLOv8-seg model for pothole detection and segmentation, followed by depth-informed characterization. Their study introduces a novel dataset comprising RGB images and corresponding depth maps collected from Saudi roads. The YOLOv8nseg model achieves high performance for both box and mask predictions, with mAP50 values near 0.96. More significantly, segmented masks enable direct calculation of pothole area and percentage of road damage. The study then merges segmented pothole regions with depth maps to estimate relative pothole depth, moving from simple visual recognition toward comprehensive hazard characterization. This work illustrates a major conceptual shift: pothole systems are no longer judged solely by detection accuracy but by their ability to support pavement condition assessment. By combining area and relative depth, Baroudi et al. offer a richer basis for severity estimation than classification confidence or box size alone. Their findings also underline a persistent research gap: publicly available pothole datasets rarely include trustworthy depth ground truth.

VI. RGB-DEPTH AND POINT CLOUD FUSION

The fusion of 2D vision with 3D geometry represents another major research direction. Zhong et al. [5] propose a method integrating YOLOv8 detection with point cloud analysis from an RGB-D depth camera. Their system first identifies candidate pothole regions using YOLOv8, then extracts corresponding 3D point clouds, analyzes surface smoothness, identifies boundary contours, and filters false positives through elevation thresholds. This approach directly addresses a well-known limitation of image-only detection: stains, patches, shadows, or surface discolorations may visually resemble potholes. The multimodal framework improves precision from 89.3recall at 93.3multimodal systems can simultaneously improve detection reliability and geometric measurement accuracy. The study emphasizes that false positive reduction is central to practical deployment, especially in well-maintained roads where visual anomalies are subtle. Compared with the segmentation-plus-depth approach of Baroudi et al., Zhong et al. provide a more explicit geometric pipeline grounded in point cloud processing. Together, these studies signal the field's gradual movement from detection to measurement.

VII. DEPLOYMENT, LIGHTWEIGHT DESIGN, AND EDGE-CLOUD INTELLIGENCE

A recurring concern in the literature is the tension between model accuracy and deployment feasibility. Highperforming models often remain computationally expensive, limiting real-world applicability on embedded devices or moving vehicles. Liu et al. [6] directly address this issue by proposing a real-time pavement distress detection system based on edge-cloud collaborative computing. Their YOLO-LFE architecture replaces heavier backbones with MobileNetV3 and introduces ESPP and improved feature fusion mechanisms. Relative to YOLOv8, the model reduces parameters by 32.5This work is particularly important because it reframes pothole detection as a systems engineering problem rather than a pure vision benchmark. By combining edge inference, cloud storage, and ByteTrack-based tracking, the study moves toward continuous, scalable road monitoring. It reflects an emerging consensus that future pothole systems must be lightweight, network-aware, and operationally integrated. Chandana B. D. [1], Reddy et al. [7], and Vinoth et al. [8] also contribute to deployment-oriented literature through Flask dashboards, real-time uploads, PDF reporting, and web interfaces. Although these are simpler implementations, they demonstrate demand for complete application pipelines rather than isolated model evaluations.

VIII. RESEARCH GAPS AND CHALLENGES

Despite strong progress, several unresolved issues recur throughout the literature. Dataset limitations remain severe, particularly for depth information, segmentation masks, and diverse environmental conditions. Generalization across lighting, weather, road texture, and camera perspective remains difficult. Severity assessment lacks standardization; some studies use confidence scores, others surface area, relative depth, or heuristic thresholds. False positives from patches, stains, shadows, and non-road objects continue to challenge image-only models. Deployment constraints on embedded devices necessitate continued research on lightweight architectures and edge-cloud collaboration. A further challenge is the lack of unified benchmarking. Reported metrics vary widely across studies, datasets, and task definitions, making direct comparison difficult. Some studies report classification accuracy, others mAP, and still others geometric error in perimeter or depth. The field would benefit from standardized evaluation protocols spanning detection, segmentation, quantification, and runtime performance.

IX. CONCLUSION

The literature on road hazard detection reveals a rapidly maturing field shaped by advances in deep learning, realtime object detection, segmentation, and multimodal sensing. CNN-based classifiers established the feasibility of automated pothole recognition, but YOLOv8-based detectors now dominate due to their balance of speed and accuracy. Segmentation models have extended the task from detection to delineation, while RGB-depth and point cloud fusion methods have enabled more credible quantification of pothole area and depth. Simultaneously, edgecloud systems and lightweight architectures are pushing these methods toward deployable smart-city applications. Taken together, the reviewed studies suggest that the future of road hazard detection lies in integrated systems combining fast visual detection, precise segmentation, geometric characterization, GPS-aware mapping, and scalable deployment infrastructure. The key research frontier is no longer simply whether a pothole can be detected, but whether it can be reliably localized, physically characterized, prioritized, and communicated in real time for road maintenance and transportation safety

REFERENCES

[1] C. B. D., ”Deep learning-based pothole detection and severity classification with location mapping,” International Advanced Research Journal in Science, Engineering and Technology, vol. 12, no. 8, 2025, doi: 10.17148/IARJSET.2025.12829. [2] N. Bhavana, M. M. Kodabagi, B. Muthu Kumar, P. Ajay, N. Muthukumaran, and A. Ahilan, ”POT-YOLO: Real-time road potholes detection using edge segmentation-based YOLO v8 network,” IEEE Sensors Journal, vol. 24, no. 15, 2024, doi: 10.1109/JSEN.2024.3399008. [3] U. Baroudi, A. BaHamid, Y. Elalfy, and Z. A. Alami, ”Enhancing pothole detection and characterization: Integrated segmentation and depth estimation in road anomaly systems,” 2025. [4] C. F. Ahmad, A. T. Al-Sayegh, A. Cheema, W. Qayyum, R. Ehtisham, S. Saghir, and A. Ahmad, ”Classification of different size of potholes based on surface area using convolutional neural network,” Discover Applied Sciences, vol. 6, p. 492, 2024, doi: 10.1007/s42452-024-06207-3. [5] J. Zhong, D. Kong, Y. Wei, and B. Pan, ”YOLOv8 and point cloud fusion for enhanced road pothole detection and quantification,” Scientific Reports, vol. 15, p. 11260, 2025, doi: 10.1038/s41598-025-94993-0. [6] Y. Liu, F. Liu, Y. Huang, J. Hu, W. Zhang, and Y. Hou, ”The real-time pavement distress detection system based on edge-cloud collaborative computing,” IEEE Transactions on Intelligent Transportation Systems, vol. 26, no. 7, pp. 10512–10522, Jul. 2025, doi: 10.1109/TITS.2025.3544240. [7] C. S. Reddy, M. Anitha,

and G. Sravani, ”Road pothole detection using deep learning with YOLOv8 model,” 2025. [8] V. M., V. K., S. R., and S. G., ”Deep learning approach for road pothole detection using YOLOv8,” International Journal of Research Publication and Reviews, vol. 6, no. 11, pp. 7332–7336, 2025.

