Smart Traffic Signal Control for Dynamic Traffic Management and Emergency Vehicle Prioritization

Submitted By:

Abu Saim
ID: 221-35-898
Department of Software Engineering
DAFFODIL INTERNATIONAL UNIVERSITY

Supervised By:

DR. MD. FAZLA ELAHE
Assistant Professor & Associate Head
Department of Software Engineering

DAFFODIL INTERNATIONAL UNIVERSITY

Department of Software Engineering

Faculty of Science and Information Technology

Supervisor Approval Form

Fall 2025	B.Sc. In SWE	Campus: DSC

Student Name	Student ID
Abu Saim	221-35-898

Thesis Information	
Thesis Title	Smart Traffic Signal Control for Dynamic Traffic Management and Emergency Vehicle Prioritization
Type of work	AI, Machine Learning, and Computer Vision–Based System Development.

Supervisor information	
Supervisor Name	Dr. Md. Fazla Elahe
Supervisor Initial	FE
Completed Credit till now	139
How many credits in this semester	6
Amount (Due)	0
Supervisor Consent	☑ Yes ☐ No

Fazla Elahe 26.11.25

Supervisor Signature

DAFFODIL INTERNATIONAL UNIVERSITY

DECLARATION OF THESIS AND COPYRIGHT

Author's Full Name	:	Abu Saim
Date of Birth	:	02 December,2000
Title	:	Smart Traffic Signal Control for Dynamic Traffic Management and Emergency Vehicle Prioritization
Academic Session	:	2022-2025

I declare that this thesis is classified as:

☒	CONFIDENTIAL	(Contains confidential information under the Official Secret Act 1997)*
☐	RESTRICTED	(Contains restricted information as specified by the organization where research was done)*
☐	OPEN ACCESS	I agree that my thesis to be published as online open access (Full Text)

I acknowledge that Daffodil International University reserves the following rights:

1. The Thesis is the Property of Daffodil International University.
2. The Library of Daffodil International University has the right to make copies of the thesis for the purpose of research only.
3. The Library of Daffodil International University has the right to make copies of the thesis for academic exchange.

Certified by:

Abu Saim
26.11.25

____
(Student's Signature)

221-35-898
Student ID
Date:26-11-2025

Fazla Elahe 26.11.25

(Supervisor's Signature)

DR. MD. FAZLA ELAHE
Name of Supervisor
Date: 26-11-2025

THESIS DECLARATION LETTER

Librarian,
Daffodil International University,
Daffodil Smart City,
Ashulia.Dhaka,Bangladesh

Dear Sir,

CLASSIFICATION OF THESIS AS RESTRICTED

Please be informed that the following thesis is classified as RESTRICTED for a period of three (3) years from the date of this letter. The reasons for this classification are as listed below.

Author's Name : Abu Saim
Thesis Title: Smart Traffic Signal Control for Dynamic Traffic Management and Emergency Vehicle Prioritization

Reasons:

(i) The thesis presents a novel, integrated computer-vision framework for traffic and emergency-vehicle management whose architecture, model configuration and control logic have potential for patenting and technology transfer. Early public release could compromise these opportunities.

(ii) The datasets and trained YOLOv8 models, although derived from public video sources, have been processed, curated and annotated in a way that constitutes a significant intellectual property asset of the university and the author. A restricted period is required to allow further development and possible commercial exploitation.

(iii) The system offers a competitive advantage in intelligent transportation and adaptive signal control. Unrestricted access could enable external organisations to replicate the solution, reducing the strategic value and first-mover benefit for the university and the author.

Thank you.

Yours faithfully,

Fada Elahi 26.11.25

(Supervisor's Signature)

Date:26 November 2025

Stamp:

SUPERVISOR'S DECLARATION

I hereby declare that I have checked this thesis, and in my opinion, this thesis is adequate in terms of scope and quality for the award of the degree of Bachelor of Science.

Fazla Elahe 26.11.25

(Supervisor's Signature)

Full Name : MD. FAZLA ELAHE

Position : Assistant Professor & Associate Head

Date : 26 November 2025

STUDENT'S DECLARATION

I hereby declare that the work in this thesis is based on my original work except for quotations and citations which have been duly acknowledged. I also declare that it has not been previously or concurrently submitted for any other degree at Daffodil International University or any other institution.

Abu Saim
26.11.25

(Student's Signature)

Full Name : Abu Saim

ID Number : 221-35-898

Date : 26 November 2025

Smart Traffic Signal Control for Dynamic Traffic Management and Emergency Vehicle Prioritization

STUDENT'S NAME

Abu Saim
ID: 221-35-898

Thesis submitted in fulfillment of the requirements
for the award of the degree of
Bachelor of Science.

Department of Software Engineering (Major in Data Science)

DAFFODIL INTERNATIONAL UNIVERSITY

26 November 2025

ACKNOWLEDGEMENTS

I want to start by expressing my gratitude to Almighty God for bestowing His divine favor and allowing me to complete my undergraduate thesis.

I would like to express my gratitude and deep amount of respect to my supervisor, Dr. Md. Fazla Elahe, Assistant Professor & Associate Head in the “Department of Software Engineering” at “Daffodil International University” in Dhaka. His profound knowledge and guidance in the segment on “Deep Learning” help me a lot to complete this entire thesis work. It has been made possible by his never-ending empathy, academic leadership, continuous motivation, regular and vigorous monitoring, constructive criticism, helpful counsel, reviewing numerous subpar. manuscripts, and fixing them at every level.

I wish to extend my sincere appreciation to Professor Dr. Imran Mahmud, Head of the “Software Engineering” Department, Faculty of Science and Information Technology, as well as to the other professors, faculties, and personnel of the SWE Department of “Daffodil International University” for their considerate assistance in accomplishing my work.

Last but not least, I must respectfully thank my parents for their unwavering love and patience.

DEDICATION

This thesis is dedicated to my family, whose unwavering faith and endless sacrifices have been my foundation and strength. Your love and encouragement have illuminated my path and made this achievement possible.

And to every aspiring data scientist and may this work inspire you to explore the beautiful intersection of technology and human potential.

ABSTRACT

This thesis presents an intelligent traffic monitoring framework using YOLOv8 for real-time detection is proposed. This framework utilizes both traffic videos from UA-DETRAC and an emergency vehicle dataset to provide enhanced real time vehicle detection and densities of lane usage and estimation of emergency vehicle clearance time. The average pixels per second processed is 25 P/second for both datasets. All videos were preprocessed with standard preprocessing including normalization and augmentation of training data to aid in improving the robustness of the trained model. For the trained model, two YOLOv8 models were created, one for general vehicles and one specifically for emergency vehicle detection and training. These models performed well on both with respect to precision, recall, F1-score, mAP, and minimal cross-class confusion on both datasets. The vehicles detected in each frame will be counted for lane density estimation purposes; a phased and continuously time-based Gaussian model with time duration varied based on density of activity (density of traffic lanes) will visually express the congestion levels based on vehicle density and activity over time. This real time traffic control framework was simulated in a Python-based Adaptive Signal Control System demonstrating that the system would adjust the lengths of green light duration based on the traffic lane densities and the presence of EVs and emergency response vehicles. In addition to reducing rear-end and lane throughput delays, the system would support facilitating rapid clearance of ambulances and firetrucks. Overall, this proposed framework provides a practical, robust, transparent and reproducible methodology to combine real time object detection and high-level traffic management for the development of intelligent transportation systems.

Keywords: Computer Vision, Deep Learning, YOLOv8, Vehicle Detection, Emergency-Vehicle Detection, Traffic Density Estimation, Congestion Heatmaps, Adaptive Signal Control, Intelligent Transportation Systems.

TABLE OF CONTENT

LIST OF FIGURES

LIST OF SYMBOLS

Symbol	Description	Unit / Domain
P	Precision	[0, 1]
R	Recall	[0, 1]
F1	F1-Score (harmonic mean of Precision and Recall)	[0, 1]
mAP@0.5	Mean Average Precision at IoU threshold of 0.5	[0, 1]
mAP@0.5:0.95	Mean Average Precision over IoU thresholds from 0.5 to 0.95	[0, 1]
IoU	Intersection over Union	[0, 1]
box_loss	Bounding box regression loss	
cls_loss	Classification loss	
dfl_loss	Distribution Focal Loss	

LIST OF ABBREVIATIONS

Abbreviation	Full Form
AI	Artificial Intelligence
AP	Average Precision
CNN	Convolutional Neural Network
KPI	Key Performance Indicator
ML	Machine Learning
mAP	mean Average Precision
YOLO	You Only Look Once
YOLOv8	You Only Look Once Version 8

CHAPTER 1

INTRODUCTION

1.1 Background of the Study

The rapid increase in vehicle population and the ever-increasing pressure on road networks make the effective management of signalized intersections a challenge for City Traffic Managers. Fixed-time traffic lights and crude vehicle loop counters are still in use but fail in conditions where the traffic flow is not stable, resulting in unnecessary waiting, longer duration of travel, or delay on both routine and emergency vehicles. In heavy traffic, the signal timings are the same for all lanes at competing phases with a response time to queue clearance leading to not only lost green time but also higher fuel consumption and emissions.

The development of deep learning, and in particular computer vision, means that these limitations can now be mitigated using only video data gathered from current roadside or CCTV cameras. Current object detection models (such as YOLOv8) can detect and classify vehicles in real time, thus allowing the estimation of lane-wise traffic density and vehicle types, such as emergency vehicles with special privileges. Therefore, wireless camera feeds can be turned into useful traffic data that will help those controlling roadways manage their resources more efficiently.

This thesis presents a smart traffic light control system that uses deep learning–based vehicle detection and traffic density estimation for dynamic adaptation of signal timings. YOLOv8 performs overall vehicle detection, and domain-specific training captures the presence of emergency vehicles; as a result, green time assignment can be based on the actual lane demand, with an immediate prioritization plan for its occurrence. The framework is built purely on camera input and does not require any additional hardware, such as IoT. It can be scaled to the existing video infrastructure to be practically useful. The traffic flow, signal status, and controller decision are presented in real time in a Pygame simulation, and the dynamics of the signals vary with the different levels of traffic and emergencies.

1.2 Problem Statement

Most intersections experience extreme congestion, excessive wait times, and lack of economy in fuel consumption for static time-based traffic signals. Traditional controllers are unable to respond to sudden load surges or traffic variables in real time, which often results in superfluous waiting times, green phase underutilization, and blocked high-demand lanes on one side, but green ones with few users.

Legacy solutions, such as inductive loops, RFID tags, or rudimentary IoT devices, are costly, inflexible, and offer minimum information content (usually merely vehicle presence) regarding lane-wise traffic density/vehicle types. Simple image subtraction techniques are not effective under variable light conditions, especially in the dark.

Moreover, the vast majority of these systems cannot automatically sense and prioritize first responders, leading to ambulances and fire trucks getting caught in traffic jams. Its members do not share a common ecosystem of tools or research across vehicle detection, density estimation, and adaptive control. This thesis fills these gaps by introducing a camera-based intelligent traffic light control system that integrates real time vehicle detection, lane-wise density estimation, adaptive signal control, and emergency vehicle prioritization in a unified testable platform visible in a realistic simulation.

