Creating Intelligent Devices To Improve The Commutation Sector

Dr. Aravindhan k
School of Computer Science and Engineering
Reva University
Bangalore,India
aravindhan.kurunthachalam@reva.edu.in

D.Abinav Datha
School of Computer Science and Engineering
Reva University
Bangalore,India
dornadulaabhinav@gmail.com

Y.M.Hanish Kumar Reddy
School of Computer Science and Engineering
Reva University
Bangalore,India
hanishym9000@gmail.com

Y.Someswar Naidu
School of Computer Science and Engineering
Reva University
Bangalore,India
someswarnaidu58@gmail.com

G.Chakravardhan Reddy
School of Computer Science and Engineering
Reva University
Bangalore,India
gaddamchakri21@gmail.com

Abstract— Traffic in megacities is an interesting problem. On paper its straightforward, yet in practice understanding how to effectively manage congestion is hard. Foreseeable solutions include slot-based traffic systems or, more innovatively, adaptive machine learning based systems in which the cities adapt to patterns instead of the reverse. Today, a majority of cities still run using an outdated variant of fixed-clock systems in which traffic lights switch with a predefined time scale based on timers set years ago, denying the system the ability to adapt to real-time events. We have a system in which congestion occurs, ambulances are forced to wait behind the congested traffic, and significant amounts of fuel are wasted idling, all while the lights are switching according to a set timer. This paper develops TrafficAI, a system in which cameras feed real time video analysis, vehicle detection, and counting to inform the timers of live, dynamic data in an effort to seek congestion reduction. To avoid overreacting, we parameterized the inputs using an Exponential Moving Average (EMA) approach to smooth out the timeseries. To mitigate robotic flickering, and using an hysteresis approach, we used a decision engine that deemed when to switch at an appropriate rate. Our system also implemented a Finite State Machine (FSM) to guarantee that state transitions occurred safely (transition events came in order with the colors). Finally, our system detected the presence of an emergency vehicle and activated a signal override to ensure the firetruck and ambulance can pass unencumbered. Our experiments demonstrated that we could achieve measurable improvements in traffic flow..

Keywords— Traffic Management System, Artificial Intelligence, Computer Vision, YOLOv8, Smart Traffic Signals, Exponential Moving Average, Finite State Machine, Emergency Vehicle Detection, Intelligent Transportation Systems

I. INTRODUCTION

Any driver who has stared at a red light at an empty intersection has experienced the lesson that something is wrong with the way we operate. Our traffic systems have not caught up to the huge and expanding demands placed by our burgeoning cities. Fixed-timer signals were a practical way to handle traffic in the 50s. They are not such in 21st century urban centers.

This is an issue of adaptability. Conventional systems either timer-based (box), or using buried inductive loop detectors do not have a clue what they are doing. They cannot tell between a light-trafic Sunday morning or a heavy-traffic Friday night. They cannot identify an approaching emergency vehicle. They just countdown and go, blind to their surroundings.

The advancements in recent year in both Artificial Intelligence and Computer Vision can provide another path. Deep learning solutions in the form of convolutional neural network including YOLO (You Only Look Once) are capable of real time classification of vehicles in an accurate manner. Combine that with intelligent decision making and you have the beginning for a very reactive traffic system.

Of course this is more difficult to deploy in the real world than it is to demonstrate in lab. Detection is noisy partially occluded car, a shadow, a reflection can all cause detection issues. Switching signals from raw detections, with no temporal smoothing, results in unreliable behavior. Switching a signal without a safe transition sequence, can create a dangerous situation for drivers. These are non-trivial engineering issues.

TrafficAI will be used as a case study to introduce a system that takes all of these constraints into account. TrafficAI uses YOLOv8 based vehicle detection, EMA based density smoothing, hysteresis based decision logic and a Finite State Machine to achieve safe signal control. An emergency vehicle detection application is also built that uses HSV color heuristics to identify the ambulances and fire trucks, thus allowing signal priority to be gained without having to train a dedicated model. I wanted to design a system that will work reliably in the real world and not just on a benchmark dataset.

The main goal of the present work is the development of a real-time intelligent traffic management system capable of providing efficient traffic flow, alleviating traffic congestion and providing safe and adaptive signal control based on an intelligent system.

II. Literature Survey

Intelligent traffic management has been evolving over the years. Researchers have been dealing with a wide variety of subjects, such as signal control based on deep learning or reinforcement leaning. Some of these works are presented below and research gaps that TrafficAI will fill are highlighted.

