

Attention-Enhanced YOLO Framework for Robust Real-Time Object Detection in Autonomous Driving Environments

Dasam Srinivas^{1,*} A.V.S. Pawan Kumar²

Abstract

Accurate object detection is essential for autonomous driving systems, where vehicles must continuously perceive surrounding road users and infrastructure in real time. Although the You Only Look Once (YOLO) family of deep learning detectors has achieved significant success due to its high speed and efficiency, detecting small or partially occluded objects in complex traffic environments remains challenging. This paper presents an attention-enhanced YOLO framework designed to improve feature representation and detection reliability for autonomous driving applications. The proposed method integrates spatial and channel attention mechanisms into the detection architecture, enabling the network to focus on informative regions of the scene while suppressing irrelevant features. The model is evaluated using representative autonomous driving datasets, including KITTI and BDD100K, and compared with existing YOLO variants. Experimental results demonstrate that the proposed framework achieves precision of 0.97, recall of 0.92, and mean average precision (mAP@0.5) of 0.982, outperforming baseline models such as YOLOv5 and YOLOv8. In addition, the system maintains real-time performance of approximately 94 frames per second, making it suitable for deployment in practical autonomous vehicle perception systems. The results indicate that incorporating attention mechanisms significantly enhances detection accuracy, particularly for small and partially occluded objects, while preserving the computational efficiency required for real-time driving environments.

Keywords: Autonomous vehicles, object detection, YOLO, attention mechanism, deep learning, real-time detection, computer vision, autonomous driving perception

*Author for Correspondence

Dasam Srinivas
E-mail: 24mtcse039.srinivasdasam@giet.edu

¹PG Scholar, Department of CSE Gandhi Institute of Engineering and Technology University, Gunupur Odisha, India

²Associate Professor, Department of CSE Gandhi Institute of Engineering and Technology University, Gunupur Odisha, India

Received Date: July 02, 2026

Accepted Date: August 11, 2026

Published Date: August 16, 2026

Citation: Dasam Srinivas, A V S Pawan Kumar. Attention-Enhanced YOLO Framework for Robust Real-Time Object Detection in Autonomous Driving Environments. Journal of Multimedia Technology & Recent Advancements. 2026; 13(2): 41–52p.

INTRODUCTION

Autonomous vehicles depend on reliable perception systems to interpret their surroundings and make safe driving decisions. Among the various components of the perception stack, object detection plays a crucial role because it enables vehicles to identify and localize important road elements, such as cars, pedestrians, cyclists, traffic signs, and other obstacles. Accurate detection allows the planning and control modules of an autonomous system to respond appropriately, thereby improving road safety and navigation efficiency. However, detecting objects in real-world driving environments is challenging due to factors such as varying object sizes, occlusions, lighting variations, and complex traffic scenes.

Earlier object detection approaches relied on handcrafted feature extraction techniques, such as histogram of oriented gradients (HOG), haar-like features, and scale-invariant feature descriptors, combined with traditional classifiers. Although these techniques have made initial progress in visual recognition tasks, they often struggle to generalize under diverse environmental conditions and require extensive manual feature engineering. The emergence of deep learning, particularly convolutional neural networks (CNNs), has significantly transformed the field of computer vision by enabling automatic feature learning directly from large datasets.

Modern deep learning based object detection algorithms can be broadly classified into two categories: two-stage and single-stage detectors. Two-stage detectors, such as Faster R-CNN, first generate candidate object regions and then classify them, typically achieving high accuracy but at the cost of increased computational complexity. In contrast, single-stage detectors such as the You Only Look Once (YOLO) framework perform detection and classification simultaneously within a single network pass. This architecture enables significantly faster inference, making YOLO particularly suitable for real-time applications, such as autonomous driving.

Despite the impressive speed and performance of YOLO-based detectors, several challenges remain in their application to autonomous vehicle perception. In real traffic environments, objects often appear at different scales and may be partially occluded by other vehicles or infrastructure elements. Small or distant objects, such as pedestrians crossing at a distance or cyclists partially hidden behind vehicles, are particularly difficult to detect reliably. Additionally, real-world driving conditions frequently involve varying illumination, weather changes, and complex backgrounds that can degrade the detection performance.

