Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Edge AI Enabled Road Inspection Vehicle for On Field Road Damage Assessment and Traffic Data Acquisition

Dr. Yogesh S. Khandekar

Associate Professor, Department of Civil Engineering
Sipna College of Engineering and Technology, Amravati, India

Abstract
Real-time road monitoring systems play a vital role in enhancing transportation safety, traffic efficiency, and infrastructure maintenance by continuously assessing road surface conditions and traffic dynamics; however, conventional stationary deployments significantly restrict spatial coverage and limit scalability across large road networks. As a continuation of our previous work on an edge-AI- based intelligent road monitoring system using YOLOv8, this paper presents a hardware-driven, field-deployable road inspection vehicle designed for continuous on-road assessment. The proposed system integrates a Raspberry Pi–based edge computing unit with a vision module, GPS receiver, and a motorized inspection platform to enable real- time detection of potholes, cracks, and vehicles during active field traversal. The YOLOv8 model is executed locally on the embedded platform to ensure low-latency inference, while road damage severity is estimated using a hybrid metric that combines detection confidence scores and bounding-box area measurements. Each detected event is geotagged and time-stamped, enabling spatial mapping and historical analysis, and an event-driven communication strategy is employed to selectively transmit critical data to a cloud dashboard, thereby minimizing bandwidth usage. Field experiments conducted on urban and semi-urban roads demonstrate reliable detection accuracy, stable real- time performance, and effective road coverage under varying traffic density and lighting conditions, validating the feasibility of extending edge-AI road monitoring systems into real-world, field operational deployments for intelligent transportation infrastructure.

Keywords: Edge Artificial Intelligence, Road Inspection System, YOLOv8, Embedded Vision, Field Deployment, Intelligent Transportation Systems, Road Damage Detection, Traffic Monitoring

1. Introduction

To extend the proposed edge-AI road monitoring framework into real-world operation, this work further develops a hardware-driven, field-deployable road inspection system capable of performing continuous on-road assessment. The system integrates the Raspberry Pi–based YOLOv8 inference module with a motorized inspection platform, on- board camera, GPS unit, and power management circuitry to enable real-time road damage detection during field traversal. Unlike stationary monitoring setups, the hardware platform allows active coverage of extended road segments, enabling the collection of geotagged damage and traffic data under varying surface, lighting, and traffic conditions. Detected potholes, cracks, and traffic events are processed locally on the embedded device to ensure low-latency inference, while a severity-aware, event-driven communication strategy selectively uploads only critical alerts to the cloud, thereby minimizing bandwidth requirements. Field trials conducted on urban and semi-urban roads demonstrate that the deployed system maintains stable real-time performance and detection accuracy during motion, while significantly improving spatial coverage compared to fixed installations. This hardware and field validation confirm the feasibility of transitioning edge-AI road monitoring systems from laboratory prototypes to scalable, real-world infrastructure inspection solutions.

2. Literature Review

Building upon recent advances in lightweight object detection and edge-based road monitoring, this work moves beyond model-centric evaluations toward a hardware-driven and field-operational implementation. Earlier studies on road anomaly detection relied on traditional vision techniques such as disparity transformation and surface modeling, which demonstrated effectiveness under constrained conditions but lacked scalability and robustness for large-scale deployment [1]. With the advent of deep learning, several works have reported improved accuracy in pothole and crack detection using convolutional neural networks and YOLO-based architectures, primarily evaluated on benchmark datasets or controlled environments [3], [11], [14]. However, the transition from dataset-driven validation to real-world deployment remains limited.

To address these challenges, the proposed system deploys a YOLOv8-based object detection model on a low-cost embedded computing platform integrated with an on-board camera, GPS module, and power management unit. YOLOv8 has been demonstrated to provide an effective trade-off between accuracy and inference speed, making it suitable for real-time embedded vision applications [5], [11]. Prior work on lightweight deep learning models confirms that optimized architectures can achieve stable inference on resource-constrained devices when carefully integrated with hardware-aware design principles [6], [16]. By executing all detection and severity estimation tasks locally, the proposed system eliminates dependency on continuous cloud connectivity and significantly reduces end-to-end latency, which is a key requirement for edge-based transportation monitoring systems [4], [20].

