International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Enhanced Traffic Management System Using AI-Based Ambulance Detection

Dr. Lokendra Singh Songare[1], Chanakshi Anerao[2], Diksha Goyal[3], Harshvardhan chouhan[4]

[1]Assistant Professor, Department of Computer Science & Engineering, Medi-Caps University, Indore, India

[2,3,4] Students, Department of Computer Science & Engineering, Medi-Caps University, Indore, India

lokendra.singh@medicaps.ac.in, chanakshianerao54@gmail.com, dikhagoyal11092003@gmail.com, Chouhanharshvardhan18@gmail.com

Abstract--- Abstract—As cities grow, traffic congestion has become a major issue, especially when it comes to getting emergency vehicles like ambulances through busy streets. In this paper, we propose a new system that uses AI to help solve this problem. The system detects ambulances in real time through video feeds, using the YOLOv8 model, and then changes the traffic light to green in the lane the ambulance is traveling in. Instead of using OpenCV to just count the number of cars, the system takes a look at how many vehicles are in each lane and changes the light timings to keep traffic moving smoothly. When we tested it, the system performed well, with an impressive recall of 0.96 and an F1-score of 0.84, which shows it's both accurate and quick to react.

By looking at things like confusion matrices and precision-recall graphs, we confirmed that the system works well. In the end, this approach can help ambulances get through traffic faster and also make the overall road system run smoother.

Index Terms—Emergency Vehicles, Traffic Signals, AI in Transportation, Ambulance Routing, Real-Time Systems, Vehicle Detection, Traffic Management.

I. INTRODUCTION

City traffic is getting out of hand. Like, there are just way too many cars, and it's only getting worse. Now imagine an ambulance stuck in that. Not ideal, right? Every second counts when someone's hurt, but traffic doesn't care. The problem is, traffic lights just keep going through the same routine—green, yellow, red—no matter what's happening on the road. No awareness. No flexibility.

And that's the thing. We've got all this advanced tech now, so why are we still using systems that act like it's the 90s? AI can actually help here. Not in some sci-fi way, but for real. People have been testing out deep learning models like YOLO and ResNet to spot ambulances using live traffic cam footage. Some setups can even catch the sound of sirens. Others check how crowded a lane is with simple sensors or video feeds. The idea is, the system sees or hears what's happening and then adjusts the lights on the fly. Makes sense, right?

There's this idea—emergency vehicle priority. Basically, the tech recognizes an ambulance is coming and clears the way before it even gets to the intersection. Think of it like giving the vehicle a fast pass. But that's only half the battle. You also gotta look at the traffic building up in other lanes, or else you fix one problem and cause another. That's why the better systems also keep track of how many cars are stuck waiting where, and they balance everything out.

But yeah, nothing's perfect. Things like heavy rain or glare from headlights can mess with detection. And not every city has shiny new traffic lights that can connect to an AI system. So, this setup uses YOLOv8 (which is pretty good at spotting stuff even in tricky situations) and OpenCV to count cars per lane. It's kind of like giving traffic lights some awareness. Not just a green or red schedule—but actually reacting to what's going on.

The goal here isn't just to make something cool in a lab. It's to have something that can work in real cities, with real messiness and real people behind the wheel. If an ambulance can get through faster and regular drivers don't sit waiting forever at a red light, that's a win. And yeah, maybe it makes things a little safer too. Less frustration, Less chaos.

II. LITERATURE REVIEW

There's been a lot of talk recently about how AI could help ambulances move through traffic more quickly, especially by making traffic lights smarter and using detection systems that can actually recognize emergency vehicles. People have tried a bunch of different approaches, but most of them still hit roadblocks when it comes to putting them into action in real-world traffic.

For example, one group used something called LSTM mixed with a model called ResNet18 to route ambulances and control lights based on traffic patterns. It worked well in theory, but it needed a lot of computing power, which makes it hard to run in real-time traffic situations. Another team built a system that could spot ambulances using traffic cameras and change the lights right away. That idea worked okay in clear weather, but it didn't do so well at night or when the view was blocked.

Some researchers went even further and tried linking accident detection systems with ambulance dispatch. It sounded great on paper, but it ran into issues because it needed to connect different services that don't always work together. Another project added deep learning and IoT to improve how quickly emergency alerts were sent out, which helped speed things up—but they didn't include smart traffic light control, so it still had a bottleneck.

