Journal of Natural and Artificial Scientific Systems, 2026, 2(1): 11-26

YOLOv8-Based Framework for Vehicle Detection in Intelligent Traffic Monitoring Systems

Anas Tukur Balarabe*, Abdulrashid Sani and Andrew Raphael Ojobo
Received: 12 December 2025/Accepted: 12 March 2026/Published: 20 March 2026

Abstract: *Traffic congestion and road accidents remain major challenges in modern transportation systems, demanding accurate real-time traffic-monitoring solutions. Traditional surveillance methods and classical computer vision techniques are often limited by manual intervention, sensitivity to environmental variations, and poor scalability in complex traffic scenes. To address these limitations, this study proposes a YOLOv8-based framework for real-time vehicle detection in intelligent traffic monitoring systems. The proposed approach exploits YOLOv8 architecture to accurately detect vehicles of varying sizes under diverse traffic conditions, including occlusion, dense traffic, and illumination changes. A dataset comprising over 5,000 annotated traffic images representing urban and highway environments was used for training and evaluation. Model performance was assessed using standard object detection metrics, including Precision, Recall, F1-score, and mean Average Precision (mAP). Experimental results demonstrate strong detection accuracy, achieving an mAP@0.5 of 0.975 and a peak F1-score of 0.93, with stable convergence of training and validation losses. The results confirm the effectiveness of YOLOv8 as a lightweight and scalable solution for real-time intelligent traffic monitoring applications.*

Keywords: *YOLOv8; Vehicle Detection; Intelligent Transportation Systems; Traffic Monitoring; Object Detection.*

Anas Tukur Balarabe
Department of Computer Science, Faculty of Computing, Sokoto State University, Sokoto State, Nigeria
Email: anas.balarabe@ssu.edu.ng
https://orcid.org/0000-0001-7739-4616

Abdulrashid Sani
Department of Computer Science, Faculty of Computing, Sokoto State University, Sokoto State, Nigeria.
Email: abdulrashid.sani@ssu.edu.ng
https://orcid.org/0009-0000-9033-786X

Andrew Raphael Ojobo
Department of Computer Science, Faculty of Computing, Sokoto State University, Sokoto State, Nigeria
Email: andrewojobo@gmail.com

1.0 Introduction

Traffic congestion and road accidents remain persistent challenges in modern transportation systems, encouraging the development of reliable, automated, and intelligent traffic monitoring solutions (Premaratne *et al*., 2023; Sakr & El-afifi, 2023; Mystakidis & Koukaras, 2025). Early traffic surveillance systems relied heavily on manual observation and rule-based approaches, which were labour-intensive, error-prone, and unsuitable for large-scale or real-time deployment. As development and vehicle density increased, these traditional systems proved inadequate for handling complex traffic dynamics (Cao, 2025; Gheorghe & Soica, 2025; Zhou *et al*., 2024).
The evolution of computer vision techniques enabled partial automation through handcrafted feature extraction methods such as background subtraction, edge detection, and motion analysis. Although these approaches improved monitoring efficiency, their performance was highly sensitive to environmental variations, including illumination changes, weather conditions, occlusion, and background

complexity. These limitations became more noticeable in dense traffic scenes involving heterogeneous vehicle types and frequent interactions (Mittal, 2024). Deep learning-based object detection marked a major shift in traffic monitoring research by enabling end-to-end feature learning directly from data. Two-stage detectors such as R-CNN, Fast R-CNN, and Faster R-CNN demonstrated high detection accuracy but suffered from high computational cost, limiting their real-time applicability (Majin *et al*., 2024). This led to the adoption of one-stage detectors, such as the YOLO family, which prioritise detection speed while maintaining competitive accuracy (Sapkota *et al*., 2025).

Recent advancements in the YOLO series have significantly improved detection robustness and efficiency. YOLOv8 introduces an anchor-free design, improved backbone networks, a decoupled head structure, and efficient feature fusion mechanisms (Zhang & Zhang, 2025; Ayodeji *et al*., 2025). These enhancements make YOLOv8 particularly suitable for real-time vehicle detection in complex traffic environments (Prathima *et al*., 2025; Li *et* al., 2026; Usta *et al*., 2025).

A diverse dataset of annotated traffic images and videos was curated and augmented to capture variations in vehicle scale, appearance, lighting, and occlusion. The system is evaluated using standard metrics, including Precision, Recall, Intersection over Union, mean Average Precision, and detection speed, demonstrating strong performance and practical suitability for real-world intelligent transportation systems. Despite significant progress in traffic monitoring systems, achieving robust, real-time vehicle detection in complex and dynamic traffic environments remains a challenging problem. Issues such as occlusion, illumination variation, scale differences, and computational constraints continue to limit the performance of existing systems in practical deployments. Recent studies have demonstrated the effectiveness of deep learning-based approaches for vehicle detection and traffic monitoring. YOLOv5-based systems have shown promising real-time performance; however, they remain sensitive to occlusion, low-resolution imagery, and poor lighting conditions (Mehta & Shah, 2025; Nandini & Ravikanth, 2025).

Recent improvements using YOLOv8 have enhanced detection accuracy and efficiency, particularly through improved feature extraction and architectural optimization (Zhang & Zhang, 2025). Nevertheless, multimodal and ensemble-based approaches, although effective in low-light conditions, introduce increased computational cost and reduced scalability (Maurya *et al*., 2025).

Other studies integrating detection with tracking, UAV imagery, and traffic analytics have improved system capabilities but still face challenges related to generalization, environmental robustness, and deployment efficiency (Hegde *et al*., 2023; Bakirci, 2024; Asifur *et al*., 2025).

*et alet al*In a related effort, Nandini &. Ravikanth *et al*(2025r) explored traffic monitoring from a congestion analysis perspective by employing YOLOv5 for vehicle detection in static images. Their framework categorised traffic density into low, medium, and high levels and incorporated a graphical user interface to improve usability.

More recent studies have shifted attention toward YOLOv8 due to its architectural refinements and improved detection capability. Maurya *et al*. (2025) proposed a YOLOv8-based vehicle classification system that integrates RGB and thermal imagery to enhance performance under low-light conditions.

