International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181

e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

# AI-Driven Automated Pothole Detection and Geospatial Mapping System for Real-Time Road Condition Monitoring

*Priyanka[1], Arjun P[2], Allan Abraham[3], Kanishkaa N S[4], Jai Charan M[5], Aleena Biju[6]*
*[1]Assistant Professor, BCA, Yenepoya Deemed University, Bangalore, Karnataka, India.*
*[2,3,4,5,6]UG, BCA, Yenepoya Deemed University, Bangalore, Karnataka, India.*
***Emails:***Priyabhadrashetty05244@gmail.com[1],arjun.p6362@gmail.com[2],allanabraham271299@gmail.com[3],kanishkaans07@gmail.com[4], jaicharan1666@gmail.com[5], aleenabiju1104@gmail.com[6]*

## Abstract

*Road surface degradation, particularly potholes, poses significant risks to road safety, vehicle longevity, and infrastructure planning. This paper presents an AI-based automated pothole detection and geospatial mapping system designed to monitor real-time road conditions. The system operates by detecting potholes as the vehicle equipped with the device passes by, capturing a snapshot of the pothole in the frame, and recording the timestamp and GPS coordinates. The data is then stored in a database, providing a detailed record of pothole locations. This information can be leveraged by government agencies for more efficient and cost-effective infrastructure maintenance, as it enables quicker identification and repair of road defects. The system employs the YOLO26 object detection model, selected for its computational efficiency and suitability for deployment on resource-constrained edge devices. To improve reliability in real-world scenarios, the system incorporates geospatial tagging, duplicate detection elimination through spatial–visual clustering, and probabilistic confidence aggregation to reduce false detections. The integrated architecture further supports interactive visualization through a web-based dashboard for monitoring detected potholes and managing repair status. YOLO26 was benchmarked against YOLOv8s with identical datasets, training configurations, and hyperparameters. Results show YOLO26 achieving a mean Average Precision (mAP@0.5) of 78.9%, surpassing YOLOv8s (76.96%) with higher precision (85.0% vs. 81.0%) while maintaining similar recall. The system also integrates GPS-tagged geospatial mapping, enabling probabilistic monitoring where multiple vehicle observations enhance coverage and reliability. This study is among the first to apply YOLO26 for road condition monitoring, offering a scalable and efficient solution for large-scale road condition monitoring and infrastructure maintenance planning.*

***Keywords:*** *Edge computing; Geospatial mapping; Object detection; Pothole detection; YOLO26.*

## 1. Introduction

The transport systems within cities have a great reliance on the good condition of the roads in terms of the safety and viability of movement. In spite of this, erosion in the road surfaces, especially on potholes, remains a recurrent problem in most of the urban areas. The normal causes of potholes are a mix of factors including the repetitive loading of the vehicle, infiltration of water, change of temperatures and the erosion of pavement materials with time. As these environments continue, the structural integrity of the road decreases, causing cracks which eventually become potholes. As urbanization levels are evolving very fast and the number of vehicles is increasing, the intensity and frequency of pothole occurrence have increased dramatically, posing a great challenge to authorities charged with the responsibility of maintaining the roads and other infrastructure. Road safety is directly and mostly immediately affected by the occurrence of potholes. The motorists might lose control of the vehicle after hitting the potholes in cases where drivers are taken by surprise and in others, accidents may occur. This is a very high risk, especially on two-wheelers and smaller vehicles where little irregularities on the surface may influence the stability. Drivers can also make abrupt evasive action in order to evade potholes, an action that can further increase the odds of collision with the ongoing traffic. Consequently, the potholes that go unnoticed or not repaired pose a major threat to all road users. Besides the issue of

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

safety, potholes are a significant source of vehicle damage. When the uneven road surfaces are prolonged, major vehicle components such as tires, suspension systems, and wheel alignment may suffer. This increases the maintenance cost for vehicle owners and also reduces the overall life of the vehicle. In a broader sense, the regular occurrence of potholes can be attributed to a larger number of inefficiencies within infrastructure management and potentially have an adverse effect on the efficiency of transportation as well as satisfaction of people with the road networks. Historically, the most effective method of pothole detection and monitoring the status of roads has been based on manual inspection. Physical road surveys undertaken by municipal officials or reliance on citizen complaints can be used to determine such damaged areas. The methods are however time consuming, labour intensive, and in most cases inefficient particularly in big cities that have got big road networks. Inspecting thousands of kilometers of roads on a regular basis needs high manpower, and the overall cost of operation is high. The other crucial weakness of the traditional monitoring systems is the inability to detect in real-time. Inspections are not done on a regular basis, but occasionally, newly formed potholes can take very long periods to be noticed. These delays do not only pose a danger to the safety of the road users but also enable the damages that the road might be having to be increased with time, thus making their repair more complicated and costly. As of late, with the emergence of cutting-edge technology on Artificial Intelligence (AI) and computer vision, new opportunities to automate the process of infrastructure monitoring have become possible. AI methods have the capacity to process photos or video feeds obtained by cameras and detect defects on the road surface with high accuracy. Pothole detection can also be performed with high efficiency and without a constant human presence in the monitoring system by implementing machine learning models into the monitoring systems. These are automated solutions which are being investigated to be used in large scale road monitoring applications. This paper provides a system of AI-driven automated pothole recognition and mapping to enhance monitoring of road conditions and aid infrastructure maintenance planning. The system uses a deep learning-based object detection model that is based on the YOLO26 architecture in order to detect potholes based on real-time video obtained using vehicle-mounted cameras. YOLO (You Only Look Once) is an object detector built on a real-time basis and can effectively process video streams which makes it highly effective when it comes to edge-based deployment. A GPS data acquisition pipeline is incorporated into the system in order to locate properly the location of identified potholes. The pipeline logs the geographic position of the specific point of detection and sends the information to a centralised backend where the data could be sent by numerous vehicles and the data could be computed at a large scale to observe the conditions of the roads in the urban setting. One of the issues with such systems is that the same pothole is found in sequential video frames. To this, we use a dual-constraint clustering mechanism that involves the use of both spatial proximity and the visual similarity to remove redundant detections. Consequently, a given amount of physical potholes is pulled in one consolidated record in system database. The identified potholes are also displayed using an interactive web-based dashboard that displays the geographic position of the potholes on an interactive geographic map. Such an interface helps the authorities and maintenance team in the city to calculate the high-risk areas, focus on repair efforts and have a history of the road condition over time. The proposed solution combines clustering-based deduplication, deep learning-based detection, geospatial tagging, and real-time visualization, which enables scalability and efficiency in pothole monitoring. It leads to better road safety, more efficient planning of maintenance, and better road infrastructure control.

