ISSN: 2582-7219

International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Impact Factor: 8.206

Volume 9, Issue 4, April 2026

www.ijmrset.com ijmrset@gmail.com +91 6381 907 438 +91 63819 07438

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Road pothole detection and automated reporting system using YOLOv8 and IoT integration

Hafil A[1], Mohamed Arfath M[2], Mohamed Mufrih A[3], and Kiruthiga R[4]

Student, Department of Artificial Intelligence and Data Science, Aalim Muhammed Salegh College of Engineering, Avadi, Chennai, Tamil Nadu, India[1]

Student, Department of Artificial Intelligence and Data Science, Aalim Muhammed Salegh College of Engineering, Avadi, Chennai, Tamil Nadu, India[2]

Student, Department of Artificial Intelligence and Data Science, Aalim Muhammed Salegh College of Engineering, Avadi, Chennai, Tamil Nadu, India[3]

Assistant Professor, Department of Artificial Intelligence and Data Science, Aalim Muhammed Salegh College of Engineering, Chennai, Tamil Nadu, India[4]

ABSTRACT: Road infrastructure degradation, particularly pothole formation, poses significant safety hazards and economic burdens globally. Traditional manual inspection methods are slow, infrequent, and labor-intensive, leading to delayed maintenance and increased accident rates. This paper presents an automated road pothole detection and reporting system leveraging the YOLOv8 deep learning architecture for real-time visual detection, integrated with Python-based processing pipelines, OpenCV for image preprocessing, and IoT sensor networks for location-aware data collection. Vehicle-mounted cameras and sensors capture road surface data, processed by YOLOv8 to detect and classify potholes by severity. Detected events are geo-tagged and uploaded to a centralized pothole database accessible via a real-time monitoring dashboard. Maintenance teams can visualize pothole maps, priorities repairs, and dispatch field crews with GPS- guided routing. Post-repair updates are recorded and exposed through an OpenAI interface for integration with smart city systems. Experimental evaluation demonstrates high detection accuracy (mAP@0.5 of 87.3%) with real-time inference, making the system suitable for urban deployment.

KEYWORDS: Pothole Detection, Road Infrastructure Monitoring, YOLOv8, Deep Learning, Computer Vision, OpenCV, IoT Sensors, Smart City Systems, Real-Time Detection, Geo-Tagging, GPS Routing, Image Processing

I. INTRODUCTION

Road potholes are one of the most pervasive hazards in modern urban and rural transport infrastructure. They arise due to repeated traffic loading, material fatigue, water infiltration, and inadequate maintenance cycles. The consequences are multifaceted: increased vehicular damage, elevated risk of road accidents including fatalities, disruption to traffic flow, and long-term degradation of road networks [1]. In developing and developed nations alike, the economic cost of pothole-related vehicle damage alone runs into billions of dollars annually.

Conventional approaches to pothole identification rely on scheduled manual road surveys conducted by trained inspectors. These methods are inherently reactive: surveys occur at fixed intervals, meaning newly formed or rapidly worsening potholes may go undetected for extended periods. The labor and logistics costs of manual inspection are substantial, and the resulting data is often fragmented and difficult to integrate into digital maintenance management systems [2].

The rapid maturation of computer vision, deep learning, and Internet of Things (IoT) technologies presents an unprecedented opportunity to automate and continuously improve road condition monitoring. Object detection models, particularly the YOLO (You Only Look Once) family, have demonstrated remarkable capability in detecting small, irregular objects in real-time video streams with high accuracy and efficiency [3]. When combined with geospatial tagging through GPS modules and connected to cloud-based reporting infrastructure, such systems can provide city administrators with live road condition data, enabling proactive and targeted maintenance.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

This paper presents the design, implementation, and evaluation of an end-to-end Road Pothole Detection and Automated Reporting System. The system employs YOLOv8 as its core detection engine, processes video input using OpenCV, and utilizes Python for orchestrating data pipelines, severity classification, and communication with a centralized database and monitoring dashboard. The remainder of the paper is organized as follows: Section 2 reviews related work, Section 3 describes the system architecture, Section 4 explains the methodology, Section 5 presents experimental results, Section 6 discusses future scope, and Section 7 concludes the paper.