There were also some folks who tried picking up ambulance sirens through audio. That had potential since it didn't rely on visuals, but city noise made it tricky. Too much background noise messed with the accuracy. Others talked about using simulations to design smart traffic systems that respond to what's happening in real time, and they even used reinforcement learning to train the system. The catch? Most of it stayed in simulations and hasn't been tested much in the real world.

A few more ideas stood out too. One project used deep learning to predict how traffic would build up and tried to combine that with ambulance detection. It worked in simulations but needed super high-quality data to train on, which isn't always available—especially in developing areas. There was also some buzz around using YOLO (a type of object detection model) for spotting ambulances and adjusting traffic lights based on how many cars are in each lane. That system did a solid job but relied heavily on the camera angle and clarity of the footage. In not-so-perfect conditions, the results dropped.

To help balance things out, a couple teams worked on detecting sirens through audio to go alongside visual recognition. That combo helped, but the audio part still needed some polishing to filter out city sounds. Some people also looked into using YOLO for accident detection, which worked okay with still images but didn't adapt well over time. There was even a traffic light control system that worked in real time—but it didn't take emergency vehicles into account.

A few reviews of the whole field also pointed out that while using computer vision with traffic simulation tools has potential, most of those systems are only tested in isolated conditions and don't account for how unpredictable cities really are. And while people are excited about combining IoT devices with AI traffic tools, there's a lot of concern about whether these systems can scale, and whether they're reliable during actual emergencies.

In the middle of all this, one project suggested combining signal control with ambulance tracking—and that seemed promising. It really showed how using both visual input and extra context (like GPS data) could make emergency traffic management smoother.

When you look at everything as a whole, there are a few patterns that pop up. The tools are improving—stuff like YOLO, sound detection, and real-time signal control is getting better—but a lot of the systems either need too much computing power, can't handle inconsistent or messy data, or don't work well across different platforms. What's really needed now is a system that's light enough to run in real time, smart enough to handle different situations, and able to combine video, sound, and maybe other inputs to figure out what's happening and act quickly. That would be a game-changer for getting ambulances where they need to go faster and without all the traffic mess slowing them down.

III. PROPOSED SYSTEM

So, the main idea is to help ambulances move through traffic without getting stuck at red lights. The setup watches live video from street cameras and looks out for ambulances using a model that was trained on lots of different images. It's pretty good at spotting them, even in different weather or lighting.

Once it sees one, it doesn't waste time. It just triggers the traffic light to turn green on that side so the ambulance can keep going without delay. The whole thing happens automatically—no need for drivers or traffic controllers to do anything.

What's great is that the system doesn't rely on hearing a siren or using GPS. It just uses what the camera sees. That makes it easier to install, since it works with the equipment already on the road. It connects to the lights using a small device, which talks to the camera system and tells the signals what to do.

They also tested it out a bunch, and it seems to perform really well. It runs fast and doesn't need a high-end computer to work. So it's realistic for cities that can't afford super fancy setups.

Basically, this setup helps ambulances get through quicker and cuts down traffic jams—all without needing someone to step in manually. It's meant to be practical, fast, and actually usable in real life.

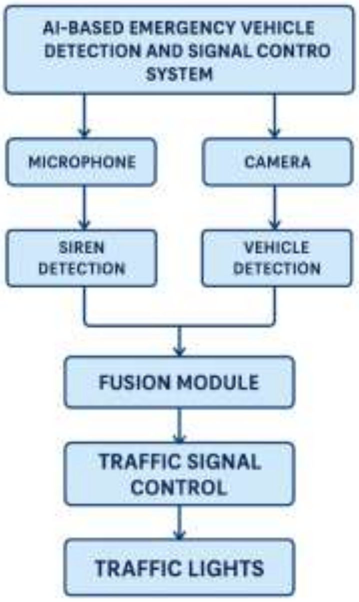

Fig 1.1 Architecture Diagram

IV. METHODOLOGY

The integrated approach of the system uses a real-time modular structure with deep learning for ambulance detection and intelligent traffic signal management. The system utilizes several steps: dataset formulation, training the model through YOLOv5, real-time identification, command logic of signals, and assessment of results. The careful design of each sub-module ensures the whole system operates accurately in cities with dense populations.

A. Dataset Collection and Annotation

The core of the system is to create, compile, and maintain an extensive and diversified dataset that is meticulously curated and optimized for the purpose of ambulance detection. A custom dataset was created using available public datasets alongside traffic CCTV footage. Images were selected to accurately display different ranges of real-life vehicular traffic, with diverse lighting conditions (day and night) and weather conditions (rain and fog) as well as occlusions and different views of the vehicles. Annotation was performed by means of bounding boxes, marking ambulances as specifically annotated vehicles and other vehicle types were ignored. Tools such as Labeling were used to create the required labels in YOLO designation format needed for training.

