International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI:10.55041/ISJEM05843
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

An AI-Based Real-Time Road Safety Monitoring System for Smart Cities

Dr. K. Satyam[1], R JayaChandra Reddy[2]

[1]Associate Professor, Department of MCA, Annamacharya Institute of Technology & Sciences, Tirupati, Andhra Pradesh, India.

[2] Post Graduate, Department of MCA, Annamacharya Institute of Technology & Sciences, Tirupati, Andhra Pradesh, India.

Abstract

Due mostly to slow emergency response and delayed hazard detection, traffic accidents continue to rank among the world's leading causes of mortality. Conventional road monitoring systems frequently fall short of offering real-time and intelligent analysis of dynamic traffic conditions since they mainly rely on manual surveillance and static rule-based methods. Vigilant RoadShield, an AI-driven real-time road hazard identification framework developed utilising the YOLOv8 deep learning model combined with a Django-based web monitoring dashboard, is presented in this study in order to overcome these constraints.

Using live video or picture inputs, the suggested system automatically identifies road dangers like cars, obstructions, and possible collision situations by utilising computer vision algorithms. For high-speed object identification with optimal accuracy, YOLOv8 Nano is used, guaranteeing low latency performance appropriate for real-time deployment.

An interactive online interface that allows for ongoing monitoring and visualisation of risks that have been discovered is easily integrated with the detection outputs.

Strong detection capability with high precision and recall while keeping effective frame processing speed is demonstrated via experimental evaluation. The technology offers a workable and scalable solution for smart city infrastructures and intelligent transportation systems. Vigilant RoadShield uses real-time AI-powered monitoring to improve road safety and lower accident-related risks by facilitating automatic hazard identification and quicker response mechanisms.

Keywords

Road Hazard Detection; YOLOv8; Real-Time Object Detection; Intelligent Transportation Systems; Deep Learning; Computer Vision; Smart City Surveillance; Road Safety Analytics.

I. Introduction

Because it facilitates mobility, trade, and economic growth, road transport is a crucial part of contemporary society. However, the likelihood of traffic accidents and dangerous circumstances has increased dramatically due to the growing number of vehicles on the road. Road traffic injuries continue to rank among the world's leading causes of mortality, according to the World Health Organization, underscoring the critical need for more efficient safety measures. In addition to human mistake, delayed hazard detection and a lack of intelligent monitoring systems that can react instantly are also major causes of accidents.

CCTV cameras under human supervision are the main component of traditional road surveillance systems. Although these systems offer visual coverage, they are frequently constrained by human tiredness, delayed response, and scalability limitations, and they lack automated decision-making capabilities. Conversely, rule-based automated systems find it difficult to adjust to dynamic traffic conditions where erratic events are common. Adopting cutting-edge technical solutions that can automatically assess road conditions and identify possible hazards without human intervention is crucial as metropolitan areas continue to grow and traffic congestion rises.

New avenues for intelligent transportation systems have been made possible by recent advancements in computer vision and deep learning. Real-time object identification models have proven to be highly effective at quickly identifying items in complicated images, especially those belonging to the YOLO (You Only Look Once) family. Among these, YOLOv8 outperforms its predecessors in terms of detection accuracy, computational complexity, and inference speed. Because of these features, it is ideal for real-time road monitoring applications where speed and accuracy are crucial.

In this regard, the suggested solution, Vigilant RoadShield, presents a framework for real-time road hazard monitoring and identification powered by AI.

To enable continuous surveillance and visualisation of recognised risks, the system combines a web dashboard powered by Django with an object identification engine based on YOLOv8. The framework facilitates quicker awareness and prompt response by identifying vehicles, obstructions, and other risk-related items on the road by processing live video streams or picture inputs. The suggested method seeks to improve road safety, promote smart city projects, and offer a scalable basis for next intelligent transportation systems by integrating deep learning and web-based monitoring.

II. Objectives of the Study

Designing and creating an intelligent real-time road hazard detection system that can increase traffic safety through automated visual analysis is the main goal of this project. The goal of the project is to use deep learning methods to recognise road hazards from live video or image inputs, including cars, obstructions, and potentially hazardous circumstances. The system aims to achieve high detection accuracy while keeping low latency appropriate for real-time deployment by incorporating a sophisticated object detection model, notably YOLOv8.

Integrating the detection engine with a Django-based web dashboard to create a scalable and user-friendly monitoring framework is another important goal. Continuous surveillance and usefulness are made possible by this integration, which guarantees that identified hazards are effectively and clearly visualised.

In order to guarantee that the suggested system can function well in dynamic traffic settings, the research also focuses on performance optimisation in terms of precision, recall, and frame processing speed.

