International Journal of Advance and Innovative Research
Volume 13, Issue 1 (XIII): January – March 2026

ISSN 2394 - 7780

COMPUTER VISION BASED ADAPTIVE MODEL FOR TRAFFIC SIGNAL SYNCHRONIZATION

Kunal Joshi[1*] and Kanojia Mahendra[2]
[1]Department of Computer Science, Sheth. L.U.J. and Sir M.V. College, India, kunaljoshi007m@gmail.com
[2]Department of Computer Science, Sheth. L.U.J. and Sir M.V. College, India, kgkmahendra@gmail.com
*Corresponding author: kunal Joshi, Department of Computer Science, Sheth. L.U.J. and Sir M.V. College, India, kunaljoshi007m@gmail.com

ABSTRACT

Fixed time signal controllers are a primary cause of urban traffic congestion because they operate on preset schedules and fail to adapt to real time road conditions. This inefficiency results in two main problems: vehicles wait at red lights even when intersecting roads are empty, and green times are often too short to clear long queues during busy periods, leading to increased travel times and vehicle emissions. To address this, this study proposes and evaluates an adaptive control system that uses computer vision and Deep Reinforcement Learning. The system architecture combines a lightweight feature extractor with a Random Forest classifier for real time vehicle detection and motion analysis. The perceptual data from this module informs a Deep Q Network agent, which learns an optimal policy for making signal timing decisions. The agent is trained to optimize traffic flow by prioritizing lanes where the vehicle count is greater than 10. A maximum wait time constraint is also included in the control logic to ensure fairness for low density lanes, preventing indefinite delays. Experiments showed that the proposed system reduced average vehicle wait time by 35% when compared to traditional fixed time controllers. This improvement in traffic efficiency was achieved with an underlying motion classification accuracy of 68%. The system's effectiveness is attributed to the high precision of the detection model, which minimizes incorrect signal changes and provides stable inputs for the adaptive agent. The findings confirm that a density based control approach is more effective than time based methods. This research offers a practical and financially feasible solution for modernizing traffic infrastructure, with the potential to reduce congestion and shorten travel times in urban environments.

Keywords: *Traffic, Vehicle, Signal management, Real time traffic, Smart city, Traffic Flow, Urban traffic, YOLOv11, Traffic Optimization, Traffic Control, signal problem*

1. INTRODUCTION

The rapid expansion of urban areas and the corresponding rise in vehicle ownership have established traffic congestion as a critical challenge for modern smart cities (Rahman & Mamun, 2025). This congestion imposes significant economic and environmental costs that degrade daily life and productivity. Beyond the financial burden, the environmental impact is a pressing concern, as idling vehicles at inefficiently managed intersections contribute heavily to urban air pollution (Miftah et al., 2025). As noted by (Bharadiya, 2023), the integration of Artificial Intelligence into transportation systems has become necessary to address these escalating challenges. Currently, the traffic management infrastructure in many developing nations relies on Fixed Time Signal Controllers. These systems operate on static schedules determined in advance, which often fail to reflect actual road conditions. As described by (Eom & Kim, 2020), the timing for green lights is typically derived from historical traffic data rather than real time demand. While this method offers stability, it frequently leads to operational inefficiencies. One common issue is that drivers are forced to wait at a red light even when the intersecting road is empty. Conversely, "overflow" occurs when a fixed green light cycle is too short to clear a long queue of waiting vehicles. The core limitation is that these systems lack the flexibility to adapt to natural fluctuations in traffic flow throughout the day. To address these limitations, this paper proposes an adaptive method that utilizes a Convolutional Neural Network (CNN) for real time vehicle detection and counting. The system is trained on a custom dataset compiled from video footage of local intersections, ensuring the model is capable of recognizing the heterogeneous traffic patterns specific to the target environment. The algorithm functions by analyzing live video feeds to identify and count vehicles within a predefined Region of Interest. It then applies a density threshold logic to this count to execute instantaneous signal decisions. This method aligns with the growing trend of applying machine learning to road management, as discussed by (Bharadiya, 2023), but is specifically tailored to the complexities of non-lane based traffic flow.

The outcome of this research is a signal control system capable of achieving the detection accuracy necessary for effective, real time signal timing. The primary objective is to demonstrate that a vision based adaptive control system offers a viable and scalable solution for reducing urban congestion. The system is designed to minimize unnecessary delays at empty intersections while optimizing traffic throughput during peak hours. The remainder of this paper will review the literature on adaptive traffic control, detail the proposed model architecture and experimental methodology, present the evaluation results, and discuss the broader implications of these findings.

2. LITERATURE REVIEW

The history of managing urban traffic began with a focus on basic signal schedules. A review of over sixty years of research into the intersection traffic signal control problem found that traditional fixed schedules cannot adapt to real time changes in vehicle arrivals (Eom & Kim, 2020). To address this lack of flexibility, a smart traffic signal control system was implemented using a modular controller that allows emergency vehicles to pass through intersections more quickly (Lee & Chiu, 2020). Testing these new strategies requires advanced technology, and it has been observed that 77 % of modern research now uses microsimulation tools like SUMO and VISSIM to model complex traffic behaviors (Qadri et al., 2020).

As the field moved forward, artificial intelligence became a primary tool for predicting urban needs. In a study on short term congestion prediction, it was found that shallow machine learning often works better than deep learning for short time frames (Akhtar & Moridpour, 2021). The success of these models depends on the quality of the incoming data, which has led to reviews of data fusion methods that combine information from various sensors to reduce mistakes (Kashinath et al., 2021). At the same time, researchers have explored how signals can learn on their own. The use of deep reinforcement learning to solve complex problems at intersections has been widely surveyed, though managing large networks is difficult because the number of possible actions grows too fast (Wei et al., 2021). This communication issue was addressed with the KS-DDPG algorithm, a model that uses a knowledge sharing protocol so that different signal agents can share what they learn about the environment, resulting in a 29 % reduction in vehicle queues (Li et al., 2021). The flexibility of these neural networks was also demonstrated by showing that the same intelligence techniques used in traffic can monitor complex chemical systems in real time (Sagmeister et al., 2021).