A. Object Detection and YOLO-Based Traffic Monitoring

Object detection is a core component of any vision-based traffic system. Amongst the various object detection architectures available, the YOLO family has arguably become the dominant choice for its real-time applications, mainly due to its single-pass architecture which trade some accuracy for vast speed improvements. Zhang et al. [1] have shown that an improved YOLOv8 architecture, with multi-scale feature fusion, can process the overhead views seen by drone-assisted urban monitoring. Li et al. [2] have improved detection further in cluttered urban scenes with the use of a SE attention mechanism and deformable convolutions. Carrasco et al. [3] systematically surveyed all YOLO variants used in traffic sign detection and attribute YOLOv8 with the most promising speed and accuracy for real-time use.

B. Adaptive Traffic Signal Control Systems

The drawbacks of fixed-timer systems has been studied, and adaptive-signal control developed, for nearly 50 years. Shah et al. [4] examined machine learning centered solutions, concluding that dynamic detection of demands and their allocation to signals would consistently outperform using fixed timers, in wait time for vehicles, and noted that this was especially true during rush hours. Reinforcement learning has shown drivers of promise; for example, Polydoros et al. [5] cited is state-based decision mechanisms, such as hysteresis thresholds, as an effective way to avoid the oscillating effects associated with naively adaptive signals, which can drive them to be worse than fixed timers. Li et al. [6] fused vision-based AI and large language models in an innovative hybrid AI approach, which could be more widely adaptable to different real-world conditions.

C. Emergency Vehicle Detection and Prioritization

Vehicles' prioritization is not an uncommon matter and has been pondered by radio based or GPS based traffic control that necessitates into-vehicle and at-intersection additional equipment. The recent research has considered vision based approaches for that matter. Alaoui et al.[7] implemented a YOLOv8 model compatible with pan-tilt-zoom cameras for the indication of emergency vehicles and modification of the signals for their priority access; the initial results are encouraging when used in smart cities context. Johny and Sharma [8] tested several detection platforms and discovered the fact that YOLOv8 was the most robust one against different light and environment conditions is an important aspect in practice..

D. Traffic Density Estimation and Smoothing Techniques

Determining the true density of a lane is slightly more complex than a glance at the number of cars. Putra et al [9] performed one of the first reviews of vision-based density estimates; they found that hybrid methods combining detection with a temporal smoothing function always worked better than an instant-by-instant counts. Sankaranarayanan et al [10] convincingly argued that temporal smoothing is essential for steady traffic measurements, as counting one frame at a time results in too much noise to make a reliable signal decision. Kumar et al [11] proved machine learning based density estimation can optimally program traffic signals in the absence of any physical sensors invaluable knowledge for cost cutting.

E. Intelligent Traffic Management Systems

Multiple research teams have gone the route of complete system design over individual parts in order to create integrated systems with detection, estimation and control. Abbas et al. [12] developed a vision based traffic light control system using a Faster R-CNN that allocated lanes the appropriate amount of time based on the number of vehicles in them, and showed that it was possible to create an end-to-end integrated system in a controlled environment, an example of the holistic approach that new urban AI systems should take.

F. Research Gaps and Motivation

As one can read through this corpus, there are several consistent issues that appear to not have been solved. Detection instability, even for established methods, is an ongoing issue, and most methods don't have a means of dealing with noisy or unstable detection outputs. For some systems, the signal switching logic is hardly more sophisticated than simple flip-flops and thus oscillates or spends too much green-time. Emergency vehicle detection systems often require special hardware and/or separate trained models. And for those systems that rely on expensive infrastructures, scalability can be an issue.

To overcome these limitations, TrafficAI was created. Integrating YOLOv8 detection with EMA smoothing, hysteresis decision logic, and a light-weight HSV based emergency detection heuristic, ensures the system to be reliable and practical for real-world application.

III. METHODOLOGY

A. System Description

The principle of the TrafficAI system is a simple one: "Don't use a clock, look at what is happening, and then act". The route is determined using the live video streams from a camera at an intersection as input. This is then fed into a pipeline of artificial intelligence and signal control algorithms which predict load, count traffic, allocate lanes for the green, and so forth. The intention is to move away from the dead, unresponsive behavior of the conventional timers.