II. RELATED WORK

Prior research on automated pothole detection spans several technological generations. Early works employed accelerometer-based and vibration sensing approaches, where anomalous vehicle vibration patterns were used as proxy indicators of road surface defects [4]. While computationally lightweight, these methods conflate potholes with other road features such as speed bumps and drain covers, leading to significant false positive rates.

Image-processing-based approaches using traditional computer vision techniques such as edge detection, texture analysis, and background subtraction have also been explored [5]. These methods perform adequately under controlled lighting and camera positioning but degrade significantly under real-world conditions including varying illumination, shadow, occlusion by vehicles, and surface wetness.

The advent of convolutional neural networks (CNNs) substantially improved detection robustness. Architectures such as VGG, ResNet, and early YOLO variants were applied to pothole classification with improved accuracy [6]. The introduction of YOLOv5 and subsequently YOLOv8 brought further improvements in speed-accuracy trade-offs, enabling real- time inference on embedded hardware. Several studies have employed YOLOv5 for pothole detection from dashcam footage, reporting meanaverage precision values exceeding 80% [7]. The integration of detection systems with complete reporting and maintenance workflows remains comparatively underexplored, and our work directly addresses this gap.

III. SYSTEM ARCHITECTURE AND TECHNOLOGIES

YOLOv8 object detection model

YOLOv8, developed by Ultralights, represents the current state of the art in single-stage real-time object detection. Unlike two-stage detectors such as Faster R-CNN, YOLOv8 processes the entire image in a single forward pass, producing bounding box predictions and class probabilities simultaneously. This architectural choice results in inference speeds suitable for real-time deployment [8]. YOLOv8 introduces an anchor-free detection head, decoupled classification and regression branches, and a C2f bottleneck module that improves gradient flow and feature reuse across scales. In the proposed system, YOLOv8s (small) is selected after benchmarking for its balance between detection accuracy and computational efficiency on edge hardware. The model is fine-tuned on a custom pothole dataset collected from Indian road conditions, ensuring domain specificity.

OpenCV for image and video preprocessing

OpenCV provides the image acquisition, preprocessing, and visualization layer of the system [9]. Frames captured from vehicle-mounted cameras are decoded using OpenCV's Video Capture interface. Preprocessing steps include Gaussian blurring for noise suppression, adaptive histogram equalization (CLAHE) for contrast normalization under varying illumination, and frame resizing to the YOLOv8 input resolution of 640 x 640 pixels. Detected pothole bounding boxes are rendered onto frames and logged for storage. OpenCV's multi-threading support ensures that frame acquisition and inference occur concurrently without pipeline bottlenecks.

Python-based processing pipeline

Python serves as the integration language for the entire system. The processing pipeline coordinates frame ingestion from OpenCV, inference through the YOLOv8 API, severity classification based on bounding box area relative to the frame, GPS coordinate extraction from connected IoT modules, payload construction, and data transmission to the backend server. Libraries including NumPy, Pandas, and Requests are employed for numerical processing, data formatting, and HTTP communicationrespectively. The modular architecture ensures each component can be independently updated or replaced without disrupting the overall workflow.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

IoT sensors and GPS integration

The hardware layer comprises vehicle-mounted 1080p dashcams, GPS receivers for precise geolocation, and Raspberry Pi 4B microcontrollers running the detection pipeline. An optional accelerometer module provides complementary vibration data to corroborate visual detections. Sensor data is transmitted over 4G LTE to the cloud backend. The system also interfaces with roadside infrastructure-mounted cameras for stationary monitoring of high-traffic intersections and historically pothole-prone road segments.

Centralized database and monitoring dashboard

Detected pothole records, each containing GPS coordinates, timestamp, severity classification (minor, moderate, or severe), thumbnail image, and vehicle identifier, are persisted in a PostgreSQL database hosted on a cloud server. A web-based monitoring dashboard built with Flask and Leaflet.js renders an interactive real-time map overlaid with pothole markers color- coded by severity. Maintenance supervisors can filter by region, severity, and date range, assign work orders to repair crews, and mark potholes as repaired. All data is exposed through a RESTful open API endpoint enabling integration with third-party road management and smart city platforms.