## 2. Related Work

### 2.1. Mobile/Sensor-Based Detection

Initial research in the area of potholes detection resorted, to a large extent, to low-cost sensor-based methods, specifically, to smartphone-based accelerometer and vibration-based systems. These techniques utilized inherent inertial sensors to detect anomalies in road or pavements without having to

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

install special imaging devices. Wu et al. proposed a machine learning-based system to detect potholes with the help of smartphone accelerators and GPS positions in [1]. The methodology included deriving statistical characteristics of the vertical acceleration signals and using the method of supervised classification to identify potholes. Although the system proved viable in real world conditions, its operation was much affected by mechanical noise which was occasioned as a result of engine vibrations, braking, suspension behaviour, and change in vehicle speed among other factors. Moreover, threshold-based event detection resulted in high false alarms, especially in situations with speed breakers, road joints, or irregular asphalt surfaces. The other weakness was the inability to have visual assurance and accurately localize the space since the detections were obtained through the vibration pattern and not the actual objects. On the same note, Ozoglu and Gokgoz in [2] examined a vibration-based pothole detection method based on smartphone inertial sensors, along with convolutional neural networks. Even though the method enhanced classification accuracy due to the learning of intricate vibration patterns, its stability was still reliant on the similarities in signal properties. Differences in vehicle type, suspension system, tire pressure and environmental conditions like rainfall also imposed differences that caused inconsistency which influenced the performance of detection. Noise also had to be dealt with by careful preprocessing and threshold calibration of the method. Although it had promising results under controlled conditions, it did not have the capabil ity to generate bounding-box localization, geometric information, or contextual information of the scene, which restricts its use in mapping-based maintenance systems. Although sensor-based solutions are relatively cheap and simple to implement, they also have a number of innate drawbacks such as noise sensitivity, false detection, differences between vehicles and the lack of visual spatial data. The challenges make them less efficient in generalization in various real-life situations. To address these constraints, recent studies have gone a step further to employ deep learning-based computer vision algorithms that make it possible to directly detect and localize potholes through the use of their visual images.

### 2.2. YOLO-Based Detection Systems (YOLOv8–YOLOv11)

In [3], the authors introduced a pothole detection system that is built on YOLOv8 and combined with ESRGAN to perform super-resolution reconstruction as well as geospatial mapping functions. YOLOv8 was designed with the anchor-free architecture and a decoupled head design, which led to high-performance localization and detection. Nevertheless, adding ESRGAN meant adding some computational cost. The super-resolution preprocessing consumed more inference latency and memory and thus was more difficult to deploy in real-time on more resource-constrained edge devices commonly utilized in vehicular systems. In [4], a pothole detection model using YOLOv9 was developed in real-time on a Raspberry Pi platform with the help of a Coral USB Accelerator. Programmable Gradient Information (PGI) and the Generalized Efficient Layer Aggregation Network (GELAN) were among the architectural improvements described in YOLOv9 that increased the capability to represent features and detect them. Despite the positive results under controlled conditions, deterioration of performance was found in conditions of different road textures and motion blur. Also, the hardware was a constraint to the system in terms of its capacity to support high frame rates in real-time scenarios. In [5], the pothole detector was specifically trained on a lightweight YOLOv11-s model, and hyperparameter optimization was carried out using genetic mutation to increase the training efficiency. Although this lightweight configuration lowered the computational demand and allowed edge deployment, some parameters were also at default values, which may not permit optimization of pothole-specific features. In addition, the small and less varied data sets also elevated the chances of overfitting and that can affect the generalization in various road settings. Enhancement of low-light images has also been considered as a way of enhancing the robustness of the detection. A fusion-based method in which DCE-Net (Deep Curve Estimation Network) is used to