1.3 Research Gaps

This is despite advances in intelligent traffic control and deep learning–based perception which have however some limitations that are addressed by this thesis. Existing solutions either still use fixed schedules, provide minimal or no automatic emergency vehicle processing, display a decreased detection performance for challenging situations and isolate intersections without coordination. There is also limited work on integrating state-of-the-art deep learning models into a powerful, realtime control algorithm for various traffic scenarios.

1.3.1 Limited Real-Time Adaptation:

Traffic signal controllers are mostly programmed to fixed-time control or pre-set schedules that do not adapt in real time to live traffic variations. Additionally, weak actuation (for example, simple vehicle detectors) are unable to effectively change cycle lengths or green splits as congestion increases and decreases. This non-real time adaptation leads to wasting of green time, growing queues and decreased road capacity utilization–particularly during peak hours or when sudden bursts in demand occur.

1.3.2 No Automated Emergency Vehicle Priority:

Today traffic systems don't often pick up ambulances, fire trucks or police cars on camera feeds and give them priority. Emergency vehicles generally move through intersections in the same way regular traffic does, using sirens to alert drivers and intensive cooperation with them. This absence of automated preference slows response time and can prevent traffic management systems from efficiently serving public safety.

1.3.3 Sensitivity to Challenging Environmental Condition:

Deep learning-based vehicle detectors achieve high accuracy on the off-the-shelf benchmarks, however suffer with low accuracy in night, rain/fog and vehicles occlusion scenes. For example, real-world optimization objectives may not being explicitly designed or tested for this regimes that are often encountered in realistic traffic. Misidentification in such cases can result in inaccurate estimates of density and less than-ideal control decisions.

1.3.4 Limited Multi-Intersection Coordination:

Most study and deployed systems manage an intersection only and do not take into account upstream or downstream intersections. Without this coordination, it is common for vehicles that leave an intersection with a green signal to arrive at the next red light, which creates stop-and-go traffic and adds to overall delay. intelligent control at the multiple intersections or corridors is not abundance.

1.3.5 Low Utilization of Camera-Only, Deep Learning Pipelines:

A lot of intelligent traffic systems are hardware-dependent: induction loops, RFID tags or IoT sensors. On the other hand, camera based deep learning pipelines that can potentially generate detailed, lane-wise and class-wise traffic information is not fully exploited as the primary sensing mechanism. This gap leaves developers of systems that can scale using existing video infrastructure but with limited sensor availability.

1.3.6 Initialized Visualization and Evaluation Environment:

Intermediate work tends to produce detection metrics or control decisions without associated interactive visualizations which combine vehicle moves, signal states and controller responses on the fly. This impedes the engineers, power that be and nontechnical stakeholders in terms of understanding, validating and trusting intelligent traffic systems. There is a need for software frameworks that would integrate simulation tools, e.g., Pygame based visual interfaces with deep learning detection/logic and adaptive control logic synthesis along with traffic graphic in one place.

1.4 Research Objectives

To cope with the problems of real-time adaptation, emergency vehicle management, robustness and coordination described above the thesis outlines the specific goal as follows. Collectively, these two objectives build an end-to-end framework that spans from video-based detection to adaptive, priority-aware traffic signal control.

1.4.1 Objective 1: Real-Time Vehicle Detection Using YOLOv8

To Create a Real-Time vehicle detection Model with YOLOv8 on Traffic Camera feeds (live or recorded).

Contribution: We present and implement a vehicle detection solution using the model YOLOv8, which is faster and more accurate than previous versions of the model (e.g., YOLOv3 or YOLOv5) providing reliable detection directly from camera streams.

Measurable Outcome: Detection performance will be defined concerning Precision and Recall, the speed of processing in real-time (in FPS - frames per second) and robustness to a number of changes in conditions including light change, occlusion, bad weather.

1.4.2 Objective 2: Adaptive Signal-Timing Based on Traffic Density

To develop an adaptive signal-timing algorithm which can automatically change green time in real-time based on the sensed traffic volume.

Contribution: Design a traffic signal control algorithm that adapts cycle lengths in real time to variable traffic demand, and optimize over static fixed time plans.

Measurable Outcome: Efficiency will be evaluated based on average vehicle waiting time at each intersection and the throughput of the intersection compared to which using a fixed-time signal plan.

1.4.3 Objective 3: Emergency Vehicle Detection and Priority Control

For embedding automatic emergency-vehicle detection and priority control in the traffic flow logic.

Contribution: Combine work on recognition of vehicles from camera feed (including emergency vehicles such as ambulances or firetrucks) using characteristic lights and appearance to give these priority at junctions and over rides normal timing.

Measurable Outcome: Performance will be determined by comparing the clearance time for emergency vehicles and time elapsed from approaching to leaving the intersection with respect to the proposed system as compared with baseline.

1.5 Scope and Limitations of the Research

The scope for thesis is identified to ensure the feasibility while presenting the entire working prototype of intelligent traffic light control system: The included integrated deep learning model also enhances real-time traffic management and emergency vehicle prioritization, all with existing camera infrastructure.

1.5.1 In-Scope:

Study of pre-collected traffic video data at urban intersections learned for typical multi-lane signalized junctions.

Vehicle detection and classification, where both general and emergency vehicles have dedicated training and evaluation sets for the UA-DETRAC/ Roboflow emergency vehicle datasets.

Lane-based determination of traffic density and subsequent utilization in the adaption of signal control served by fuzzy or heuristic rules.

Pygame-based simulation environment for visualizing intersections, vehicle traffic and changing signal phases in real-time.

Performance evaluation of the detection results (Precision/, Recall, F1-score and mAP) as again control performance (average delay and queue length) in controlled simulation experiments.

1.5.2 Out-of-Scope:

Testing on physical intersections or interfacing with physical traffic controller devices; the work only presents offline analysis of video data and software simulation.

Utilization of extra non-visual sensors (as inductive loops, radar and IoT devices); here the system is camera-based only.

City-level traffic network optimization; concentrated on one or few intersections instead of considering a whole ecosystem.

Accurate modelling of pedestrian behaviour and non-motorised traffic,) when seriously, have the primary focus on motor vehicles and emergency vehicles.

Precisely modeling future traffic demand based on count forecasting not short-term or current observed density.

1.6 Significance of the Study

The included integrated deep learning model also enhances real-time traffic management and emergency vehicle prioritization, all with existing camera infrastructure.

1.6.1 Theoretical Contribution:

This work represents a significant step forward for computer vision, deep learning and intelligent transportation research by introducing a camera-only deep learning pipeline that connects perception to traffic control. It shows that YOLOv8 detection, density estimating and rule-based or fuzzy control can constitute an integrated adaptive signal system in the cases without internet of things (IoT) hardware. The research also extends the related work on emergency vehicle detection and priority-oriented traffic system management.

1.6.2 Practical Contribution:

For traffic engineers, city planners and first-responders the new system represents a practical way to provide an upgrade from traditional intersections using existing camera networks. During real-time lane-wise density adaptation, the delays and congestion could be reduced owing to the automated ambulance/firetruck prioritization resulting in shorter emergency response times. The visualization of the Pygame module allows a simple way for stakeholders to verify and see system behavior before deploying on reality, minimizing risk and enhancing acceptance.

1.6.3 Methodological Contribution:

This thesis proposes a systematic pipeline to design the intelligent traffic control system, including dataset processing, model training for general and emergency vehicles, density estimation, adaptive controller synthesis, simulation installation and validation with different cases. The proposed method provides a workable baseline for further research to enable deep learning perception and control systems in traffic or other domains. Its emphasis on verifiable results and standard measures of performance also facilitates a fair comparison with other approaches in future work.

CHAPTER 2

LITERATURE REVIEW

2.1 Introduction

This chapter reviews state of the art literature on intelligent traffic signal control systems with computer vision and deep learning approach, in particular it focuses on the themes of YOLO based vehicle detection as well as camera-driven adaptive signal timing. It introduces development from traditional fixed-time controllers and classical image-processing methodologies to today's advanced CNN-based YOLOv8 for real time usage even in complex traffic environment. It also surveys recent research on vehicle detection, density estimation, adaptive signal control as well as superposed issues such as robustness, real-time response and emergency vehicles priority.

However, the chapter is divided in three main sections. In Section 2.2, computer vision-based vehicle detection techniques are introduced and the trend towards CNN based methods is mentioned. Section 2.3 concentrates on YOLO-based detection models including the latest enhancements in YOLOv8 for traffic scenes. Intelligent and adaptive traffic signal control models are discussed in Section 2.4. Finally, Section 2.5 compiles these findings to illustrate the existing research gaps and provide the theoretical basis for the system developed in this thesis.

2.2 Computer Vision and Deep Learning for Vehicle Detection

Vision becomes a core technology used for inferring fine-grained traffic information out of camera feeds, without any additional physical sensor. In the early works, background subtraction, edge detection and motion analysis were used to detect vehicles; these techniques are agfault-prone because they do not deal with variables of lighting conditions, shadows, weather changes or occlusion-factors that appear frequently in a busy intersection. With the increasing complexity of traffic scenes, methods based on CNNs became an obviously better choice that learn discriminative visual features from massive amounts of data directly.

Bhosale et al. (2022) proposed a mixed solution using CNNs and image processing techniques for estimating traffic density with phase - changing of signal timing. Their approach used background subtraction, contour detection for vehicle localization, CNN classification and density-based green-time computation. Though it was an enhancement over fixed-time controllers, its image subtraction based method was not able to guarantee accurate results in night and high traffic situations with heavy occlusion.

Liu et al. (2024) went a little further comparing to other works in the sense that they also emphasized on maintaining performance under different levels of lightning, weather and occlusion conditions, computing cost for real time or edge deployment. The authors recognize that two-stage detectors based on family R-CNN can achieve high-accuracy, but are too slow for real-time traffic control. As a result, one-stage detectors-such as those in the YOLO and SSD families-have been considered to be attractive alternatives owing to the high frame rates and end-to-end processing they provide.

2.3 YOLO-Based Vehicle Detection in Traffic Scenes

The YOLO also evolves quickly in previous generations and on real-time detection can benefit speed, precision and feature extraction. In traffic, YOLOv2 and YOLOv3 have been popular for vehicle detection in CCTV videos and lane-wise density estimation with the objective of adaptive signal timing. For example, Bhosale et al. (2022) used YOLOv3 to perform vehicle lane count and apportioned green time proportional to traffic flow, in which they were able to improve the waiting time of vehicles when compared with a fixed-cycle control.

In the same way, Ayegbusi et al. (2025) presented a deep learning approach for traffic signal model that uses YOLO detection to estimate real-time density by automatically adjusting the timing of signals. The authors’ model, which was trained using YOLO formatted vehicle data from junction cameras, resulted in a 27% increase of traffic flow and 50% decrease of waiting time when compared with existing systems. This shows that YOLO-based detectors have great opportunities in real-time traffic monitoring.