To address these challenges, recent research has explored the integration of attention mechanisms into deep neural networks. Attention modules enable a network to focus on the most relevant spatial and channel features while suppressing less informative information. By enhancing feature representation, attention mechanisms can improve the detection of challenging objects without significantly increasing the computational overhead.

In this study, an attention-enhanced YOLO framework was proposed to improve object detection performance in autonomous driving environments. The proposed approach incorporates attention mechanisms within the YOLO architecture to strengthen feature extraction and enhance the detection of small and partially occluded objects. The effectiveness of the framework is evaluated using widely used autonomous driving datasets, and the experimental results demonstrate improved detection accuracy while maintaining real-time inference capability, as shown in Figure 1.




Figure 1. Block diagram of the R²M²-GridFM framework.

LITERATURE REVIEW

Recent studies have greatly enhanced object detection algorithms for self-driving cars, particularly in terms of efficiency and real-time detection accuracy. YOLO object detection algorithms are popular because of their efficient single-stage detection and fast detection speed. However, recent studies conducted between 2024 and 2026 have focused on improving the detection of smaller objects, robust detection in complex environments, and the efficiency of algorithms.

In 2024, YOLOv10 was developed as an end-to-end object detection algorithm that eliminated the dependency on traditional non-maximum suppression techniques during training and improved the efficiency during inference with high accuracy [1]. The object detection algorithm used architectural design and dual assignment strategies to improve efficiency and faster detection performance for real-time applications like self-driving cars.

Researchers have also worked on improving the detection of small objects and their features. Enhanced versions of the YOLO algorithm introduced multi-scale feature aggregation and neck structures to better detect fine spatial features that are otherwise lost in deep learning models. This improvement enables better detection of distant pedestrians, cyclists, and other small objects on the road, which are essential for safe driving systems [2].

Another major area of research for improving object detection models is enhancing object detection under adverse weather conditions. It is common for self-driving cars to drive through rainy weather, fog, or at night when the weather conditions are poor. Several studies have proposed weather-robust versions of the YOLO algorithm using attention models and preprocessing techniques to ensure the stability of object detection models under adverse weather conditions [3].

Attention-based object detection models are another trend in recent studies. By using spatial and channel attention models, object detection models are now capable of focusing on more informative regions of an image. This has enabled object detection models to better detect occluded objects on roads [4].

Recent studies have explored domain adaptation and multimodal perception methods to improve detection reliability. These methods focus on improving detection generalization across various driving scenarios and datasets, as well as maintaining computational efficiency for embedded autonomous systems [5].

However, some limitations remain in achieving reliable detection results for small or occluded objects in real time. These limitations encourage researchers to propose better frameworks for detecting objects using attention-based methods and efficient architectures [6–8].

PROPOSED ARCHITECTURE AND METHODOLOGY

In this study, an attention-enhanced YOLO network architecture is introduced to improve the performance of object detection in an autonomous driving environment. The introduced framework includes an attention mechanism in the object detection process to improve feature representation for the detection of small, distant, and occluded objects while maintaining real-time performance [9–10].

Overall Architecture

The architecture has four major stages: input preprocessing, backbone feature extraction, attention-based feature enhancement, and detection head prediction. The entire architecture processes the image and predicts the object locations along with the class probabilities in a single forward pass [11].

For an image I , the detection network predicts a list of bounding boxes and class scores as follows:

$$D = f(I; \theta)$$

where $f(\cdot)$ represents the detection network and θ denotes the learned model parameters.

Feature Extraction Backbone

The backbone network employs hierarchical features from the input image based on convolutional operations.

Each convolutional layer in the network produces feature maps that contain spatial information, such as edges, textures, and objects.

The convolution operation can be expressed as

$$F_l = \sigma(W_l * F_{l-1} + b_l)$$

Where

F_l = feature map at layer l W_l = convolution kernel weights b_l = bias term $*$ = convolution operation σ = activation function. These feature maps encode the spatial and semantic information necessary for object detection.

