Journal of Science, Computing and Engineering Research (JSCER)
Volume-9, Issue-3, March 2026.
DOI: https:jscer.org/doi.org/10.46379/jscer.2026

CITY GUARDIAN: AN AI-POWERED URBAN SAFETY COMPANION

[1]J.R. Arun Kumar, [2]Lakshya Sharma, [3]Akshat Sharma, [4]Deepak Singh, [5]Aditya Sharma, [6] Himek Saini

[1]Professor, Department of CSE, Modern Institute of Technology and Research Centre, Rajasthan, India.

[2,3,4,5,6]UG Student, Department of CSE, Modern Institute of Technology and Research Centre, Rajasthan, India

Article Information

Received : 25 March 2026

Revised : 25 March 2026

Accepted : 26 March 2026

Published : 28 March 2026

Corresponding Author:

Deepak Singh

Abstract— Rapid urbanization has placed increasing pressure on city administrations to manage road infrastructure and public safety effectively. Existing systems rely on manual inspections and reactive complaint portals that lack real-time automation and structured data collection. This paper presents City Guardian, an AI-powered urban monitoring platform that integrates deep learning-based hazard detection, citizen crowdsourcing, geolocation mapping, and an administrative analytics dashboard. The system employs a YOLOv8 object detection model to identify road anomalies — including potholes and open manholes — from smartphone images, achieving a detection confidence of 0.95. Geotagged reports are stored in a Firebase cloud backend and visualized on interactive maps, enabling municipal authorities to monitor and prioritize repairs in real time. A gamified Guardian Points mechanism incentivizes citizen participation while crowd-based verification maintains data integrity. The cross-platform mobile application, built with Flutter, delivers proximity-based hazard alerts and an invalid-report penalty system to prevent misuse. Experimental results demonstrate high detection accuracy and practical feasibility across real urban environments. City Guardian provides a scalable and proactive framework for smart city governance, aligned with national initiatives such as Smart Cities Mission and Digital India.

Keywords: ***Computer Vision, Crowdsourcing, Flutter, Machine Learning, Pothole Detection, Smart Cities, Urban Safety, YOLOv8***

Citation: J. R. Arun Kumar, Lakshya Sharma, Akshat Sharma, Deepak Singh, Aditya Sharma, Himek Saini, "CITY GUARDIAN: AN AI-POWERED URBAN SAFETY COMPANION", Journal of Science, Computing and Engineering Research, 9(03), Mar 2026.

I. INTRODUCTION

In the era of smart cities and digital transformation, efficient monitoring of urban infrastructure and public safety has become increasingly important. Rapid urbanization has increased the pressure on city authorities to manage roads, public spaces, and safety conditions effectively. Traditional inspection mechanisms rely on manual surveys, delayed complaint portals, and siloed data systems that are neither scalable nor responsive to the pace of modern city growth.

City Guardian is an AI-enabled urban infrastructure and safety monitoring platform designed to modernize civic issue detection and reporting mechanisms. It serves as an integrated ecosystem connecting citizens, artificial intelligence models, and municipal authorities within a unified digital framework. The primary goal of the system is to enable structured detection, reporting, visualization, and resolution of urban hazards through real-time technological integration.

The application empowers users to capture or upload images of damaged roads, open manholes, or other hazards using their smartphones. These images are processed by a deep learning model trained on the YOLOv8 architecture, which identifies anomalies, generates bounding boxes, estimates severity, and classifies hazard categories. Each detection is automatically mapped using GPS coordinates to ensure location accuracy and administrative traceability.

The platform also includes a reward-based participation model where verified contributors earn digital Guardian Points. This gamification mechanism not only increases engagement but also strengthens data authenticity through collaborative validation. The administrative dashboard offers real-time visualization, filtering, and analytics tools that support data-driven governance and resource optimization.

II. PROBLEM STATEMENT

Urban cities face persistent challenges related to infrastructure degradation and public safety management. Road surface damages such as potholes and open manholes contribute significantly to accidents, traffic congestion, and economic losses. Despite these issues, many municipal systems rely heavily on manual inspections or outdated complaint portals that lack automation and real-time responsiveness.