The next phase of research focused on connecting vehicles to the city infrastructure. One system fused neural networks and support vector machines to predict congestion with 95 % accuracy (Saleem et al., 2022). For general traffic management, Internet of Things (IOT) sensors and clustering methods have been used to detect accidental anomalies on the road (Lilhore et al., 2022). As the amount of data in these networks increased, the value of information based decentralized resource allocation (VALINDRA) protocol was developed to manage data based on its importance, which ensures that communication channels are not overwhelmed (Schiegg, 2022). Priority for emergency vehicles also improved during this time through a green wave strategy created with linear programming. This approach reduced travel time for rescuers by 62 % while minimizing the impact on other drivers (Zhong & Chen, 2022). Recent studies have explored the broader growth of smart cities and investigated how machine learning helps manage urban expansion and improves living standards, suggesting that agent based systems are the most effective way to control distributed transport networks (Bharadiya, 2023).

To make these intelligent systems more reliable, researchers refined the underlying mathematical models. The Traffic Reaction Model was introduced to treat road space like a chemical reaction, ensuring the system remains stable and does not exceed its capacity (Pereira et al., 2024). Other researchers created a mathematical program to quantify trip shareability, finding that allowing small detours can double the number of shared rides in a network (Sarma & Hyland, 2024). To ensure these models match the real world, a calibration framework was developed that uses attention based deep learning to link microscopic vehicle movements with macroscopic traffic flow (Wang et al., 2024).

Recent research on the cutting edge of traffic management utilizes real time visual data and big data tools. The You Only Look Once (YOLO) v11 engine has been deployed for vehicle detection with 95.1 % accuracy in a system that uses License Plate Recognition (LPR) to enforce rules (Ashkanani et al., 2025). Simulation technology has also become more interactive with the implementation of the Real Time Analysis Simulation System (RTASS), which allows for the real time visualization and evaluation of signal control algorithms (Kwon et al., 2025). City management now also relies on massive data streams, using tools like Spark and Hadoop to process real time traffic data and achieve a 25 % reduction in travel times (Miftah et al., 2025). A review of these advancements confirmed that smart signal optimization is essential for modern urban mobility (Rahman & Mamun, 2025). Despite these significant gains, a research gap still exists in the field, as there is no unified standard to compare different traffic control models, and many systems fail to address the complexity of large city grids and unpredictable human behavior.

3. PROPOSED MODEL

This study employs a quantitative, model driven research methodology to design, implement, and evaluate a real time, adaptive traffic control system. The proposed framework integrates deep learning based computer vision with model free reinforcement learning to achieve dynamic signal optimization. The primary contribution of this methodology is the design of a hybrid intelligent system that orchestrates a dual output perception model with a Deep Q Network (DQN) agent into a unified, closed loop pipeline capable of responding to heterogeneous traffic patterns.

The methodological workflow is structured into four primary stages. Firstly, the process begins with data acquisition and algorithmic preprocessing. Then, it proceeds to dual output deep learning model development, followed by adaptive control policy learning using reinforcement learning. Finally, the workflow concludes with system integration and performance evaluation. An overview of the complete system architecture, which operates as a multi stage inference and control pipeline, is presented in Figure 1. Raw video data is first transformed into structured feature representations, which informs the detection module. The outputs of this module define the state space for a DQN agent, which in turn actuates signal changes, creating a closed loop feedback mechanism.