By offering a foundation that may be expanded for smart city applications, the study also seeks to support intelligent transportation systems. Vigilant RoadShield aims to improve situational awareness, facilitate quicker response mechanisms, and eventually contribute to safer road conditions by automating hazard detection and lowering dependency on manual monitoring.

III. Problem Statement

Because cognitive real-time analysis is lacking, road accidents and hazardous incidents continue to rise despite the widespread deployment of surveillance cameras and traffic monitoring systems. Human supervision, which is frequently ineffective, uneven, and prone to delayed response, is a major component of traditional monitoring frameworks. Observing several video streams by hand may cause important events to be missed, particularly in crowded cities where constant attention is necessary. Furthermore, rule-based automated systems are not flexible enough to react to changing and uncertain road conditions.

Instead of proactive danger detection, the majority of current transportation safety systems concentrate on post-incident analysis. The severity of accidents increases and emergency response times are greatly impacted by this delay in recognising risky situations. Additionally, a lot of traditional technologies cannot effectively handle massive amounts of live video data in real time and are not scalable for smart city infrastructure. An AI-driven system that can automatically identify traffic dangers with high accuracy and low latency is therefore desperately needed. The system must intelligently evaluate visual inputs, recognise any hazards like cars, obstructions, or unusual road conditions, and display the findings via a centralised monitoring interface. The proposed Vigilant RoadShield system, which seeks to

ISSN: 2583-6129
DOI:10.55041/ISJEM05843

improve road safety through real-time deep learning-based hazard detection and intelligent monitoring capabilities, is built on filling this gap.

IV. System Architecture

Vigilant RoadShield's architecture uses an integrated deep learning and web-based monitoring system to enable real-time road hazard detection. Input acquisition, preprocessing, detection engine, danger interpretation, backend processing, and dashboard visualisation are the six main layers that make up the system's modular architecture. Scalability, effectiveness, and smooth interaction between the AI model and the web application are guaranteed by this tiered method.
The first part is the input acquisition layer, which records live video streams or image inputs from uploaded sources or security cameras. The preprocessing module receives these visual inputs and applies formatting, normalisation, and scaling to conform to the input specifications needed by the detection model. This stage guarantees optimal processing speed and consistent inference performance.

The YOLOv8-powered object detection engine is the system's key component. Because of its high-speed inference capability and lightweight design, the YOLOv8 Nano model is used for real-time applications. After processing each frame, the model produces bounding boxes and confidence scores for any items it detects, including cars, obstructions, and other road-related dangers. To reduce redundant detections and increase accuracy, non-maximum suppression is used.
The hazard interpretation layer classifies detected objects and assesses their significance for traffic safety after detection. The Django framework-implemented backend layer then receives the processed detection findings.

Django controls request-response cycles, keeps pertinent data in the SQLite database for logging and monitoring, and facilitates communication between the detection module and the user interface.
Lastly, the detection results are shown in an interactive online interface via the dashboard visualisation layer. Users can monitor threats in real time by using the dashboard, which shows annotated photos or video frames with highlighted bounding boxes. A comprehensive end-to-end solution that can assist smart city safety programs and intelligent transportation is created by integrating deep learning-based detection with a web-based monitoring system.

V. Output

Fig: Output-1

This picture shows how the system can identify speedbreakers and other road infrastructure features in real time. The dashboard is updated and an automated alert is generated by the system upon detection. This attests to the framework's ability to warn drivers about potential hazards before they happen.

Fig: Output-2

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI:10.55041/ISJEM05843
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

The suggested framework's real-time pothole detecting capability is demonstrated in this graphic. When the technology detects abnormalities in the road surface, it promptly sounds a safety alert. Particularly in low visibility situations, this kind of identification is essential for preventing vehicle damage and lowering the danger of accidents.

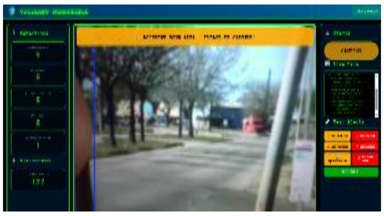

Fig: Output-3

The system's capacity to recognise contextual hazards is shown in this screenshot. The framework can identify accident-prone areas and alert users accordingly, rather than only identifying physical things. In high-risk areas, this function facilitates safer navigation and improves situational awareness.