Unlike stationary camera installations or cloud- dependent monitoring setups, the developed platform operates directly on urban and semi-urban road segments, capturing real-world variations in lighting, traffic density, and surface degradation. Mobile and embedded vision systems have been shown to enhance spatial coverage and adaptability compared to fixed infrastructure, particularly in intelligent transportation and smart city applications [9], [19]. Detected potholes, cracks, and vehicle counts are geotagged and time-stamped to support spatial mapping, historical trend analysis, and maintenance planning, aligning with AI-driven priority-based road maintenance strategies discussed in recent studies [18]. To ensure efficient bandwidth usage, a priority-aware, event-driven communication mechanism is employed, where only high-severity or safety-critical events are transmitted to a centralized dashboard, while non-critical data is stored locally for deferred synchronization [8], [13].

Field experiments conducted under diverse environmental and traffic conditions demonstrate that the proposed hardware-driven system maintains reliable detection accuracy and stable inference speed during continuous operation. The results confirm that edge-based intelligence enhances responsiveness and robustness compared to cloud- dependent approaches, enabling scalable and practical road infrastructure monitoring aligned with smart mobility initiatives [4], [12], [15], [17], [20].

3. Mechanical Setup

The mechanical design of the road inspection vehicle focuses on three primary objectives: structural stability, modular integration, and reliable motion control under on-road conditions. The mechanical architecture is organized into three major subsystems, namely the chassis and structural framework, the drive and locomotion system, and the sensor mounting and power distribution unit, which together support continuous field deployment and real-time data acquisition [9], [16], [20].

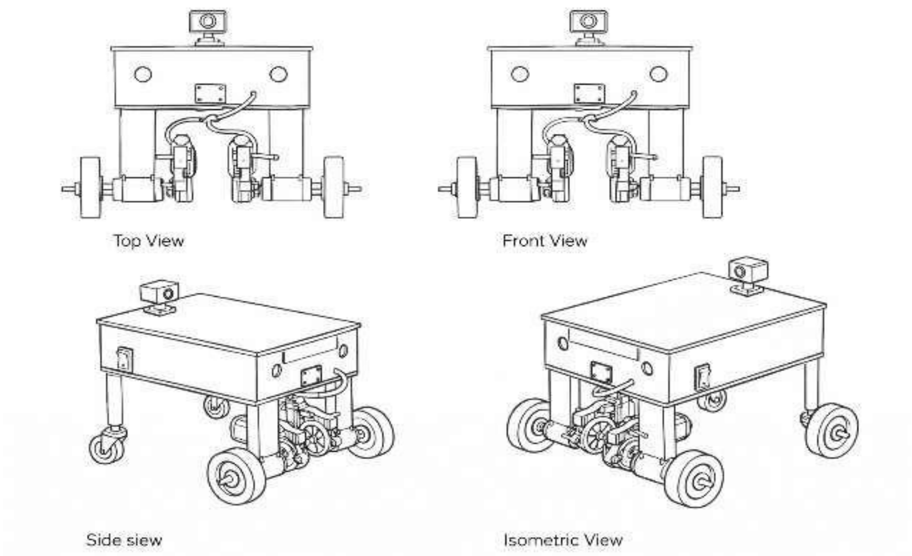

Fig. 1. Multi-view CAD representation

A. Chassis and Structural Framework:

The inspection vehicle is built on a rigid metal chassis designed to withstand vibrations, uneven road surfaces, and prolonged field operation [16], [20]. The chassis provides sufficient load-bearing capacity to support the embedded edge computing unit, camera module, GPS receiver, battery pack, and motor drivers [9]. A low center of gravity is maintained to ensure stability during motion and sudden speed variations [17]. The camera mounting frame is positioned to provide an unobstructed forward-facing view of the road surface, enabling accurate road damage detection and traffic monitoring while the vehicle is in motion [9], [16].