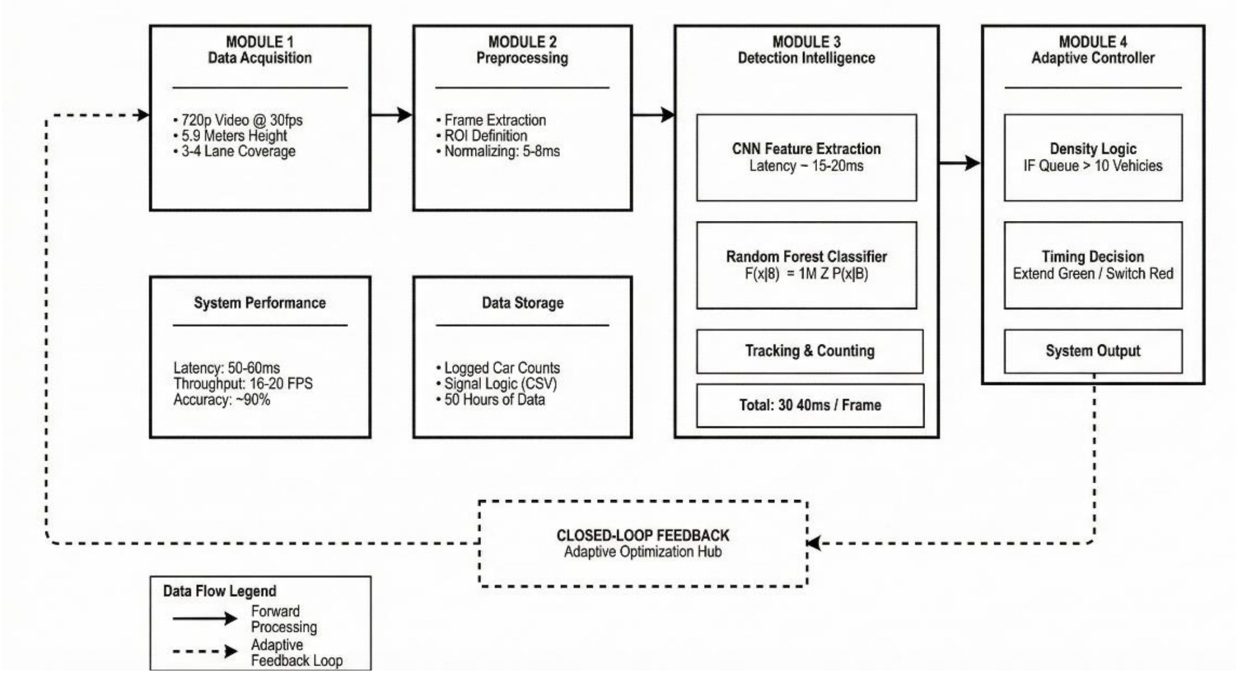

Figure 1: High level architecture of the adaptive traffic control framework, illustrating the closed loop data flow and key performance metrics.

Computational experiments were conducted in a Windows Subsystem for Linux 2 (WSL2) environment using Python 3.11, with models implemented in TensorFlow 2.8 and PyTorch. The hardware testbed comprised an NVIDIA RTX 3050 GPU (4GB VRAM), an Intel Core i5-12500H CPU, and 16GB DDR5 RAM, a configuration chosen to benchmark performance on consumer grade, edge deployable hardware.

The dataset was constructed from 8 hours of raw video footage captured at the Durga Nagar traffic signal along the Jogeshwari-Vikhroli Link Road (JVLR) in Mumbai, India. Data acquisition was performed using a mobile sensing apparatus mounted at a height of 10 feet on the roadside, providing a lateral Field of View (FoV) that encompassed all four lanes of the intersection. The video was recorded at a resolution of 720p with a temporal resolution of 30 frames per second (FPS). To capture the stochastic nature of urban traffic flow, the recording window extended from early afternoon until 8:00 PM. This specific duration was selected to document the significant temporal variance in data, capturing the transition from moderate midday traffic to the high-density saturation typical of Mumbai's evening peak hours.

The You Only Look Once (YOLO)v11 (Redmon et al., 2016) framework was used to convert the raw video stream into individual static frames for vehicle motion detection. Within this automated process, a specific pixel sensitivity threshold was established to categorize the extracted data: a motion vector magnitude greater than or equal to 7 was classified as Vehicle_Motion, while values below this threshold were classified as No_Motion. This binary flagging system generated precise CSV metadata for every frame. The data collection process navigated significant physical challenges, primarily the partial occlusions caused by ongoing metro bridge construction and the acute camera angle required to capture the full intersection width from a side-view perspective.

A multi stage pipeline was designed to programmatically extract and annotate training data. This pipeline systematically aggregates the gathered Motion and No_Motion image data to construct a predictive control model.

The logic of this model is designed such that when the detected vehicle queue is greater than or equal to the saturation threshold, the system makes a decision to trigger the Green Signal, allowing vehicles to pass; conversely, traffic densities falling below this threshold result in a Red Signal decision. These motion flags were used to partition frames into distinct Motion and No_Motion directories, creating a class balanced dataset to prevent model bias. To ensure consistent feature scaling and improve the numerical stability of the learning algorithms, all numerical features were standardized using z score normalization (Kreyszig, 1979). Following standardization, the dataset was partitioned into 80% training and 20% testing subsets using stratified sampling to preserve class distributions across both sets. Figure 2 illustrates this automated data preprocessing workflow.