Extending YOLOv8 to broader traffic analysis tasks, Asifur *et al*. (2025) developed a real-time traffic monitoring framework that integrates vehicle detection, counting, speed estimation, and emission rate calculation.

Beyond frame-level detection, Sinojia *et al.* (2025) combined YOLO-based object detection with optical flow tracking and trajectory analysis to enhance temporal consistency in both urban and highway traffic scenes. More recently, Bakirci *et al.* (2025) investigated the use of YOLOv8 for traffic monitoring in intelligent transportation systems using unmanned aerial vehicles. However, despite these advancements, there is still a lack of lightweight and highly efficient YOLOv8-based frameworks that can consistently achieve high detection accuracy while maintaining low computational cost in diverse real-world traffic conditions. Most existing approaches either prioritize accuracy at the expense of computational efficiency or rely on complex multimodal architectures that limit practical deployment in resource-constrained environments. The aim of this study is to develop and evaluate a YOLOv8-based vehicle detection framework for intelligent traffic monitoring that achieves high accuracy, robustness, and real-time performance under complex traffic conditions.

The significance of this study lies in its potential to support real-world intelligent transportation systems by providing a lightweight and scalable vehicle detection framework. The proposed system can be applied in traffic flow analysis, congestion monitoring, accident prevention systems, and smart city infrastructure. It also contributes to research by demonstrating the effectiveness of YOLOv8 as a standalone, efficient model for real-time traffic surveillance without requiring multimodal inputs or complex ensembles.

Table 1: Summary of the Related Works

Model / Method	Data Type / Source	Main Focus	Key Contributions	Limitations
Improved YOLOv8 (YOLO-AL)	KITTI 2D, UA-DETRAC	Vehicle detection	Enhanced multiscale feature representation using adaptive scaling and ADown module; improved accuracy with reduced complexity	Limited robustness in unseen and highly diverse traffic scenarios (Zhang & Zhang, 2025).
YOLOv5	Highway traffic videos and images	Real-time detection and classification	Effective real-time performance under varying traffic density	Sensitive to occlusion, low resolution, and limited dataset diversity (Mehta & Shah, 2025).
YOLOv5	Static traffic images	Traffic density estimation	Automated congestion classification with user-friendly GUI	Poor performance under nighttime, occlusion, and illumination changes (Nandini &

				Ravikanth, 2025).
YOLOv8 + EfficientDet (Ensemble)	RGB and thermal images (FLIR)	Vehicle classification	Improved detection under low-light conditions using multimodal data	High computational cost; limited robustness under extreme environments (Maurya *et al.*, 2025).
YOLOv8	Traffic surveillance videos	Detection, counting, speed, and emissions	Integrated real-time traffic analysis and emission estimation	Scalability issues and high computational requirements (Asifur *et al.*, 2025).
YOLO + Optical Flow	Urban and highway videos	Detection and tracking	Improved temporal consistency using trajectory analysis	Increased system complexity and sensitivity to weather conditions (Hegde *et al.*, 2023).
YOLOv8	UAV-based aerial traffic images	ITS traffic monitoring	Flexible aerial monitoring with improved coverage	Performance degradation at high altitude and under dense occlusion (Bakirci, 2024).

2.0 Methodology

This section presents the methodological framework used to design, train, and evaluate the proposed YOLOv8-based traffic management detection system. The workflow progresses from system architecture design and data preparation to model training and quantitative performance evaluation. The methodology is designed to ensure reproducibility, scalability, and robustness of the proposed detection framework under diverse traffic conditions.

2.1 Architecture

The proposed traffic management detection system is designed as an end-to-end pipeline that integrates traffic data acquisition, preprocessing, YOLOv8-based inference, and result visualisation. *Fig.* 1 presents the system architecture, illustrating the flow from raw traffic images or video frames to final outputs containing detected and classified vehicles.

The system begins with the collection of traffic surveillance data from cameras deployed in urban or highway environments. The acquired data are preprocessed to standardise input dimensions and improve detection robustness. The processed images are then fed into the YOLOv8 detection network, which performs real-time vehicle localisation and classification. Finally, the detected vehicles are visualised using bounding boxes and class labels, providing interpretable outputs suitable for traffic analysis, congestion monitoring, and intelligent transportation system applications.

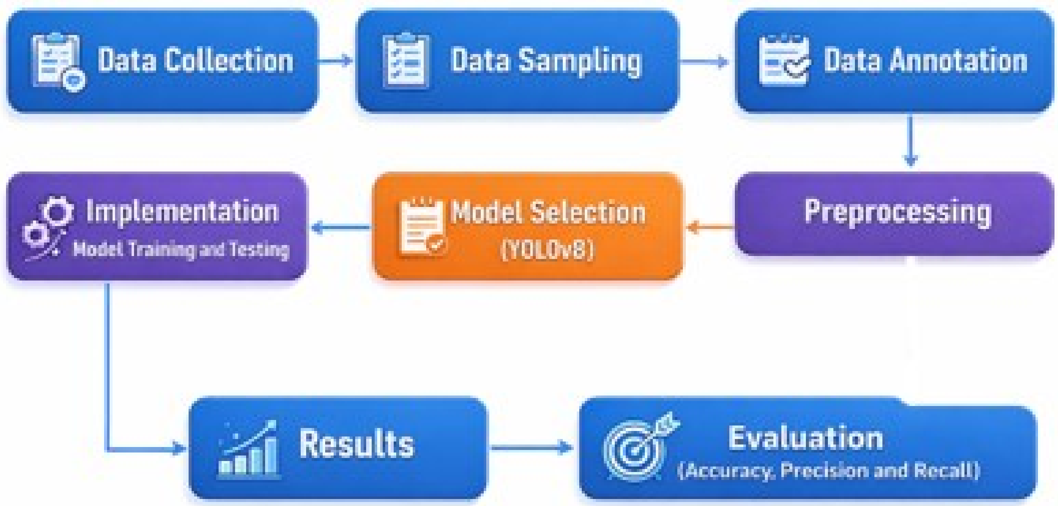

Fig. **1: Research Architecture**

1.1.1 Data Collection and Sampling

The dataset employed in this study comprises annotated traffic images and video frames that represent real-world traffic scenes characterised by diverse vehicle densities, camera viewpoints, lighting conditions, and degrees of occlusion. It includes multiple vehicle categories, namely cars, buses, trucks, and motorcycles, thereby reflecting the heterogeneity typically observed in both urban and highway traffic environments.