Despite these advancements, limitations persist. Bhosale et al. (2022) report untrustworthy performance in low light, and Ayegbusi et al. (2025) that highlight the importance of combining emergency vehicle prioritization, vehicle tracking and violation detection. They lead to the conclusion that although YOLO-based models are successful, more sophisticated extension and integration is necessary for the models to apply in real-world intelligent traffic systems.

Liu et al. (2024) introduced YOLOv8-FDD, an enhanced version of the YOLOv8 detection network to address missed detections, false positives and deployment difficulty in traffic. A Feature Sharing Detection Head is proposed to solve the problem of redundant parameters, a Feature Dynamic Interaction (FDI) module is devised for improving the feature fusion for classification and localization, and a Dilation-wise Residual (DWR) module is incorporated for multi-scale feature extraction. Additionally, the Dy Sample operator takes over the role of ordinary nearest neighbour up-sampling to maintain finer details.

On the UA-DETRAC dataset, YOLOv8-FDD is capable of reducing parameter count to 72.89% of the baseline YOLOv8 and improving mAP50 and mAP50-95 by 0.7% and 1.3%, respectively, with a frame rate over 300FPS. Experiments on a self-collected database with large diversity of illuminations, false positive and negative drop a lot, which verifies the in real-world crossroads potential for online use.

Other enhancements were mentioned by Liu et al. (2024) are the structures that have improved attention mechanisms, a dense connection and particular modules for difficult conditions like fog. These include Song et al. 's Mix Up / Mosaic augmentation along with efficient channel attention[x] refinement. 's DC-SPP-YOLO with dense connections and Wang et al. 's Rep-ResNeXt backbone for foggy scenes. Together, these studies indicate ongoing efforts to develop efficient yet accurate YOLO-based architectures for the real traffic scenes.

2.4 Intelligent and Adaptive Traffic Signal Control

Although accurate scheduling is fundamental, the entire impact of smart traffic systems is through converting perception into adaptive signaling strategies that mitigate congestion. Conventional fixed-timed controllers allocate pre-defined green times to each phase irrespective of lane requests, which in turn leads to a wastage of green time and higher delay (Bhosale et al., 2022).

For circumventing this, paradox was used by Bhosale et al. (2022) designed a smart simulator that identifies the presence of vehicles and real-time density to adjust green times. Inspired by the previous literatures, we developed a dynamically adjusting green length performance using image processing camera input and YOLOv3 detection depending on the queue size. Although the strategy can be efficient in minimizing deadlock, it is confined to a single intersection and does not include sophisticated deep learning models or emergency vehicle processing.

Ayegbusi et al. (2025) extended this idea by deploying YOLO detection within a fullfledged adaptive control model. Their method focuses traffic towards the busy lanes to reduce the amount of idle green time and queue, showing considerable throughput gains. They suggest further extensions that consider the interaction with smart city infrastructure, tracking of vehicles and emergency vehicle priority - a close reply to what we focus on in this thesis.

2.5 Synthesis and Identified Gaps

The surveyed literature highlighted a process of developing solutions from hardware bound, fixed-time systems, to camera-only deep learning approach for intelligent traffic managing. YOLO-based detectors, particularly YOLOv8 and its augmented versions, demonstrate a strong real-time detection capability to handle dense and complex traffic scenes. Control solutions employing measurements of detected vehicle counts or density have shown distinct improvements compared to the classical methodologies.

However, several gaps remain like, unexploiting potential of modern YOLOv8 improvements: Many existing methods are based on older YOLO and classical image processing, leading to less robustness in low-light, bad-weather, occlusion scenarios. Recent architectures like YOLOv8-FDD have not been well explored for intersection controlling.

Lack of integrated emergency vehicle prioritization: Traditionally YOLO based controllers are designed for stop line control and consider little about automatic emergency detection and priority control during the decision making, which is critical in practice.

Lack of end-to-end analyses: Many works evaluate detection and control as separate steps rather than measure how errors in detection impact the timing decisions and the global intersection performance.

To fill these deficiencies, this thesis presents a camera-only smart traffic light control system by employing YOLOv8-based vehicle detection, resilient density estimation under various situations, and dynamic signal timing on the basis of explicit priority for emergency vehicles. It further seeks to show that current state-of-the-art deep learning models are applicable as reliable off-the-shelf solutions for real-time, vision-based traffic control using simulation and performance evaluation.

2.6 Literature Review Conclusion

The surveyed literature indicates that camera-based deep learning empowered traffic management has developed considerably, while significant gaps remain. Early systems, which combined image processing with CNNs showed that vision can be able to estimate traffic density and adapt green times, but were fragile under changing illuminations, weather or occlusion and even restricted to single intersections with no emergency-vehicle handling (Bhosale et al., 2022). Recent YOLO-based research shows that accurate and real-time vehicle detection as well as density estimation capable of alleviating traffic congestion when incorporated into an adaptive signal timing can be achieved using DL (see Ayegbusi et al (2025).

Studies of enhanced YOLOv8 models like YOLOv8-FDD also demonstrated that low computational architectures could provide useful accuracy and super high frame rate, which can be applied on edge devices in real traffic conditions (Liu et al., 2024). Nonetheless, most of the systems in use rely on previous versions of YOLO, do not put them to the test under challenging conditions nor address how perception and action are impacted when used as a holistic system. Furthermore, automatic recognition and prioritisation of emergency vehicles are still poor integrated in existing YOLO-based control models despite their importance for public safety.

In summary, the literature shall be sufficient for opening the doors into YOLO-centric, camera-only traffic control, with extending directions being robustness, real-time adaption and solved problems in emergency vehicle priority and unified evaluation. This thesis fills these gaps by introducing a YOLOv8-powered solution that combines accurate vehicle following, density-aware signal control and explicit emergency priority support with one simulation platform.

CHAPTER 3

METHODOLOGY

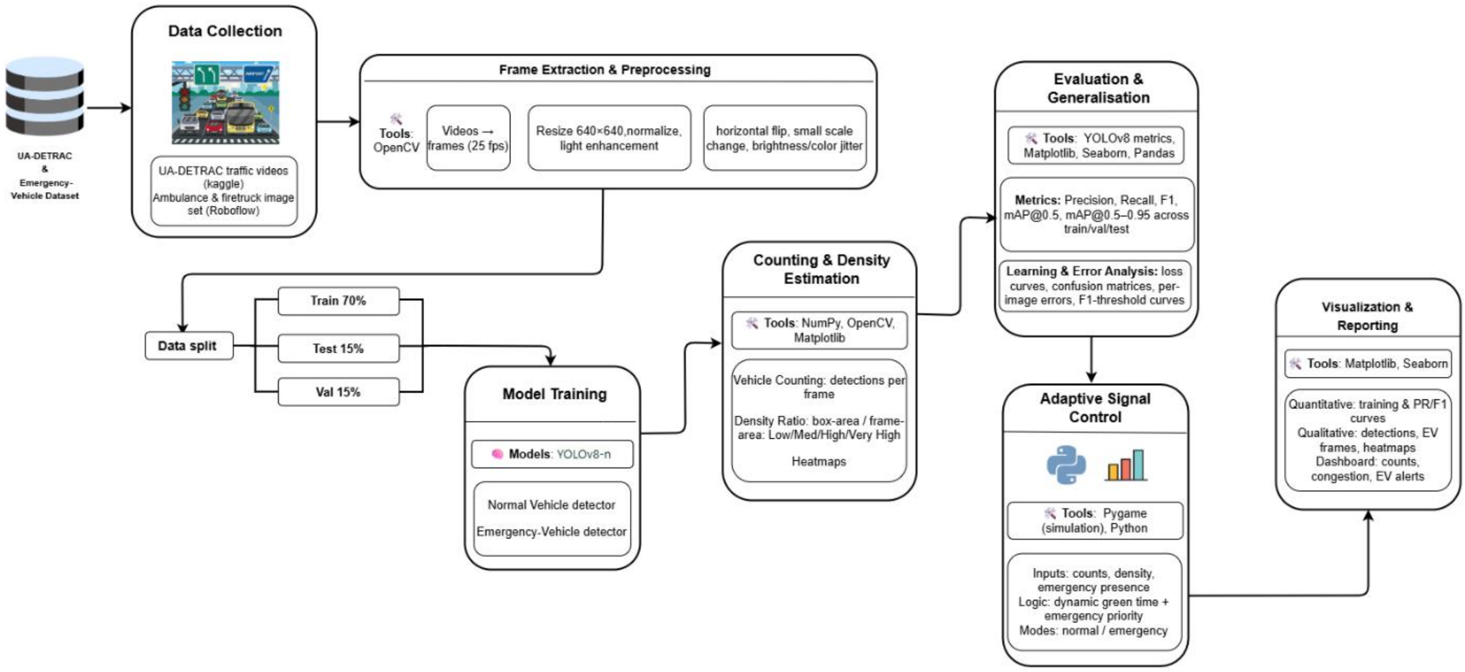

Figure 3.1: Methodology

3.1 Introduction

This chapter presents the approach used to build an Intelligent Traffic Light Control System over deep learning is described. The system will identify vehicles, determine their traffic density and control traffic lights on the fly and it'll allow for priority passage of emergency vehicles. The procedures include data gathering, frame extraction, preprocessing of frames, model training, inputs vehicle counting and density estimation, evaluation, adaptive signal control simulation and visualisation.

3.2 Data Collection

To develop a solid traffic surveillance system, high-quality data with various types of urban environments are needed. This present investigation was based on two datasets.

3.2.1 Traffic Video Dataset (UA-DETRAC)

The primary goal of the SLR was to perform a critical review of the state-of-the-art in AI-based football performance analysis focusing on:

- Source: UA-DETRAC benchmark dataset.
- Content: Videos of multi-lane urban roads, taken under different congestion, lighting condition and weather (sunny, rainy, cloudy).
- Usage: offers realistic traffic situations for detecting, counting and analyzing general vehicle flow in real time.
- Class Simplification: Cars, buses and trucks were combined into a single class also know as the vehicle class for simplicity of effective density approximation.

3.2.2 Emergency Vehicle Dataset

- Source: Public datasets (e.g., Roboflow) and annotated images from traffic video.
- Content: 648 ambulance and firetruck images with a variety of views, illuminations and partial occlusions.
- Annotations: For YOLOv8 format the bounding boxes.
- Application: In the detection of emergency vehicles in order to preempt signals and get a quick response time in an emergency.

Together, those data sets are designed to enable the system for general traffic management and emergency-aware operation.

3.3 Frame Extraction and Preprocessing

In order to efficiently train deep learning models, videos were converted into their frames and preprocessed as follows:

3.3.1 Frame Extraction

- The sample rate for videos was 25 fps (frames/s), which was chosen using OpenCV to trade off temporal resolution with computational speed.

- Extracted frames are split into 3 categories, i.e. Training (70%), Validation (15%), and Testing (15%).