Attention Enhancement Module

To improve feature representation, an attention mechanism was introduced after the backbone stage. In this attention block, relevant object information is emphasized, and irrelevant background information is suppressed [12].

Channel Attention

In channel attention, the importance of each channel in the feature map is determined. Global average pooling is performed on the feature map to obtain a channel descriptor.

$$z_c = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W F(i, j, c)$$

Where

H and W represent the spatial dimensions of the feature map. The attention weights are computed as follows:

$$A_c = \sigma(W_2 \cdot \text{ReLU}(W_1 \cdot z))$$

The feature map is then refined using channel attention:

$$F_c = A_c \odot F$$

where $\odot$ denotes element-wise multiplication.

Spatial Attention

Spatial attention highlights the most informative spatial regions of the feature map.

$$A_s = \sigma(\text{Conv}(F_c))$$

The refined spatial feature map becomes

$$F' = A_s \odot F_c$$

This process allows the network to focus on relevant object regions in complex traffic scenes.

Detection Head

The detection head predicts the bounding boxes, class probabilities, and confidence scores from the enhanced feature maps.

For each predicted bounding box, the model estimates

$$B1 = (x, y, w, h)$$

Where

x, y represent box center coordinates
w, h represent width and height.

The object confidence score is defined as

$$\text{Conff} = P(\text{object}) \times \text{IoU}$$

where IoU represents the intersection over union between the predicted and ground-truth bounding boxes.

Loss Function

The training objective minimizes the combined detection loss, which consists of localization, confidence, and classification components:

$$L = L_{\text{box}} + L_{\text{conf}} + L_{\text{class}}$$

where

- L_{box} measures bounding box regression error
- L_{conf} penalizes incorrect objectless prediction
- L_{class} evaluates classification error.

Minimizing this loss enables the network to accurately localize and classify objects in driving scenarios.

Training Strategy

The model was trained using annotated autonomous driving datasets. The input images were resized and augmented to improve generalization. During training, the network learns the optimal parameters θ through gradient-based optimization [13].

The integration of attention-guided feature refinement and multi-scale detection enables the proposed framework to improve the detection performance while maintaining the real-time processing capability required for autonomous vehicle perception systems.

Algorithm 1: Attention-Enhanced YOLO Training and Inference

- *Input:* Training dataset $D=(I_i, T_i)$, where I_i is the input image, and T_i represents the ground-truth annotations. Learning rate η
- *Output:* Trained Attention-Enhanced YOLO detection model

Training Phase

1. Initialize network parameters θ of the backbone, attention module, and detection head
2. for epoch = 1 to E do
3. for each mini-batch (images, targets) in dataset D do
4. Preprocess input images (resize, normalization)
5. Apply data augmentation (flip, scale, brightness variation)
7. # Feature Extraction
8. $F \leftarrow \text{BackboneNetwork}(\text{images})$
- 9.
10. # Attention Enhancement
11. $F_c \leftarrow \text{ChannelAttention}(F)$
12. $F' \leftarrow \text{SpatialAttention}(F_c)$
14. # Object Detection
15. $P \leftarrow \text{DetectionHead}(F')$
17. # Compute Detection Loss

18. $L_{box} \leftarrow \text{BoundingBoxLoss}(P, \text{targets})$
19. $L_{conf} \leftarrow \text{ConfidenceLoss}(P, \text{targets})$
20. $L_{class} \leftarrow \text{ClassificationLoss}(P, \text{targets})$
21. $L_{total} \leftarrow L_{box} + L_{conf} + L_{class}$
23. # Update Model Parameters
24. Compute gradients $\nabla\theta(L_{total})$
25. $\theta \leftarrow \theta - \eta * \nabla\theta(L_{total})$
26. end for
27. end for

Inference Phase

1. Input test image I
2. Preprocess image (resize and normalization)
3. Extract features using backbone network
4. Apply channel attention and spatial attention modules
5. Generate predictions using detection head
6. Apply non-maximum suppression (NMS) to remove overlapping boxes
7. Output final bounding boxes, object classes, and confidence scores (Figure 2).