IV. METHODOLOGY

Dataset preparation

A custom dataset was assembled comprising images and video frames captured from roads in Chennai, Tamil Nadu and surrounding regions. The dataset contains approximately 4,200 annotated images with bounding box labels for pothole instances, spanning diverse conditions including dry, wet, dawn, and night-time lighting. Annotations were created using Labeling in YOLO format. The dataset was partitioned into training (70%), validation (15%), and test (15%) splits. To address class imbalance, oversampling of underrepresented severity classes was performed, and synthetic images were generated using mosaic and copy-paste augmentation strategies.

Model training

YOLOv8s was selected as the primary model after preliminary benchmarking demonstrated superior accuracy relative to YOLOv8n withacceptable latency for the target hardware. Training was conducted for 100 epochs on a system equipped with an NVIDIA RTX 3060 GPU (12 GB VRAM) using the Ultralights training API. The initial learning rate was set to 0.01 with cosine annealing scheduling. Stochastic gradient descent with momentum (0.937) was employed as the optimizer. Early stopping with a patience of 15 epochs was applied to prevent overfitting, with a fixed batch size of 16 images per step.

System workflow

The operational workflow proceeds as follows. Vehicle-mounted sensors continuously capture road footage. Each frame is preprocessed by the OpenCV module and passed to the YOLOv8 inference engine. Frames containing detected potholes with confidence scores above a threshold of0.45 are flagged. Severity is determined by the ratio of bounding box area to total frame area: minor (below 3%), moderate (3-8%), and severe (above 8%). The flagged frame, GPS coordinates, timestamp, severity label, and vehicle ID are packaged into a JSON payload and transmitted via HTTP POST to the backend API. The backend validates and stores the record in the database and updates the live dashboard. Duplicate submissions within a 50-metre radius and 10-minute window are deduplicated to prevent database pollution. Upon receiving a new record, the dashboard alerts maintenance supervisors via automated email and SMS notifications. Repair crews are provided GPS-guided navigation to the pothole location via a mobile companion application, and resolved potholes trigger an automated re- inspection scheduled after 30 days.

V. RESULTS AND DISCUSSION

Detection performance

The YOLOv8s model achieved a mean average precision at IoU threshold
0.50 (mAP@0.5) of 87.3% on the held-out test set, with precision and recall values of 89.1% and 84.7% respectively. Inference speed on the Raspberry Pi 4B hardware averaged 11 frames per second (FPS) after ONNX model export and INT8 quantization, sufficient for real-time detection at typical urban driving speeds. On the development GPU system, inference achieved 42 FPS. False positives were most commonly observed on painted road markings, manhole covers, and shadow-cast irregular patches under strong sunlight, representing known challenges for texture-based defect detection systems.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

System performance

End-to-end latency from frame capture to database record creation averaged2.1 seconds under normal 4G LTE network conditions. The deduplication mechanism successfully suppressed 94% of redundant submissions during field trials covering a 15-kilometre test route in Chennai. The dashboard reliably updated with new pothole markers within 5 seconds of detection, providing near real-time situational awareness. System uptime during a two- week field evaluation period was 99.2%, with brief interruptions attributable to network connectivity losses in tunnel sections.

Comparison with existing approaches

The proposed system compares favorably against manual inspection methods and earlier sensor-only approaches. Manual inspection typically covers a given road segment at most once per quarter, whereas the proposed system provides continuous monitoring. Accelerometer-only systems reported in the literature achieve precision values of approximately 70-75% [4], compared to 89.1% for the proposed vision-based approach. The addition of the automated reporting and dispatch pipeline represents a qualitative improvement over detection-only systems, translating detections directly into actionable maintenance workflows.