Traditional inspection models are labor-intensive and unable to scale efficiently across rapidly expanding urban territories. Absence of structured data collection and centralized analytics prevents authorities from identifying high-risk clusters or prioritizing repairs effectively. Delays in addressing minor damages often escalate into serious infrastructure failures, resulting in significant costs to both citizens and municipalities.

Another major issue is the communication gap between citizens and local authorities. Many individuals encounter hazards daily but lack a transparent and technology-driven platform for structured reporting. Existing complaint systems rarely provide feedback loops, reducing citizen trust and long-term participation. Additionally, there is limited availability of geospatial visualization tools that combine infrastructure hazards with perceived safety risks, preventing commuters from making informed route decisions.

The core challenge lies in the absence of an integrated AI-powered system capable of detecting hazards automatically, enabling structured citizen participation, analyzing real-time geospatial data, and supporting data-driven governance at scale.

III. RELATED WORK

Urban infrastructure monitoring has emerged as a critical research domain due to increasing population density and rapid city expansion. Traditional inspection mechanisms rely on manual surveys and complaint-based reporting systems, which are often time-consuming and inefficient. Early studies focused on sensor-based detection systems where accelerometers and vibration sensors were installed in vehicles to identify road irregularities [1]. While these systems provided measurable accuracy, their large-scale deployment required significant infrastructure investment, limiting scalability in developing regions.

With the advancement of image processing techniques, researchers shifted toward vision-based pothole detection methods. Classical computer vision algorithms such as Canny edge detection, threshold segmentation, and morphological filtering were applied to identify irregular road textures [2]. However, these methods were highly sensitive to environmental factors such as lighting variations, shadows, and camera angles, resulting in inconsistent performance across diverse urban settings.

The emergence of deep learning significantly transformed object detection capabilities. Convolutional Neural Networks (CNNs) demonstrated superior performance in feature extraction and classification tasks. Models such as Faster R-CNN, SSD (Single Shot Detector), and YOLO (You Only Look Once) became widely adopted for real-time object detection applications [3]. Among these, YOLO-based architectures achieved a strong balance between detection speed and accuracy, making them suitable for mobile and edge-device deployment.

Recent research utilizing YOLOv5 and YOLOv8 architectures has shown improved bounding box precision and reduced inference latency [4]. These improvements make real-time road anomaly detection feasible on smartphones.

Additionally, TensorFlow Lite and model quantization techniques have enabled on-device inference, reducing dependency on high-bandwidth cloud connectivity.

Beyond road damage detection, studies on participatory sensing and crowdsourcing platforms have highlighted the importance of citizen engagement in urban governance [6]. Applications like Safetipin integrate geolocation tagging and crowd-based validation to generate safety indices [8]. However, most existing systems either focus solely on reporting or solely on AI detection, lacking integration between automated analysis and civic engagement. City Guardian addresses this gap by combining YOLOv8-based detection, Firebase cloud services, geospatial visualization, and a reward-driven citizen participation model into a single scalable framework.

IV. SYSTEM ARCHITECTURE

The City Guardian system architecture is designed using a modular and scalable approach to ensure efficient integration between mobile interfaces, AI models, and cloud infrastructure. The system consists of four primary layers: User Interaction Layer, AI Processing Layer, Cloud Backend Layer, and Administrative Visualization Layer.

A. *User Interaction Layer*

A Flutter-based cross-platform mobile application provides the citizen-facing interface. Users can capture hazard images, report unsafe areas, view interactive hazard maps, and receive real-time proximity alerts. GPS services are integrated to capture precise location data during reporting. The interface is designed to minimize input complexity while ensuring structured, high-quality data collection from non-technical users.

B. *AI Processing Layer*

A YOLOv8 deep learning model handles object detection within the AI Processing Layer. The model is trained on labeled datasets containing potholes, open manholes, and related road anomalies. During inference, uploaded images are processed to detect hazards, generate bounding boxes, calculate confidence scores, and classify severity levels as Low, Medium, or High. For mobile optimization, TensorFlow Lite enables lightweight on-device inference, reducing latency and energy consumption without requiring constant server connectivity.