EXPERIMENTAL SETUP AND DATASET

The proposed framework for attention enhancement in the YOLO framework is evaluated using benchmark datasets for object detection in autonomous driving scenarios. The evaluation was performed considering real-world scenarios for object detection in autonomous driving. All the evaluation procedures are carried out using images with annotations for various road objects such as vehicles, pedestrians, and road signs [14].

Datasets

Two datasets, KITTI and BDD100K, were selected for this study. These datasets were selected because they offer diverse driving scenes in urban areas. They are also commonly used to test perception systems in autonomous vehicles.

The KITTI dataset comprises images captured on roads using a camera fixed on a vehicle. These images were captured under various conditions, such as urban, suburban, and highway. It has various objects such as cars, pedestrians, and bicycles marked, making it suitable for testing object detection models in driving scenarios. For this study, the dataset was divided into various sets for testing and training.




Figure 2. Attention-enhanced YOLO architecture.

The BDD100K dataset comprises images captured on roads under various conditions. These images were taken under various conditions, such as during the day, night, cloudy, and rainy. These conditions are important for testing the robustness of the proposed object detectors. A subset of this dataset was selected for testing [15].

Data Preprocessing

Before the start of the training process, all images were normalized to a fixed resolution of 640×640 pixels. Pixel value normalization was used to enhance the stability of the learning process. The bounding boxes were converted to the format required by the YOLO algorithm.

To enhance the robustness of the proposed model, different augmentation strategies were used during the training process. This included random horizontal flipping, scaling, translation, and brightness. This will allow the proposed detector to learn invariant features, which will be helpful in avoiding overfitting in different driving scenarios.

Training Configuration

The suggested model was trained in a supervised learning environment using images and bounding box annotations. In this model, transfer learning was used by utilizing the weights of the pretrained YOLO model. This accelerated the training process and aided in the learning of features.

The parameters used to train the model were as follows: batch size = 16, learning rate = 0.001, training duration = 50 epochs, and the optimization algorithm used was stochastic gradient descent. In addition, the momentum term was set to 0.9, and the weight decay was set to 0.0005.

Evaluation Metrics

For evaluation of detection performance, standard evaluation metrics for object detection were used, such as precision, recall, F1-score, mAP@0.5, and mAP@0.5:0.95. Frames per second (FPS) was used to measure the inference speed, which is an important criterion for real-time processing in autonomous driving systems.

Experimental Platform

The experiments were performed on a deep learning platform that has support for GPUs, and PyTorch was used as the framework. This facilitated the efficient execution of the training and testing of the developed model, as well as comparisons with other baseline YOLO models under similar experimental conditions.

RESULTS AND DISCUSSION

This section discusses the experimental results of the proposed attention-enhanced YOLO framework for object detection. The results are compared with those of other commonly used YOLO variants. The experimental results demonstrate the accuracy, robustness, and real-time inference ability of the framework in autonomous driving scenarios.

All YOLO variants were trained on the same data split and training conditions for a fair comparison. Object detection metrics, such as precision, recall, F1-score, mean average precision (MAP), and FPS, were used to evaluate the models.

Overall Detection Performance

Table 1 summarizes the overall detection performance of the evaluated models.

The results show that the proposed architecture has the best detection accuracy across almost all metrics. The improvement in the mAP indicates that the proposed architecture improves object localization and classification.

Detection Accuracy by Object Category

The detection accuracy was further analyzed for important road objects. Table 2 presents the category-wise performance.

Performance gains are particularly noticeable for pedestrians and cyclists, which are typically harder to detect because they occupy smaller regions in the image.

Small-Object Detection Analysis

The accurate detection of small objects is critical for the safety of autonomous driving. Table 3 shows a comparison of the small-object detection performance.

The proposed framework demonstrated a clear improvement in detecting small objects, indicating that the attention mechanism enhances spatial feature representation.

Detection Accuracy Graph

Figure 3 illustrates the detection accuracy comparison among the evaluated models.

Graph Values

The graph shows a steady increase in detection accuracy across YOLO versions, with the proposed method achieving the highest performance (Table 4).