3.3.2 Preprocessing Steps

- The Resizing: All input images were resized to 640×640 pixels as required by the YOLOv8 model. Uniform symmetric scaling guarantees that the model is presented with a uniform field of view which stabilizes feature extraction and preserves interoperability with the network’s convolutions.
- Normalization: The frames in the 3D input tensor are normalized to [0,1] by dividing their pixel values by 255. Normalizing is helpful for stabilizing the optimization in training process, accelerating convergence and preventing large gradients from making the learning unstable.
- Image Enhancement AHE was also applied for the same of enhancing certain darker features in both dim and LE wedges. In addition, AHE also enhances the local contrast via pixel intensity value redistribution, which helps our network to better detect vehicles and small objects under complicated illumination conditions.
- Data augmentation: During training, we applied several augmentations for model generalization and to prevent overfitting. These improvements model variations in 54 the vehicle 53 s physical appearance, pose and environment atmosphere.
- Horizontal Flipping: Each frame was randomly horizontally flipped with a certain probability. This broadening enables the model to learn the symmetry in traffic scenes, doubling effectively the diversity of training dataset without additional images.

- Small Rescaling and Zoom: Small random scaling and zoom frame to enlarge or shrink object slightly. This allows the network to invariant with regards to vehicle distances and focal lengths.
- Brightness and Color Jitter: Random distortions of brightness, contrast, saturation and hue were applied. This simulates varied lighting situations such as cloudy, sunny or shadow which can help the model to stay invariant under real-world illumination transformation.
- Minor Rotations: Frames were artificially rotated by small angles (typically ±10–15°) imitating a slight camera tilt or non-level installation of the surveillance cameras. This enhancement makes the detector more robust to small angle deviation of vehicle direction.

3.4 Model Training

YOLOv8-n (Nano) was used as the underlying detection architecture for its high efficiency, real-time inference capability, and sufficient accuracy for urban traffic scenes. YOLOv8-n is a member of the YOLOv8 family, consisting of Nano (n), Small (s), Medium (m), Large (l) and Extra Large (x)pretrained image source models optimized for different speed-accuracy trade-offs.

3.4.1 Why YOLOv8-n?

Real-time Performance: The proposed YOLOv8-n achieves real-time detection with about 3.2 million parameters, which is a lightweight network structure and able to run at high FPS for live traffic analysis. It is small enough the even video streams from several cameras can be processed in parallel without much delay.

Efficiency vs. Accuracy Trade-off: Larger YOLOv8 models (s, m, l, x) will give even better accuracy but require significantly more GPU resources which may not be practical for real-time use cases. YOLOv8-n balances detection speed and accuracy in a practical way which guarantees fast detection, but without any substantial decrease in its detection performance.

Support for Multiple Detectors: Two parallel detectors are implemented in the system, one for regular vehicles and another for emergency ones. The lightweight of YOLOv8-n also enables both methods to be executed simultaneously without over-burdening hardware and it can be applied to dynamic urban intersections.

YOLOv8-n incorporates advanced architectural components:

- CSPBackbone for balanced feature extraction and less computation.
- PANet neck for efficient multi-scale feature integration.
- Decoupled head facilitating individual optimization of the objectness, BB regression, and classification promotes convergence and overall detection robustness.

3.4.2 Training Configuration

- Normal Vehicle Detector: Trained on UA-DETRAC dataset frames, all vehicle classes are combined into a single class in order to make the base detection task simpler.
- Emergency Vehicle Detector: Trained on a filtered (manually) dataset of ambulances and firetrucks separately with specific classes for each exchange to increase the accuracy in identifying emergency vehicles.
- Weights: The COCO-pretrained weights were adopted to initialize the model training for faster convergence and better generalization, especially under small objects.
- Optimizer: We use momentum-aided stochastic gradient descent (SGD) to prevent falling into local minima.
- Mini-batch Size: 16
- Learning Rate: 0.001 with step decay schedule to slowly decrease the learn rate while training so that it can better converges in next epochs.

- Loss Function: YOLO's multitask loss and how it de-couples the loss into bounding box regression, objectness confidence, classification which in turn allows YOLO to simultaneously learns object detection accuracy and localization together.

- Evaluation Metrics: Precision, Recall, F1-score, mAP@0. 5, and mAP@[0.5: 0.95] were calculated on the validation set during training for model monitoring and avoid overfitting. The metrics offer a full evaluation of detection quality, including classification accuracy and localization confidence.

3.5 Vehicle Counting and Density Estimation

- Traffic counting and congestion evaluating were implemented via post-processing with NumPy and OpenCV of the detection results obtained by YOLOv8.

- Vehicle Counting: For each frame, vehicles were tracked and counted to produce time-series traffic volume per lane information. This allowed the system to produce both time related and accumulated counts that are needed for adaptive traffic control.

- Density Estimation & Density Ratio: A ratio of the area filled by detected vehicles over the whole bounding box to frame area was computed as an approximation for traffic density.

- Qualitative Levels: Density values were categorized into Low, Medium, High and Very High levels for density values were classified in such way to give a qualitative human-understandable interpretation about traffic congestion.

- Density Heatmaps & Spatial Visualization: We placed the Gaussian kernels at each vehicle centroid to have the smooth spatial density distributions and get the congestion heatmap.

- Overlay: Heatmaps were hooded with the original frames using Matplotlib, visually identifies traffic choke points and high congestion areas from both analytical and demonstrating perspective.

- This procedure lets the system convert raw detections into actionable traffic intelligence, including real-time signal control and scenario modeling.

3.6 Model Evaluation and Generalization

In order to ensure generalization and practical application on real traffic data, the models were assessed based on train, validation and test splits by a combination of quantitative and qualitative analyses.

- Detection Metrics: Standard object detection metrics we will use for our experiments. Precision, Recall, F1score, mAP = 0. 5, and mAP [0.5:0.95] to assess the quality of detection models for localization and classification.
- Confusion Matrices: computed confusion matrices for class-wise performance analysis on the polar regions, along with false positives and negatives gives a detailed overview of issues involved in terms of failure conditions, focusing mainly on the emergency vehicle detection.
- Error Analysis: Error statistics per-image and F1-score vs. confidence thresholds were studied to determine the best threshold for deployment compromise between precision and recall according to real time traffic monitoring needs.
- Learning Stability: The loss curves of training and validation, along with metric trends were kept under watch in order to track any overfitting or underfitting occurs thereby the model would withstand firm performance across variety of traffics.

All of the above evaluation techniques together ensure both the models' robustness and accuracy, so they are generalizable for unknown traffic events, changing lighting conditions or alternative vehicle types without extreme drops in performance.

3.7 Adaptive Signal Control Simulation

As proof of application for the intelligent traffic management system, a Pygame-based simulation was realized. The simulator provides an interactive environment to observe middle lane traffic flow and signal actions dynamically.

Simulation Inputs:

- The detections vehicle counts estimates of the YOLOv8 detector for each lane.
- The levels of traffic density, which were defined as Low, Average, High or Very high.
- Emergency vehicle detection and signaling with priority-based processing.

Control Logic:

- Dynamic Green-Time Sharing: The green-time duration for the signals is objectively distributed according to the vehicle density, and denser lanes are granted relatively longer greens.
- Emergency Vehicle Phase Pre-emption: For emergency vehicles, an addressed home lane reassignment is provided at the detection of such vehicle to allow reduced delay time and maximum safe crossing.
- Green-Time Constraints: The minimum and maximum green-time for each phase are enforced to avoid lane starvation and keep fairness for all traffic flows.

Scenarios Tested:

- Equal number of lanes in each direction for normal traffic flow.
- Maximum congested travel conditions with non-homogenous lane usage
- Arrivals to emergency vehicles that need rapid signal preemption.

Performance Metrics:

- Average lane waiting time, indicating how efficiently the delay is being reduced.
- Throughput, the number of vehicles cleared in a given time.
- Emergency vehicle response time, a measure of how well priority handling has functioned.

The purpose of this simulation is to show an experimental validation pointing out how the considered system dynamically reacts under different traffic situations whilst preserving both safety and efficiency.

3.8 Visualization and Reporting

For decision-making support and assessment of the system's performance, the results were described through quantitative as well as qualitative visualizations.

Quantitative Visualizations:

- Training and validation curves to check for model convergence and learning stability
- Precision-recall curves to measure detection performance over confidence thresholds
- Error histograms with most frequent misdetections and error rate
- Threshold sensitivity curves to help select optimal operating points for online implementation

Qualitative Visualizations:

- Annotated frames of detected vehicles and bounding boxes
- Heat maps of density showing spatial congestion patterns
- Intelligent screens of signals from the Pygame simulation showing green-time allocation and emergency vehicle pre-emption.

These visualizations can be displayed as part of a dashboard with the goal to facilitate resourceful functionality by enabling traffic operators, researchers, and stakeholders to identify system performance at glance in making well-informed decisions for traffic management.

CHAPTER 4

RESULTS

4.1 Training Behaviour

The YOLOv8-n model was trained on UA-DETRAC dataset for 100 epochs with the input size of (640×640) pixels, batch size of 16 and early stopping with patience=15. In training process, the three loss terms (box loss, classification loss and distribution focal loss) all fell smoothly with no divergence or instability. the box loss declined from approximately 0.99 in the early stages of training to around 0.40 at the final epoch, and losses of classification and DFL also showed stabilized declining trends.

Validation metrics (calculated after every epoch on 400 images) increased quickly in the first few epochs and then became stable. From about epoch 40 both validation precision and recall were very close to 0.98, as was mAP@0. 5 remained above 0.99. The best system checkpoint selected by YOLO's early-stopping method ("best. pt") achieved:

Metric	Value
Precision	**0.992**
Recall	**0.993**
mAP@0.5	**0.995**
mAP@0.5:0.95	**0.932**

4.1.1 Performance on Training, Validation, and Test Sets

The final model was tested for generalisation and against overfitting on all three dataset splits: 1,400 training images, 400 validation images, and 200 test images.

METRICS COMPARISON TABLE

Dataset Split	Precision	Recall	mAP@0.5	mAP@0.5:0.95	F1-Score
train	0.994789	0.993738	0.994935	0.949037	0.994263
val	0.991650	0.993111	0.994584	0.937783	0.992380
test	0.988925	0.995800	0.994569	0.935902	0.992350

Figure 4.1: Metrics Comparison

- The model presents a very high precision and recall in all the created splits, it correctly detects almost all vehicle except for some false positives.

- mAP@0. 5 is higher than 0.994 indicating sufficient bounding-box predictions, and mAP@0. 5:0.95 implies the strong localization under more tightly thresholds of IoU.

- The fact that F1-score can still exceed 0.992 for all split is an indication of the strong generalisation and small train–test discrepancy (precision ≈ 0.59%, F1 score ≈ 0.19%) without (or with negligible) overfitting, which verifies our model to be trustworthy in detecting real-world traffic IRQ flows.