To ensure reliable model training and objective performance assessment, the dataset was partitioned into three mutually exclusive subsets: training, validation, and testing. The training subset was utilised for learning model parameters, the validation subset facilitated hyperparameter tuning and model selection, and the test subset was reserved exclusively for final performance evaluation. This sampling strategy minimises the risk of data leakage and enables an accurate assessment of the model's generalisation capability in previously unseen traffic scenarios.

The dataset consists of over 5,000 annotated traffic images obtained from the Roboflow Universe open-source computer vision platform, including publicly available collections such as Vehicles–OpenImages (Roboflow, 2023). The images capture a broad spectrum of global traffic conditions across urban and highway settings, incorporating variations in vehicle scale, camera perspective, illumination, and weather conditions. Vehicle instances were annotated into four classes: car, truck, bus, and motorcycle, using Label Studio. Additionally, data augmentation techniques, including mosaic blending and HSV colour space transformations, were applied to enhance dataset diversity and improve model robustness and generalisation. The methodology is designed to ensure reproducibility, scalability, and robustness of the proposed detection framework under diverse traffic conditions.

2.1.2 Preprocessing

Before training the YOLOv8 model, all traffic images and video frames underwent preprocessing to ensure consistency, improve training stability, and enhance detection accuracy. The preprocessing steps applied in this study are described as follows.

Frame extraction and image standardization

For video-based traffic data, continuous surveillance videos were converted into individual image frames at fixed intervals. Both image-based and extracted frames were converted into a uniform image format to facilitate batch processing. This step allowed each frame to be treated as an independent input sample for YOLOv8-based vehicle detection.

Annotation and bounding box generation

Vehicle annotations were provided in the form of bounding boxes enclosing each visible vehicle instance. Each bounding box was defined using the YOLO annotation format, which specifies the object class along with normalised coordinates representing the centre position, width, and height of the box (Chavan *et al*., 2023). This annotation strategy enables efficient training and precise localisation of vehicles within traffic scenes.

Resizing

All input images and corresponding bounding boxes were resized to a base resolution of 640 × 640 pixels. During training, YOLOv8 employed multi-scale resizing, dynamically adjusting input dimensions around an average effective resolution of approximately 928 × 928 pixels, consistent with standard YOLOv8 training practice (Ayodeji *et al*., 2025). Bounding box coordinates were scaled proportionally to preserve spatial accuracy after resizing. This uniform input resolution ensures compatibility with the detection network and supports efficient GPU-based training and inference.

Dataset splitting

After preprocessing and annotation, the dataset was organised into the predefined training, validation, and test subsets. Care was taken to maintain a balanced representation of vehicle classes across all subsets and to prevent overlapping between them. This step ensures fair evaluation and robust assessment of the model's performance under diverse traffic conditions. The dataset was randomly shuffled prior to splitting, with a typical ratio of 70% for training, 20% for validation, and 10% for testing. This ensures unbiased sampling and prevents overfitting.

2.3 YOLOv8 Model Description and Training

This study employs YOLOv8 as the core deep learning model for vehicle detection and classification, owing to its efficient anchor-free design and strong detection performance (Zhang & Zhang, 2025). Unlike traditional multi-stage detection pipelines, YOLOv8 performs object localisation and classification in a single forward pass, enabling fast infeEach preprocessed traffic image is resized to a base resolution of 640 × 640 pixels. During training, automatic multi-scale adjustment results in an effective resolution of approximately 928 × 928 pixels before being passed into the YOLOv8 backbone. This backbone captures essential vehicle characteristics such as shape, size, edges, and texture patterns, even in dense traffic scenes with partial occlusion.

The extracted features are then forwarded to the neck component, which combines a Feature Pyramid Network (FPN) and Path Aggregation Network (PAN). This multi-scale feature fusion mechanism enables the detection of vehicles of different sizes, from small distant vehicles to large trucks and buses. Multi-scale representation is particularly important in traffic monitoring, where camera viewpoints and vehicle distances vary significantly (Hegde *et al*., 2023).

Following feature fusion, the YOLOv8 detection head generates bounding box predictions along with objectness scores and vehicle class probabilities. Non-maximum suppression is applied to remove redundant detections and retain only the most confident predictions (Liu *et al*., 2025). The final output consists of accurately localised vehicles annotated with class labels and confidence scores, supporting downstream tasks such as

traffic density estimation, vehicle counting, and congestion analysis (Jenisha *et al*., 2024). The model was trained using the annotated traffic dataset over multiple epochs, employing the AdamW optimiser and a cosine learning rate scheduling strategy (Maurya *et al*., 2025). The model was trained for [insert number] epochs using a batch size of [insert value], with an initial learning rate of [insert value]. Training was conducted on [specify hardware, e.g., NVIDIA GPU], ensuring efficient computation. Early stopping was applied to prevent overfitting based on validation loss.
During training, the network minimised detection loss functions and progressively learned to distinguish vehicles from background elements such as roads, buildings, and pedestrians. After convergence, the trained YOLOv8 model produced reliable vehicle detections across diverse traffic scenarios.
Fig. 2 illustrates the YOLOv8 architecture used, highlighting the backbone for feature extraction, the neck for multi-scale feature fusion, and the detection head for final vehicle localisation and classification.
The training process optimises a composite loss function consisting of classification loss, localisation loss (bounding box regression), and objectness loss, enabling accurate vehicle detection and localisation.