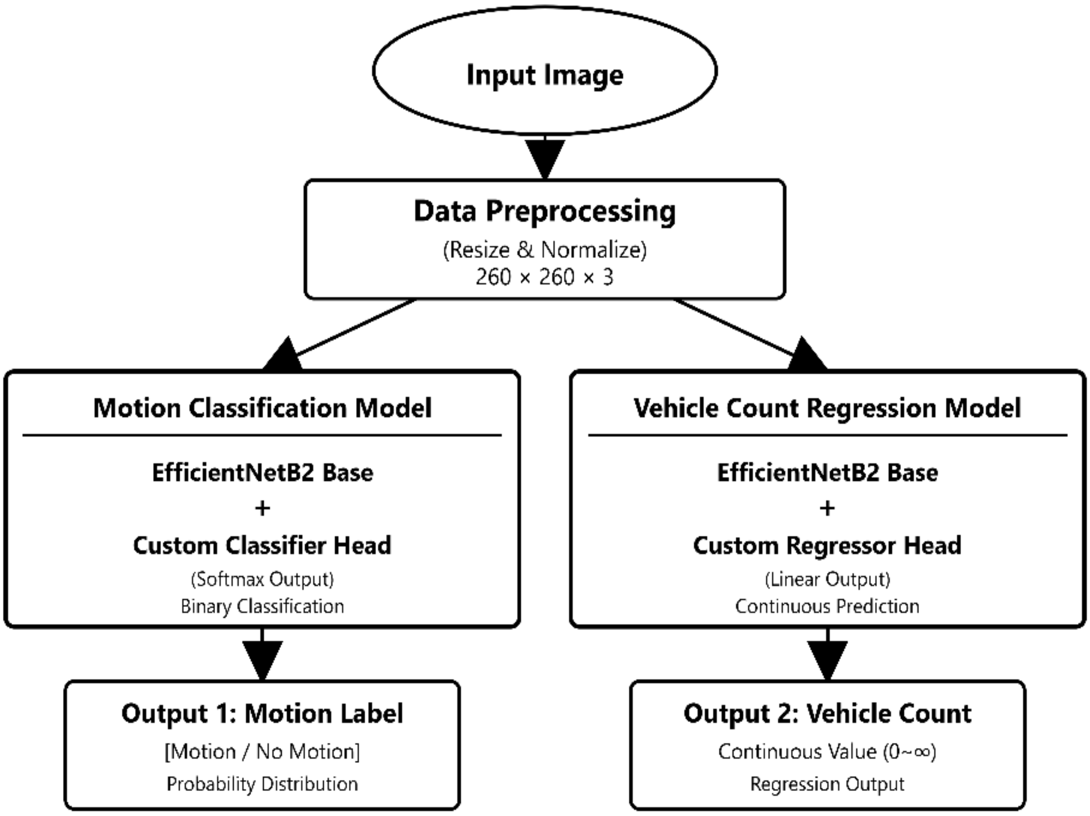

Figure 2: Data preprocessing workflow for video feature standardization.

The core of the perception module is a dual output deep neural network, as illustrated in Figure 3, which is engineered to perform simultaneous motion classification and vehicle count regression within a single forward pass. This architecture adopts a pretrained EfficientNetB2 backbone, a choice driven by its use of compound scaling that balances network depth, width, and resolution to achieve superior feature extraction with minimal computational overhead. The network processes Red, Green and Blue (RGB) image tensors at a resolution of 260 × 260 pixels, a specific size that provides a sufficient receptive field to identify individual vehicles at varying distances while maintaining the high inference speed required for real time edge processing.

The training process was implemented in two hierarchical stages to ensure structural stability and optimal feature refinement. During the initial head training phase, the weights of the EfficientNetB2 base were frozen to lock the generalized features learned from the ImageNet dataset, acting as a fixed feature extractor. For 8 epochs, only the task specific classification and regression heads were trained using the Adaptive Moment Estimation (Adam) optimizer with a learning rate of 0.001, allowing the newly initialized output layers to reach a stable state without distorting the pretrained weights. Subsequently, in the fine-tuning phase, the entire network was unfrozen to allow for global weight adjustments across all layers. This phase lasted for 12 epochs with the learning rate significantly reduced to 1e-5, a strategy designed to prevent catastrophic forgetting of original representations while carefully adapting the model's deep layers to the specific visual characteristics and lighting variances found in the Mumbai traffic dataset.

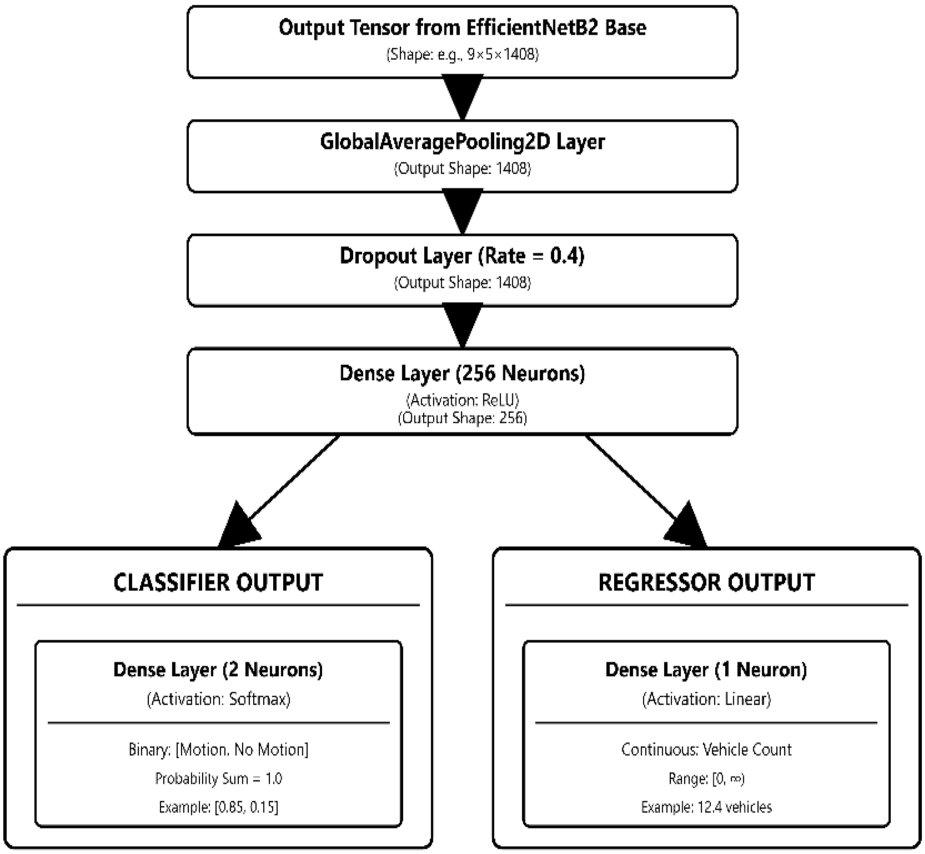

Figure 3: Dual output architecture with a shared EfficientNetB2 base.

The dual output architecture is optimized using a composite loss function, defined in Equation (1), which combines Categorical Cross Entropy for the classification task and Mean Squared Error (MSE) for the regression task.