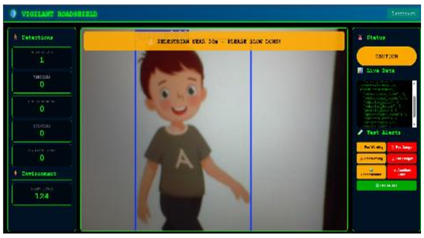

Fig: Output-4

The YOLOv8 model's pedestrian detection functionality is shown in this graphic. The accuracy of the model's detection is validated by the bounding box and confidence score. Because they allow for early driver response, real-time pedestrian alerts greatly increase road safety.

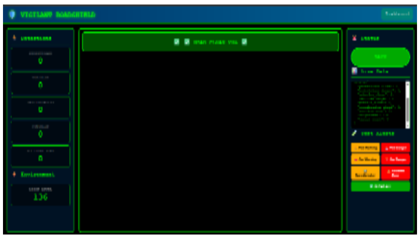

Fig: Output-5

The system's capacity to discern between dangerous and safe road conditions is seen in this screenshot. The dashboard verifies a safe driving condition when no dangers are identified. By doing this, false warnings are avoided and alerts are only activated when required.

VI. Results and Analysis

Real-time testing with live image and video inputs was used to assess the effectiveness of the suggested Vigilant RoadShield framework. The system's capacity to produce contextual alerts, dynamically update the monitoring dashboard, and reliably identify road risks was evaluated. The YOLOv8 Nano model's integration with the Django-based interface made it possible to detect and visualise dangers in a variety of circumstances with ease. The system effectively recognised a number of road-related dangers during experimental evaluation, such as potholes, speedbreakers, pedestrian presence, and accident-prone areas. The model produced bounding boxes with confidence scores and sent out the proper alert messages upon detection.

For example, the system indicated the presence of a speedbreaker and suggested slowing down when it identified one. In a similar vein, the system's capacity to identify surface imperfections in real time was demonstrated by the prompt caution alerts that followed pothole identification.

ISSN: 2583-6129
DOI:10.55041/ISJEM05843

Experiments on pedestrian detection confirmed the responsiveness of the model. The system updated the detection counter and produced a proximity-based alert message whenever a pedestrian came within close proximity. In urban traffic situations where pedestrian safety is a major concern, this functionality is especially crucial. The dashboard showed the quantity of entities found and the associated threat level, while the bounding box visualisation verified precise object localisation.

The approach showed contextual awareness by detecting accident-prone regions in addition to object-based risks. The system raised the threat count and sent out a warning when it recognised such a situation. This multi-class detection capabilities demonstrates the deep learning model's resilience and flexibility in handling a variety of traffic situations.

The dashboard interface was crucial in improving monitoring effectiveness and usability. Counts of detected cars, pedestrians, speedbreakers, potholes, and accident zones were dynamically shown in the left panel. In order to provide more contextual information that could affect detection reliability, the environmental module also tracked light intensity levels. Structured detection logs, containing hazard kind, brightness levels, and calculated threat count, were displayed in the live data section on the right panel.

The system's capacity to integrate visual detection with analytical interpretation is validated by its organised output. According to performance studies, the YOLOv8 Nano model preserved good detection accuracy while maintaining a steady inference speed appropriate for real-time operation. The algorithm successfully distinguished between conditions that were dangerous and those that weren't. The dashboard ensured that notifications were generated only when necessary by displaying a "SAFE" status and a "ROAD CLEAR" indication when there were no dangers. By doing this, false positives are decreased and needless driver distraction is avoided. Overall, the experimental findings support the suggested framework's ability to provide intelligent, real-time road hazard monitoring and detection. Proactive safety management and enhanced situational awareness in dynamic traffic settings are two benefits of the system's shown practical capability for incorporation into smart transportation infrastructures.

References

[1] World Health Organization, *Global Status Report on Road Safety 2023*. Geneva, Switzerland: WHO, 2023.

[2] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.

[3] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," *arXiv preprint arXiv:2004.10934*, 2020.

[4] G. Jocher et al., "Ultralytics YOLOv8: State-of-the-Art Real-Time Object Detection," Ultralytics, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[5] S. S. Sengar, R. K. Gupta, and A. Kumar, "Real-Time Road Accident Detection Using Deep Learning," *IEEE Access*, vol. 9, pp. 123456–123468, 2021.

[6] M. Teichmann, M. Weber, M. Zoellner, R. Cipolla, and R. Urtasun, "Multinet: Real-time Joint Semantic Reasoning for Autonomous Driving," in *Proc. IEEE Intelligent Vehicles Symposium*, 2018.

[7] Z. Zhang, Q. Liu, and Y. Wang, "Road Surface Damage Detection Based on Deep Convolutional Neural Networks," *Computer-Aided Civil and Infrastructure Engineering*, vol. 35, no. 7, pp. 680–698, 2020.