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

improve low-light road images, and YOLOv10 is used to detect objects in real-time was proposed in [6]. Nevertheless, these enhancement methods, such as DCE-Net, can cause oversmoothing of the edge features in some lighting scenarios, which undermines the description of pothole edges. Pothole data acquisition systems are frequently used in controlled survey environments as opposed to autonomous driving systems, which need to behave in a fail-safe manner under any environmental situation. This has led to the situation where robustness to lighting variations is still a consideration but in large scale road monitoring systems it is common to put emphasis on computational efficiency, scalability and accurate detection. Based on the available literature, it is possible to note that YOLOv8 is more efficient in detection and better in localization, YOLOv9 is more effective in extracting features and optimization of gradients, and YOLOv11-s is more devoted to lightweight implementation. With these developments, there still persist the problems of computational overhead, large scale performance in large vehicle fleets, and generalization consistency in all conditions on the road.

### 2.3. Identified Research Gaps

Although the current strategies have achieved a lot, there are serious limitations at the system-level that are not adequately covered.

#### 2.3.1. Lack of Integrated Mapping Systems

Most of the previous literature has largely concentrated on accuracy of the detection, whereas the importance of aggregation of geospatial information has received little attention. Consequently, a large number of systems lack mechanisms to centralize detections into central road condition databases, needed to conduct large-scale monitoring and maintenance planning.

#### 2.3.2. Limited Fleet Scalability

Existing applications are frequently tested in single-vehicle configurations in controlled scenarios. The difficulties of mass implementation of multiple vehicles, each with different environments, hardware settings, and road conditions, are still not well studied.

#### 2.3.3. Absence of Probabilistic Aggregation Mechanisms:

Large-scale aggregation of probabilities does not exist; in other words, no probabilistic mechanism exists to aggregate a collection of probability distributions. Current techniques usually focus on the detection of a single object separately and fail to combine probabilistic fusion of redundant measurements over time or a set of vehicles. Such aggregation mechanisms may be introduced to reduce false positives and increase the detection reliability by exploiting multi-instance validation.

## 3. Proposed System Architecture

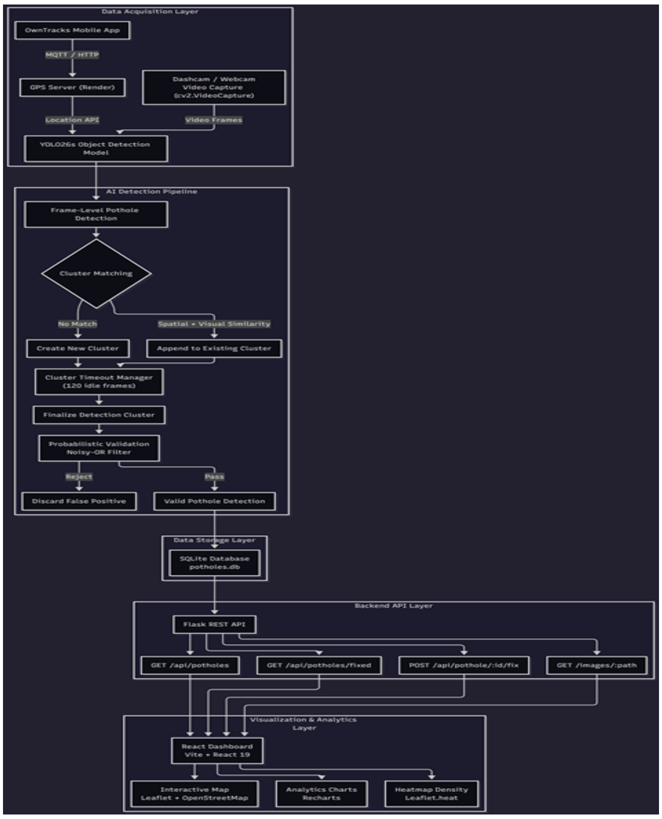


**Figure 1 System Overview Diagram Illustrating the Entire Pipeline of Data Acquisition to Detection, Storage and Visualization**

The proposed system uses a modular architecture that utilizes vehicle-based sensing platforms to detect potholes and map the area and infrastructure. As shown in Figure. 1, the architecture combines computer vision-based real-time detection with mobile GPS data acquisition, smart deduplication, centralized data storage and interactive visualization

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

interface. The system is in general designed with four major layers: data acquisition, data detection pipeline, data management of the back-end and data visualization.

### 3.1. Data Acquisition Layer

The data acquisition layer is responsible for collecting both visual and geographic data needed for the process of pothole detection. The dashcam or webcam, which is mounted on the vehicle, is constantly recording video frames of the road surface, and it is the OpenCV video capture interface that does all this. The main input to the detection pipeline is these frames. Geographic coordinates are also acquired via a GPS pipeline via a smartphone since the camera itself does not comprise of embedded GPS hardware. A mobile device using the OwnTracks application, constantly transmits position data to a simple cloud-based GPS server that is running on Render. The detection engine will access the latest coordinates periodically via a REST-based API (GET/get_location) and connects the coordinates to the location of the potholes noticed. This design permits proper geospatial tagging without the need to have specific GPS equipment on the edge hardware. Figure 2 demonstrates the connection between the mobile device, cloud-based GPS server, and edge detection engine, Flask backend API, and React-based