Fig. 2. Chassis and Structural Framework

B. Drive and Locomotion System

The locomotion system employs a four-wheel configuration driven by two high-torque 12 V DC geared motors connected to the rear wheels, while the front wheels function as free-rotating caster units to facilitate smooth turning [16]. This differential drive mechanism enables stable low-speed traversal, which is essential for capturing clear road surface imagery and ensuring reliable edge-AI inference [9], [17]. The system allows tight turning radii, making it suitable for navigating narrow road segments and curved paths [20]. Vehicle speed is regulated through motor driver control to synchronize motion with real-time image acquisition and processing [4], [16].

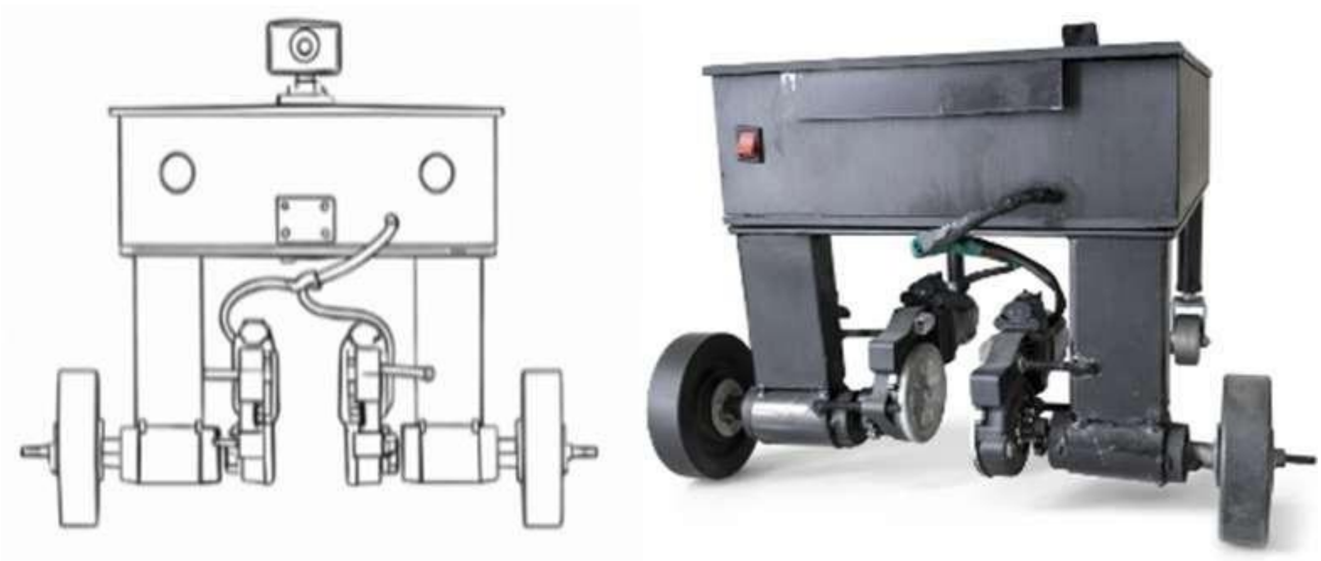

Fig. 3. Drive and Locomotion System

C. Sensor Mounting and Power Distribution Unit:

The sensing and power subsystem integrates a high- resolution camera for road surface imaging and traffic monitoring [9] and a GPS module for geotagging detected events to support spatial mapping and analysis [14]. Autonomous operation is enabled through a rechargeable battery pack designed for prolonged field deployment [16]. To maintain image quality during motion, the camera mount is vibration-isolated, improving the robustness of edge-AI inference under real-world conditions [17]. Stable and regulated power delivery to the Raspberry Pi, motor drivers, and peripherals ensures reliable embedded operation, while the modular design allows future integration of additional sensors for system enhancement [6], [12], [20].