C. *Cloud Backend Layer*

The Cloud Backend Layer is powered by Firebase services: Firestore provides a NoSQL cloud database for structured report storage with real-time synchronization; Firebase Storage handles image uploads; Firebase Authentication manages user login via Google OAuth and email; and Firebase Cloud Messaging delivers push notifications to users and administrators when new hazards are reported or proximity thresholds are crossed.

D. Administrative Visualization Layer

A React.js-based web dashboard provides municipal authorities with a comprehensive Control Center. The dashboard displays live metrics including total users, total reports, pending reports, and fixed potholes. It provides report trend charts, system health analytics, and a severity-filtered Pothole Reports management table. Administrators can verify, update, or remove reports directly through the portal, supporting transparent and data-driven governance workflows.

V. METHODOLOGY

The development of City Guardian involves a multi-layered methodology integrating real-time data processing, AI-powered hazard detection, cloud-based services, and community-driven feedback. The following subsections detail each major functional component.

A. Data Acquisition and User Interaction

Users interact with the system through the mobile application, providing camera input, GPS location, and optional voice or text descriptions. The app fetches the user's real-time GPS coordinates at the moment of capture, ensuring automatic geotagging without manual input. Users can either view the current Danger Map displaying live hazard overlays, or actively report a hazard by submitting a captured image through the Report Incident interface.

B. On-Device AI Hazard Detection Pipeline

The detection pipeline follows six steps: (1) the user opens the app and captures an image via the camera interface; (2) the image is passed to the on-device YOLOv8 model via TensorFlow Lite; (3) the model outputs bounding box coordinates, class label (pothole or manhole), and a confidence score; (4) severity is estimated from the bounding box area ratio relative to the total image area — small bounding boxes indicate Low severity, medium indicate Medium, and large bounding boxes indicate High severity; (5) the report is submitted with GPS coordinates, timestamp, and severity label; (6) the backend stores the report in Firestore and triggers proximity alerts for nearby users via Firebase Cloud Messaging.

C. Invalid Report Handling and Data Quality

To maintain data integrity, City Guardian implements a three strike penalty system. If the YOLOv8 model detects no pothole or manhole in a submitted image, an Invalid Report Warning dialog is displayed, informing the user that no hazard was detected and warning that a third consecutive invalid submission will result in a 50-point Guardian Points deduction. This mechanism discourages careless or fraudulent reporting while allowing genuine users to correct mistakes without immediate penalty.

D. Safe Route and Proximity Alert System

The Safe Route Engine combines graph-based routing algorithms with hazard weighting to generate alternate navigation paths that avoid known hazard clusters. Push notifications are triggered automatically when a user approaches within a configurable radius of a verified hazard. The Hazard Nearby alert displays the severity level and distance, enabling drivers and pedestrians to take immediate precautions. Historical hazard data is also used to generate visual heatmaps on the interactive map interface.

E. Community Verification and Reward System

Reported hazards can be verified or dismissed by other users through crowd-voting logic. Verified reports gain trust weight, improving model feedback and long-term data quality. Repeated dismissals result in automatic removal of low-confidence reports. Verified submissions earn +10 Guardian Points and crowd verifications earn +5 points, tracked in real time via Firestore. Points are displayed on user profiles and power a global leaderboard, promoting sustained community competition and civic engagement.

F. Technology Stack

The complete technology stack spans mobile, AI, cloud, and web layers. Flutter (Dart) enables cross-platform Android and iOS deployment from a single codebase. YOLOv8 (Ultralytics) provides state-of-the-art real-time object detection, with TensorFlow Lite handling on-device inference. The backend infrastructure runs on Firebase: Firestore for database, Storage for images, Authentication for user management, and Cloud Messaging for notifications. The administrative portal is built with React.js and the Firebase Admin SDK, while Google Maps API handles geolocation and map rendering across all layers.

VI. RESULTS AND DISCUSSION

The system was evaluated through real-world testing across urban road environments. The YOLOv8 model demonstrated strong detection performance, achieving a maximum confidence score of 0.95 on pothole images under standard conditions. The model correctly classified potholes across varying lighting, road surface types, and image angles, demonstrating robustness suitable for practical deployment.