B. System Architecture

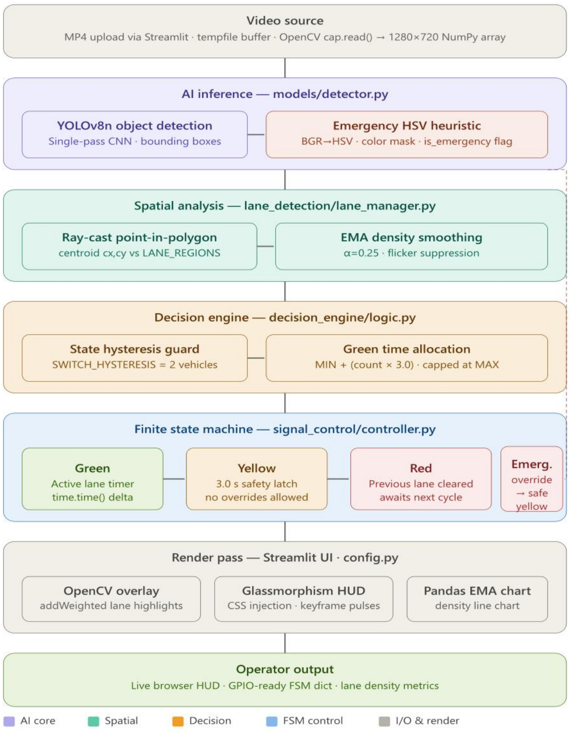

Figure 1: TrafficAI — Vision-Based Intelligent Traffic Signal Optimization with Emergency Vehicle Prioritization

The system is organized into five modules, each handling a distinct step in the processing pipeline:

- Vehicle Detection Module (YOLOv8)– Performs detection and classification of vehicles frame by frame. It is a pretrained deep learning network.
- Lane Mapping Module (OpenCV + Geometry) - Finds the corresponding lane each vehicle is in by geometric centroid calculus and polygon based lane regions.
- Traffic Density Estimation Module (EMA) Transforms raw vehicle counts into stabilized lane-wise densities.
- Decision Engine(Hysteresis Logic) Determines which lane should be given a green light, using hysteresis thresholds to prevent frequent changes.
- Signal Control Module (FSM) Responsible for the state transition from Green to Yellow to Red, to make sure no transition is performed until it can be achieved safely and with proper order.

It is desirable to keep this modules separate for possible maintainability, testability and ability to extend. For example, we can just replace the yolov8 detection module by a better detection model.

C. Core Algorithms

1. Exponential Moving Average (EMA)

Raw detection counts are noisy to begin with. Just a headlight reflection can cause a lane‘s apparent density to spike or dip one frame. A passing vehicle trip using a detection model that is a pixel off, a blind spot in the camera all of these can cause the raw lane density to have false positives.

What they do instead is take a simple exponentional moving average (EMA) of each lane, which takes the average of the new value and the previous time step‘s. This combines the new value and old smoothed value into the new value, returning a smooth series that represents the real variance and not the frame noise. Instead of being noisy, the values "hold steady" when things really haven‘t changed.

2. Time Allocation Strategy

In TrafficAI, Signal durations are not predetermined. Since the higher EMA values indicate greater densities of vehicles, those lanes are allocated longer green times, while the lighter lanes are allocated shorter times to clear the accumulated vehicles. This intelligently balances the need to move vehicles out of jammed Lanes quickly, with the desire to allocate a fair amount of time to the other Lanes. This value is recomputed at every decision step based on current values.

3. Hysteresis-Based Decision Logic

Another less obvious but more critical choice in the design of TrafficAI is the use of hysteresis in the criteria used for switching signals. Without the hysteresis, a system that always switches to the densest lane at each decision point may oscillate between two lanes of near-equal density, never giving a green light long enough to clear any cars, or worse, back-and-forth between two further-apart lanes leaving no lane green long enough to clear any cars. Hysteresis prevents this with a threshold that must be reached requiring the competing lane to be a certain amount denser than the present lane before switching to retain the yellow state and not slip back and forth.

4. Emergency Vehicle Detection (HSV Heuristic)

Emergency vehicle detection in TrafficAI doesn‘t use a separate trained classifier. It simply uses HSV color-space filtering to look for the general color signatures of emergeny vehicles--the red and blue flashing lights of ambulances and fire trucks and police cars. When these signature colors become visible within a lane region, an emergency override is triggered immediately grabbing that lane at high priority. This is a cheap, no retraining, no additional data solution that provides a first line of defense that works under urban intersection lighting conditions

D. Finite State Machine (FSM)

The signal control logic is implemented as a FS: this guarantees that transitions between signals states always follow a safe, predictable sequence. The FSM operates across three primary states:

- Green — Active flow state for the selected lane.
- Yellow — A mandatory intermediate phase that warns drivers before the signal changes, using a fixed delay.
- Red — Inactive state for lanes not currently selected.

