

PAVE: Pothole Avoidance and Vision Engine

Bhadra M U
Dept. of Computer Science and
Engineering
Vidya Academy of Science and
Technology
Thrissur, Kerala, India
bhadrannim@gmail.com

Harikrishna K S
Dept. of Computer Science and
Engineering
Vidya Academy of Science and
Technology
Thrissur, Kerala, India
harikrishnaa9241@gmail.com

Harikrishnan M A
Dept. of Computer Science and
Engineering
Vidya Academy of Science and
Technology
Thrissur, Kerala, India
harikuttan9306@gmail.com

Jerin V Brijesh
Dept. of Computer Science and
Engineering
Vidya Academy of Science and
Technology
Thrissur, Kerala, India
jerinvbrijesh11@gmail.com

Rejusha T R
Dept. of Computer Science and
Engineering
Vidya Academy of Science and
Technology
Thrissur, Kerala, India
rejusha821@vidyaacademy.ac.in

Abstract—The rapid growth of transportation systems has increased the demand for safer and smarter road infrastructure. Road defects such as potholes and improper recognition of traffic signals are major causes of accidents and traffic inefficiencies. This paper presents PAVE (Pothole Avoidance and Vision Engine), an intelligent system that integrates pothole detection and traffic signal recognition using computer vision and deep learning techniques. The proposed system utilizes a YOLOv8-based model to detect potholes and identify traffic signals such as red and green lights from images and real-time video streams. A dataset consisting of road images is used to train the model for accurate classification and detection. The system enables efficient processing and real-time performance. The integrated approach improves driver awareness by providing timely alerts for road hazards and traffic signals. This enhances road safety, reduces accidents, and supports advanced driver assistance systems (ADAS). Experimental results demonstrate that the system achieves high accuracy and performs effectively in real-time environments.

Keywords—pothole detection, traffic signal recognition, YOLOv8, computer vision, deep learning, ADAS

I. INTRODUCTION

In the modern era, road transportation has become one of the most essential components of daily life, facilitating the movement of people, goods, and services across cities and countries. With rapid urbanization and economic growth, the number of vehicles on the road has increased significantly, leading to higher traffic density and increased demand for efficient transportation systems. While these developments have improved connectivity and convenience, they have also introduced serious challenges related to road safety, traffic management, and infrastructure maintenance.

One of the major concerns affecting road safety is the poor condition of road infrastructure, particularly the presence of potholes. Potholes are surface defects that occur due to environmental factors such as heavy rainfall, water infiltration, temperature variations, and continuous vehicular

stress. Over time, these factors weaken the road surface, leading to cracks and eventually forming potholes. These defects can cause severe damage to vehicles, including tire punctures, suspension failure, and loss of vehicle control. In many cases, potholes are difficult to detect in advance, especially under low-light conditions, high-speed driving, or during adverse weather, making them a significant hazard for drivers.

In addition to road surface issues, improper recognition and interpretation of traffic signals also contribute significantly to road accidents. Traffic signals are fundamental components of road safety systems, as they regulate vehicle movement and ensure smooth traffic flow at intersections and pedestrian crossings. However, drivers may sometimes fail to notice or correctly interpret these signals due to distractions, poor visibility, unfavorable weather conditions, or lack of attention. This can result in violations such as running red lights, sudden braking, or collisions, thereby increasing the likelihood of accidents.

According to recent transportation studies, a considerable percentage of road accidents in developing countries are caused by infrastructure-related issues and human error. The combination of undetected road hazards and missed traffic signals creates a dangerous driving environment, particularly in densely populated urban areas. Traditional methods of road inspection and traffic monitoring are often manual, time-consuming, and inefficient, making it difficult to ensure real-time safety for drivers.

In recent years, advancements in Artificial Intelligence (AI), Machine Learning (ML), and Computer Vision have provided new opportunities to address these challenges. Computer vision techniques enable machines to interpret and analyze visual data from images and videos, making it possible to detect objects and patterns automatically. Deep learning models, especially Convolutional Neural Networks (CNNs), have demonstrated exceptional performance in tasks

such as image classification, object detection, and scene understanding.

Among various object detection algorithms, the YOLO (You Only Look Once) family of models has gained significant attention due to its ability to perform real-time detection with high accuracy. The latest version, YOLOv8, offers improved performance, faster processing speed, and better generalization capabilities compared to its predecessors. These features make it highly suitable for applications that require real-time decision-making, such as autonomous driving and driver assistance systems.