Precision–Recall Curve

The precision–recall curve is commonly used to analyze object detection performance. Figure 4 shows a comparison between the proposed model and the baseline detectors.

Trend:

- The proposed model maintains higher precision at higher recall levels
- YOLOv5 curve drops earlier due to lower detection confidence
- YOLOv8 and YOLOv11 show improvement but are still slightly below the proposed model

Speed–Accuracy Comparison

Real-time performance is essential for autonomous driving perception. Figure 5 illustrates the relationship between inference speed and detection accuracy (Table 5).

Observation:

- YOLOv8 achieves high speed but slightly lower accuracy
- YOLOv11 improves accuracy but reduces speed slightly
- The proposed model achieves both higher accuracy and higher FPS

Table 1. Overall detection performance comparison.

Model	Precision	Recall	F1-Score	mAP@0.5	mAP@0.5:0.95	FPS
YOLOv5	0.91	0.86	0.88	0.92	0.67	78
YOLOv8	0.94	0.88	0.91	0.96	0.74	92
YOLOv11	0.95	0.89	0.92	0.97	0.76	85
Proposed Model	0.97	0.92	0.94	0.982	0.79	94

Table 2. Detection accuracy for different object categories.

Model	Vehicle	Pedestrian	Cyclist	Traffic light
YOLOv5	0.94	0.88	0.83	0.81
YOLOv8	0.96	0.90	0.86	0.84
YOLOv11	0.97	0.91	0.88	0.86
Proposed Model	0.98	0.94	0.91	0.89

Table 3. Small-object detection performance.

Model	Small-object mAP
YOLOv5	0.71
YOLOv8	0.75
YOLOv11	0.77
Proposed Model	0.84




Figure 3. Mean average precision comparison.




Figure 4. Precision–recall curve comparison.

This demonstrates that the attention module improves feature learning without increasing computational overhead significantly.

Training Convergence Analysis

Figure 6 illustrates the training convergence of the proposed model compared with that of YOLOv8.

Trend

The proposed architecture converged faster and reached lower loss values, indicating more effective feature learning (Table 6).

Confusion Matrix Analysis

The confusion matrix provides insight into classification accuracy across object categories Table 7.




Figure 5. Accuracy versus FPS.




Figure 6. Training loss versus epoch.

Table 4. Model of graph values.

Model	mAP
YOLOv5	0.92
YOLOv8	0.96
YOLOv11	0.97
Proposed	0.982

Table 5. Speed–accuracy comparison.

Model	mAP	FPS
YOLOv5	0.92	78
YOLOv8	0.96	92
YOLOv11	0.97	85
Proposed	0.982	94

Table 6. Lower loss values.

Epoch	YOLOv8 loss	Proposed loss
10	1.12	1.05
20	0.86	0.73
30	0.63	0.52
40	0.48	0.39
50	0.41	0.33

Table 7. Confusion matrix.

Actual / predicted	Vehicle	Pedestrian	Cyclist
Vehicle	482	9	6
Pedestrian	11	453	12
Cyclist	8	15	436

Most predictions appear along the diagonal elements, which indicates strong classification performance.

DISCUSSION

From the experimental evaluation, it is confirmed that by incorporating attention mechanisms in the YOLO model, detection performance is enhanced in autonomous driving applications. The attention module in the model enhances feature discrimination by focusing on important information and suppressing background information.

Another interesting fact observed in this model is the improvement in the detection of smaller objects in real-world driving scenarios. In real-world driving conditions, pedestrians and cyclists may be distant from the vehicle and occluded by other vehicles in the scene. The proposed model enhances the detection of such objects by improving the feature maps using attention mechanisms.

Another advantage of the proposed model is its real-time detection performance. The model achieved over 90 FPS, which meets the real-time detection requirements for autonomous driving applications.

However, some issues may be resolved in future work by integrating this model with multimodal perception techniques such as LiDAR and radar sensor integration.