The FSM also manages emergency interrupts while in this framework, so even high-priority transitions experienced the Yellow warning period first before transition occurs. This prevents harsher movements which could lead to hazardous conditions at interections.

E. Implementation Tools

The system was implemented using the following tools and libraries:

- YOLOv8 — Real-time vehicle detection.
- OpenCV — Image processing, geometric lane mapping, and HSV color filtering.
- NumPy — Efficient numerical computation for EMA calculations and array operations.
- Streamlit — Interactive real-time dashboard for monitoring signal states and lane densities.

IV. RESULTS AND ANALYSIS

A. *Real-Time Vehicle Detection and Lane Mapping*

Figure 1: Real-Time Vehicle Detection and Lane Segmentation

B. The very first thing that the system needs to get right is seeing vehicles and understanding which lane they are in. Tests showed that it did this very well. The YOLOv8 model monitored all four lanes at all times and reliably drew bounding boxes, and correctly classified all vehicles in real time. The fire engines and police cars were easily distinguished from most other vehicles.

C. Lane assignment which is performed using geometric centroid computation and segment (cell) based lane regions functioned well at all traffic levels. Switched vehicles were well tracked with the lanes updated smoothly as the centroids traveled through the cells. The detection and mapping layer is the foundation upon which future algorithms are built upon.*Emergency Event Logging and System State Transition*

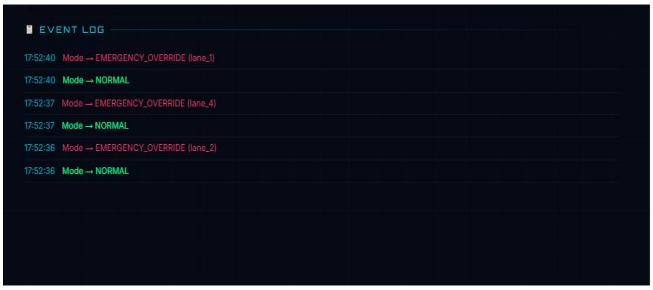

Figure 2: Event Log Showing Emergency Override Transitions

The event log showed the way the system responded when a fire engine entered the system. In several tests the system correctly responded to the appearance of a fire engine in the lanes in various positions creating the correct EMERGENCY_OVERRIDE state. Whether the fire engine was in lane 1, lane 4 or lane 2 the system came back to the normal operation cleanly after the fire engine had left. There was no additional overriding not required.

This is critical. Whilst a system in override state for too long is almost unacceptable, one that never reacts to an emergency situation is just as bad. Logs reveal that TraffiC AI handles both entry into and exit from emergency mode correctly.

D. *Signal Control and Lane Density Visualization*

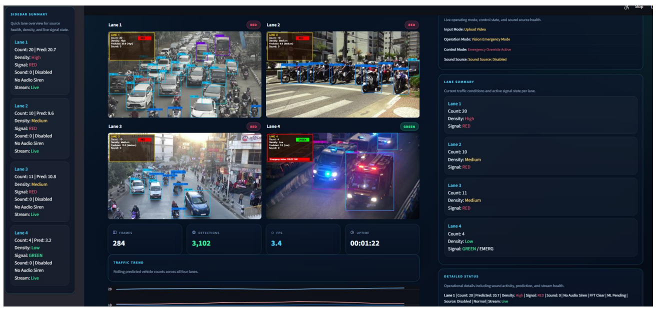

Figure 3: Signal Grid and Lane Density Monitoring Interface

E. . Intelligent Traffic Optimization and System Behavior

The dashboard also provides a screen that exposes the system‘s internal state in real time. Using the signal grid the test server always revealed how much the system‘s decisions are ever changing: the system‘s states transiting from red to yellow to green from lane to lane as the flow of traffic was like, all of those transitions being in the true sequence order required by the FSM.

One particular order this occurred was with lane 4 in EMERGENCY_OVERRIDE and lane 1 in Yellow completely realistic situation, two things being done at once, and technically there‘s nothing wrong with that! System behaved well, with safe state management throughout all lanes, lane 4 onto EMA and lane 1 in the Yellow Phase, no conflicts The EMA density readings proved to be pretty representative of actual traffic load: there were substantial changes both across lanes and over time. You could see the smoothing quite clearly: the density values changed very smoothly from frame to frame, as the density signal control system should allow.***D. System Performance and Traffic Trend Analysis***

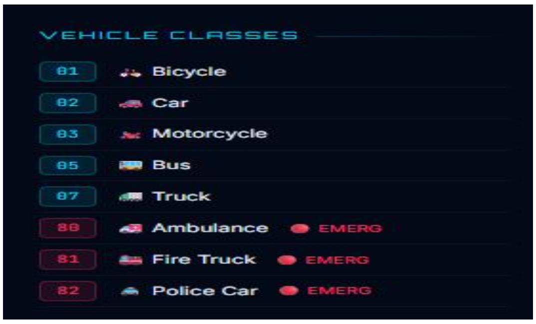

Figure 4: Vehicle Detection Metrics and Lane Density Trends

Frames per second demonstrated a speed of above 15 fps, suitable for real-time operation on an off-the-shelf hardware platform. Performance data collected during the test session indicated the system has been able to sustain a throughput of detecting 576 vehicles within a period of 35 minutes of 186 frames.