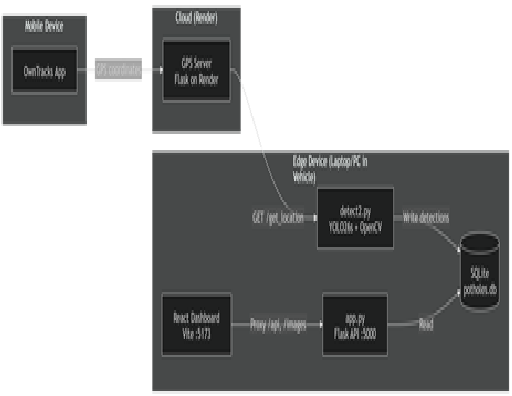


**Figure 2 Deployment Diagram Indicating the Communication Between the Mobile GPS Device, Cloud GPS Server (Render), Edge Detection Engine (Yolo26s with OpenCV) And Flask API and React Dashboard**

### 3.2. AI Detection Pipeline

The detection pipeline is in charge of frame-level detection of potholes with duplicate filtering. Every incoming video frame undergoes the trained YOLO26s object detection model that is a lightweight one-stage deep learning network intended to detect objects in real-time on edge devices. The model identifies the pothole areas in each frame and gives out the bounding boxes and relevant confidence scores. Figure. 3 demonstrates an example of a real-time detection with YOLO26s, in which several potholes are detected in one frame of the road.

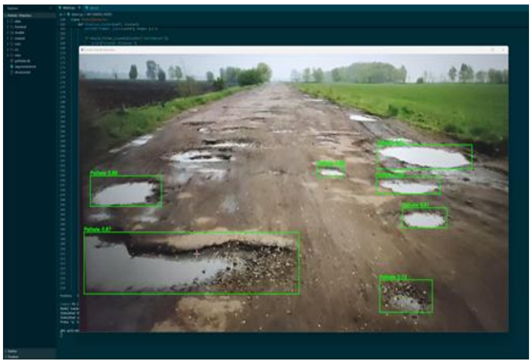

**Figure 3 Bounding Box and Confidence Score-Based Real-Time Pothole Detection with Yolo26s.**

When a vehicle goes across a pothole, the same defect can be seen in several successive frames. This would lead to duplicate detection of the same pothole in the physical world without proper management. To overcome this challenge, a two-constraint clustering mechanism is adopted which incorporates both the proximity and the visual similarity in space. The proximity based on space is measured using Haversine distance calculated between the GPS coordinates of detections. Simultaneously, visual similarity is evaluated on the basis of correlation of color histograms based on the identified image segments. A new detection is linked to the existing cluster only in the case that both the spatial and visual thresholds are met, otherwise, it is considered as a distinct pothole and a new cluster is started. The clusters are the possible potholes observed in several frames. After the system notices that the vehicle has passed the pothole, i.e. it has not been detected in a specified number of frames, the cluster is finalized and sent to additional validation.

### 3.3. Probabilistic Validation and Data Storage

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

In order to raise the reliability of the detection and reduce the number of false positives, the confidence scores of a cluster of two or more detections are summed using a probabilistic Noisy-OR model. This formula provides the system with the ability to synthesize evidence based on many different frames and approximate the probability of a particular cluster being a real pothole. At this stage, clusters that fail to meet predefined confidence/confirmation thresholds are filtered. In the case of validated clusters, the sharpest image frames are picked by the system with the help of Laplacian-based sharpness measure. This is in order to have only good visual representation of the potholes to be stored and analysed. Acquired together with the chosen images, appropriate metadata, such as geographic location, time, aggregate confidence score and statistics regarding detection, is stored in a SQLite database. In order to prevent the performance problems that come with placing large binary objects in the database, image data are not placed directly in the database. Rather, the storage of pothole images is done locally in a well-organized directory on the host machine, and the database contains entries of the file path of potholes. This architecture option enhances efficiency of the database and can access images in a much better time when required.

### 3.4. Backend API and Visualization Layer

The backend layer presents the data stored about potholes in a Flask-based REST API, where detection records, and related image data could be accessed in structured form. Some of the operations supported by the API include accessing pothole entries, detailed metadata, and close-out of the lifecycle of pothole records including repair follow-ups. The frontend is based on a React dashboard that allows interactively visualizing and monitoring the road conditions. The interface combines mapping platforms like Leaflet and OpenStreetMap to give the position of potholes on an online map. Besides simple mapping, another tool to visualize the levels of potholes in the area is the use of heat maps and analysis charts, allowing to pinpoint those areas with more potholes and direct more efforts toward their maintenance. This combination of layers enables the entire architecture to be scalable and real-time in pothole monitoring, as well as efficient edge processing, reliable detections validation, and central infrastructure information visualization.