The mobile application successfully demonstrated all core features during live testing: proximity hazard alerts were triggered accurately within configured distance thresholds; the invalid report penalty system correctly identified non-pothole images and displayed the warning dialog with the Guardian Points balance; and user Guardian Points were tracked and updated in real time. The admin dashboard provided live visibility into report trends, severity distributions, and system health metrics across the testing period.

A. Mobile Application Screenshots

Fig. 1 shows the Login screen featuring Google OAuth and email-based authentication with the City Guardian branding. Fig. 2 illustrates the home map screen showing the interactive Google Maps interface with a Hazard Nearby alert triggered when a low-severity pothole was detected within 3 meters of the user's current location. Fig. 3 shows the Invalid Report Warning dialog displayed when no pothole was detected in the submitted image, warning the user of the pending penalty. Fig. 4 shows another instance of the Hazard Nearby proximity alert on the map interface.

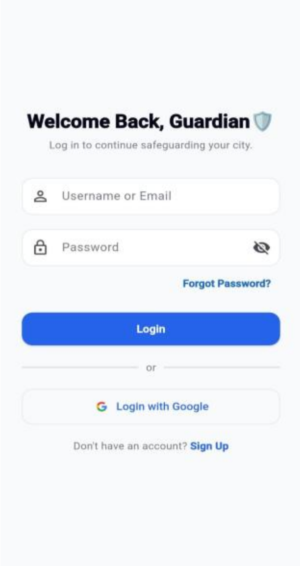

Fig. 1. Login Screen

Fig. 2. Map with Hazards

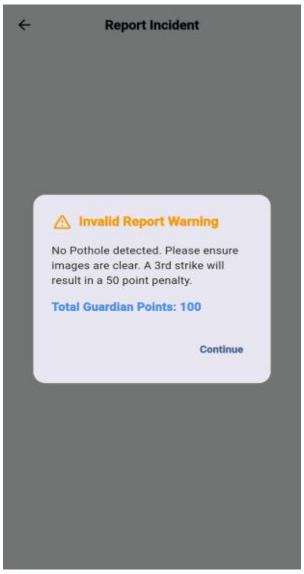

Fig. 3. Invalid Report Warning

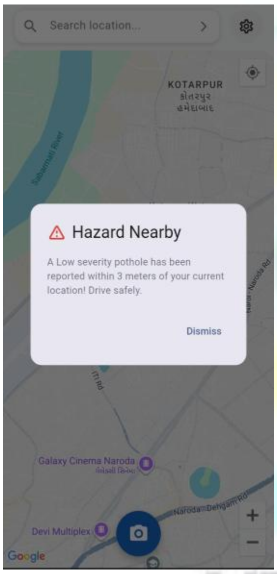

Fig. 4. Hazard Nearby Alert

B. *Admin Dashboard Screenshots*

The administrative portal provides municipal officers a comprehensive real-time Control Center. Fig. 5 shows the main dashboard with live statistics: 5 total users, 12 total reports, 0 pending reports, and 1 fixed pothole, with percentage change indicators. Fig. 6 illustrates the report trend line chart tracking cumulative reports from October to March, alongside the System Health donut chart showing 8% resolved versus 92% rejected or pending. Fig. 7 displays the Pothole Reports management table with severity badges (Low, Medium, High), reporter details, dates, verification status, and action controls.