Motivated by these technological advancements, this paper proposes PAVE (Pothole Avoidance and Vision Engine), an intelligent vision-based system designed to enhance road safety by integrating pothole detection and traffic signal recognition into a single framework. The system utilizes a YOLOv8-based deep learning model to detect potholes and classify traffic signals from real-time video input captured by a camera. By combining these functionalities, the system provides a comprehensive solution for identifying both road hazards and traffic conditions simultaneously.

The proposed system operates by capturing video frames, processing them using computer vision techniques, and applying a trained deep learning model to detect objects of interest. Once a pothole or traffic signal is detected, the system generates appropriate outputs such as visual alerts or driving commands. This enables drivers or autonomous systems to take necessary actions, such as slowing down, stopping, or steering away from hazards.

The primary objective of this work is to develop a reliable and efficient system that can assist drivers in real-time decision-making and reduce the risk of accidents caused by road hazards and traffic signal misinterpretation. Furthermore, the system aims to contribute to the development of Advanced Driver Assistance Systems (ADAS) and support the implementation of intelligent transportation systems in smart cities.

In addition to improving driver safety, the proposed system can also assist government authorities and infrastructure agencies in monitoring road conditions and identifying damaged areas that require maintenance. By automating the detection process, the system reduces the dependency on manual inspection and enables faster response to infrastructure issues.

In conclusion, the integration of deep learning and computer vision techniques in transportation systems has the potential to significantly improve road safety and efficiency. The PAVE system represents a step towards achieving this goal by providing a practical and scalable solution for real-time pothole detection and traffic signal recognition. The following sections of this paper discuss the related work, methodology, implementation, results, and future scope of the proposed system in detail.

II. RELATED WORKS AND CRITICAL ANALYSIS

A comprehensive literature survey was conducted focusing on recent advancements in computer vision for autonomous

driving. The survey highlights a growing, yet highly compartmentalized, adoption of YOLO-based architectures for identifying either road anomalies or regulatory infrastructure.

Recent studies have heavily explored YOLOv8 for infrastructure monitoring. For example, Wijaya et al. (2024) [5] explored infrastructure monitoring through pothole detection using YOLOv8 trained on an open dataset, proving the model's capability to identify pavement defects. Similarly, Zhang et al. (2024) [1] introduced Improved YOLOv8-CBAM, utilizing a Convolutional Block Attention Module to dynamically adjust the model's focus to critical spatial regions of the road. While these architectural modifications demonstrate high recognition accuracies (exceeding 99%), they introduce significant computational overhead. The reliance on heavy attention modules makes these models difficult to deploy on ultra-constrained edge devices operating without cloud connectivity. Furthermore, these studies relegate detection strictly to road surface anomalies, ignoring the broader driving environment.

Conversely, parallel research has focused entirely on traffic signal perception. Liu et al. (2024) [2] introduced RCW-YOLO, utilizing Receptive-Field Convolutional Blocks to boost TT100K traffic sign dataset accuracy by 3.5%, while Biswas et al. (2023) [4] successfully validated standard YOLOv8's efficacy for detecting erratic traffic signals on Indian roads. However, a critical limitation persists across this domain: these frameworks are inherently single-task oriented. An autonomous vehicle or ADAS relying on these models would require running multiple, heavy neural networks in parallel—one for potholes and one for traffic lights—which drastically increases inference latency and hardware costs.

The critical gap identified in current literature is the absence of a cohesive, multi-task perception engine optimized for edge-device deployment. Previous works treat infrastructure monitoring (potholes) and regulatory perception (traffic signals) as mutually exclusive domains.

The novelty of the proposed PAVE (Pothole Avoidance and Vision Engine) framework lies in bridging this gap. Rather than heavily modifying the network with computationally expensive attention modules, PAVE utilizes the unmodified, highly efficient YOLOv8 Nano architecture to unify both hazard detection and signal recognition into a single, concurrent pipeline. By achieving this dual-task perception within a lightweight framework, PAVE directly transitions academic infrastructure monitoring into a practical, real-time safety protocol for Advanced Driver Assistance Systems (ADAS).

III. METHODOLOGY

The proposed system, PAVE (Pothole Avoidance and Vision Engine), is designed to detect potholes and recognize traffic signals using deep learning and computer vision techniques. The system follows a structured pipeline consisting of data collection, preprocessing, model training,

object detection, and real-time decision-making. Each stage is carefully designed to ensure accuracy, efficiency, and real-time performance in road environments.

A. Overview of the System

The PAVE system operates as an intelligent vision-based pipeline that processes road images and video streams to detect hazards and traffic signals. The workflow begins with dataset preparation and ends with real-time detection and driving assistance outputs.