The lane density trend graphs were very valuable. Instead of the high variance, jagged lines that would have resulted from raw frame-by-frame counts, the trends averaged over adjacent frames provided smooth, relatively realistic changes in congestion over time. The signals and decisions based on these smoothened values were very reliable and predictable, without the oscillation that would occur if each those values were based on each tiny change in flow.

Figure 5: Intelligent Traffic Optimization and System Behavior

The whole system view rather than individual componentsshowed the overall system behaved correctly confirming that the structure was the correct one to address the sets of requirements: Signal priority showed the cues to always be correct with regard to lane. Congestion, the hysteresis constraint prevented rapid switching between congested lanes having high density, and the emergency overrides behaved correctly and cleared quickly.

The thing that finally makes the system “robust” enough tohave it running on real roads is that it is operationally predictable and not only under best conditions and that it behaves controllably in the mess that reality throws at intersections. Moreover, the twin EMA smoothings and the hysteresis logic serve to eliminate many a tremor in the traffic control decisions.
As results, obviously found that the proposed system will intelligente traffic optimization according to circumstance.

V. REFERENCES

[1] Zhang *et al.*, "Analysis of Vehicle and Pedestrian Detection Effects of Improved YOLOv8 Model in Drone-Assisted Urban Traffic Monitoring System," *PLOS ONE*, vol. 20, no. 3, p. e0314817, Mar. 2025.

[2] Li *et al.*, "Enhanced Multi-Target Detection in Complex Traffic Using an Improved YOLOv8 with SE Attention, DCN_C2f, and SIoU," *World Electric Vehicle Journal*, vol. 15, no. 12, p. 586, Dec. 2024.

[3] M. Carrasco *et al.*, "Traffic Sign Detection and Recognition Using YOLO Object Detection Algorithm: A Systematic Review," *Mathematics*, vol. 12, no. 2, p. 297, Jan. 2024.

[4] R. Shah *et al.*, "Machine Learning Solutions for Adaptive Traffic Signal Control: A Review of Image-Based Approaches," *World Journal of Advanced Engineering Technology and Sciences*, vol. 13, no. 1, pp. 476–481, Sep. 2024.

[5] A. Polydoros *et al.*, "Traffic Signal Control via Reinforcement Learning: A Review on Applications and Innovations," *Infrastructures*, vol. 10, no. 5, p. 114, May 2025.

[6] Y. Li *et al.*, "Incorporating Vision-Based Artificial Intelligence and Large Language Model for Smart Traffic Light Control," *Engineering Applications of Artificial Intelligence*, Elsevier, 2025.

[7] A. Omari Alaoui *et al.*, "Real-Time Traffic Signal Adjustment Using YOLOv8 for Improved Integration of Emergency Vehicles in Smart Traffic Systems," *Signal, Image and Video Processing*, Springer Nature, 2025.

[8] C. Johny and A. Sharma, "Emergency Vehicle Detection for Signal Optimization Using YOLOv8," *Journal of Electrical Systems*, vol. 20, no. 3, pp. 6952–6960, 2024.

[9] R. Putra *et al.*, "A Systematic Review on Vision-Based Traffic Density Estimation for Intelligent Transportation Systems," *IET Intelligent Transport Systems*, Wiley, 2025.

[10] M. Sankaranarayanan, C. Mala, and S. Jain, "Traffic Density Estimation for Traffic Management Applications Using Neural Networks," *International Journal of Intelligent Information Technologies*, vol. 20, no. 1, pp. 1–19, 2024.

[11] V. Kumar *et al.*, "Machine Learning Based Adaptive Traffic Prediction and Control Using Edge Impulse Platform," *Scientific Reports*, Nature Portfolio, May 2025.

[12] S. Abbas *et al.*, "Vision Based Intelligent Traffic Light Management System Using Faster R-CNN," *CAAI Transactions on Intelligence Technology*, IET/Wiley, Apr. 2024.