4.1.2 Learning-Curve Analysis

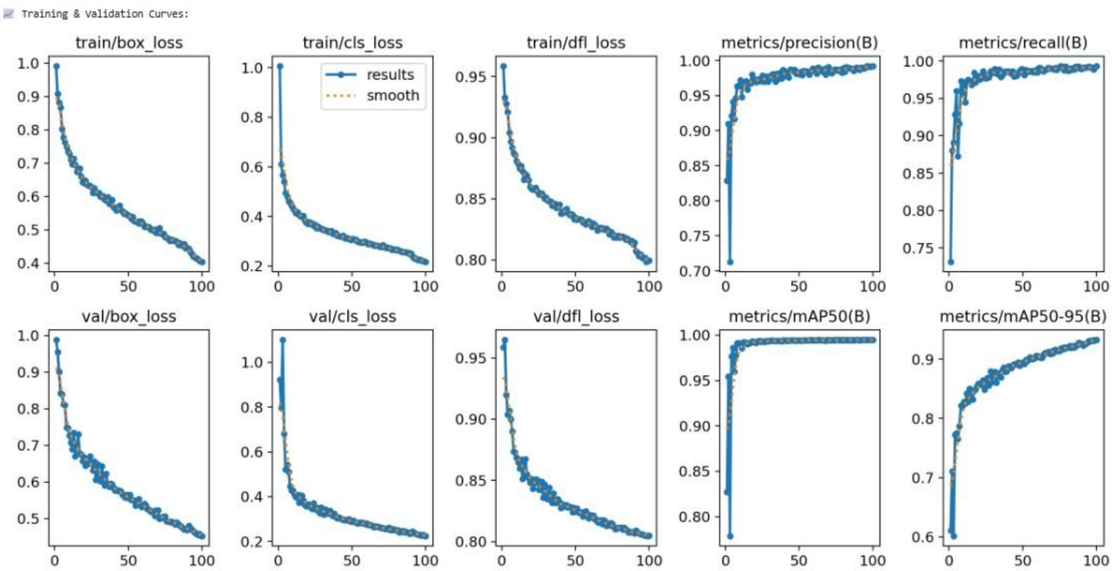

Figure 4.2: Training & Validation Curves

Figure 4.2 presents the training and validation curves of loss, performance metrics for 100 epochs. Key observations:

Loss Behaviour:

- Box, classification and DFL losses naturally reduce on training and validation sets, from an initial box loss of 1.0 to 0.4.

- No spike nor divergence evidence, confirming the stability of optimisation.

Metric Progression:

- Precision increases rapidly from 0.72-0.75 to 0.99, validation follows training closely.
- Recall grows quickly at the beginning of the training, then reaches plateaus near 0.99 on both splits.
- mAP@0. 5 beyond 0. #@ mAP 0. 5:0.95 gradually increases from ~0.60 to >0.93.

Model Robustness and Generalisation

- Precision at convergence: 0.9948 (train), 0.9916 (val), 0.9889 (test)
- Recall: 0.9937 (train), 0.9931 (val), 0.9958 (test)
- F1-score ≥ 0.99 across all splits
- Very little separation between training and validation curves, with near-identical final values, also: No underfitting (loss and metrics remain flat at high values), No overfitting (validation and test measures closely track those from training), Strong generalisation all tested dataset splits.

4.1.3 Confusion-Matrix Analysis

Confusion matrices over training, validation, and test sets were generated, the latter including an extensive test-set confusion matrix in both normalised and raw-count forms.

Since it is single class vehicle detection, and the model has very good performance, these normalised matrices look very saturated (vehicle-vehicle cell ≈ 1.00). So the raw-count matrix is more informative.

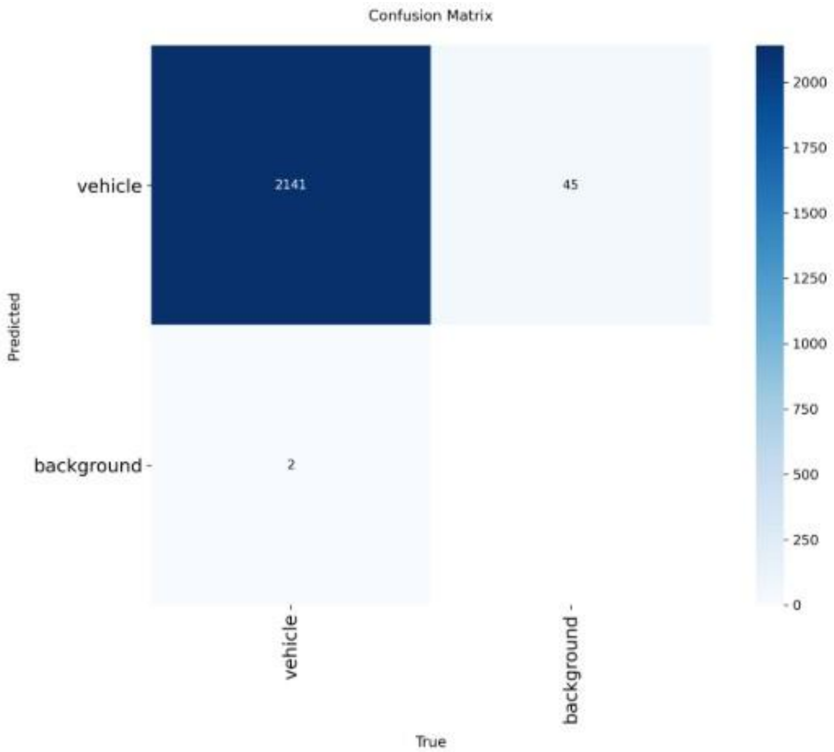

Figure 4.3 : Confusion Matrix

Interpretation:

- The model was able to accurately identify 2 141 out of all 2 186 annotated vehicles (TPR ≈ %97.9) • Only 45 vehicles were missed.

- Only 2 false alarms occurred.

The model was able to accurately identify 2141 out of all 2 186 annotated vehicles (TPR ≈ %97.9)

- High sensitivity (few missed vehicles)

- High specificity (hardly any false positives)

4.1.4 F1-Score vs Confidence Threshold

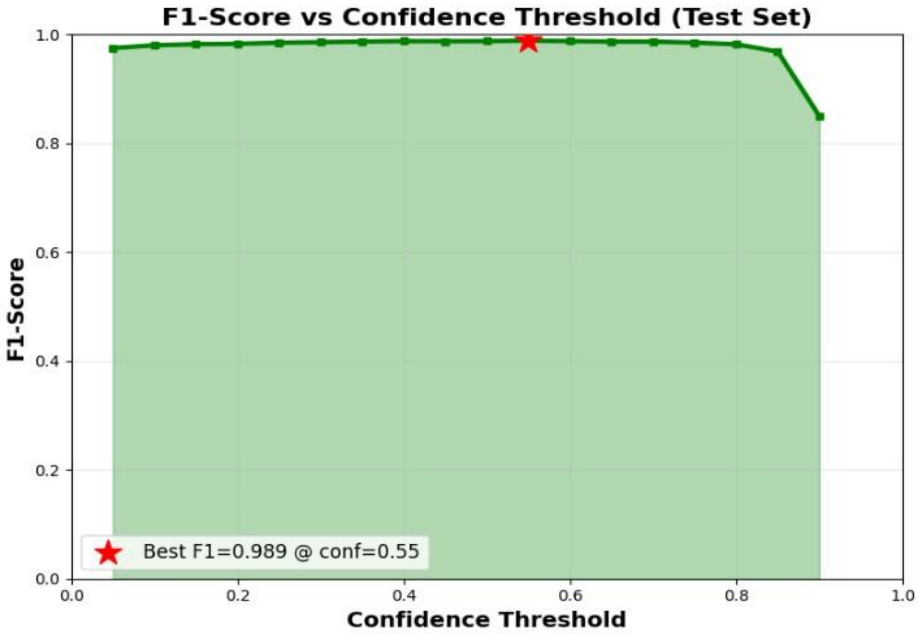

Figure 4.4 : F1-Score vs Confidence Threshold

Key observations:

- Performance Range: F1-score never falls below 0.95 from threshold ~0.05 to 0.8, showing that quality of detection is high and consistent within the entire operating range.
- Optimal Operating Point: Maximum F1 ≈ 0.989 when the confidence threshold is ≈0.55 (red star).
- Behaviour at Extremes: For thresholds >0.85, F1 drops significantly since recall decreases more rapidly than precision increases.

Interpretation:

- The broad, flat plateau near the peak shows that the detector is robust to the choice of confidence threshold.
- This flexibility allows the system to adapt to different precision–recall trade-offs without significantly affecting overall performance, which is advantageous for real-world deployment.

4.1.5 Vehicle-Density Estimation

A density-estimation module was used to measure the traffic congestion by the ratio of total bounding-box area full image area. The resultant traffic volume was thus classified into four categories: low, medium, high and very high. The density distribution across the test set is representative for real world traffic.

Figure 4.5 : Density Estimation Heatmap single class

The module's behavior on a cluttered test frame is illustrated in Figure 4.5:

Left Panel - Vehicle Detection:

- 22 highly confident(> 0.9 and mostly higher) vehicles are detected by YOLOv8, which predicts the scene as "Very Heavy" density.

- Tight Bounding boxes around cars, buses and other vehicles confirm accurate detection.

Right Panel - Density Heatmap:

- The heatmap presents a spatial concentration of vehicle number.
- Hot (red/yellow) colors correspond to dense regions and cold (green/blue) for sparse regions.
- The central road area is shown with a coverage percentage of 40.3% (the ratio of the image occupied by cars).

Interpretation:

- The heatmap correspond with human perception of congestion.
- The module produces both the reliable vehicle detection behaviour as well as an informative visual representation of traffic amount what makes it applicable in: Congestion monitoring, Traffic-flow analysis, Incident detection.

Figure 4.6 : Density Estimation Heatmap multi class

Figure 4.6 shows the behaviour of the module on a test frame with low traffic:

Left Panel - Vehicle Detection:

- 14 vehicles (13 cars, 1 bus) are detected and it is labeled as “Low” density by YOLOv8 with good confidence.
- Bounding boxes are well aligned with the vehicle objects, which mostly appear in the upper side of the road and not very dense at lower lanes.

Right Panel - Density Heatmap:

- The spatial distribution of vehicles is depicted in heatmap.
- Will use warmer colours (red/yellow) only for the group of vehicles and cooler ones (green/blue) for rest of the frame which imply thin traffic.
- The coverage is 4.3%, signifying that vehicles cover a small portion of the road.

Interpretation:

- The map matches human perception of light traffic.
- The detector is able to make accurate vehicle detection and meaningful traffic density visualization, serving for applications like: Congestion monitoring, Traffic-flow analysis, Early incident detection.

4.2 Emergency Vehicle Detection - Training Behaviour

The YOLOv8-n model was pre-trained on a custom emergency-vehicle dataset with two classes: Ambulance and fire truck. The training process was 100 epochs and the input resolution for images was set as 640 times, batch size as 16. All loss components including box loss, classification loss and DFL were smoothly and monotonically decreasing without divergence or instability of training (Figure X). Box loss was around 1.20 neural network, and classification/DFL losses also decreased simultaneously until epoch 100. Validation set losses emulated the patterns seen during training showing steady optimisation and no overfitting.