The system is divided into three major phases:

- Dataset preparation and preprocessing
- Model training using YOLOv8 Nano
- Real-time detection and control using video input

This modular architecture ensures scalability and easy integration with autonomous driving systems.

B. Data Collection

The primary dataset was aggregated from publicly available Kaggle repositories containing road potholes [7] and traffic signals [8]. These initially separate datasets were merged to create a unified training environment. The finalized data contains extensive variations in lighting, road textures, and object scales to ensure model robustness.

C. Data Preprocessing

To ensure high data quality and model consistency, several preprocessing steps are implemented:

- Image Resizing: All images are uniformly resized to a standard resolution of 640x640 pixels.
- Label Remapping: Class IDs are systematically remapped to prevent semantic collisions, establishing non-overlapping labels for 'potholes', 'green lights', and 'red lights'.
- Data Augmentation: Techniques such as mosaic augmentation, horizontal flipping, and brightness adjustments are used to improve the model's ability to generalize across environmental conditions.
- Data Splitting: The finalized dataset of 1,218 images is partitioned into an 80:20 ratio for training and validation.

These preprocessing techniques help improve generalization and reduce overfitting during training.

D. Justification of Model Selection

The selection of the YOLOv8 Nano (yolov8n.pt) architecture is primarily driven by the stringent computational constraints inherent to real-time ADAS. Traditional two-stage object detectors, such as Faster R-CNN, offer high localization accuracy but incur significant inference delays. Conversely, while single-stage detectors like SSD achieve real-time processing, they often exhibit degraded mean Average Precision (mAP) when detecting irregularly shaped hazards like potholes. YOLOv8 Nano resolves this tradeoff via an anchor-free detection head and an optimized C2f module. This significantly reduces the parameter count and computational footprint while maintaining robust gradient flow, providing

the equilibrium of ultra-low latency inference and high accuracy required for a dual-task engine

E. Model Training

The system uses the YOLOv8 Nano (yolov8n.pt) model for object detection due to its lightweight architecture and real-time performance capability.

Key training details include:

- Epochs: 100 training epochs
- Optimizer: The AdamW optimizer is used to tune hyperparameters like the learning rate.
- Loss Functions: A composite loss function including CIoU and Distribution Focal Loss (DFL) handles bounding box regression, while Varifocal Loss (VFL) is used for classification.

Training is performed using GPU acceleration (e.g., Google Colab) to speed up computation.

The model learns to identify patterns and features corresponding to potholes and traffic signals from the training dataset.

F. Object Detection and Classification

In this system, object detection and classification are performed using the YOLOv8 Nano model. The model identifies objects and categorizes them into three classes: Potholes, Red Lights, and Green Lights.

The anchor-free detection head directly predicts object locations and class labels, improving detection speed and efficiency. For each detected object, the model generates bounding boxes along with confidence scores, enabling accurate real-time detection.

G. Evaluation Metrics

To evaluate the performance of the models, several metrics are used:

- Precision: Measures the correctness of positive predictions.
- Recall: Measures the ability to identify all relevant instances.
- F1 Score: Harmonic mean of precision and recall.
- Mean Average Precision (mAP): Evaluates detection accuracy across different IoU thresholds

H. System Implementation

The system is implemented using Python and integrates the following technologies:

- Ultralytics/PyTorch: For building and training the YOLOv8 model.
- OpenCV: For real-time video processing, frame extraction, and visual rendering of bounding boxes.
- Matplotlib: For generating static and interactive data visualizations.

During execution, video frames are captured, processed by the trained model, and annotated with bounding boxes and labels.

I. Advantages of the Proposed Methodology

- Real-time detection using a lightweight model
- High accuracy in detecting traffic signals and potholes
- Improved generalization through data augmentation
- Scalable and modular design
- Suitable for deployment in edge devices and ADAS systems

J. Limitations

Despite its advantages, the system has certain limitations:

- Performance may degrade under poor lighting or extreme weather conditions
- Pothole detection is challenging due to irregular shapes and textures
- The system relies only on camera input; sensor fusion is not included
- The decision-making module is heuristic-based and can be further improved

K. System Architecture

The architectural flow of the PAVE framework is illustrated in Fig. 1. The system follows a modular architecture where input images are first preprocessed and then passed to the YOLOv8 Nano model for detection. The trained model processes real-time video frames and outputs detected objects along with their locations.

These outputs are further used to generate driving assistance actions such as stopping or steering adjustments, ensuring safer navigation in real-world conditions.