VI. FUTURE SCOPE

Several directions are identified for future enhancement of the system. First, real-time 3D depth estimation using stereo cameras or LiDAR sensors would enable volumetric pothole characterization, providing maintenance teams with quantitative data for material and cost estimation. Second, integration with municipal Geographic Information Systems (GIS) and existing road asset management platforms would streamline operational workflows. Third, the deployment of federated learning protocols would allow the model to be continuously refined using data from the entire fleet of participating vehicles without centralizing sensitive GPS trajectory data, improving both model accuracy and user privacy.Fourth, autonomous vehicle support represents a high-value application domain. By publishing real-time pothole location data through standardized APIs, the system can feed hazard avoidance modules in autonomous and advanced driver-assistance systems, enabling path planning adjustments to avoid pothole-dense segments. Fifth, predictive analytics using historical pothole frequency and environmental data such as rainfall intensity and temperature cycling could enable proactive maintenance scheduling, addressing road degradation before it reaches the pothole stage.

Finally, expansion to detect other road surface defects including cracks, rutting, and delamination would broaden the system's utility as a comprehensive road health monitoring platform aligned with smart city infrastructure objectives.

VII. CONCLUSION

This paper has presented a comprehensive automated road pothole detection and reporting system integrating YOLOv8 deep learning-based object detection, OpenCV video processing, IoT sensor infrastructure, and a cloud- based monitoring and maintenance dispatch platform. The system achieves high detection accuracy (mAP@0.5 of 87.3%) with real-time inference capability on edge hardware, overcoming key limitations of manual inspection and earlier sensor-based approaches. The end-to-end pipeline from detection to maintenance dispatch represents a holistic solution to the road pothole management challenge, offering significant potential for deployment within smart city frameworks. The proposed architecture is scalable, hardware-agnostic, and extensible to broader road surface monitoring tasks. Future work will focus on 3D defect characterization, federated model improvement, and integration with autonomous vehicle safety systems.

REFERENCES

1. S. Ryu, J. Lim, J.H. Cha, Pothole detection and recognition using deep learning. KSCE J. Civ. Eng. 24, 3193-3205 (2020)
2. A. Eriksson, M. Girod, B. Hull, R. Newton, S. Madden, H. Balakrishnan, The pothole patrol: using a mobile sensor network for road surface monitoring, in Proc. ACM MobiSys, Breckenridge, USA (2008), pp. 29-39
3. G. Jocher, A. Chaurasia, J. Qiu, Ultralytics YOLOv8, version 8.0 (2023). https://github.com/ultralytics/ultralytics
4. P. Mohan, V.N. Padmanabhan, R. Ramjee, Nericell: rich monitoring of road and traffic conditions using smartphones, in Proc. ACM SenSys, Raleigh, USA (2008), pp. 323-336
5. Y.C. Lin, K.C. Liu, A pothole detection system using image processing, in Proc. IEEE ICSIPA, Kuching, Malaysia

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

(2013), pp. 293-297

6. M. Maeda, Y. Sekimoto, T. Seto, T. Kashiyama, H. Omata, Road damage detection and classification using deep neural networks with smartphone images. Comput.-Aided Civ. Infrastruct. Eng. 33, 1127-1141 (2018)
7. F. Naddaf-Sh, M.M. Naddaf-Sh, H. Kashani, S. Zargarzadeh, An efficient and scalable deep learning approach for road damage detection, in Proc. IEEE Big Data, Atlanta, USA (2020), pp. 5602-5608
8. C.Y. Wang, A. Bochkovskiy, H.Y.M. Liao, YOLOv7: trainable bag-of- freebies sets new state-of-the-art for real-time object detectors, in Proc. IEEE CVPR, Vancouver, Canada (2023), pp. 7464-7475
9. G. Bradski, The OpenCV library. Dr. Dobb's J. Softw. Tools 25, 120-125 (2000)