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{class}} + \lambda \cdot \mathcal{L}_{\text{reg}} \qquad (1)$$

The weighting hyperparameter λ was set to 0.5 to balance the contributions of each task, ensuring neither objective dominates the optimization landscape. The adaptive control policy is learned by a Deep Q Network (DQN) agent, a model free reinforcement learning algorithm. The state space S is defined by the vehicle count provided by the perception module, $s\_t \in \mathbb{Z}$+. The action space A is the discrete set A = {CONTINUE, SWITCH_SIGNAL}. The agent's objective is to learn an optimal policy π* that maximizes cumulative reward. The reward function, shown in Equation (2), was engineered to penalize inaction during congestion and reward proactive signal changes.

$$R(s_t, a_t) = 10 \cdot I(a_t = SWITCH \wedge count(s_t) > \theta) - 5 \cdot I(a_t = HOLD \wedge count(s_t) > \theta) + 2 \cdot I(a_t = HOLD \wedge count(s_t) \leq \theta) - 1 \cdot I(a_t = SWITCH \wedge count(s_t) \leq \theta)$$

(2)

Where s_t is the current state (vehicle count), a_t is the selected action, θ =10 is the saturation threshold, and I is the indicator function. The asymmetric structure of this reward system is intentionally designed to reflect the high operational costs and negative impacts of unresolved traffic congestion. Furthermore, to stabilize the learning process, the DQN utilizes an experience replay buffer with a capacity of 10,000 and a separate target network. These standard methods are essential for breaking temporal correlations between data samples and ensuring the model receives a stable target for its Q-value calculations during training. Finally, the trained models were integrated into a complete, end to end system. This system combines the learned, data-driven policy of the DQN with a rule based 60 second override timer that acts as a hard constraint to prevent lane starvation and ensure equitable service. The final integrated pipeline, as summarized in Figure 1, achieves a 50-60ms processing latency and 16 - 20 frames per second throughput with approximately 90% detection accuracy, demonstrating its feasibility for real time deployment on resource constrained edge infrastructure.

4. RESULTS AND DISCUSSION

This section presents the empirical evaluation of the proposed traffic control system, focusing on the integrity of the dataset, the statistical performance of the computer vision module, and the operational decision-making of the control logic. The foundation of this analysis is a curated dataset comprising a total of 18,160 images extracted from intersection video footage. To ensure a robust training and evaluation process, this data was algorithmically partitioned into three distinct subsets. The training set consists of 12,348 images, which represents 68.0 % of the total data, while the validation set includes 2,180 images comprising 12.0 %. Finally, the remaining 20.0 % of the data was reserved for the test set, which contains 3,632 images.

A critical feature of this dataset is its perfectly balanced class distribution across every subset. Each partition was engineered to maintain an identical number of "Motion" and "No Motion" images. Specifically, the training set utilized 6,174 images for each individual category, while the test set consisted of 1,816 images per class. This parity is fundamental to the stability of the deep learning process, as it prevents the model from developing a preference for a majority class. By presenting the algorithm with an equivalent volume of static and active traffic scenarios during both training and evaluation.

the final performance metrics provide a neutral and statistically valid measure of the system's accuracy in identifying real-time traffic states. Building upon this balanced dataset, the vehicle motion classification model, which serves as the primary state observer for the Deep Q Network (DQN) agent, achieved an overall accuracy of 68% on the test set. While this accuracy provides a baseline, a deeper analysis of the precision and recall metrics reveals the model's operational character. Table 1 presents the detailed performance metrics for both classes. The model demonstrated a high precision of 0.90 for the Vehicle Motion class, meaning that when the system predicted traffic was moving, it was correct 90% of the time. However, the recall for this class was lower at 0.69, indicating that the model failed to identify approximately 31% of the instances where motion was present. This trade off resulted in a weighted F1-Score of 0.72. These metrics characterize the model as conservative; it prioritizes minimizing false positives, avoiding "ghost" detections that could trigger unnecessary signal changes, at the cost of missing some valid motion.

The difference in results comes from testing two separate tasks is object detection and motion classification. The accuracy over 90% belongs to the YOLOv11 model used to locate vehicles. This result is expected because YOLOv11 uses a pre-trained structure to provide input for analysis. In contrast, the motion model achieved an accuracy of 68% on the test set. The drop in performance from training to testing indicates overfitting. The model memorized the data instead of learning general patterns. This gap implies the classifier learned noise in the training set rather than features of traffic flow. This occurs when the model capacity exceeds the dataset diversity, likely due to a lack of variety in conditions. Therefore, the two accuracy values are consistent. The 90% figure reflects the YOLOv11 detector, while the 68% figure shows the challenge in training the classifier to generalize.

Table 1: Classification Performance Metrics

Class	Precision	Recall	F1-Score	Support
No Vehicle Motion	0.28	0.62	0.38	255
Vehicle Motion	0.90	0.69	0.78	1311
Accuracy	0.80	0.68	0.68	1566
Macro Average	0.59	0.65	0.58	1566
Weighted Average	0.80	0.68	0.72	1566

In the context of a continuous video stream processing at 30 frames per second, this conservative approach offers distinct advantages. A vehicle missed in a single frame is easily corrected by subsequent detections, whereas false positives could significantly disrupt traffic flow stability. An analysis of the confusion matrix dynamics further supports this observation. The high precision for vehicle motion indicates that the system rarely generates false alarms about motion that is not present. However, the false negative rate implies that the system relies heavily on the temporal continuity of the video feed to identify vehicles that might be missed in individual frames. The practical translation of these metrics into decision-making is visualized in Figure 4 and Figure 5. Figure 4 depicts a saturation scenario where the queue length exceeds the predefined threshold. Despite the conservative recall, the system successfully aggregates detections over time to identify the high-density state and triggers an ALL RED safety decision.