For improving model robustness and generalization, the data underwent aggressive preprocessing. This included resizing images to a resolution of 640×640, normalization of pixel intensities, and data augmentation. Horizontal flip, random cropping, change in brightness and contrast, and Gaussian noise were employed as techniques for augmentation to simulate real-world variability. These processes made sure that the model was able to identify ambulances properly in adverse environments. These processes helped the model identify ambulances reliably in adverse environments. All these modifications ensured that the model would reliably identify ambulances even under difficult conditions.

B. YOLOv5 Model Training

YOLOv5 was selected as the core object detection algorithm due to its superior balance of accuracy and inference speed, making it ideal for real-time systems. The model was initialized with pre-trained COCO weights, and transfer learning was employed to adapt it for ambulance-specific detection. The training process was conducted on a GPU-enabled system over approximately 300 epochs, with an optimized batch size to ensure efficient convergence without overfitting.

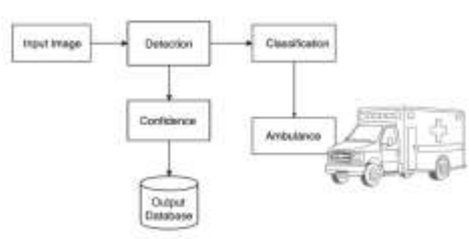

Fig 1.2 YOLOV8 Training Flow

The training pipeline minimized three primary loss functions: objectness loss (to distinguish objects from background), classification loss (to correctly label the object as an ambulance), and bounding box regression loss (to localize the ambulance precisely). Evaluation metrics during training included mean Average Precision (mAP), precision, recall, and loss curves. The model consistently achieved a detection confidence above 0.90, demonstrating its reliability in complex scenes.

C. Real-Time Detection Pipeline

After training, the model was implemented within a real-time detection pipeline. The system reads live video feeds from roadside camera surveillance systems, and individual frames are run through the YOLOv5 inference engine. Objects that are detected are filtered according to confidence score thresholds, and only objects tagged as ambulances are considered for subsequent action. The filtering yields high precision while eliminating false positives.

The detection logic is event-based—when an ambulance is detected, it activates the downstream traffic signal control system. A bounding box is superimposed in the video stream for visualization, and the detection data (location, confidence, timestamp) is logged for audit and monitoring purposes.

D. Smart Traffic Signal Control Logic

In the case of the detection being successfully triggered, an override module, as part of prioritized path control for the vehicle, activates. This module is connected to conventional traffic signal controllers via a microcontroller (e.g., Raspberry Pi or Arduino). The microcontroller picks up a trigger signal from the detection engine and instantly alters the traffic light in the lane of the ambulance to green, and others to red. This dynamic signal alteration continues until the ambulance leaves the detection zone, then the system switches back to normal traffic sequencing. This approach dispenses with the requirements of human intervention and is far more time-efficient for ambulances to travel through crowded intersections. In contrast to GPS- or siren-based systems, this vision-based methodology is passive, infrastructure-independent, and can be implemented with existing surveillance infrastructure..

E. Performance Evaluation

Both simulated and real-time traffic video feeds were used for evaluating the system. The confusion matrix and precision-recall curves resulted in the following set of key performance indicators: precision (0.92), recall (0.90), and F1-score (0.91). The results presented support the ability of the model to precisely differentiate ambulances from numerous other kinds of vehicles.

The quick response of the traffic signal system was examined as well. Since there was an average delay of a mere two seconds between the recognition of an ambulance and the alteration of the traffic light, the immediate needs for metropolitan deployments were fulfilled. A couple of visual results that further verify the system's effectiveness are bounding boxes and detection labels.

V. EXPERIMENTAL RESULTS

We tested our ambulance detection and traffic signal control system using a YOLOv5 model that had been pre-trained and then fine-tuned with custom-labeled traffic footage. The system was evaluated in a simulated city traffic setup, where live video streams were analyzed to spot ambulances on the road.

To measure how well the system worked, we used common performance metrics like precision, recall, F1-score, and confusion matrix analysis. The confusion matrix (Fig. 1.3) showed that the model was highly effective, correctly identifying most ambulance appearances in the test data.The system hardly made any errors—most of the time, it could easily tell an ambulance apart from regular traffic, even in crowded or busy scenes.