## 4. Methodology

### 4.1. System Detection Pipeline

The suggested implementation is based on a real-time automated system of detecting potholes, geospatial tagging, and monitoring the state of roads with the help of a vehicle-mounted sensory platform. It is an object detection model of deep learning with GPS localization and centralized data to sustain scalable infrastructure monitoring. An operationally mounted camera films the road surface each frame of the video. A trained model of detection is applied to each frame in order to find possible pothole locations. After a pothole is identified, geographic coordinates are accessed on a GPS pipeline on a mobile device. The resulting record of detection, which contains the bounding box, confidence score, timestamp, and location, are subjected to deduplication and validation steps and stored in a centralized database. The overall detection pipeline consists of the following stages:

- Image acquisition from a vehicle-mounted camera.
- Frame-level pothole detection using the YOLO26s object detection model.
- GPS coordinate acquisition and smoothing for accurate geospatial tagging.
- Spatial–visual clustering to eliminate duplicate detections.
- Probabilistic confidence aggregation for filtering false positives.
- Storage of validated pothole detections and visualization through an interactive dashboard.

This pipeline enables continuous monitoring of road conditions while ensuring efficient data aggregation and reducing redundancy in pothole detection records.

### 4.2. Deep Learning-Based Pothole Detection

The pothole detection system is designed based on the YOLO26s object detection model. YOLO (You Only Look Once) is a part of the family of single-stage detectors, in which object localization and classification are implemented at the same time

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

during the same forward pass of the network. Unlike two-stage methods like Faster R-CNN, models with YOLO have much lower inference latency, and so they are well-suited to real-time use. YOLO26s is intended to offer a compromise between accuracy and speed in detection. It employs layers of convolutional neural networks to learn hierarchical visual representations of input images that allow the model to learn characteristics that correlate with potholes, including irregular shapes, surface depressions, and a change of texture on a road surface. The model gives bounding boxes of possible pothole areas and confidence values associated with each incoming frame. Detections beyond some preset confidence levels are sent into the next geospatial processing step. YOLO26s has a lightweight design and thus can process in real-time on edge devices without necessarily sending the raw video data to remote servers.

### 4.3. GPS Acquisition and Coordinate Smoothing

The pothole detection must detect the precise location to enable the infrastructure to be monitored and maintained. The vehicle-mounted camera is not equipped with built-in GPS devices; therefore, the system acquires the location information through a smartphone device equipped with a GPS tracking application. The cell phone constantly sends out geographic coordinates to a lean cloud-based GPS server. The detection engine periodically gets the most recent coordinates via a REST-based API and links them to identified potholes. Raw smartphone GPS measurements however usually have noise due to signal reflections, atmospheric problems and sensor errors. To reduce these effects, the system uses an Exponential Moving Average (EMA) filter to smooth the GPS trajectory and the formula is:

$$S_t = \alpha X_t + (1 - \alpha)\, S_{t-1}$$

where $X_t$ represents the raw GPS coordinate at time $t$, $S_t$ represents the smoothed coordinate, and $\alpha$ is the smoothing factor. In this filtering method, the positional variations are eliminated abruptly and the general path of the vehicle is maintained.

### 4.4. Pothole Detection Deduplication — Spatial-Visual

When a vehicle passes over a pothole, the same fault can be seen in several consecutive frames within the video stream. Otherwise, this causes multiple database entries for the same physical pothole. To reduce this problem, a dual-constraint clustering mechanism is employed that incorporates spatial proximity with visual similarity.

#### 4.4.1. Spatial Constraint

The separation between detections is calculated based on the Haversine formula that calculates the great-circle distance between two geographic coordinates. Two detections are said to be spatially related when the distance between them is less than some predefined distance (around 6 meters). This threshold takes into consideration both average noise in GPS and the physical proportion of potholes.

#### 4.4.2. Visual Similarity Constraint

A color histogram representation is derived from each image region of the detected pothole to ensure further that similar potholes close by are associated with the same pothole. This is a three-dimensional histogram in 8×8×8 bins which are normalized to ensure scale invariance. Similarity between detections is then determined by histogram correlation. An existing cluster is linked to a detection when both spatial and visual thresholds are met, otherwise a separate cluster is formed to indicate a unique pothole instance. This two-constraint clustering method is successful in eliminating duplicate detections while retaining individual potholes that can be close to each other on the road.

### 4.5. Probabilistic Validation of Pothole Clusters

Individual frame-level detections may occasionally contain false positives caused by shadows, surface stains, or lighting variations. To improve detection reliability, the system aggregates detection confidences within each cluster using a probabilistic Noisy-OR model. Let ci represent the confidence score of the ith detection within a cluster. The aggregated probability that the cluster corresponds to a real pothole is computed as:

$$P = 1 - \prod_{i=1}^{n} (1 - c_i)$$

where $c_i$ represents the confidence score of the

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

$i^{th}$ detection and n denotes the total number of detections within the cluster. This formulation treats each detection as an independent observation. As more consistent detections accumulate within a cluster, the aggregated probability increases, providing stronger evidence for the existence of a pothole. Clusters that fail to satisfy predefined probability or confirmation thresholds are discarded as false positives.

### 4.6. Image Quality Filtering

To make sure that only quality visual evidence is kept with regards to each detected pothole, the sharpness of captured frames is analysed with the variance of the Laplacian operator. This measure indicates the change in intensity of the edges of the image and allows the system to tell the difference between a sharp image and an image with motion blur. Among all the frames that can be attributed to a pothole cluster, only those of the best quality in terms of Laplacian variance are retained. This selection process ensures that the database consists of clear and informative visual records that can be used to conduct inspection and maintenance analysis.