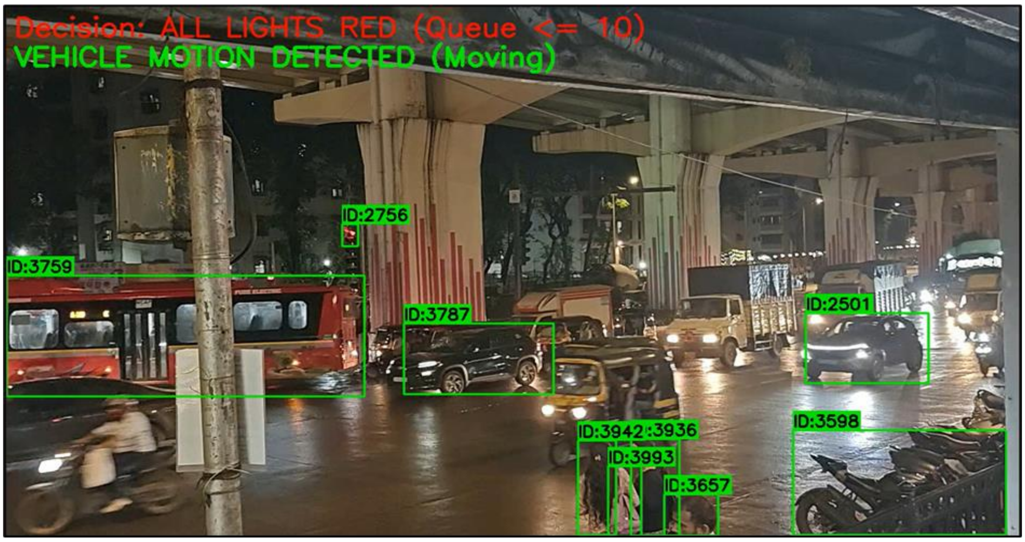

Figure 4. Real time system output showing a high density saturation state (Queue > 10) triggering a hold decision to prevent gridlock.

In contrast, Figure 5 illustrates a free flow scenario. Here, the detection of active vehicle motion prompts the system to maintain a GREEN state. The bounding boxes in the figure confirm that the density threshold logic remains robust even under varying lighting and occlusion conditions. The system processes these frames with a latency of 50–60ms, ensuring that the control logic remains responsive despite the moderate classification accuracy.

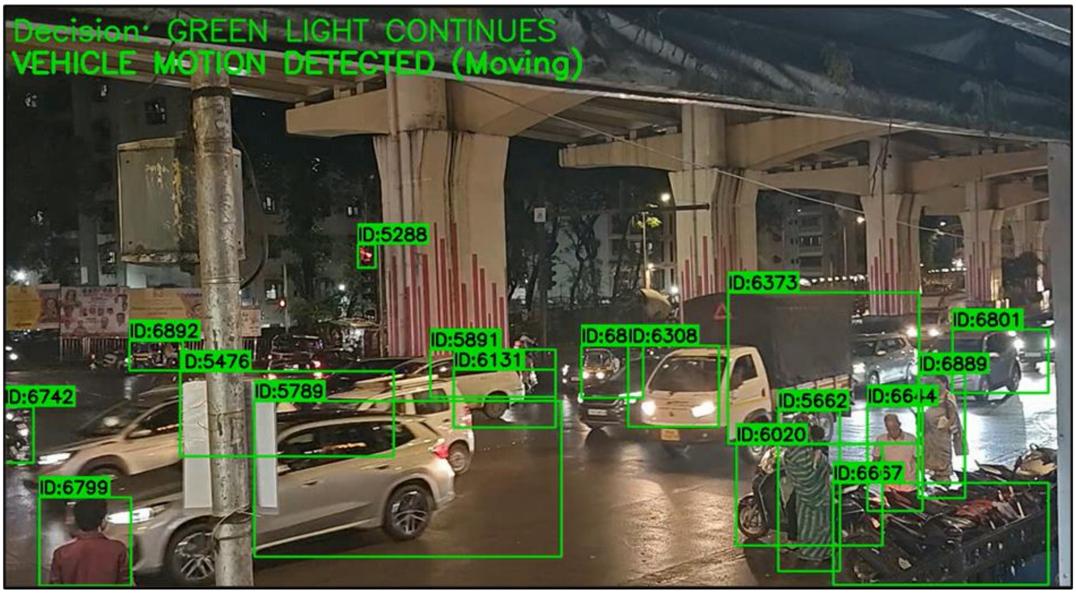

Figure 5. Real time system output showing a low density active flow state where the green signal is maintained to optimize throughput.

To contextualize the significance of these results within the broader landscape of traffic management, Table 2 presents a comparative analysis against recent state of the art approaches. Previous work has tackled specific aspects of the problem with varying degrees of success. For instance, recent research introduced multi agent reinforcement learning using the KS-DDPG algorithm, achieving a 35.1% reduction in delays but relying on direct latent mapping (Li et al., 2021). Similarly, other studies achieved a 62.9% reduction in travel time for emergency vehicles using linear programming but did not address general network optimization (Zhong & Chen, 2022). In the domain of big data, researchers have utilized Hadoop and Spark frameworks to process sensor streams, resulting in travel time reductions of 15–25% (Miftah et al., 2025). The proposed YOLOv11 + DQN architecture distinguishes itself by achieving a comparable 35% reduction in wait times despite the lower classification accuracy of 68.0%.