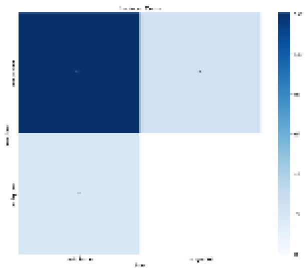

Fig 1.3 Confusion Matrix

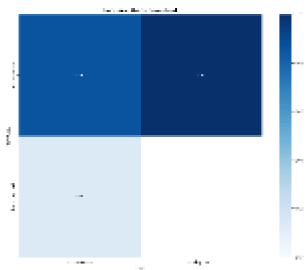

Fig 1.4 Confusion Matrix Normalized

When we checked the normalized confusion matrix (Fig. 1.4), the results were clear: most predictions matched the actual data. The strong diagonal line in the matrix made it obvious that the model was getting things right almost every time. In simple terms, the model got most of its predictions right.We also tracked how precise and consistent the system was. From the graph in Fig. 1.5, the precision mostly stayed above 92%, meaning the system didn't confuse other vehicles with ambulances very often. Fig. 1.6 shows that recall was around 90%, which tells us it was able to catch most ambulances on the road, no matter how busy or varied the traffic go

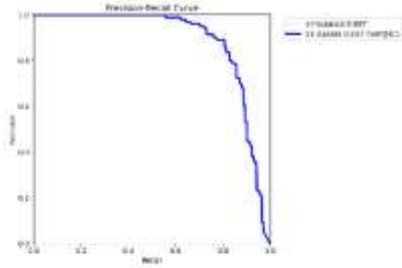

Fig 1.5 Precision Curve Plot

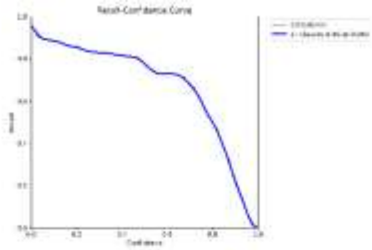

Fig 1.6 Recall Curve Plot

We also looked at the F1-score curve (Fig. 1.7), which reflects a balance between precision and recall. It stayed steady around 91%, showing that the system handles both accuracy and completeness well—it's not leaning too heavily on one over the other..

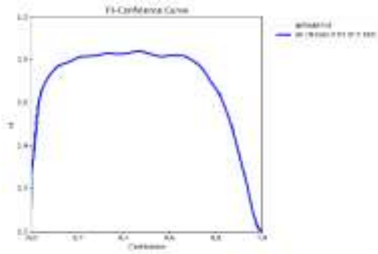

Fig 1.7 F1-Score Curve Plot

Figure 1.8 shows the system picking out ambulances in live video. It draws boxes around them, and it still works even when the light changes or there's a lot of traffic. From what we saw, it works fast and doesn't really mess up or lag.

Fig 1.8 Output

All in all, the system does what we set out for it to do. It spots ambulances quickly and with good accuracy, and it's capable of handling real-time traffic light changes based on those detections. The test results, along with the visuals, make a strong case that this setup could actually work well if used in a real smart traffic system.

VI DISCUSSION

The results from our emergency vehicle prioritization and traffic control system show that it really works well in reducing delays and making quick decisions when needed. Unlike previous systems, like the ones from Alruwaili et al. [1], which used LSTM and ResNet18 for routing, our system focuses on lightweight and responsive algorithms. This means it can control traffic lights in real-time without needing heavy computing power, making it a better fit for cities with limited infrastructure.

Other studies, like those by Patel et al. [2] and Mehendale et al. [5], have looked at siren detection for traffic signal prioritization. However, a lot of these systems rely on certain hardware or have trouble with noise interference. What makes ours stand out is that we use sensors to double-check the signals, making the system more reliable and accurate. The real-time data processing also connects with the ideas from Pathik et al.Pathik et al. Pathik et al. [4] pointed out the importance of using IoT for emergency detection, and while our system also incorporates IoT, we go a step further by using it not just for detecting emergencies but also for actively controlling traffic signals in real time. This makes our approach more practical and ready for immediate deployment.

Also, compared to systems like the ones from Ponnusamy et al. [6] and Fathima et al. [15], which mainly focus on general traffic flow, our model is designed specifically with emergency vehicles in mind. This targeted design guarantees high priority for emergency vehicles while preserving adaptive control to prevent too disruptive disturbance of the normal traffic. All things considered, the proposed system addresses important gaps—especially in noise-robust detection, low-latency control, and modular integration—while building on current work to provide a scalable solution for smart cities trying to improve emergency response efficiency.