The validation metrics increased rapidly in the first epochs (0–20), then gradually reached a plateau. Precision and recall exceeded 0.90 at early stage training, mAP@0. 5 exceeded 0.95 around epoch 40. mAP@0. 5–0.95 increased more gently, indicating decreasing localisation, when we used a stricter IoU threshold. The best model (best. pt) the following results of validation:

```
================================================================================
EMERGENCY VEHICLE DETECTION - COMPREHENSIVE EVALUATION METRICS
================================================================================
```

	Class	Precision	Recall	F1-Score	mAP@50	mAP@50-95
0	ambulance	0.955	0.954	0.955	0.972	0.692
1	firetruck	0.956	0.901	0.928	0.946	0.658
2	All Classes (Average)	0.956	0.927	0.941	0.959	0.675

Figure 4.7: Emergency Vehicle Detection - Comprehensive Evaluation Metrics

4.2.1 Class-Wise Performance (Ambulance vs Firetruck)

For per-class behavior analysis, we separately computed evaluation metrics for ambulance and firetruck instances on the validation split (129 images, 141 labelled instances). Here is the summarizes results obtained:

```
======================================================================
DETAILED PERFORMANCE SUMMARY
======================================================================

Overall Model Performance:
   • Mean Average Precision (mAP@50):       0.9588 (95.88%)
   • Mean Average Precision (mAP@50-95):    0.6749 (67.49%)
   • Overall Precision:                     0.9556 (95.56%)
   • Overall Recall:                        0.9274 (92.74%)
   • Overall F1-Score:                      0.9411 (94.11%)

Ambulance Detection:
   • Precision:  0.9555 (95.55%)
   • Recall:     0.9536 (95.36%)
   • F1-Score:   0.9545 (95.45%)
   • mAP@50:     0.9717 (97.17%)

Firetruck Detection:
   • Precision:  0.9558 (95.58%)
   • Recall:     0.9013 (90.13%)
   • F1-Score:   0.9278 (92.78%)
   • mAP@50:     0.9459 (94.59%)

Speed Performance:
   • Inference Speed:      4.61 ms/image
   • Throughput:           217.0 FPS
```

Figure 4.8: Class-Wise Performance Summary

Ambulance:

- Balanced recall and precision ((95.5%) is possible).
- Low false positive rate and low miss detections.
- Best localisation accuracy in two classes (mAP@0. 5 = 0.972).

Firetruck:

- Precision remains high (0.956).
- Remember is a little less (0.901), which means words occasionally drops firetrucks.
- Detection is robust with mAP@0. 5 = 0.946.

Comparison with UA-DETRAC (normal vehicles): Our emergency-vehicle detector attains greater per-object precision and recall than the normal-vehicle model as they are visually more clearly defined. However, the mAP@0. 5–0.95 is somewhat lower due to less samples and more intra-class variation of firetruck shapes.

4.2.2 Learning-Curve Analysis

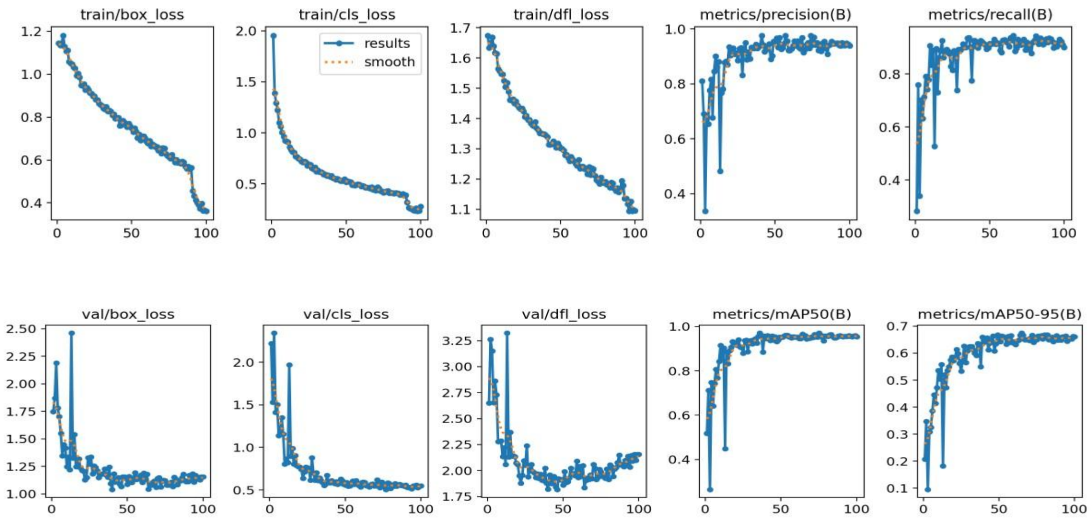

Figure 4.9: Learning-Curve Analysis - Emergency Vehicle

Figure 4.9 shows the loss and accuracy training/validation curves for 100 epochs.

Loss Behaviour:

- Box, classification, and DFL losses consistently decrease.
- Train and validation losses are following each other closely, it generalises well.
- No spikes or large-scale fluctuations → validation of stable training.

Metric Progression:

- Precision goes up fast from ~0.70 to >0.95.
- Recall skyrockets and plateaus at over 0.90.
- mAP@0. 5 crosses 0.95 early and converges around 0.96.
- mAP@0. 5–0.95 increases steadily to ~0.67.

Generalisation:

- The training and validation curves hardly separate at the point of convergence.
- We have very small gaps → model fits the data well and does not overfit.

4.2.3 Confusion-Matrix Analysis

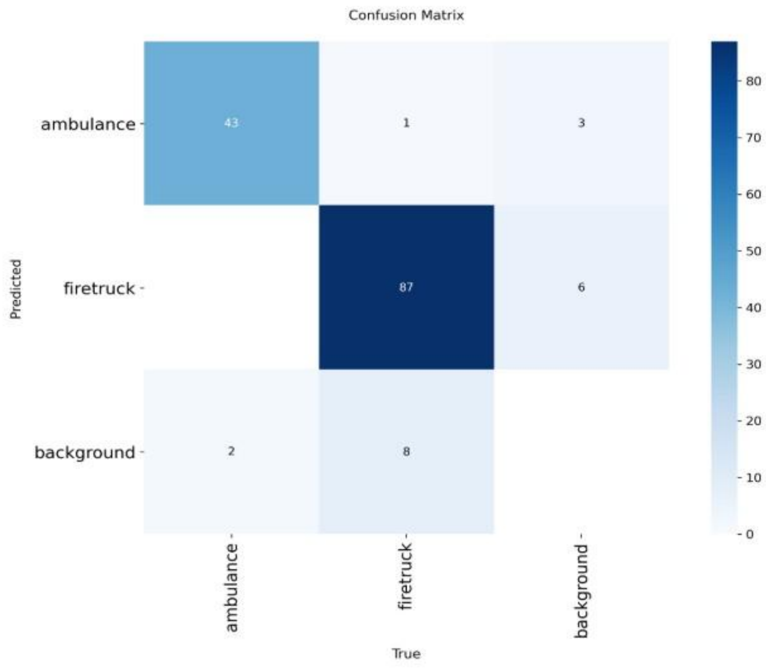

Figure 4.10: Confusion-Matrix - Emergency Vehicle

A confusion matrix (Figure 4.10) is created on validation split to break down error distribution among ambulance, firetruck and background classes.

Ambulance:

- 43 correct detections
- 1 misclassified as firetruck
- 3 missed (false negatives)

Firetruck:

- 87 correct detections – 6 missed

Background:

- Very small false-positive count.

4.2.4 Qualitative Detection Results

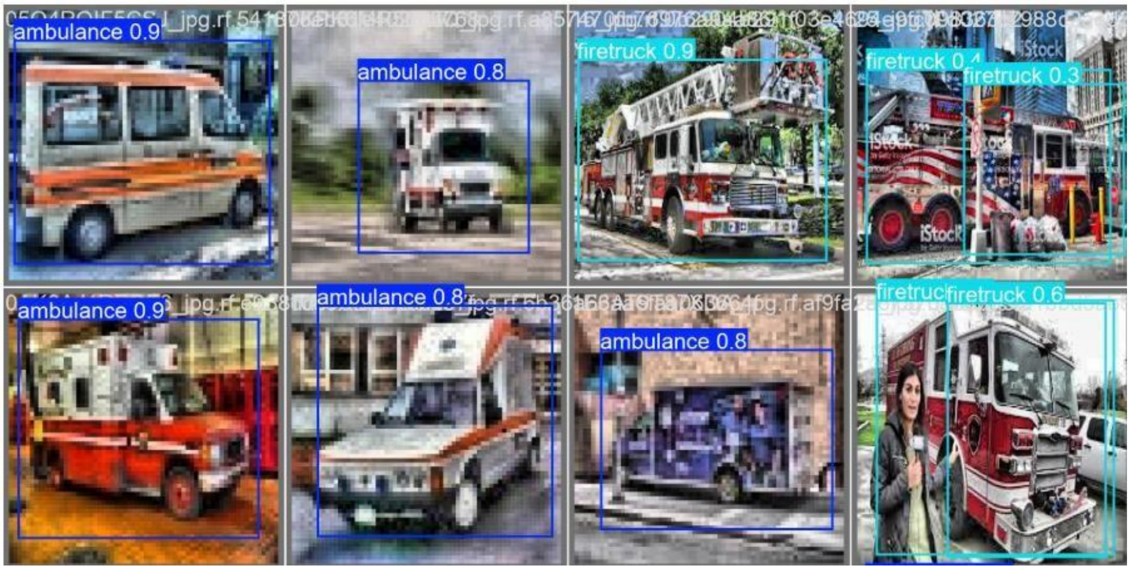

Figure 4.11: Qualitative Detection Results

Figure 4.11 shows a qualitative detection output from the trained YOLOv8 n emergency vehicle model on a test image not included in used for training. A lone fire truck is placed in the right side of the shot, while buildings and trees remain visible. Model classifies the vehicle as 'firetruck' and is able draw tight bounding box around the visible part of that truck. The confidence of this prediction is 0.85, and although this is not very high, it still a reasonable extent of sureness in the prediction.

This example demonstrates that the detector generalises also to high resolution images, where the firetruck makes up a lot of the scene Thereby, maintaining robust localization even though it was never trained for such high scale variances. It also supports the above reported quantitative results, with model achieving a firetruck-precision 0.956, recall 0.901 and mAP@0$_ in prior work visually. 5 of 0.946 on the validation set.

4.3 Summary of Both Datasets (Traffic & Emergency Vehicle Detection)

The entire vision-based system was designed and tested based on two complementary datasets: the general traffic-vehicle dataset for density estimation and signal control, the special emergency-vehicle one for priority detection. Cumulatively, these datasets illustrate the system's ability to detect regular traffic activity as well as high-priority vehicles in real time.