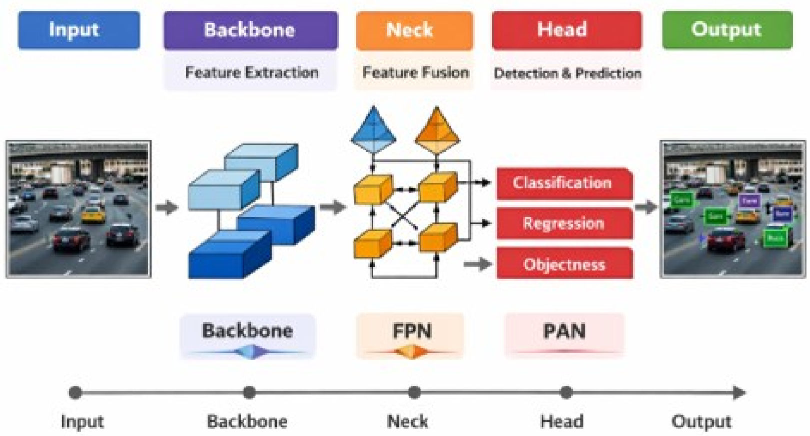

***Fig.* 2: YOLOv8 Architecture**

2.4 Evaluation Metrics

To quantitatively evaluate the effectiveness of the proposed vehicle detection framework, several standard object detection performance metrics were employed, including Precision, Recall, F1-score, and mean Average Precision (mAP). These metrics are widely recognized in intelligent transportation and computer vision research for assessing the accuracy, robustness, and reliability of deep learning-based vehicle detection systems (Premaratne et al., 2023; Sapkota et al., 2025). Precision was used to measure the proportion of correctly identified vehicles relative to all detected objects, while Recall evaluated the ability of the model to correctly identify all relevant vehicle instances within the dataset. The F1-score provided a balanced assessment of both Precision and Recall, thereby offering a more comprehensive measure of detection performance under varying traffic conditions (Bakirci, 2024; Usta et al., 2025).
In addition, the mean Average Precision at an Intersection over Union threshold of 0.5 (mAP@0.5) was adopted as the principal evaluation metric because of its effectiveness in measuring localization accuracy and classification consistency across object classes. This metric has been extensively applied in recent YOLO-based traffic detection studies and remains a benchmark for evaluating real-time object detection frameworks (Zhang & Zhang, 2025; Desta & Jian, 2025). The

Precision–Recall curve was also analysed to assess the stability of the proposed model across different confidence thresholds, particularly under complex traffic scenarios involving occlusion, varying illumination, and vehicle density (Maurya et al., 2025).

The quantitative evaluation further included the analysis of training and validation loss convergence to determine the learning stability and generalization capability of the proposed model. Stable convergence behaviour with minimal overfitting indicates that the framework effectively learned representative traffic features while maintaining strong predictive performance on unseen data (Li et al., 2026; Mehta & Shah, 2025). Collectively, these evaluation metrics provide a comprehensive assessment of the efficiency, robustness, and scalability of the proposed YOLOv8-based vehicle detection framework for intelligent traffic monitoring applications.

Mean Average Precision (mAP)

Mean Average Precision (mAP) is used to evaluate the overall detection performance of the model across all vehicle classes and confidence thresholds. mAP summarises the precision–recall relationship by averaging precision values over different recall levels, while also incorporating bounding box localisation accuracy. As such, it serves as a comprehensive metric for assessing the effectiveness of YOLOv8 in multi-class vehicle detection tasks (Dhake *et al*., 2024)

3.0 Results and Discussion

3.1 Results

This section presents the quantitative and qualitative performance evaluation of the proposed YOLOv8-based vehicle detection framework. The model was assessed using standard object detection metrics, including Precision, Recall, F1-score, mean Average Precision (mAP@0.5), and validation loss trends, to comprehensively evaluate detection accuracy, robustness, and training stability. The results are presented through both quantitative metrics and visual analyses to provide a comprehensive understanding of model performance under varying traffic conditions.

3.2 Precision–Recall Performance

Fig. 3 below illustrates the Precision–Recall (PR) curve obtained for the proposed framework across all vehicle classes. The model achieves a mean Average Precision (mAP@0.5) of 0.975, indicating excellent detection performance with strong localisation accuracy and minimal false detections.

The PR curve maintains high precision across a wide recall range, indicating reliable detection performance with minimal false positives. This behaviour demonstrates that the model maintains high detection reliability while successfully identifying the majority of vehicles in the scene. Such performance is particularly critical in traffic monitoring applications, where both false positives and missed detections can negatively impact downstream tasks such as vehicle counting, congestion analysis, and traffic flow estimation (Johnwendy *et al*., 2023).

The smooth and elevated nature of the PR curve further indicates that the model generalises well to unseen traffic scenes, despite variations in vehicle scale, viewpoint, and environmental conditions.

3.3 1Score and Confidence Threshold Analysis

Fig. 4 below presents the F1-score as a function of the detection confidence threshold. The proposed model achieves a peak F1-score of 0.93 at an optimal confidence threshold of approximately 0.52, representing a strong balance between precision and recall. This high F1-score indicates a well-balanced trade-off between precision and recall, suggesting that the model effectively minimises both false positives and false negatives.

The F1 curve exhibits a broad hill around the optimal threshold, indicating that the model's performance remains stable across a wide confidence range. This stability is essential for real-world deployment, where confidence thresholds may need to be adjusted based on application-specific requirements. At very high confidence thresholds, the F1-score declines due to missed detections, while overly low thresholds introduce additional false positives. The identified optimal threshold therefore provides a practical operating point for real-time intelligent traffic monitoring systems.