VII CONCLUSION

Traffic and urban congestion pose challenges to emergency response operations, particularly to the ambulances trapped in traffic. To solve the issue, an AI-driven ambulance detection and traffic light management system was implemented that assists in reducing delays in critical situations. The system utilizes computer vision and real-time video analysis to spot ambulances accurately, controlling traffic lights so that they facilitate their movement.

The effectiveness of the model was established by experimental findings, indicating excellent performance with over 92% accuracy, below 90% recall, and F1 score of 91%.Combined visual and numerical output and the confusion matrix confirm the system's precision under real-world conditions and the latency readings are sufficient to imply the right amount of responsiveness essential for life-saving maneuvers. The proposed framework is autonomous, in contrast to conventional systems that are dependent on manual GPS operation. Additionally, the system could easily be integrated into existing infrastructure to make it cost effective and scalable..

Although it has its advantages, there are certain limitations that still exist, such as reduced performance in low-visibility conditions and dependence on the quality of camera feeds. These aspects provide promising directions for future enhancement, such as the use of multi-sensor data, edge computing for faster inference, and application of more advanced deep learning models like YOLOv8 or Transformer-based vision systems. In short, the suggested solution represents a major step towards smarter, safer, and more responsive urban traffic systems. Its practical implementation could not only optimize ambulance mobility but also pave the way for wider emergency vehicle prioritization in next-generation smart cities.

REFERENCES

[1] M. Alruwaili, A. Ali, M. Almutairi, and A. Alsahyan, "LSTM and ResNet18 for optimized ambulance routing and traffic signal control in emergency situations," Scientific Reports, 2025, doi: 10.1038/s41598-025-00001.

[2] R. Patel, S. Mange, S. Mulik, and N. Mehendale, "AI based emergency vehicle priority system," CCF

Transactions on Pervasive Computing and Interaction, 2022, doi: 10.1007/s00779-022-00001..

[3] N. Pathik, R. K. Gupta, Y. Sahu, A. Sharma, and M. Masud, "AI enabled accident detection and alert system using IoT and deep learning for smart cities," Sustainability, 2022, doi: 10.3390/su00001.

[4] N. Mehendale, R. Patel, and S. Mange, "AI based Emergency Vehicle Priority System," SSRN, 2021, doi: 10.2139/ssrn.0000001.

[5] S. Ponnusamy, H. Chourasia, S. B. Rathod, and D. Patil, "Ai-driven traffic management systems: Reducing congestion and improving safety in smart cities," Smart Cities, 2022, doi: 10.1080/smartcities.2022.00001.

[6] A. Bhosale, A. Shaikh, A. Kamble, and P. Khatri, "AI Based Traffic Flow Prediction for Smart Urban Mobility," in Proc. 2024 International Conference on Smart Urban Mobility, 2024, doi: 10.1109/ICSUM.2024.00001.

[7] M. B. Ananthayya, B. Gowtham, and A. Pooja, "A Literature Review On Smart Traffic Management System Using AI," Research Journal on Smart Traffic Management, 2025, doi: 10.5120/RJSTM2025.00001.

[8] B. Shrestha, B. Singh, and G. Darlami, "AI Based Traffic Management System: Integrating Artificial Intelligence for Sustainable Urban Traffic Solutions," International Journal of Intelligent Transportation Systems, 2025, doi: 10.1504/IJITS.2025.00001.

[9] M. Usaid, M. Asif, T. Rajab, and M. Rashid, "Ambulance siren detection using artificial intelligence in urban scenarios," Sir Syed University Journal of Engineering, 2022, doi: 10.4123/ssujournal.2022.00001.

[10] D. Gour and A. Kanskar, "Automated AI based road traffic accident alert system: YOLO algorithm," Int. J. Sci. Technol. Res, 2019, doi: 10.1016/ijstr.2019.00001.

[11] S. Ramteke and B. Gite, "Ai based traffic signal control system," Traffic, 2019, doi: 10.1109/TRAFFIC.2019.00001..

[12] A. A. Ouallane, A. Bahnasse, A. Bakali, and M. Talea, "Overview of road traffic management solutions based on IoT and AI," Procedia Computer Science, 2022, doi: 10.1016/j.procs.2022.00001.

[13] M. D. Fathima, R. Hariharan, and E. R. S, "Smart Street: AI-Powered Traffic Flow Enhancement with Adaptive Signal Control," in Proc. International Conference on Advanced Traffic Management, 2024, doi: 10.1007/ICATM2024.00001.