AIJFR
Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

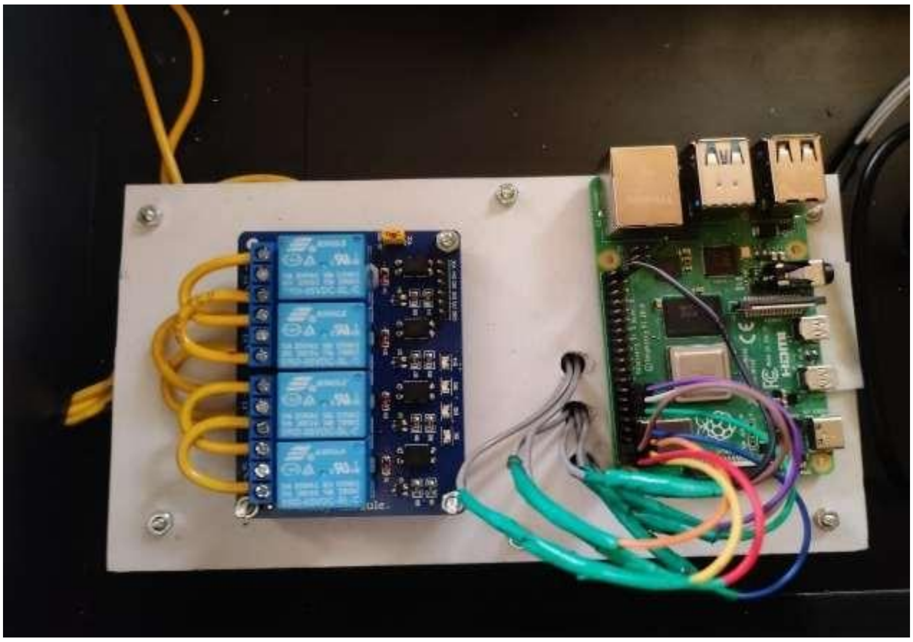

Fig. 4. Sensor Mounting and Power Distribution Unit

4. Methodology

The proposed road inspection system is implemented as a hardware-driven, field-operational platform that integrates edge-based intelligence, real-time sensing, and cloud-assisted event management. The overall design incorporates embedded hardware configuration, on-device deep learning inference using YOLOv8, real-time image acquisition and preprocessing, traffic and damage analysis, and cloud-enabled alert dissemination. The complete operational workflow of the system is illustrated in Fig 5.\

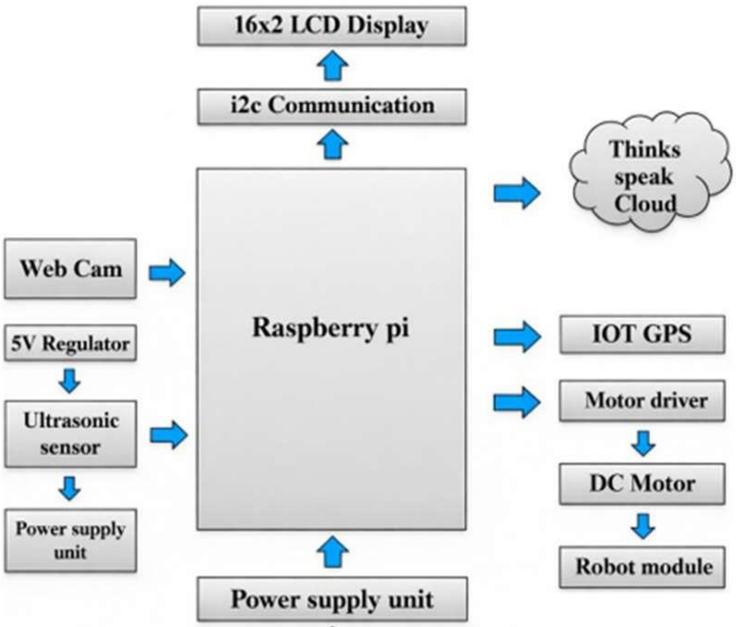

Fig. 5. Block Diagram

A. Hardware Configuration and Edge Processing Unit

The core processing unit of the system is a Raspberry Pi 4, selected for its balance between computational capability, power efficiency, and cost-effectiveness for edge-based intelligent transportation and embedded vision applications [6], [16]. A high-resolution Pi camera module is mounted on the inspection vehicle to enable continuous road image capture during traversal, providing visual input for road damage detection and traffic analysis [9], [17]. A GPS module is integrated to provide precise geotagging of detected road anomalies, enabling accurate spatial mapping and location-aware analysis of road conditions [14], [18]. Local alerts are generated using an on-board buzzer to indicate critical detections in real time, supporting immediate situational awareness during field operation [20]. The system is powered using a regulated DC power source suitable for extended field deployment, ensuring reliable operation of the processing unit, sensors, and peripherals [16]. To enhance deep learning inference performance, an optional Coral TPU accelerator is incorporated, enabling faster execution of YOLOv8 models while maintaining low power consumption, which aligns with recent approaches for accelerating edge-AI workloads on embedded platforms [6], [11].