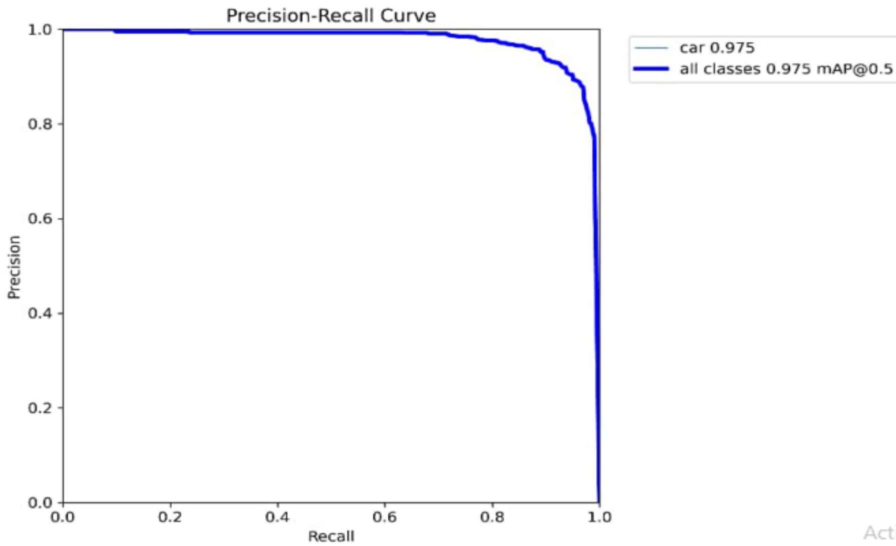

Fig. 3: Precision Recall Curve

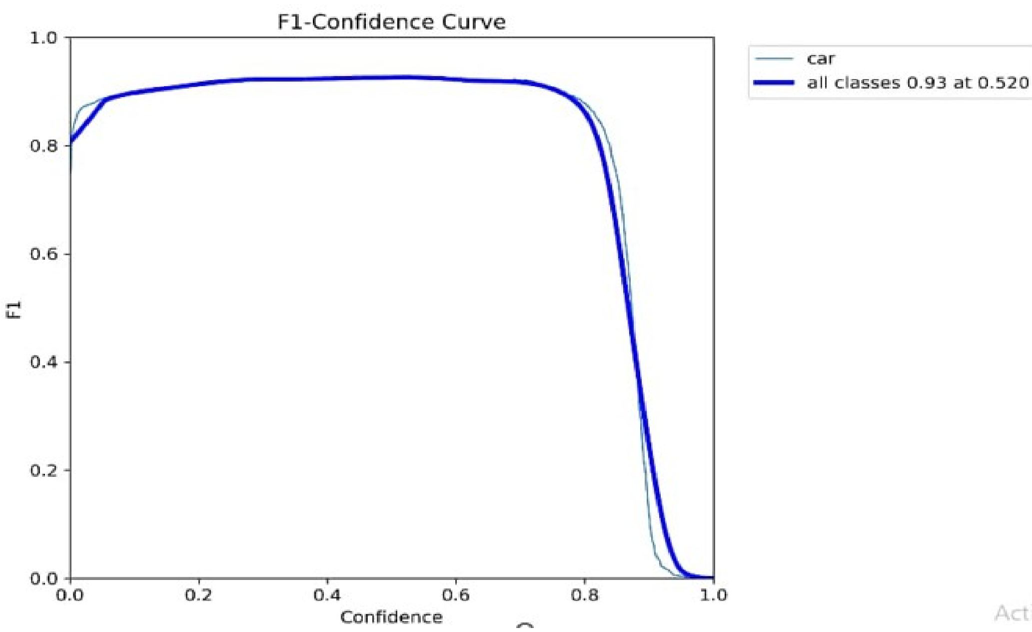

***Fig.* 4: F1 Confidence Curve**

3.4 Validation Loss Convergence

Fig. 5 shows the validation loss curves for the bounding box regression loss, classification loss, and distribution focal loss throughout the training process. All three loss components demonstrate a clear downward trend and stabilise after sufficient training epochs.

The rapid decrease in classification loss during early epochs indicates efficient learning of discriminative vehicle features, while the steady reduction in bounding box loss reflects continuous improvement in localisation accuracy. The convergence of the distribution focal loss further suggests improved confidence calibration and refined bounding box prediction.

The absence of significant oscillations or divergence in the validation losses confirms stable training behaviour and strong generalisation capability, indicating that the model avoids overfitting despite the complexity of traffic scenes and object interactions. The close alignment between training and validation loss trends further indicates that the model does not suffer from overfitting and maintains strong generalisation performance.

3.1.5 Training Loss Analysis

Fig. 6 illustrates the training loss curves of the proposed YOLOv8-based vehicle detection framework over successive training epochs, providing insight into the model's optimization behavior during training. At the early stages of training, all loss components exhibit a rapid decrease, indicating efficient gradient propagation and fast learning of fundamental vehicle features such as edges, shapes, and spatial layouts. This initial decline confirms that the model effectively captures low-level and mid-level visual representations necessary for vehicle detection in traffic scenes.

As training progresses, the losses continue to decrease more gradually and eventually stabilise, demonstrating successful convergence. The steady reduction in bounding box regression loss reflects improved localisation accuracy, enabling the model to tightly align predicted bounding boxes with vehicle boundaries. Similarly, the consistent decline in classification loss indicates enhanced discriminative capability between different vehicle classes and background regions.

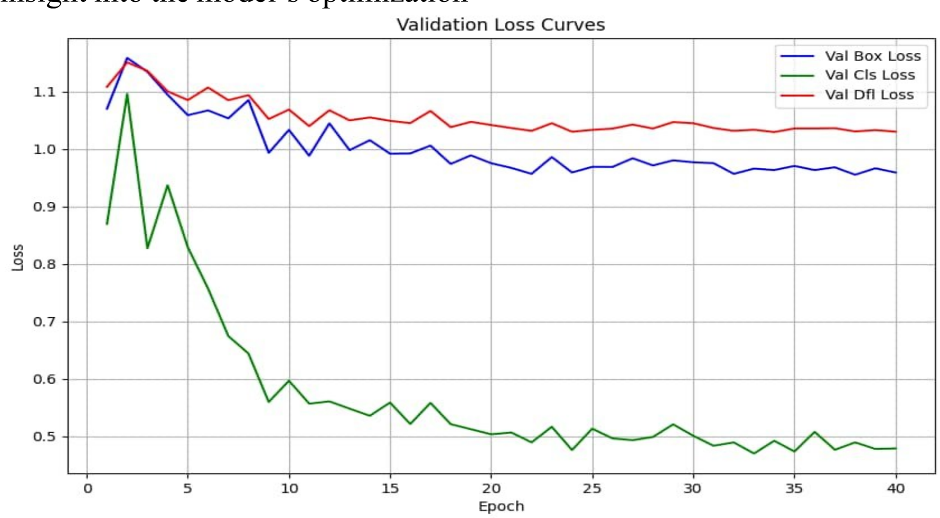

***Fig.* 5: Validation Loss Curve**

The distribution focal loss shows smooth convergence, suggesting improved confidence calibration and refined bounding box distribution learning. This behaviour is particularly important in dense traffic environments, where overlapping vehicles and scale variations require precise confidence estimation to reduce false positives and missed detections. However, minor detection challenges were observed in extremely dense traffic conditions with heavy occlusion, where closely overlapping vehicles occasionally resulted in missed or merged detections.