IV RESULTS AND EVALUATION

The performance of the proposed PAVE – Pothole Avoidance and Vision Engine is evaluated using the YOLOv8 Nano model on a structured dataset consisting of potholes and traffic signal images. The evaluation aims to measure how effectively the system can detect and classify road hazards and traffic lights in real-time scenarios. This section provides a detailed analysis of the model performance based on training behavior, evaluation metrics, and class-wise results.

A. Experimental Setup

The dataset used for training and evaluation consists of annotated images of potholes, red traffic lights, and green traffic lights collected from publicly available datasets. The dataset is divided into training and validation sets in an 80:20 ratio.

Before training, preprocessing steps such as image resizing, normalization, and data augmentation (including scaling and mosaic augmentation) are applied to improve model generalization. The YOLOv8 Nano model is trained for 100 epochs using GPU acceleration.




Fig. 1. System Architecture for PAVE

B. Evaluation Metrics

To ensure a comprehensive evaluation, multiple performance metrics are used:

- Precision: Precision indicates how many of the predicted positive instances are actually correct. High precision indicates fewer false positives.
- Recall: Recall measures the ability of the model to correctly identify all relevant instances. High recall indicates fewer missed detections.
- Mean Average Precision (mAP): Evaluates overall detection accuracy across different Intersection over Union (IoU) thresholds.
 - mAP@50: IoU threshold of 0.5
 - mAP@50–95: Average over multiple IoU thresholds
- Loss Functions:
 - Box Loss: Measures bounding box localization error
 - Class Loss: Measures classification error
- Distribution Focal Loss (DFL): Improves bounding box edge precision

These metrics together provide a complete understanding of detection accuracy and localization performance.

C. Model Performance and Analysis of Results

The YOLOv8 Nano model demonstrates robust quantitative performance across all evaluation metrics, as summarized in Table 1. By the 100th training epoch, the model achieves an overall mean Average Precision (mAP@50) of 0.826, indicating strong generalization capabilities without overfitting.

Class	Precision (P)	Recall (R)	mAP@50	mAP@50-95
Red Light	0.859	0.806	0.864	0.505
Green Light	0.804	0.779	0.829	0.478
Pothole	0.816	0.67	0.785	0.491
All Classes	0.826	0.751	0.826	0.491

Fig. 2. System Architecture for PAVE

As visualized in the training curves (Fig. 3), both Box Loss and Distribution Focal Loss (DFL) exhibit a consistent decrease, confirming accurate bounding box localization. Furthermore, the Class Loss demonstrates rapid convergence toward zero, which aligns with the high classification accuracy observed.

A detailed breakdown of class-wise performance reveals that traffic signal detection—particularly red lights—yields the highest accuracy. As illustrated by the confusion matrix in Fig. 4, there is an absence of misclassification between red and green lights, a critical safety parameter for autonomous control.

While pothole detection yields a comparatively lower recall (0.670) due to high physical variance in hazard shapes and textures, the precision remains strong (0.816). In safety-critical ADAS applications, a slightly higher false positive rate for hazards acts as a deliberate safety mechanism, encouraging cautious navigation rather than risking catastrophic missed detections.

D. Visualization of Results

The performance of the proposed system is visually represented using training graphs, confusion matrix, and detection outputs.

Fig. 3 shows the training curves of the model, including loss and evaluation metrics over epochs. The decreasing loss and increasing metric values indicate that the model is learning effectively and converging well.

Fig. 4 presents the confusion matrix, which illustrates the classification performance of the model. It shows that most predictions are correctly classified, with very few misclassifications between classes.

Fig. 5 demonstrates the detection of potholes in road images. The model successfully identifies potholes and draws bounding boxes around them.

Fig. 6 shows the detection of green traffic lights, where the model accurately recognizes and classifies the signal.

Fig. 7 illustrates the detection of red traffic lights, which is performed with high accuracy and reliability.

Overall, the visual results confirm that the model performs effectively in real-time detection of road hazards and traffic signals.

E. Error Analysis

Despite strong performance, some errors are observed:

- **Background Confusion:** Road textures similar to potholes lead to false positives.

- **Lighting Conditions:** Poor lighting or glare can affect traffic signal detection.
- **Small Object Size:** Distant or partially visible objects may be missed.

However, in safety-critical applications, false positives (e.g., detecting a pothole when none exists) are preferable over false negatives, as they promote cautious driving.

F. Discussion

The results demonstrate that the YOLOv8 Nano model provides an effective balance between accuracy and real-time performance. The model performs exceptionally well in detecting traffic signals, particularly red lights, which are critical for autonomous decision-making.

Although pothole detection shows slightly lower accuracy, it still maintains acceptable performance for real-world deployment. The model's lightweight architecture makes it suitable for edge devices and real-time applications such as Advanced Driver Assistance Systems (ADAS).