B. Image Acquisition and Preprocessing in Field Conditions

During field operation, the Pi camera continuously captures video frames of the road surface at predefined intervals, providing a steady stream of visual data for real-time analysis on the edge device [9]. To maintain reliable detection performance under challenging real-world conditions such as fluctuating illumination, shadows cast by vehicles or roadside objects, surface texture variations, and motion-induced noise, a comprehensive image preprocessing pipeline is executed locally on the embedded platform [7], [17]. Initially, Gaussian and bilateral filtering techniques are applied to suppress sensor noise and minor surface artifacts while preserving important structural edges relevant to road damage detection [6]. Contrast Limited Adaptive Histogram Equalization (CLAHE) is then employed to enhance local contrast and improve the visibility of cracks and potholes under uneven lighting conditions commonly encountered during outdoor operation [7]. Grayscale conversion is performed to reduce computational complexity and memory usage without significantly affecting the geometric features required for detection [3].

Furthermore, region-of-interest cropping is applied to restrict processing to the road surface area within the camera’s field of view, thereby eliminating irrelevant background information and reducing false detections [1]. Perspective correction is also incorporated to compensate for camera tilt and vehicle motion, ensuring geometric consistency across frames and improving detection robustness [14]. Collectively, these preprocessing steps enhance feature clarity, reduce false positives, and improve the overall accuracy and stability of edge- AI-based road damage detection during continuous field deployment.

Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Fig. 6. Sample Road Damage Images

Fig. 7. Image Preprocessing

C. Traffic Estimation Using Edge-Deployed YOLOv8

Traffic analysis is performed using the YOLOv8n model optimized for embedded execution through ONNX Runtime or TensorRT acceleration to enable low-latency, real-time inference on resource-constrained edge devices [5], [6], [11]. The lightweight architecture allows reliable detection of multiple vehicle classes, including cars, buses, trucks, and two-wheelers, from each processed video frame captured by the on-board camera under dynamic field conditions [4], [9]. To maintain temporal consistency and avoid duplicate counting, a centroid-based object tracking mechanism is employed, wherein the centroid

Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

of each detected bounding box is computed and assigned a unique identifier that is preserved across consecutive frames using distance-based association [3], [17]. Virtual counting lines are defined within the camera's field of view, and a vehicle is registered only when its tracked centroid crosses the line in a predefined direction, ensuring robust and direction-aware traffic measurement [14], [18].