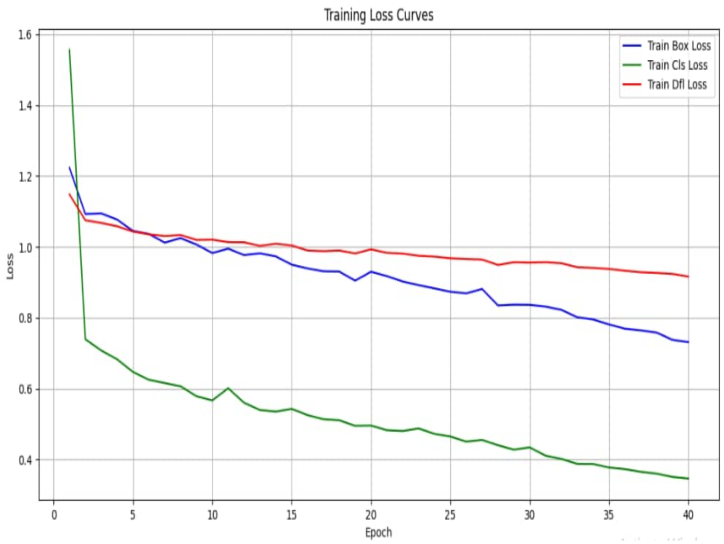

***Fig.* 6: Training Loss Curve**

3.1.6 Qualitative Detection Results

Fig. 7 provides qualitative visualisations of vehicle detection results across multiple traffic scenarios, including dense urban roads, multi-lane highways, and bridge environments. The model successfully detects vehicles of varying sizes and classes, even under conditions of partial occlusion, close vehicle proximity, and complex background structures. The detected bounding boxes are tightly aligned with vehicle boundaries, and false positives are minimal despite the presence of background clutter such as road signs, barriers, and infrastructure elements. These visual results corroborate the quantitative findings and demonstrate the framework's robustness in real-world traffic environments.

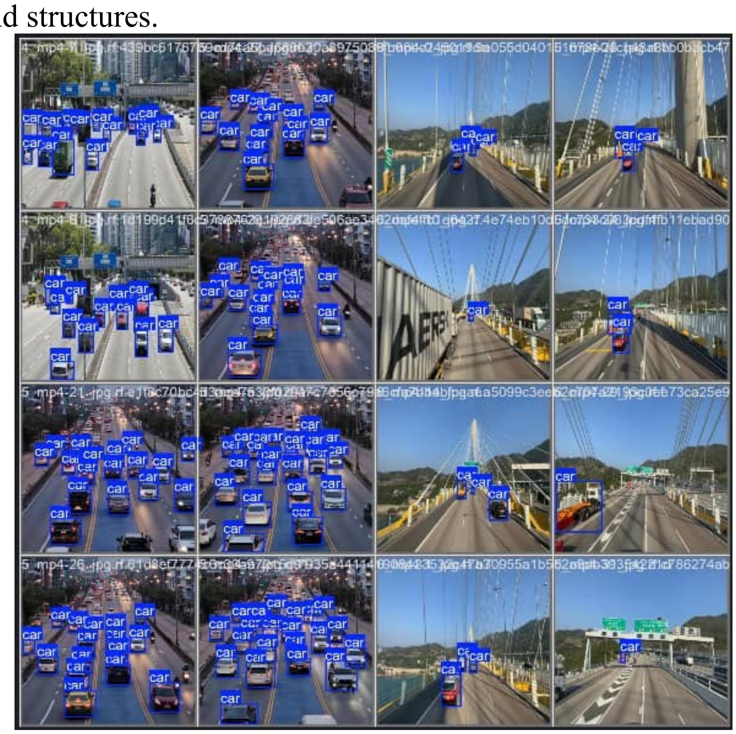

Fig. 7: Detection Results

3.2 Discussion

The experimental results demonstrate that the proposed YOLOv8-based vehicle detection framework achieves strong and competitive performance in terms of accuracy, robustness, and training stability, particularly in complex traffic scenarios.

3.2.1 *Comparison with Related Work*

Compared with earlier YOLOv5-based approaches reported in the literature, the proposed framework demonstrates substantial improvements in vehicle detection performance. Previous studies commonly reported F1-scores below 0.90 for traffic monitoring applications (Chavan et al., 2023; Nandini & Ravikanth, 2025), whereas the present framework achieved a peak F1-score of 0.93, indicating a more balanced trade-off between precision and recall. Furthermore, the proposed model attained a high mAP@0.5 value of 0.975, alongside stable Precision–Recall characteristics, demonstrating enhanced robustness under challenging traffic conditions. This performance exceeds many previously reported results, where mAP values generally ranged from 0.85 to 0.95 for comparable vehicle detection tasks (Bakirci, 2024; Usta et al., 2025).

Although Zhang and Zhang (2025) improved YOLOv8 performance through architectural modifications such as adaptive scaling and the ADown module, their method introduced additional architectural complexity. In contrast, the proposed framework achieved comparable or superior detection performance without incorporating extra modules or increasing network complexity. Instead, the framework leveraged the inherent advantages of YOLOv8, including its anchor-free detection mechanism, decoupled detection head, and efficient multi-scale feature fusion. This finding suggests that architectural efficiency can provide significant performance gains without the need for increasingly complex model structures.