Fig. 3. Training Curves




Fig. 4. Confusion Matrix




Fig. 5. Pothole Detection Output




Fig. 6. Green Light Detection Output




Fig. 7. Red Light Detection Output

G. Summary

Overall, the proposed system achieves robust and reliable performance in detecting road hazards and traffic signals. The combination of efficient preprocessing, optimized training, and advanced deep learning techniques ensures accurate real-time detection.

The evaluation confirms that the system is suitable for deployment in intelligent transportation systems, contributing to improved road safety and smarter mobility solutions.

V. PERFORMANCE COMPARISON WITH EXISTING TECHNIQUES

To contextualize the efficacy of the proposed framework, Table 2 provides a comparative analysis between PAVE and recent state-of-the-art models deployed for road infrastructure and traffic signal perception.

Reference	Model Architecture	Primary Task	mAP@50
Ziqi Zhang et al.	YOLOv8-CBAM	Potholes	0.99
Jonathan Wijaya	YOLOv8	Potholes	0.7
Supratim Biswas	YOLOv8 CNN	Traffic Signals	N/A
Proposed (PAVE)	YOLOv8 Nano	Dual-Task	0.826

Fig. 8. Performance Comparison with Recent Frameworks

Note: Direct cross-dataset metric comparisons are inherently complex due to variations in testing environments. However, this comparison highlights the architectural trade-offs.

While highly modified architectures like the YOLOv8-CBAM achieve near-perfect precision on isolated tasks, they incur a computational penalty that actively hinders real-time edge deployment. Conversely, standard YOLOv8 models applied to single tasks yield lower mean Average Precision

than the proposed system. PAVE demonstrates highly competitive accuracy (0.826 mAP@50) while achieving something the benchmarked frameworks do not: unifying two disparate safety-critical tasks into a single, ultra-lightweight Nano architecture without the latency penalty of running parallel neural networks..

VI. CONCLUSION

This paper presented PAVE, an intelligent vision-based system for pothole detection and traffic signal recognition. By efficiently applying pre-processing techniques and training an unmodified YOLOv8 Nano model on a unified dataset, the system achieves reliable, concurrent detection of road hazards and regulatory signals. The integration of object detection with real-time video processing ensures low-latency operation suitable for edge computing.

Experimental evaluations confirm that PAVE provides a highly effective balance between detection accuracy and computational efficiency, mitigating the need for parallel single-task networks. The proposed approach actively enhances driver awareness, supports autonomous decision-making, and contributes to the broader implementation of smarter mobility solutions. Future work will focus on improving small-object detection at greater distances and integrating sensor fusion to maintain optimal accuracy under adverse, low-light weather conditions.

REFERENCES

- [1] Z. Zhang, B. Shao, L. Yang, and A. Qin, "Improved YOLOv8-CBAM enhanced pothole road recognition-based model," in Proc. 2024 IEEE 4th Int. Conf. Data Science and Computer Application (ICDSCA), 2024.
- [2] Y. Liu, Y. Shi, and X. Yuan, "RCW-YOLO: Improved traffic sign detection model based on YOLOv8 for complex scenes," in Proc. 2024 9th Int. Conf. Intelligent Informatics and Biomedical Sciences (ICIIBMS), 2024.
- [3] "What is YOLOv8: An in-depth exploration of the internal features of the next-generation object detector," arXiv preprint arXiv:2408.15857v1, 2024.
- [4] S. Biswas, S. Acharjee, A. Ali, and S. S. Chaudhari, "YOLOv8 based traffic signal detection in Indian road," in Proc. 2023 7th Int. Conf. Electronics, Materials Engineering & Nano-Technology (IEMENTech), 2023.
- [5] J. Wijaya, I. K. D. W. Kusuma, M. A. Ammaar, and A. A. S. Gunawan, "Enhancing infrastructure monitoring: Pothole detection in road images using YOLOv8 and open datasets," IEEE, 2024.
- [6] "Introduction to convolutional neural networks in deep learning," Analytics Vidhya, 2022. [Online]. Available: <https://www.analyticsvidhya.com/blog/2022/03/basic-introduction-to-convolutional-neural-network-in-deep-learning/>
- [7] "Potholes-detection-YOLOv8," Kaggle. [Online]. Available: <https://www.kaggle.com/datasets/anggadwisunarto/potholes-detection-yolov8>
- [8] "Traffic signs detection using YOLOv8," Kaggle. [Online]. Available: <https://www.kaggle.com/code/pkdarabi/traffic-signs-detection-using-yolov8>