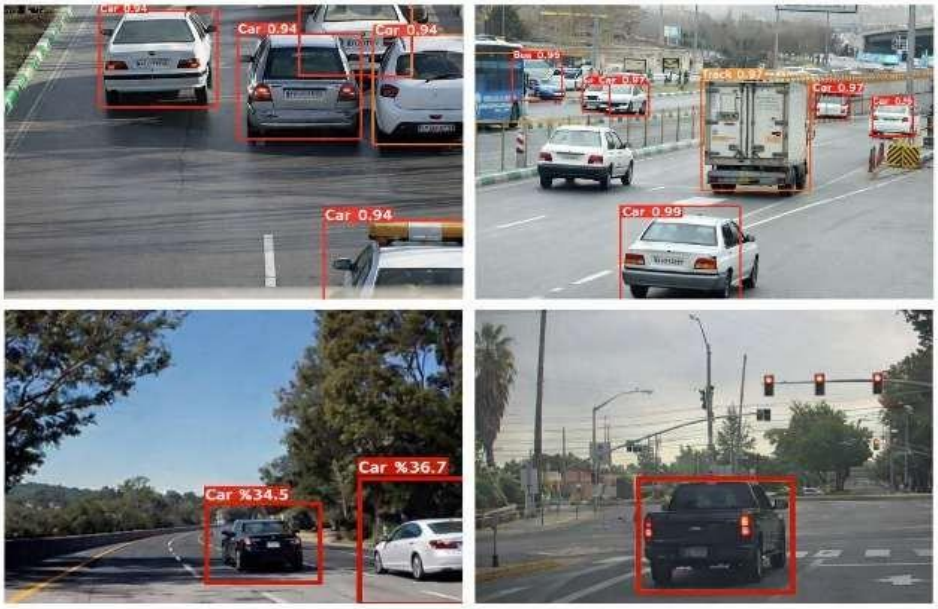

Fig. 8. Vehicle Detection and Counting

D. Road Damage Detection and Severity Assessment

Road surface anomalies such as potholes and cracks are detected using the same YOLOv8 inference framework, which has been shown to be effective for real-time road damage detection under diverse conditions [3], [5], [11]. Each detected damage instance is associated with a confidence score and bounding-box dimensions, providing both probabilistic and geometric information about the anomaly [14]. A hybrid severity estimation metric is employed by combining the model confidence values with the bounding-box area, an approach commonly adopted to approximate damage extent and visual prominence [1], [17]. Based on this combined metric, detected anomalies are classified into minor, moderate, or severe categories, enabling effective prioritization of road segments according to damage severity and potential impact on traffic safety and maintenance planning [18].

Fig. 9. Damage Severity Estimation

E.Data Fusion and Event Generation

To support intelligent decision-making, traffic density metrics and road damage severity scores are fused to compute a priority index for each detected event [18], [20]. Each event record includes time stamps, GPS coordinates, severity classification, associated traffic context, and visual evidence, enabling comprehensive situational awareness and post-analysis [14], [9]. High-priority events corresponding to severe road damages occurring on high-traffic segments are flagged for immediate reporting, while low-priority events are stored locally on the edge device for periodic synchronization [8], [13]. This priority-aware, event-driven data handling strategy ensures efficient bandwidth utilization while preserving critical information required for timely maintenance and safety interventions [12], [15].

F.Cloud Integration and Remote Monitoring

Critical event data is transmitted to a centralized cloud dashboard, where detected road damage locations, severity levels, and traffic heatmaps are visualized in real time to support situational awareness and monitoring [8], [13]. The cloud interface enables remote supervision and data- driven maintenance planning by road authorities, aligning with smart transportation infrastructure management practices [15], [18]. By adopting an event-driven communication approach, the system minimizes bandwidth usage while ensuring timely alerts for critical road conditions, a strategy widely recognized as effective for scalable edge–cloud smart city deployments [2], [12], [20].

5. Results And Discussion

The performance of the proposed Edge-AI road inspection system was evaluated through real-time experiments conducted on a Raspberry Pi 4 platform and further validated using publicly available road damage datasets [3], [10]. The evaluation focuses on key performance indicators including inference latency, detection reliability, resource utilization, and communication efficiency under realistic field-operational conditions, which are commonly adopted metrics for assessing embedded and edge- based intelligent transportation systems [4], [6],[20].

A. Inference Latency and Throughput Analysis

The system's real-time capability was evaluated by measuring the average inference latency per frame and processing throughput under different embedded hardware configurations, following standard evaluation practices for edge-based vision systems [4], [6], [20]. Initially, experiments were conducted on a Raspberry Pi 4 without any external accelerator, where the system achieved an average processing latency of approximately 140 ms per frame, corresponding to a throughput of 6–8 frames per second (FPS). This performance is consistent with CPU-only deep learning inference reported in prior embedded implementations [6], [16]. Subsequently, the system was evaluated with an integrated Coral TPU accelerator, which reduced the inference latency to nearly 80 ms per frame and enabled a sustained throughput of 12 FPS. These results demonstrate the effectiveness of hardware acceleration for real-time YOLO-based inference on edge platforms, as also observed in recent studies [5], [11], [20]. The comparative performance of both.

Configuration	Avg. Latency (ms)	Throughput (FPS)	Real Time Suitability
Raspberry Pi 4 (CPU only)	140	6-8	Moderate
Raspberry Pi 4 + Coral TPU	80	12	High

Table 1: Inference Latency and Throughput

B. Class-wise Detection Reliability

Detection reliability was analyzed separately for traffic objects and road surface defects using class- wise F1-score and average detection confidence [3], [11], [14]. The results indicate strong generalization across object categories, with vehicle detection outperforming road damage classes due to more distinct visual features and higher object contrast [7], [17]. The class-wise detection performance is summarized in table 2.