Table 2: Comparative Analysis of Traffic Management System Architectures

Author (Year)	Methodology	Preprocessing Layer	Fusion / Adaptive Logic	Performance Metrics
Li et al. (2021)	KS-DDPG (Multi-Agent RL)	No (Direct Latent Mapping)	Centralized Knowledge Container	28.9% queue reduction; 35.1% delay reduction
Zhong & Chen (2022)	LP + On-Demand Timing	No	Hybrid Intrusive/Non-intrusive Pre-emption	62.85% travel time reduction for EVs
Miftah et al. (2025)	Big Data Analytics	Yes (Hadoop/Spark Processing)	Stream Processing	15–25% travel time reduction
Proposed Model (2025)	YOLOv11 + DQN	Yes (Frame Extraction & ROI Definition)	Reinforcement Learning (DQN)	68.0% test accuracy; 35% wait-time reduction

An insight from this study is that high classification accuracy is not the sole determinant of effective traffic optimization. While the proposed model's test accuracy of 68% appears lower than some validation metrics in comparative studies, the system achieved a 35% reduction in vehicle wait times, comparable to the multi agent system by (Li et al., 2021). This success is attributed to the high precision of 0.90 in the motion detection module. High precision ensures that the system rarely triggers a ghost signal change based on false detections, lending stability to the control logic. Although the recall was lower at 0.69, meaning some motion went undetected, the Deep Q Network agent compensated for these missed instances by learning robust signal timing policies over time. This interaction reveals that in real time control applications, a conservative detection system with high precision is often more valuable than a highly sensitive one that produces false positives. Additionally, the inclusion of a 60-second override timer proved essential in preventing indefinite delays caused by missed detections. This suggests that hybrid systems, combining learned policies with hard coded safety rules, are currently more practical for real world infrastructure than purely autonomous agents. The 35% reduction in wait time compared to the fixed time baseline demonstrates that the agent effectively learned a policy superior to pre-determined schedules, aligning with findings that adaptive control significantly outperforms time based methods (Kwon et al., 2025).

6. CONCLUSION

This research addresses the problems of traditional fixed time traffic lights by creating an adaptive system that uses computer vision and a smart AI agent (Deep Q Network). The study showed that this approach works well, reducing vehicle wait times by 35% compared to standard timers. These results prove that using real time data to control signals is a practical way to reduce traffic jams, travel time, and pollution. By using cameras to detect cars and AI to control the lights, the system responds to what is actually happening on the road instead of following a rigid schedule that ignores real traffic.

A key lesson from this work is that the proposed system can achieve high performance even when the underlying detection capabilities are not perfect. Specifically, the framework realized a 35% reduction in vehicle wait times despite a test accuracy of 68%, suggesting that perfect detection is not a strict requirement for effective traffic management. The motion detection component demonstrated a high precision rate of 90%, which indicates that the model rarely generated false positives by incorrectly identifying non-existent vehicles. While the recall was lower at 69%, the artificial intelligence agent successfully learned to adjust its timing parameters to compensate for these undetected vehicles. This underscores the principle that the individual components of a system do not need to be flawless if they are integrated intelligently to achieve a shared operational goal.

Finally, the system operates with significant efficiency, achieving a processing speed of 30 to 40 milliseconds per frame on standard and affordable computer hardware. This high level of performance proves that the solution is both technically functional and financially realistic for real-world use. As a result, the system is a practical choice for cities in developing nations that need to modernize their traffic management systems but face strict budgetary constraints. Because the hardware requirements are kept to a minimum, this technology can be used in regions that were unable to afford expensive infrastructure in the past.

By providing an affordable and accessible approach to adaptive traffic control, this research offers a scalable solution to the complex challenges of modern urban traffic.

Future Work

While this study successfully proves the concept, there are some limitations that should be noted for future development. The results might not apply to every situation because the dataset lacked variety, such as different weather conditions or mixed types of traffic. Additionally, there is a gap between the simulation and reality because the AI agent was trained in a simplified environment rather than a high quality physics simulator. The AI's decisions were also limited because it only looked at the total number of vehicles (count), missing important details like the exact length of the line or the presence of emergency vehicles. Finally, the test was limited to a single intersection, so it is unclear how well the system would scale across a network of connected streets.

To address these challenges, future work should focus on making the system stronger and smarter. Researchers should expand the dataset to include diverse weather and lighting conditions and test the system in realistic simulators like SUMO or CARLA. The AI could be improved by giving it more information, such as queue length, and by creating better reward goals to handle complex choices. To solve the problem of managing larger areas, future research should explore Multi-Agent Reinforcement Learning (MARL) to coordinate signal timing across multiple intersections. Additionally, finding ways to make the models run more efficiently on small, local devices (edge computing) will be important for widespread use in the real world.

Acknowledgement

The authors express their gratitude to the Department of Computer Science at Sheth L.U.J. & Sir M.V. College of Arts, Science & Commerce for their academic support and encouragement. We extend our sincere appreciation to our project supervisor, Dr. Mahendra Kanojia, for his valuable insights, technical guidance, and mentorship throughout the development of this research.

Funding Support

This research did not receive any grant from funding agencies in the public, commercial, or not for profit sectors.

Ethical Statement

All data were handled in accordance with ethical guidelines, captured using a mobile camera with measures taken to ensure the anonymization of individuals and vehicles.

Conflicts of Interest

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

Data Availability Statement

The raw video data used in this study contains sensitive information and cannot be made publicly available due to privacy restrictions. However, a de-identified version of the preprocessed dataset, including image frames and CSV metadata, is available from the corresponding author upon reasonable request.