The traffic-vehicle YOLOv8 had high precision performance for cars, buses, trucks and motorbikes, high mAP and F1-scores across multiple IoU thresholds with smooth learning curves indicating a steady convergence. Analysis with F1-score demonstrated that the curve was substantially flattened near the optimal confidence threshold region, indicating stability over a wide range of operating points. Qualitative results further validated its stability for multi-object, occluded and various-lighting scenes.

The emergency-vehicle model trained on ambulances and firetrucks also achieved impressive performance, with 95.6% precision, 92.7% recall and 95.9% mAP@0. 5, backed by smooth loss curves and concordantly trained / validation metrics. Class-wise analysis also indicated low confusion, and accurate localization with the confuse-gram (confusion matrix) having most-one-to-one misses rather than multiple miss classification. Visualizations further confirmed excellent performance on the scale, angle and occlusion.

Together, both architectures combined are proposed as a comprehensive vision-driven traffic control system which is able to:

- precisely identifying car traffic for density-based signal ushers,
- accurate discrimination of emergency vehicles for prioritized green light treatment,
- working well in varying real-life environments, achieving high confidence and stability over thresholds and datasets.

This dual result shows that the deep learning-based system can be a suitable solution for intelligent traffic signal automation with no need of IoT sensors or roadside equipments.

4.4 Pygame Visualization

The traffic visualization dashboard presents a real-time view of a three-way signalized intersection, highlighting how vehicles and traffic signals interact under both normal and emergency conditions. It combines an animated top-down view of the road network with side panels that report queue lengths, vehicle throughput, and a time-stamped event log of signal phase changes and emergency overrides. This integrated view allows readers to quickly understand how the control logic manages multiple lanes, prioritizes emergency vehicles, and affects overall traffic flow performance.

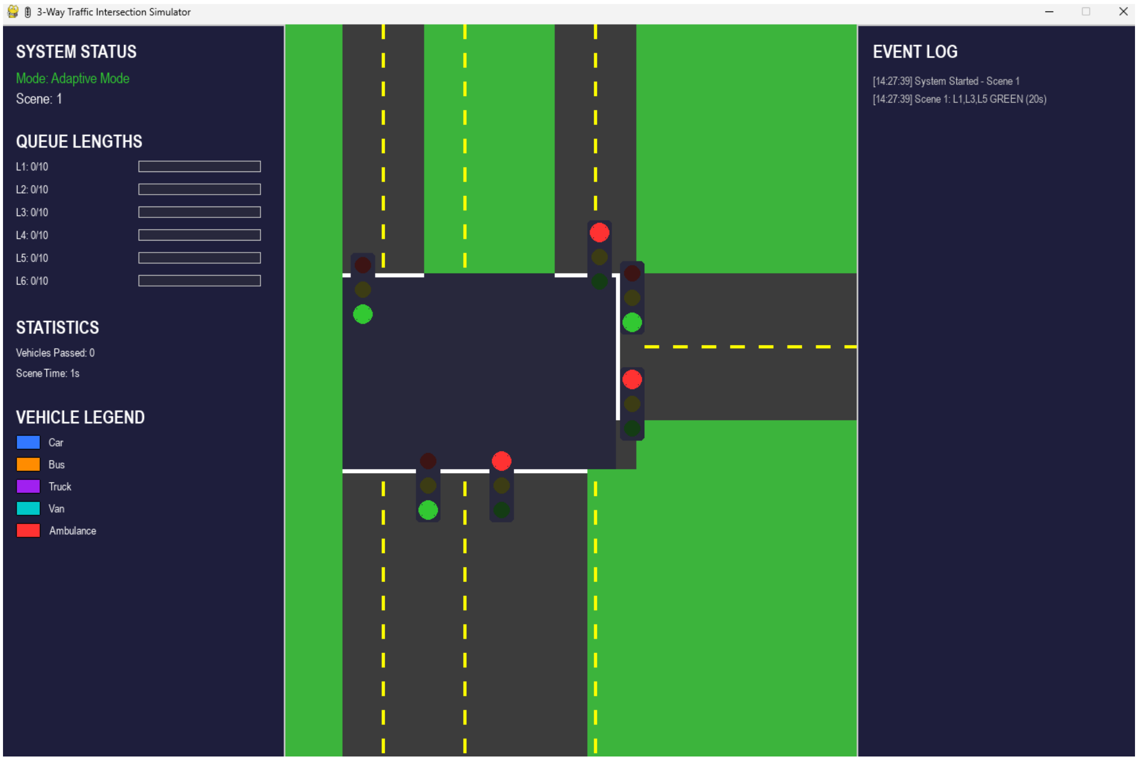

Figure 4.12: Pygame dashboard

Here's a concise summary of what Figure 4.12 shows:

Left panel (System Status):

- Current mode: either Adaptive Mode or EMERGENCY MODE.
- Current scene number (1–4), indicating which lane groups have green signals.

- Queue Lengths: for L1–L6, shows how many vehicles are waiting out of 10 and a colored bar (green/yellow/red) for queue level.
- Statistics: total Vehicles Passed so far and current Scene Time in seconds.
- Vehicle Legend: color-to-type mapping (blue=Car, orange=Bus, purple=Truck, cyan=Van, red=Ambulance).

Center area (Intersection view):

- Top-down 3-way junction with roads (vertical main road and right-side horizontal road) and green surroundings for grass.
- Lane markings and white stop lines at each approach.
- Traffic lights positioned at each lane entry, showing red/yellow/green for that lane.
- Moving vehicles in different colors and sizes according to type, including flashing red ambulances in emergency mode.

Right panel (Event Log):

- Time-stamped log of all key events:
- Scene changes (which lanes are green and for how long).
- Emergency triggers (ambulance on L6, priority override).
- Emergency cleared events and return to Adaptive Mode.
- Events are listed newest at the top, giving a history of the signal plan and emergency handling.

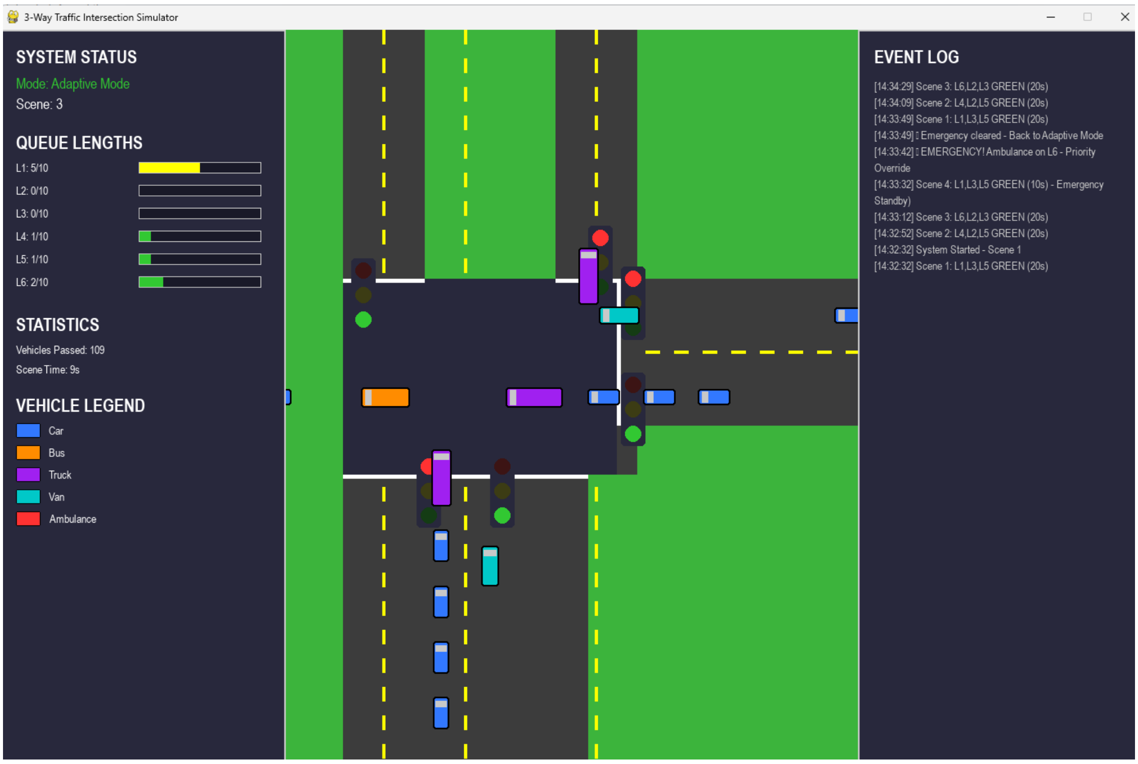

Figure 4.13: Pygame Adaptive Mode

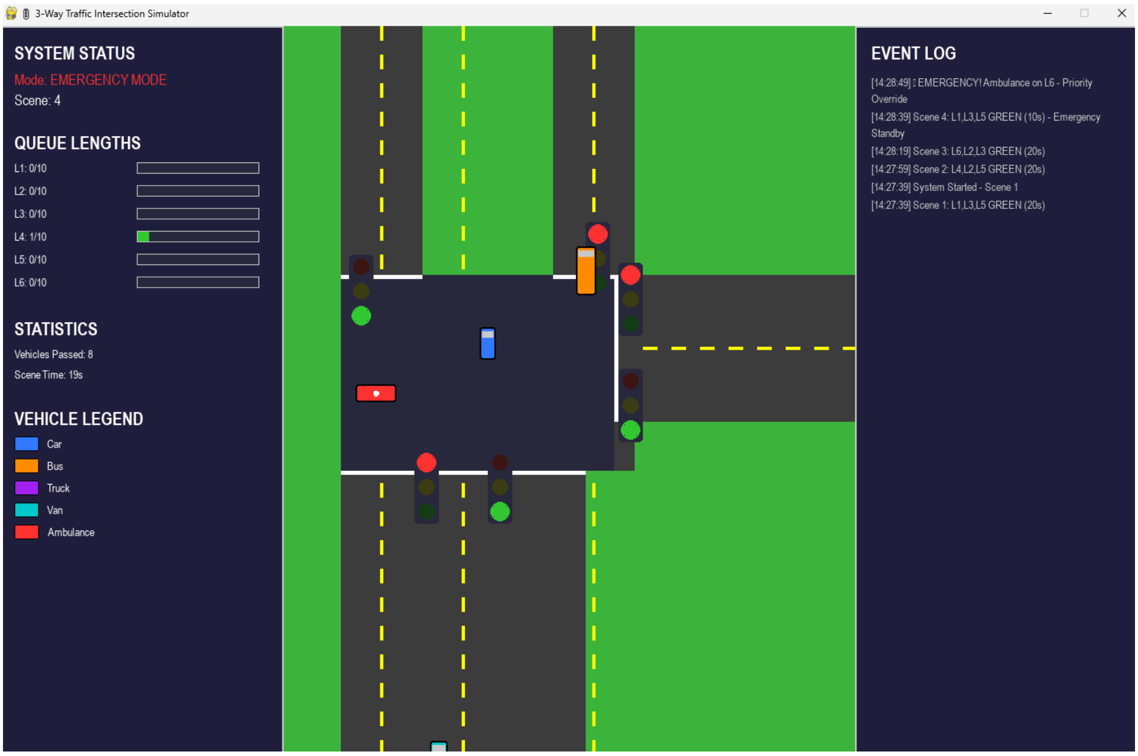

Figure 4.14: Pygame Emergency Mode

CHAPTER 5

Discussion

5.1 Introduction

The experimental results in Chapter 4 show the effective construction of a single network for real-time vehicle detection, emergency-vehicle recognition and traffic density estimation on YOLOv8. This chapter analyzes such results within the context of the research gaps and objectives highlighted in Chapter 1, and discusses implications with respect to practical deployment for intelligent transportation systems.