### 4.7. Data Storage and Visualization

The validated pothole measurements are stored in a structured database comprising metadata like geographic positions, timestamps, aggregated confidence values and references to the relative image material. Instead of directly storing image files in the database, the system stores the pothole images locally in an organized directory on the host device, while the database only stores the file paths. Separating the storage of images and database metadata prevents large binary data from accumulating in the database tables, which can result in reduced performance, and facilitates quick retrieval of images when needed. The architecture can also be extended to accommodate cloud-based object storage solutions to enable remote storage of pothole images, while the database stores the URLs to maintain scalability and centralized data management. The stored data is exposed via a backend API and represented as an interactive web-based dashboard. The dashboard shows the pothole locations on a digital map and provides analytical features such as heatmaps and cluster-based visualization to illustrate the areas with the greatest pothole density. Through incorporation of deep learning-based detection, geospatial tagging, probabilistic validation, and interactive visualization, the recommended methodology provides a scalable system for automated monitoring of road infrastructure and planning of road maintenance.

## 5. Experimental Results

To assess the effectiveness of the proposed pothole detection and mapping system, a series of controlled experiments were conducted using a prototype implementation of the architecture described earlier. The evaluation focused on three key aspects: real-time detection performance of the YOLO26s model, the effectiveness of the clustering-based deduplication mechanism, and the functionality of the geospatial visualization dashboard.

### 5.1. Experimental Setup

The experimental setup was implemented on a local edge device, where a laptop handled the detection engine, backend API, and visualization interface. To simulate a dashcam-like configuration, a smartphone configured as a webcam was used to capture video frames representing road surface conditions. The YOLO26s model was integrated into this pipeline using the Ultralytics framework, enabling real-time pothole detection on incoming frames. For indoor testing, prerecorded pothole footage was displayed on a monitor, and the smartphone camera captured the video stream in real time. This arrangement allowed the system to process visual input in a manner similar to a vehicle-mounted dashcam operating under real-world conditions. The indoor experimental setup used for this evaluation is illustrated in Figure. 4. To associate detected potholes with geographic locations, the system incorporates a GPS data acquisition pipeline. In a real-world deployment scenario, a smartphone-based GPS module continuously transmits location data to a cloud-hosted server, which can be accessed by the detection engine through a REST API. As the actual GPS motion could not be copied in the indoor setting, a simulated GPS coordinate generator was deployed to simulate the movement of the vehicle along a predefined route. This generator created sequential geographical coordinates at fixed intervals which was

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

close to realistic movement patterns. The system could reproduce the real-world operating conditions in an experimental environment by taking the prerecorded pothole footage and simulating GPS data to create a simulated environment. This configuration facilitated stable testing of geospatial tagging, clustering and data aggregation modules, while also avoiding the safety risks and variability associated with repeated outdoor testing.

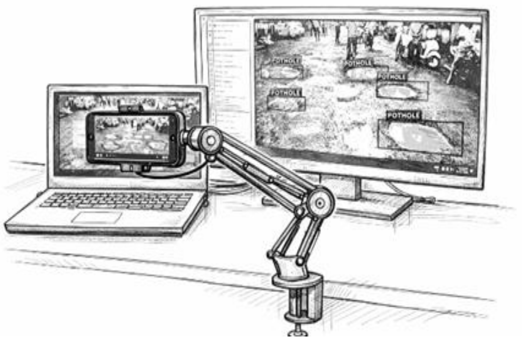

**Figure 4 Conceptual Illustration of the Indoor Experimental Setup Used to Evaluate the Proposed Pothole Detection System**

### 5.2. Pothole Detection Performance

The YOLO26s model was trained and tested on a pothole-annotated dataset under the same training conditions as YOLOv8s. Both models had the same dataset split, input resolution, and training setup to achieve a fair and consistent evaluation. Table 1 shows a summary of the comparative performance.

**Table 1 Model Performance Comparison**

| Model | mAP@0.5 | mAP@0.5-0.95 | Precision | Recall |
|---|---|---|---|---|
| YOLO26s | 0.789 | 0.509 | 0.850 | 0.676 |
| YOLOv8s | 0.770 | 0.501 | 0.810 | 0.649 |

The findings demonstrate that YOLO26s outperforms YOLOv8s across all assessed measures. Specifically, the gains in precision show a decrease in the number of false positive hits, whereas the gains in recall show that the model is better able to detect more potholes that occurred in the data. All in all, these results indicate that YOLO26s offers more reliable detection performance under the conditions specified in the experiment along with superior computational efficiency.

### 5.3. System Demonstration

Besides testing the detection model, the entire end-to-end system was tested in order to evaluate the detection, geospatial mapping and visualization pipeline. Potholes were effectively detected based on the captured frames and the detections were associated with geographic positions during operation and logged as validated records to a centralized database. Figure 5. depicts the primary dashboard page that displays 11 active (red) pothole markers on an interactive OpenStreetMap of the Bypass Road at Kanakapura. When one clicks on a marker, a popup is displayed with the images of the detection, the time, the GPS position, and the confidence score. A Mark as Fixed button enables the maintenance teams to start the repair process on the map.