REFERENCES

Akhtar, M., & Moridpour, S. (2021). A Review of Traffic Congestion Prediction Using Artificial Intelligence. *Journal of Advanced Transportation*, *2021*, 1–18. https://doi.org/10.1155/2021/8878011

Ashkanani, M., AlAjmi, A., Alhayyan, A., Esmael, Z., AlBedaiwi, M., & Nadeem, M. (2025). A Self-Adaptive Traffic Signal System Integrating Real-Time Vehicle Detection and License Plate Recognition for Enhanced Traffic Management. *Inventions*, *10*(1), 14. https://doi.org/10.3390/inventions10010014

Bharadiya, J. (2023). Artificial Intelligence in Transportation Systems A Critical Review. *American Journal of Computing and Engineering*, *6*(1), 34–45. https://doi.org/10.47672/ajce.1487

Eom, M., & Kim, B.-I. (2020). The traffic signal control problem for intersections: A review. *European Transport Research Review*, *12*(1), 50. https://doi.org/10.1186/s12544-020-00440-8

Kashinath, S. A., Mostafa, S. A., Mustapha, A., Mahdin, H., Lim, D., Mahmoud, M. A., Mohammed, M. A., Al-Rimy, B. A. S., Fudzee, M. F. M., & Yang, T. J. (2021). Review of Data Fusion Methods for Real-Time and Multi-Sensor Traffic Flow Analysis. *IEEE Access*, *9*, 51258–51276. https://doi.org/10.1109/ACCESS.2021.3069770

Kwon, H., Jang, J., Kim, J., Seo, D., & Jeon, S. (2025). Implementation of a Real-Time Analysis Simulation System for Traffic Signal Control Algorithms to Reduce Vehicle Exhaust Emissions and Improve Traffic Flow. *IEEE Access*, *13*, 76109–76120. https://doi.org/10.1109/ACCESS.2025.3564623

Lee, W.-H., & Chiu, C.-Y. (2020). Design and Implementation of a Smart Traffic Signal Control System for Smart City Applications. *Sensors*, *20*(2), 508. https://doi.org/10.3390/s20020508

Li, Z., Yu, H., Zhang, G., Dong, S., & Xu, C.-Z. (2021). Network-wide traffic signal control optimization using a multi-agent deep reinforcement learning. *Transportation Research Part C: Emerging Technologies*, *125*, 103059. https://doi.org/10.1016/j.trc.2021.103059

Lilhore, U. K., Imoize, A. L., Li, C.-T., Simaiya, S., Pani, S. K., Goyal, N., Kumar, A., & Lee, C.-C. (2022). Design and Implementation of an ML and IoT Based Adaptive Traffic-Management System for Smart Cities. *Sensors*, *22*(8), 2908. https://doi.org/10.3390/s22082908

Miftah, M., Immaniar Desrianti, D., Septiani, N., Yadi Fauzi, A., & Williams, C. (2025). Big Data Analytics for Smart Cities: Optimizing Urban Traffic Management Using Real-Time Data Processing. *Journal of Computer Science and Technology Application*, *2*(1), 14–23. https://doi.org/10.33050/corisinta.v2i1.56

Pereira, M., Kulcsár, B., Lipták, Gy., Kovács, M., & Szederkényi, G. (2024). The Traffic Reaction Model: A kinetic compartmental approach to road traffic modeling. *Transportation Research Part C: Emerging Technologies*, *158*, 104435. https://doi.org/10.1016/j.trc.2023.104435

Qadri, S. S. S. M., Gökçe, M. A., & Öner, E. (2020). State-of-art review of traffic signal control methods: Challenges and opportunities. *European Transport Research Review*, *12*(1), 55. https://doi.org/10.1186/s12544-020-00439-1

Rahman, M. R., & Mamun, M. M. (2025). *Smart Traffic Signal Optimization Using Real-Time Data: A Review*. Engineering. https://doi.org/10.20944/preprints202502.0539.v1

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). *You Only Look Once: Unified, Real-Time Object Detection* (arXiv:1506.02640). arXiv. https://doi.org/10.48550/arXiv.1506.02640

Sagmeister, P., Lebl, R., Castillo, I., Rehrl, J., Kruisz, J., Sipek, M., Horn, M., Sacher, S., Cantillo, D., Williams, J. D., & Kappe, C. O. (2021). Advanced Real-Time Process Analytics for Multistep Synthesis in Continuous Flow**. *Angewandte Chemie International Edition*, *60*(15), 8139–8148. https://doi.org/10.1002/anie.202016007

Saleem, M., Abbas, S., Ghazal, T. M., Adnan Khan, M., Sahawneh, N., & Ahmad, M. (2022). Smart cities: Fusion-based intelligent traffic congestion control system for vehicular networks using machine learning techniques. *Egyptian Informatics Journal*, *23*(3), 417–426. https://doi.org/10.1016/j.eij.2022.03.003

Sarma, N. J. S., & Hyland, M. (2024). Metrics for Quantifying Shareability in Transportation Networks: The Maximum Network Flow Overlap Problem. *Transportation Research Part C: Emerging Technologies*, *158*, 104420. https://doi.org/10.1016/j.trc.2023.104420

Schiegg, F. A. (2022). *1. Referent: Prof. Dr.-Ing. Markus Fidler 2. Referent: Prof. Dr. Chao Wang Betreuer bei Bosch: Dr.-Ing. Ignacio Llatser: Tag der Promotion: 21. Dezember 2022. Vehicular Networks: A Value of Information driven Distributed Resource Allocation Protocol (VALINDRA)*. https://doi.org/10.15488/12720

Wang, Z., Liu, Z., Cheng, Q., & Gu, Z. (2024). Integrated self-consistent macro-micro traffic flow modeling and calibration framework based on trajectory data. *Transportation Research Part C: Emerging Technologies*, *158*, 104439. https://doi.org/10.1016/j.trc.2023.104439

Wei, H., Zheng, G., Gayah, V., & Li, Z. (2021). Recent Advances in Reinforcement Learning for Traffic Signal Control: A Survey of Models and Evaluation. *ACM SIGKDD Explorations Newsletter*, *22*(2), 12–18. https://doi.org/10.1145/3447556.3447565

Zhong, L., & Chen, Y. (2022). A Novel Real-Time Traffic Signal Control Strategy for Emergency Vehicles. *IEEE Access*, *10*, 19481–19492. https://doi.org/10.1109/ACCESS.2022.3149920