Detected Class	Avg. Confidence Score	F1-Score
Vehicles	0.94	0.92
Potholes	0.91	0.89
Cracks	0.90	0.88

Table 2: Class-wise Detection Reliability

C. Resource Utilization on Embedded Platform

To assess deployment feasibility, CPU utilization, memory usage, and power consumption were measured during continuous operation [4], [6]. Despite executing deep learning inference and image preprocessing locally, the system maintained moderate resource usage, validating its suitability for long-duration field deployment [12], [16]. The embedded resource utilization observed during continuous operation is summarized in table 3.

Resource Allocation	Average Usage
CPU Utilization	68%
RAM Consumption	1.2 GB
Power Draw	6.5W
Thermal Stability	Maintained

Table 3: Embedded Resource Utilization

D. Communication Overhead and Bandwidth Efficiency

The impact of the event-driven communication mechanism was evaluated by comparing continuous video streaming with selective metadata transmission [8], [13]. Instead of uploading full video streams, only annotated images and event descriptors were transmitted for critical detections, leading to a substantial reduction in network usage [12], [20]. The comparison of data transmission modes and corresponding bandwidth utilization is summarized in table 4.

Transmission Mode	Avg. Data Rate (MB/min)	Bandwidth Reduction
Continuous Video Upload	120	75
Event-Driven Upload	35	70

Table 4: Communication Efficiency Comparison

E. System Robustness Under Field Conditions

The system was tested across multiple road segments with varying traffic density, surface conditions, and lighting environments to evaluate its robustness under real-world operational scenarios [9], [17]. These test routes included low-traffic residential roads, moderately congested urban streets, and high-traffic arterial roads, thereby exposing the system to diverse motion dynamics, occlusions, and illumination variations. Performance consistency was assessed by monitoring detection stability and false positive rates during prolonged continuous operation, ensuring that transient environmental changes did not significantly degrade system reliability [7], [20]. The evaluation results indicate that the system maintains high detection stability under low and moderate traffic conditions, while sustaining moderate-to-high stability even in dense traffic scenarios where occlusions and rapid object interactions are more frequent. The overall field robustness across different traffic scenarios, demonstrating the system's capability to operate reliably under practical deployment conditions, is summarized in table 5.

Test Condition	Detection Stability
Low Traffic	High
Moderate Traffic	High
High Traffic	Moderate–High

Table 5: Field Robustness Evaluation

F. Discussion

The results demonstrate that the proposed system achieves a robust balance between accuracy, efficiency, and resource utilization under real-world operating conditions [6], [20]. The low inference latency and stable throughput enable reliable real- time operation on embedded edge hardware [5], [11], while class-wise detection reliability confirms effective detection of both traffic objects and road surface damages [3], [14], [17]. Event-driven cloud communication significantly reduces bandwidth usage without compromising situational awareness [8], [12], [13]. Overall, the findings validate the effectiveness of deploying YOLOv8-based edge intelligence for real-world road inspection in smart city environments [2], [15], [20].

Fig. 10. Final output of the proposed Edge-AI road inspection system

6. Conclusion and Future Work

This work presented an Edge-AI–based road inspection system that leverages a lightweight YOLOv8 deep learning framework combined with optimized image preprocessing techniques to enable accurate, real-time detection of road surface damages and traffic conditions on embedded hardware [5], [6], [11]. The use of a single forward- facing camera allows continuous road surface monitoring and traffic analysis without reliance on expensive sensing infrastructure, while on-device edge inference ensures low-latency operation and reduced dependence on cloud connectivity [2], [9], [16]. By jointly analyzing road damage severity and traffic density, the system facilitates priority-based decision-making for road maintenance, ensuring that critical and high-impact road segments are identified and addressed in a timely manner [14], [18].