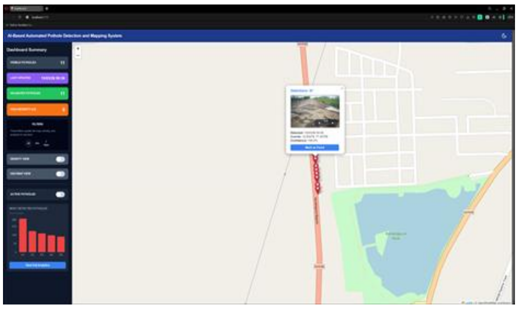

**Figure 5 There are 11 Markers on the Dashboard Indicating Active Potholes (Red) on Kanakapura Bypass with Detection Popup.**

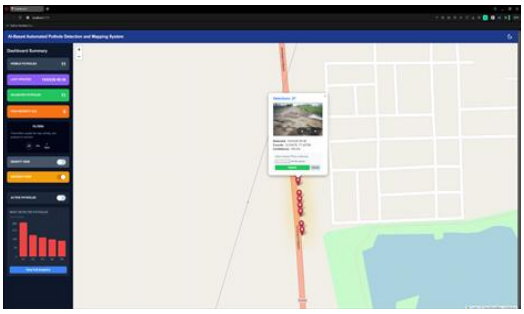

**Figure 6 Mark as Fixed Popup with the Option to Upload a Repair Photo and Submit a Confirmation.**

Figure 6. depicts the Mark as Fixed popup that enables one to add an optional post-repair photo before submitting. This repair job writes the record of

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

the potholes to the fixed potholes table and the related before and after pictures to the fixed pothole images table. Figure 7 presents the dashboard with three potholes identified as repaired. The markers are changed to green, and the popup now has the detection and repair timestamps and before and after repair images.

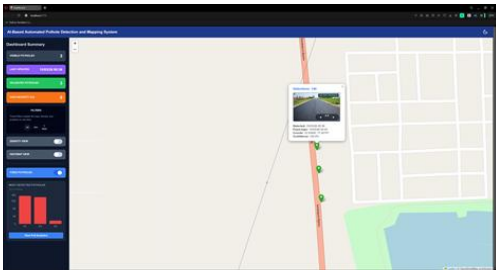

**Figure 7 Dashboard View Displaying Repaired Potholes Depicted as Green Markers, With Repair Timestamps and After-Repair Photos.**

The dashboard also has other spatial modes of analysis. In the Density View (Figure 8), the pothole clusters are represented as circles with color-coded overlays depending on the levels of severity: yellow for low density (less than 10); orange for medium density (10–19); and red for high density (20+). This visualization gives a preview of the locations of pothole concentration. Conversely, the Heatmap View (Figure 9) brings out a gradient-based depiction, with areas having a large concentration of potholes rendered with greater intensity.

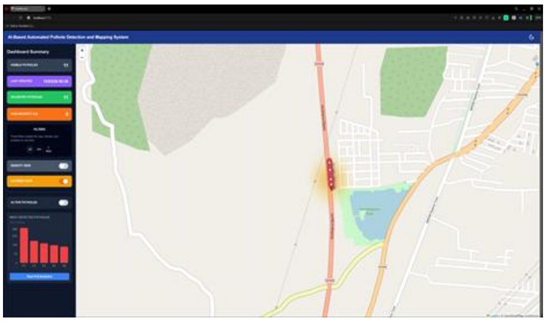

**Figure 8 Density Map Showing a Medium Concentration of Potholes On Kanakapura Bypass Using Color Coding According to Severity**

This enables easy identification of road segments that have a high concentration of damage. These findings indicate that the suggested system can incorporate real-time detection with geospatial tagging, data aggregation, and visualization in a single monitoring system.

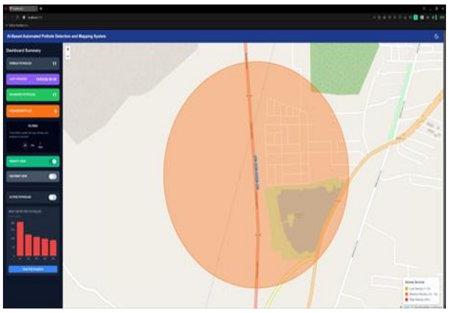

**Figure 9 Heatmap View Showing Gradient Intensities Where There is High Concentration of Potholes**

## 6. Discussion

The experimental findings reveal that the particular system can combine real-time pothole detection and geospatial mapping with central infrastructure monitoring in an integrated way. The YOLO26s model provides high detection accuracy and at the same time, it is computationally efficient such that it can be deployed in edge-based systems. Another strong side of the architecture is the spatial-visual clustering mechanism applied to duplicate elimination. When a vehicle is going over a pothole, the given defect can be presented in a series of frames. The clustering methodology is used to merge such frequent detections into one representation and, therefore, it minimizes the redundancy inside the database and enhances the consistency of the produced geospatial map. Besides that, the probabilistic Noisy-OR validation strategy also helps to enhance the reliability of detection. Overcoming the isolated false positives that can be caused by shadows, changes in lighting, or surface variation, the system can consolidate the confidence scores obtained through multiple observations. Meanwhile, some constraints should be taken into account. The current evaluation was conducted in a controlled indoor setting, the main reason being the safety limitations and the challenge of conducting repeated outdoor experiments. In practice, motion blur, different light conditions, and complicated road textures might cause additional influences on the

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181
e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

results of detection. The other weakness is associated with the utilization of GPS data on smartphones, which may introduce positional error because of signal noise and environmental interference. Even though the smoothing method based on EMA decreases sharp deviations, it is possible to do even better by using sensor fusion methods or employing more precise positioning solutions. In general, the findings indicate that the suggested system is an applicable and scalable solution when it comes to merging computer vision-based pothole detection with geospatial mapping, and also identifies the areas where real-world implementations can be improved further.