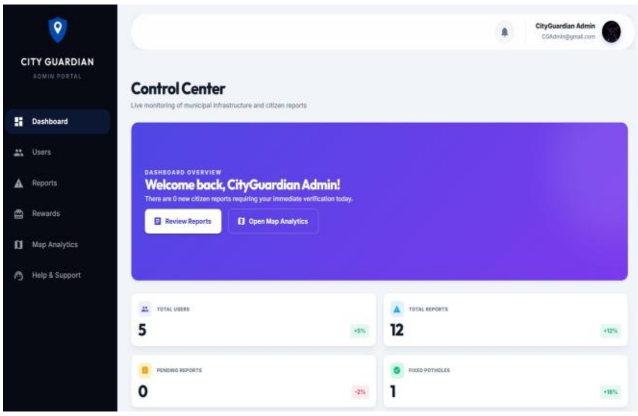

Fig. 5. Admin Control Center Dashboard

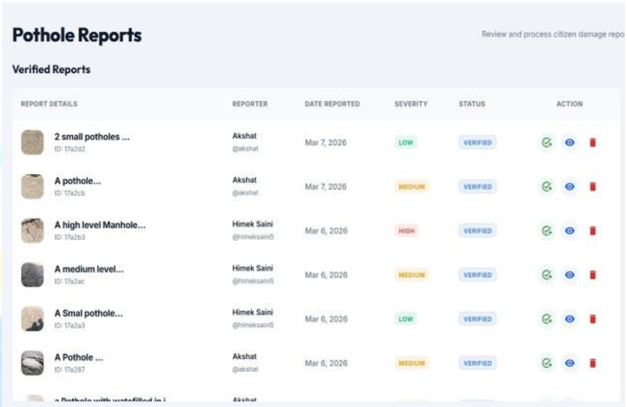

Fig. 6. Pothole Reports Management Table

VII. CONCLUSIONS AND FUTURE WORK

City Guardian successfully demonstrates how artificial intelligence and mobile-based technologies can be applied to solve practical urban infrastructure and safety challenges. The system integrates real-time hazard detection, geolocation services, cloud storage, and citizen participation into a unified digital platform designed for smart governance.

Through the implementation of a YOLOv8-based object detection model, the application is capable of identifying road anomalies such as potholes and open manholes with a confidence score of 0.95. The automated detection process minimizes dependency on manual inspections and accelerates the reporting workflow. By storing structured data with timestamps and GPS coordinates, the platform ensures better Tracking, prioritization, and transparency in maintenance operations.

The integration of unsafe area reporting further strengthens the platform's capability to address broader public safety concerns. Safety indices and heatmaps provide meaningful insights for both commuters and municipal authorities, enabling proactive intervention. The reward-based engagement mechanism promotes continuous citizen involvement and enhances the authenticity of collected data, while the administrative dashboard supports data-driven governance and resource optimization.

Several enhancements are planned for future development:

(1) Predictive Risk Modeling using historical and environmental data patterns to forecast hazard-prone zones; (2)

IoT Integration with smart streetlights, traffic sensors, and surveillance systems for automated detection; (3) Expanded Hazard Categories including road cracks, waterlogging, air pollution hotspots, and traffic congestion; (4) Advanced Route Optimization using intelligent routing algorithms considering hazard severity and user preferences; (5) Federated Learning for privacy-preserving model updates across distributed devices; (6) Automated Escalation mechanisms that forward unresolved reports to higher authorities after defined time thresholds; (7) Multi-Language and Regional Support to improve accessibility across diverse user populations; and (8) Large-Scale Pilot Deployments in multiple cities to evaluate scalability and real-world performance metrics.

REFERENCES

[1]. Al-Saad, M., et al. (2024). Vision-Based Pothole Detection Using Computer Vision and Machine Learning. Sensors, MDPI.

[2]. Kumar, A., et al. (2021). Real-Time Pothole Detection Using YOLO Architecture. arXiv preprint arXiv:2107.06356.

[3]. Zhang, Y., et al. (2022). Computer Vision Techniques for Road Damage Detection: A Comparative Review. arXiv:2204.13590.

[4]. Ultralytics. (2023). YOLOv8: State-of-the-Art Real-Time Object Detection Model. Ultralytics Documentation.

[5]. Sharma, P., et al. (2024). Smart Street Lighting Using IoT and LoRa Technology. International Journal for Research in Applied Science and Engineering Technology.

[6]. Silva, B., et al. (2022). Crowdsourcing Applications in Smart City Infrastructure Management. arXiv:2204.05421.

[7]. Verma, S., et al. (2022). Crowdsourcing for Urban Planning and Infrastructure Monitoring. Sustainability, MDPI.

[8]. Safetipin Initiative. (2022). Crowdsourcing Women's Urban Safety Experiences for Safer City Planning.

[9]. World Bank Group. (2022). Perceived Risk and Urban Safety: A Research Report on Infrastructure Challenges in Developing Cities.

[10]. Patel, K., et al. (2024). Crowdsourcing Applications for Monitoring the Urban Environment and Infrastructure. ResearchGate.