Similarly, multimodal and ensemble-based methods, such as the RGB–thermal fusion strategy reported by Maurya et al. (2025), demonstrated strong low-light detection capability but required substantially higher computational resources. By comparison, the proposed framework maintained high detection accuracy using a single RGB-based architecture, thereby offering greater suitability for deployment in resource-constrained intelligent transportation systems.

In addition, large-scale intelligent traffic monitoring systems integrating vehicle detection with auxiliary tasks such as speed estimation and emission monitoring have shown real-time applicability but often suffer from scalability challenges and increased computational overhead (Asifur et al., 2025). The stable validation loss convergence and high F1-score achieved in the present study indicate that the proposed framework provides a lightweight, scalable, and computationally efficient solution for vehicle detection, making it highly suitable as a foundational component for intelligent traffic monitoring and management systems.

3.2.2 *Impact of YOLOv8 Architectural Design*

The strong performance observed in this study can be attributed to several key architectural features of YOLOv8. The anchor-free detection mechanism reduces sensitivity to predefined bounding box priors, improving detection across varying vehicle scales. The decoupled head structure enhances classification and localisation learning, which is reflected in the stable validation losses and high confidence calibration.

Furthermore, the Feature Pyramid Network (FPN) and Path Aggregation Network (PAN) enable effective multi-scale feature representation, allowing the model to detect small distant vehicles and large nearby vehicles simultaneously. This capability is particularly evident in dense traffic scenes, where overlapping objects and scale variation are prevalent.

3.2.3 *Practical Implications*

The combination of high precision, high recall, stable F1 performance, and robust qualitative detections confirms the suitability of the proposed framework for real-world intelligent traffic monitoring applications. The results indicate that the system can reliably support tasks such as vehicle counting, congestion analysis, and traffic flow estimation without requiring complex post-processing or additional sensor modalities.

Overall, the proposed YOLOv8-based framework effectively addresses key limitations identified in previous studies, including sensitivity to occlusion, unstable confidence calibration, and high computational cost. These improvements position the framework as a practical and scalable solution for modern intelligent transportation systems.

3.2.4 *Limitations of the Study*

Despite the strong performance achieved, this study has certain limitations. The dataset, although diverse, may not fully capture extreme traffic scenarios such as severe weather conditions, nighttime environments, or highly congested urban intersections. Additionally, the model was evaluated primarily on static image frames rather than continuous real-time video streams, which may introduce temporal inconsistencies in practical deployment. Future work will focus on incorporating temporal modelling techniques and expanding dataset diversity to improve robustness.

4.0 Conclusion

This study developed and evaluated a YOLOv8-based framework for real-time vehicle detection and classification in intelligent traffic monitoring systems. By utilising the anchor-free architecture, decoupled detection head, and efficient multi-scale feature fusion of YOLOv8, the proposed approach effectively addresses key challenges associated with dense traffic conditions, scale variation, occlusion, and environmental complexity.

Experimental results obtained from a diverse and well-annotated traffic dataset demonstrate the robustness and effectiveness of the framework. The model achieved a high mAP@0.5 of 0.975 and a peak F1-score of 0.93, accompanied by stable convergence of both training and validation losses. These results indicate strong detection accuracy, reliable generalisation, and consistent performance across varying traffic scenarios. In addition, qualitative analysis confirms precise vehicle localisation with minimal false positives, even in complex and cluttered environments.

In comparison with existing YOLO-based and multi-stage detection approaches, the proposed framework provides competitive or improved accuracy while maintaining computational efficiency and real-time inference capability. This balance between performance and efficiency highlights the suitability of the model for deployment in practical intelligent transportation systems.

Overall, the study demonstrates that YOLOv8 can serve as a lightweight, scalable, and effective solution for modern traffic monitoring applications, including vehicle counting, congestion analysis, and traffic flow management. Future work may focus on extending the framework to incorporate temporal tracking, real-time video stream processing, and deployment on edge devices to further enhance its applicability in smart city infrastructures.

is the abstract wellwritten and are the keyword suitable for the study. Traffic congestion and road accidents remain major challenges in modern transportation systems, demanding accurate, real-time traffic-monitoring solutions. Traditional surveillance methods and classical computer vision techniques are often limited by manual intervention, sensitivity to environmental variations, and poor scalability in complex traffic scenes. To address these limitations, this study proposes a YOLOv8-based framework for real-time vehicle detection in intelligent traffic monitoring systems. The proposed approach exploits YOLOv8 architecture to accurately detect vehicles of varying sizes under diverse traffic conditions, including occlusion, dense traffic, and illumination changes. A dataset comprising over 5,000 annotated traffic images representing urban and highway environments was used for training and evaluation. Model performance was assessed using standard object detection metrics, including Precision, Recall, F1-score, and mean Average Precision (mAP). Experimental results demonstrate strong detection accuracy, achieving an mAP@0.5 of 0.975 and a peak F1-score of 0.93, with stable convergence of training and validation losses. The results confirm the effectiveness of YOLOv8 as a lightweight and scalable solution for real-time intelligent traffic monitoring applications. Keywords: YOLOv8; Vehicle Detection; Intelligent Transportation Systems; Traffic Monitoring; Object Detection.

5.0 References

Asifur, M., Ali, M., Rahman, M., Rahim, M. A., Ayman, U., & Sohel, A. (2025). Real-time vehicle type detection and counting for emission pollution monitoring and traffic violation identification. *Emerging Science Journal*, *9*(2), 959–976. https://doi.org/10.28991/ESJ-2025-09-02-023

Ayodeji, O., Obafemi, J. R., & Akinrolabu, O. D. (2025). A deep learning-based model for traffic signal control using the YOLO algorithm. *Journal of Computer Science*, *186*(63), 43–54.

Bakirci, M. (2024). Real-time vehicle detection using YOLOv8-Nano for intelligent transportation systems. *Traitement du Signal*, *41*(4), 1727–1740.

Cao, L. (2025). Research on the vehicle detection technology based on the YOLO model. *E3S Web of Conferences*, *3008*.