The adoption of an event-driven cloud reporting mechanism further enhances system efficiency by transmitting only essential metadata and annotated snapshots, thereby significantly reducing communication overhead while preserving situational awareness [8], [12], [13]. This data- efficient design improves system reliability and enables continuous monitoring and logging of road conditions under real-world operating environments [20]. Future work will focus on integrating inertial and vibration sensors to enhance surface anomaly detection accuracy [17], expanding training datasets to accommodate diverse road textures and climatic conditions [3], incorporating predictive maintenance analytics using historical data [18], and deploying the system at scale across multiple road networks to support comprehensive smart city infrastructure management [2], [15].

References

1. R. Fan, U. Ozgunalp, B. Hosking, M. Liu, and I. Pitas, “Pothole detection based on disparity transformation and road surface modeling,” *IEEE Transactions on Image Processing*, vol. 31, pp. 291– 305, 2023.
2. A. Zhang, K. Wang, and Y. Liu, “Edge computing–based intelligent transportation systems: Architecture, applications, and challenges,” *IEEE Internet of Things Journal*, vol. 11, no. 2, pp. 1321– 1334, 2024.
3. S. Arya, R. Sharma, and P. Gupta, “Automated road damage detection and classification using deep learning techniques,” *Journal of Intelligent Transportation Systems*, vol. 28, no. 3, pp. 211–223,2024.
4. M. K. Islam, M. Shahriar, and S. Ahmed, “Real- time traffic monitoring using deep learning on edge devices,” *IEEE Access*, vol. 12, pp. 45812–45825, 2024
5. Jocher et al., “YOLOv8 for real-time object detection: Architecture and performance analysis,” *Pattern Recognition Letters*, vol. 175, pp. 1–9, 2024.
6. Y. Li, Z. Chen, and H. Wang, “Lightweight deep learning models for real-time embedded vision applications,” *IEEE Embedded Systems Letters*, vol. 15, no. 2, pp. 89–92, 2023.
7. S. K. Sarker, M. Rahman, and A. Hossain, “Vision-based road surface anomaly detection under varying illumination conditions,” *Journal of Transportation Engineering*, vol. 150, no. 3, 2024.
8. P. Kumar and S. Banerjee, “Event-driven edge– cloud frameworks for smart city video analytics,” *IEEE Transactions on Cloud Computing*, vol. 12, no. 1, pp. 97–109, 2024.
9. J. S. Lee and V. Kumar, “Edge-AI enabled vision systems for intelligent transportation infrastructure,” *Sensors (MDPI)*, vol. 24, no. 5, p. 1837, 2024.

10. T. Lin et al., “Microsoft COCO: Common objects in context,” *International Journal of Computer Vision*, vol. 130, no. 5, pp. 1237–1256 2022.
11. A. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, “YOLO-based object detection models: Performance analysis and deployment considerations,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 9, pp. 5678–5691, 2023.
12. M. Babar and A. Akan, “Edge intelligence for smart city monitoring: A comprehensive survey,” *Sustainable Computing: Informatics and Systems*, vol. 42, 2024.
13. S. Patel and R. Mehta, “Bandwidth-efficient edge–cloud frameworks for real-time video analytics,” *IEEE Transactions on Multimedia*, vol. 26, pp. 3142–3154, 2024.
14. K. Wang, X. Liu, and Y. Zhang, “Real-time road condition monitoring using deep neural networks,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 1, pp. 876–887, 2024.
15. European Commission, “Artificial intelligence for smart and sustainable transport systems,” *EU Smart Mobility Report*, 2024.
16. N. Chen, J. Liu, and L. Zhao, “Embedded vision systems for intelligent traffic and infrastructure monitoring,” *Applied Sciences*, vol. 14, no. 7, 2024.
17. U. Sarkar and I. Ghosh, “Deep learning-based road surface condition monitoring for intelligent transportation,” *Journal of Transportation Safety & Security*, vol. 16, no. 2, pp. 215–229, 2024.
18. A. Verma and N. Singh, “Priority-based road maintenance using AI and traffic analytics,” *International Journal of Smart Cities and Applications*, vol. 5, no. 2, pp. 45–58, 2025.
19. Y. Song et al., “Embedded vision systems for real-time traffic and infrastructure monitoring,” *Applied Sciences*, vol. 14, no. 6, 2024.
20. H. Zhou, D. Zhang, and Q. Liu, “Edge AI- driven intelligent transportation systems: Design and implementation,” *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 2, pp. 412–423, 2024.