CONCLUSION

In this study, an attention-enhanced YOLO framework was developed to improve object detection in an autonomous driving environment. The developed method combines attention mechanisms into the object detection network, which can improve feature representation for object recognition, particularly for small and occluded objects. The experimental results show that the developed model has higher precision, recall, and average precision than other variants of the YOLO algorithm while maintaining real-time performance. The study suggests that the application of attention mechanisms can improve the robustness of object detection in an autonomous driving environment without significant computational overhead. Future studies should focus on incorporating multimodal information and optimizing the model for real-time performance.

REFERENCES

1. Wang CY, Bochkovski A, Liao HYM. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2023. p. 7464-75. doi:10.1109/CVPR52729.2023.00721.

2. Ultralytics. Ultralytics YOLO26, YOLO11, YOLOv8: object detection, instance segmentation, semantic segmentation, image classification, pose estimation, object tracking [Internet]. GitHub; 2017. Available from: <https://github.com/ultralytics/ultralytics>
3. Wang CY, Yeh IH, Liao HYM. YOLOv9: Learning what you want to learn using programmable gradient information. In: European Conference on Computer Vision. Cham: Springer Nature Switzerland; 2024. p. 1-21.
4. Tian Y, Ye Q, Doermann D. YOLOv12: Attention-centric real-time object detectors. *Adv Neural Inf Process Syst*. 2026;38:78433-57.
5. Carion N, Massa F, Synnaeve G, Usunier N, Kirillov A, Zagoruyko S. End-to-end object detection with transformers. In: European Conference on Computer Vision. Cham: Springer International Publishing; 2020. p. 213-29.
6. Tamilarasi T, Muthulakshmi P, Miraei Ashtiani SH. Improved YOLO-based real-time brinjal detection algorithm for vision modules in harvesting robots. *Eng Res Express*. 2025;7(3):035234. doi:10.1088/2631-8695/ade00.
7. Ganguly B, Dey D, Munshi S. An attention deep learning framework-based drowsiness detection model for intelligent transportation system. *IEEE Trans Intell Transp Syst*. 2025;26:4517-27. doi:10.1109/TITS.2025.3544138.
8. Liang E, Wei D, Li F, Lv H, Li S. Object detection model of vehicle-road cooperative autonomous driving based on improved YOLO11 algorithm. *Sci Rep*. 2025;15:32348. doi:10.1038/s41598-025-18263-9. PMID:40897806.
9. Wei F, Wang WW. SCCA-YOLO: a spatial and channel collaborative attention enhanced YOLO network for highway autonomous driving perception system. *Sci Rep*. 2025;15:6459. doi:10.1038/s41598-025-90743-4. PMID:39987326.
10. Dustali A, Hasanlou M, Azimi SM. Comparative analysis of YOLO-based algorithms for vehicle detection in aerial imagery. *Int Arch Photogramm Remote Sens Spat Inf Sci*. 2025;XLVIII-G:411-6. doi:10.5194/isprsarchives-XLVIII-G-2025-411-2025.
11. Liu J, Zhan J, Zhang J, Chen J, Song Y, Tang L, et al. Robust scale fusion and edge-aware feature attention network for remote sensing UAV road detection under harsh weather. *Results Eng*. 2025;27:106172. doi:10.1016/j.rineng.2025.106172.
12. Devgan MS, Singh G, Sharma P, Kumar T, Cheng X, Ahlawat D. Attention-guided lightweight CNN-transformer fusion for real-time traffic sign recognition in adverse environments: HACTNet. *IET Intell Transp Syst*. 2026;20:e70167. doi:10.1049/itr2.70167.
13. Hameed R, HN M, P CS, Oggu VB, B I. Domain-adaptive reinforcement learning framework for robust autonomous vehicle navigation. In: Proceedings of the 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS); 2025. p. 1-6. doi:10.1109/IACIS65746.2025.11210918.
14. Geiger A, Lenz P, Urtasun R. Are we ready for autonomous driving? The KITTI vision benchmark suite. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2012. p. 3354-61. doi:10.1109/CVPR.2012.6248074.
15. Yu F, Chen H, Wang X, Xian W, Chen Y, Liu F, et al. BDD100K: a diverse driving dataset for heterogeneous multitask learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2020. p. 2633-42. doi:10.1109/CVPR42600.2020.00271.