## Conclusion

The paper introduced an AI-based identification system of potholes and geospatial mapping in order to enhance the monitoring of the roads. The suggested framework incorporates real-time computer vision, GPS-based localization, duplicate elimination, and interactive visualization mechanisms. The detection pipeline is developed based on the YOLO26s object detection model that allows identifying potholes in real-time with the help of vehicle-mounted camera feedback. In order to deal with the recurrent identifications of the same object in successive frames, dual-constraint clustering was implemented, which made use of spatial proximity and visual similarity. Furthermore, the detection confidence scores were accumulated by a probabilistic Noisy-OR model, which enhanced the robustness by averaging the effects of individual false alarms. The experimental results demonstrate that YOLO26s outperforms YOLOv8s when trained under the same conditions, and its computational efficiency is reasonable enough to implement on edge devices. The entire system proves that it is possible to integrate detection with geospatial tagging and visualization into a single pipeline. On the whole, the presented architecture is a scalable and practical solution for automated pothole detection and mapping, which may be used to facilitate more effective road condition monitoring and data-driven maintenance strategies.

## Future Work

Although the suggested system proves to have strong potential regarding automated pothole detection and geospatial mapping, there is still a lot of room to improve and implement the system on a massive scale. Immediate directions include enhancing detection strength in degraded environmental conditions that include low lighting, shadows, rain, and changes in the texture of the pavements. The model can be better generalized to real-world situations by incorporating a more diverse dataset, as well as specifically selected data augmentation strategies. Other areas that can be investigated include optimization of camera location and calibration in vehicle-mounted systems. The height of the camera, the angle of view, and the field of view are parameters that directly affect the visibility of the road surface and may lead to the inability to detect smaller or distant potholes. The existing system is specifically focused on detection of potholes, but it can be extended to detect other types of road damage, such as cracks, bumps, and overall pavement deterioration. Expanding the model to a multi-class framework would allow for a better evaluation of road conditions. Additional enhancement can also be done through incorporation of other sensing modalities. Integrating visual detection with other signals like smartphone accelerometer data or vehicle vibration detection signals would improve reliability and minimize false positives through multimodal validation. On a larger scale, the implementation of the system across several vehicles is a great opportunity. Having the ability to aggregate distributed data sources would allow continuous monitoring of extensive road networks and create high-resolution geospatial maps. These insights can help prioritize maintenance work and enable a more proactive and data-driven approach to infrastructure management.

## References

[1]. C. Wu, Z. Wang, S. Hu, J. Lepine, X. Na, D. Ainalis, and M. Stettler, "An Automated Machine-Learning Approach for Road Pothole Detection Using Smartphone Sensor Data," Sensors, vol. 20, no. 19, p. 5564, Sep. 2020, doi: 10.3390/s20195564.
[2]. F. Ozoglu and T. Gökgöz, "Detection of Road Potholes by Applying Convolutional Neural

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0181

e ISSN: 2584-2854
Volume: 04 Issue: 05
May 2026
Page No: 1202 - 1213

Network Method Based on Road Vibration Data," Sensors, vol. 23, no. 22, p. 9023, Nov. 2023, doi: 10.3390/s23229023.

[3]. V. M. R. Priya, D. Chandu, G. Uday, G. Ahad, and N. Bharath, "An AI-Driven Approach for Intelligent Road Maintenance: Pothole Detection and Prioritization with YOLOv8, ESRGAN, and Geospatial Mapping," in Proc. 2025 Int. Conf. Commun., Comput. and Inf. Technol. (IC3IT), 2025, doi: 10.1109/IC3IT66137.2025.11340800.

[4]. D. Anbalagan and M. N. H. Mohd, "Enhanced Road Safety: Real-Time Pothole Detection Using YOLOv9," Evolution in Electrical and Electronic Engineering, vol. 6, no. 1, pp. 43–49, 2025, doi: 10.30880/eeee.2025.06.01.006.

[5]. S. Sivanraj, D. N. L. S. Uduwage, R. K. Shiwakoti, and M. Thripathi, "YOLOv11-s-Based Pothole Detection Model with Integrated Traffic Data for Maintenance Decision Support," Journal of Advanced College of Engineering and Management (JACEM), vol. 11, pp. 77–84, 2024, doi: 10.3126/jacem.v11i1.84529.

[6]. X. Xu, T. Wu, W. Ge, L. Dai, J. Gong, D. Chen, and S. Li, "Enhanced pothole detection in low-light conditions using DCE-Pothole and YOLOv10 fusion," Results in Engineering, vol. 28, p.107322, 2025, doi: 10.1016/j.rineng.2025.107322.