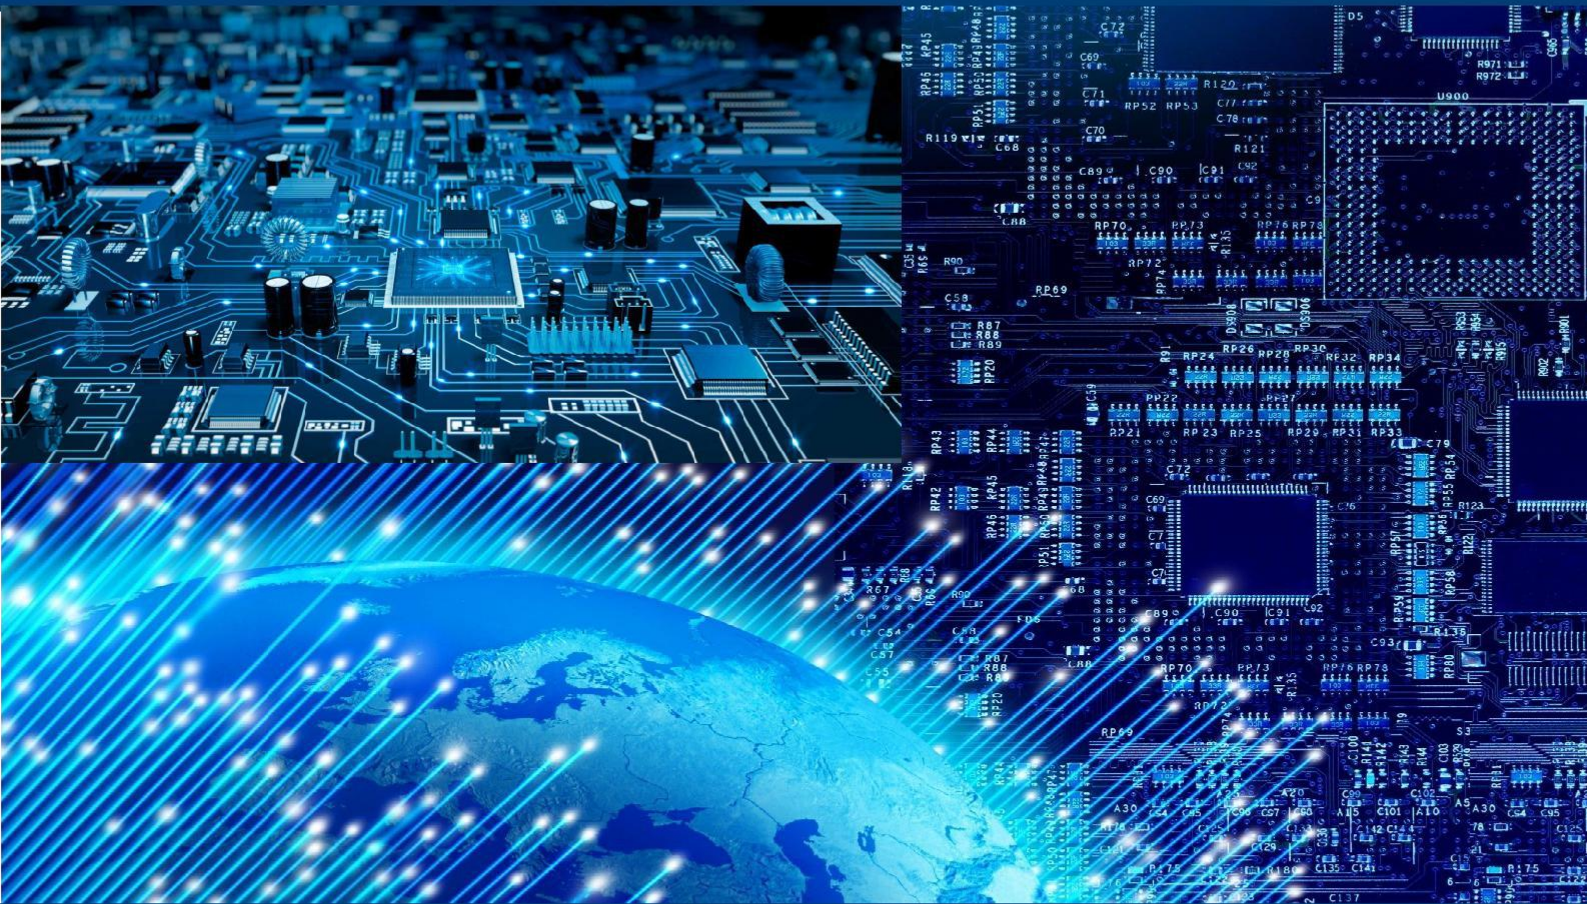

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH IN SCIENCE, ENGINEERING AND TECHNOLOGY

| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |

www.ijmrset.com