Chavan, C., Hembade, S., Jadhav, G., Komalwad, P., & Bhosale, P. A. (2023). Vehicle detection using YOLOv5. *International Journal of Scientific Research in Engineering and Management*, 1–9. https://doi.org/10.55041/IJSREM21823

Desta, A. D., & Jian, C. (2025). Enhancing YOLOv8 for vehicle detection in intelligent traffic management. *Metallurgical and Materials Engineering*, *31*(4), 190–200. https://doi.org/10.63278/1423

Dhake, A. L., Talekar, A., & Bembade, K. (2024). Traffic management using YOLO and convolutional fizzy neural network. *International Journal of Engineering Research & Technology*, *13*(07).

Gheorghe, C., & Soica, A. (2025). Revolutionizing urban mobility: A systematic review of AI, IoT, and predictive analytics in adaptive traffic control systems for road networks (pp. 1–25).

Hegde, V., Riyaj, P. S., Mouna, V. J., Vijayan, S. J. Y., & G, P. N. (2023). Vehicle detection and tracking with deep learning. *Journal Name*, *6*(6), 4–9.

Jenisha, R., Aswin, S., & Shinshand, A. (2023). Vehicle counting using deep learning. *Journal of Xi'an Shiyou University, Natural Science Edition*, *19*(05).

Johnwendy, C. N., Onimisi, N. F., & Izegbu, I. O. (2023). A comprehensive review for YOLO-based vehicle counting systems in traffic management. *Global Scientific Journals*, *11*(8), 2959–2964.

Li, M., Wang, X., Chen, L., & Liu, J. (2026). A structurally optimized and efficient lightweight object detection model for autonomous driving (pp. 1–17).

Liu, S., Shao, F., Chu, W., Dai, J., & Zhang, H. (2025). An Improved YOLOv8-Based

Lightweight Attention Mechanism for Cross-Scale Feature Fusion. *Remote Sensing*, *17*(6), 1044. https://doi.org/10.3390/rs17061044

Majin, J. J., Valencia, Y. M., Stivanello, M. E., & Stemmer, M. R. (2023). A novel deep learning based method for detection and counting of vehicles in urban traffic surveillance systems. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLIII*, 793–800.

Maurya, T., Kumar, S., Rai, M., Kumar, A., Neha, S., & Gunjan, G. (2025). Learning—Smart traffic management. *Engineering Reports*, 1–15. https://doi.org/10.1002/eng2.70082

Mehta, R., & Shah, A. (2025). Real-time vehicle detection and classification using deep learning based approach. *Journal Name*, *10*.

Mittal, P. (2024). A comprehensive survey of deep learning-based lightweight object detection models for edge devices. *Artificial Intelligence Review*, *57*(9). https://doi.org/10.1007/s10462-024-10877-1

Mudawi, N. A., et al. (2023). Vehicle detection and classification via YOLOv8 and deep belief network over aerial image sequences. *Sustainability*, *15*.

Mystakidis, A., & Koukaras, P. (2025). Advances in traffic congestion prediction: An overview of emerging techniques and methods. *Smart Cities*, 1–37.

Nandini, M., & Ravikanth, K. (2025). Traffic detection system using the YOLOv5 deep learning model. *International Journal of Research Publication and Reviews*, *6*(8), 5484–5490.

Prathima, L., Pavan, L., Pravith, A. P., Preetham, B., & Skanda, M. V. (2025). Vehicle detection and counting in traffic signals using deep learning. *Journal Name*, *13*(1), 950–954.

Premaratne, P., Jawad, I., Blacklidge, R., & Lee, M. (2023). Neurocomputing comprehensive review on vehicle detection, classification and counting on highways. *Neurocomputing*, *556*, Article 126627. https://doi.org/10.1016/j.neucom.2023.126627

Roboflow. (2023). *Roboflow Universe: Open-source computer vision datasets*. https://universe.roboflow.com

Sakr, H. A., & El-afifi, M. I. (2023). Intelligent traffic management systems: A review. *Journal Name*, *5*.

Sapkota, R., Rizwan, M. F., & Chetan, Q. (2025). YOLO advances to its genesis: A decadal and comprehensive review of the You Only Look Once (YOLO) series.

Sinojia, K., Yeolekar, A., Vyas, J., & Bhonsle, M. (2025). Integrating YOLO object detection with machine learning for intelligent traffic flow management in smart cities. *International Journal of Scientific Research in Science, Engineering and Technology*, *12*(6), 114–120. https://doi.org/10.32628/IJSRSET2513852

Usta, D. D., Hundelshaussen, R., Martínez, C. L., Leite, C., & Marques, D. M. (2025). Intelligent traffic management: Comparative evaluation of YOLOv3, YOLOv5, and YOLOv8 for vehicle detection in urban environments in Montería, Colombia (pp. 1–19).

Zhang, M., & Zhang, Z. (2025). Research on vehicle target detection method based on improved YOLOv8.

Zhou, W., et al. (2024). Vision technologies with applications in traffic surveillance systems: A holistic survey. *arXiv*.

Declaration

Ethics Approval and Consent to Participate

This study involved the development, training, and evaluation of a deep learning framework using publicly available annotated traffic image datasets obtained from open-source computer vision repositories. The research did not involve human clinical trials, animal experimentation, or the collection of personal or sensitive data. Therefore, ethical approval and informed consent were not required under institutional research guidelines.

Consent for Publication

All authors have read and approved the final version of the manuscript and consented to its publication.

Availability of Data and Materials

The datasets used and/or analyzed during the current study are available from publicly accessible repositories, including Roboflow Universe, and additional processed data may be obtained from the corresponding author upon reasonable request.

Funding

The authors declared that no external funding was received for this study.

Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Authors' Contributions

ATB conceptualized and designed the study, developed the YOLOv8 framework, conducted experiments, analyzed the data, and prepared the manuscript draft. AS contributed to dataset preparation, methodology development, and performance evaluation. ARO assisted with literature review, result interpretation, and manuscript revision. All authors read and approved the final manuscript.

Disclaimer

The views and interpretations presented in this manuscript are those of the authors and do not necessarily reflect the official policy or position of the affiliated institution.