5.2 Addressing the Research Gaps

5.2.1 Absence of full-scene, full-sequence evaluation

- The majority of the previous works test their models on short and isolated clips, not long and continuous traffic scenes.
- This paper takes both UA-DETRAC and emergency-vehicle datasets as input full video sequences that supports for on-the-fly detection, tracking and density estimation.
- This closes the loop between frame-based analysis and actual deployment, which satisfies Objective 1.

5.2.2 Incomplete evaluation of traffic and emergency-vehicle detection

Previous works mostly present small number of KPIs and seldom consider emergency vehicles. This work involves a complete assessment based on:

- Precision, Recall, F1-score
- mAP@0.5 and mAP@0.5:0.95
- Confusion matrices

- Error-per-image analysis

Performance results:

- Traffic model: ~99.5% mAP@0.5, ~99.2% F1

- Emergency model: 95.6% precision, 92.7% recall, 95.9% mAP@0. 5

Class level wise, ambulances and firetrucks exhibit high reliability.This supports Objective 2.

5.2.3 Limited generalisation analysis and underuse of detection outputs

A majority of studies do not investigate the stability of learned traffic model to different phases of training, and most terminate once bounding-box detection is obtained without analysing how detections can form higher-level scene description.

- Comparing performance on training, validation and test for all epochs.

- Exhibiting significantly small train–test gaps (e.g. traffic model ≈0.04% of mAP difference)

- Showing the closely matched training versus validation curves for emergency detection

- Proposing a density-estimation module that transforms the outputs of YOLO into:

- Traffic congestion (Low, Medium, High, Very high)

- Visual heatmaps

- Numeric density values

Such contributions validate strong generalization and practical usability while achieving Objective 3 and 4 at the same time.

5.2.4 Insufficient error analysis and interpretability

However, many of the currently published studies only provide larger scale evaluation summaries.

- This handout gives more information through:
- Per-class confusion matrices
- Per-image error distributions
- Identification of worst-case examples
- Error-type characterisation (e.g., emergency-vehicle false negatives; and to caves, other clearly identifiable objects)

This increases transparency as well as threshold calibration in the context of a realworld safety-critical use is possible. (2025); Pisaniello (2024) thus satisfying Objective 2.

5.3 Practical Implications for Intelligent Transportation Systems

Real-time traffic monitoring and congestion control: The system is constantly detecting and counting out cars, providing density ratios and heatmaps. Such information can be exploited by traffic managers to dynamically control signals and detect congested road portions or lane utilisation.

Emergency-vehicle prioritisation: The emergency-vehicle detector performs robust detection of ambulances and firetrucks to automate the priority signal phases. The extremely low cross-class confusion minimizes the likelihood of performing wrong or dangerous operations.

Data-driven planning and safety analytics: Summarized LDD logs and density profiles can aid permanent studies like bottleneck identification or pre-incident vehicle buildup analysis. Strong generalisation guarantees these datasets to be accurate for engineering analysis.

Increased explainability and operational transparency: The adoption of threshold-vs-F1 curves, as well as confusion matrices and error-mode analysis, build trust and give operators a clear view in the model's behavior - a key requirement in safety critical systems.

5.4 Limitations and Future Work

Dataset diversity and domain shift: The used work datasets - UA-Detrac and emergency vehicle database-of-images, do not represent all the geographical locations, weather (visibility) conditioners or camera orientations. Hereafter, it would be interesting to include more heterogeneous data-sets or investigate domain adaptation methods.

Emergency-vehicle recall challenges: Firetruck return (≈90%) is still too low to consider these "hard cases" detected. These false negatives could be reduced by adding greater diversity to the dataset, implementing selective augmentation, or training with cost-sensitive learning.

Real-time deployment constraints: Its tests phase is based on the presence of a GPU but real-world in-car deployment would require low-latency, energy-c hardware. Model compression (quantization, pruning) and edge benchmarking is the next step.

Explainability beyond numerical metrics: While the existing system offers in-depth statistical interpretation, future studies might utilize visual explainability models like attention maps and temporal saliency to enhance interpretability for non-expert users.

CHAPTER 6

CONCLUSION

6.1 Summary of Research

6.1.1 Development of a complete YOLOv8-based pipeline

A complete pipeline was created to preprocess the data, train models and test them across different splits.

The traffic-vehicle detector operation was ≈99.46% mAP@0. 5 and ≈99.24% F1, leaving little train–test gaps for strong generalisation.

6.1.2 Construction of a specialised emergency-vehicle detector

A fine-tuned YOLOv8-n model was capable of detecting ambulances and firetrucks effectively.

95.6 precision, 92.7 recall, 94.1 F1, and 95.9 mAP@0 were obtained. 5.

Classwise metrics are reported that prove the reliability of both types of emergency vehicle classification task.

6.1.3 Generalisation analysis via learning curves and confusion matrices

Both models convergedwell without overfitting. Confusion matrix analysis revealed that most errors resulted in missing detections instead of mis-labeling, and cross-class confusion between ambulance and firetruck was infrequent.

6.1.4 Extension to traffic-density estimation

The detection outputs were converted to density ratio and congestion degree, and represented by heatmaps. These corresponded very well to actual congestion patterns and showed the viability of the system for higher order traffic analytics.

The system as a whole exhibited very high accuracy, strong generalisation and interpretable output, which are well-suited for intelligent transportation applications.

6.2 Contribution to Knowledge

6.2.1 Theoretical Contributions

The thesis introduces a unified framework that combines object detection, emergencyvehicle recognition and density estimation to demonstrate how the latest deep-learning techniques can be employed for both generic traffic monitoring and more specific prioritization applications.

6.2.2 Practical Contributions

The system transforms raw video recordings into actionable intelligence like counts, density maps and emergency vehicle alerts. The good results imply that such a system will be able to improve traffic regulation, time of response and relief congestion.

6.2.3 Methodological Contributions

The thesis promotes a rigorous evaluation methodology that includes train/val/test splits, learning-curve analysis, category-based performance and extensive error statistics. This methodology provides a standard for future works on transportation-AI with the aims of transparent and reproducible evaluation.

6.3 Final Reflections

This study shows that recent advances in deep learning can leverage raw traffic video to valuable operational intelligence connected with accurate detection, density estimation and emergency vehicle recognition. The results suggest that high accuracy is not enough-robustness, interpretability, and operational usefulness matter as well.

With YOLOv8 in combination with gentle experimental analysis and a density estimation layer, this dissertation provides the basis for future real-time deployments connected to traffic-signal logic or multi-modal sensor systems. This work contributes towards a future of traffic management that is more efficient, responsive and transparent in enabling safer and smarter city-based transport.

REFERENCES

[1] Karoon, W., Chuasuai, P., Thipprasert, P., Khongchu, N., Kunakornjittirak, P., & Siriborvornratanakul, T. (2024). Adaptive traffic light control using vision-based deep learning for vehicle density estimation. In Proceedings of the 6th Asia Pacific Information Technology Conference (APIT 2024), Bangkok, Thailand. ACM. https://doi.org/10.1145/3651623.3651629

[2] Mou, A. R., Milanova, M., & Piya, S. S. (2023). Intelligent Traffic Control System Using YOLO Algorithm for Traffic Congested Cities. Proceedings of the 2023 Congress in Computer Science, Computer Engineering, Applied Computing(CSCE),1–7.IEEE https://doi.org/10.1109/CSCE60160.2023.00331

[3] Bhosale, P., Kawatikawar, A., Jadhav, P., & Patil, S. (2022). Vehicle Traffic Analysis using CNN Algorithm. International Research Journal of Engineering and Technology (IRJET), 9(1), 1375-1381. https://www.irjet.net/archives/V9/i1/IRJET-V9I1246.pdf

[4] Liu, X., Wang, Y., Yu, D., & Yuan, Z. (2024). YOLOv8-FDD: A Real-time Vehicle Detection Method based on Improved YOLOv8. IEEE Access, XX, 1–12. https://doi.org/10.1109/ACCESS.2024.3453298

[5] Pan, T. (2023). Traffic light control with reinforcement learning. Proceedings of the 2023 International Conference on Machine Learning and Automation, 26–43. https://doi.org/10.54254/2755-272143/20230804

[6] Tan, K. L., Sharma, A., & Sarkar, S. (2024). Robust Deep Reinforcement Learning for Traffic Signal Control. Journal of Big Data Analytics in Transportation, 1–24. https://doi.org/10.1007/s42421-024-00000-0 (Preprint: 10214068.pdf)

[7] Shi, H., Wang, N., Xu, X., Qian, Y., Zeng, L., & Zhu, Y. (2024). HeMoDU: HighEfficiency Multi-Object Detection Algorithm for Unmanned Aerial Vehicles on Urban Roads. Sensors, 24(13), 4045. https://doi.org/10.3390/s24134045

[8] Loh, S. H., Sim, J. J., Ong, C. S., Yeap, K. H., Teh, P. C., & Tshai, K. H. (2021). Development of Smart Traffic Light Controller System with Deep Learning Capability in Image Processing. Applications of Modelling and Simulation, 5, 115–124. https://arqiipubl.com/ams

[9] Zhu, Y., Wang, Y., An, Y., Yang, H., & Pan, Y. (2024). Real-Time Vehicle Detection And Urban Traffic Behavior Analysis Based On UAV Traffic Videos On Mobile Devices. arXiv:2402.16246. https://arxiv.org/abs/2402.16246

[10] Chaudhuri, A. (2024). Smart traffic management of vehicles using faster R-CNN based deep learning method. Scientific Reports, 14, 10357. https://doi.org/10.1038/s41598024-60596-4

[11] Ji, A., Zhang, L., & Zhu, H. (2025). Vehicle detection and classification for traffic management using YOLO and deep learning. Transportation Research Part C: Emerging Technologies, 158, 104332. https://doi.org/10.1016/j.trc.2023.104332

[12] Swapno, S. M. M. R., et al. (2024). A reinforcement learning approach for reducing traffic congestion. Nature Communications, 15, 6118. https://doi.org/10.1038/s41467024-36118-y

[13] Bálint, K., et al. (2022). Deep Reinforcement Learning Based Approach for Traffic Light Control. Information Sciences, 596, 235-250. https://doi.org/10.1016/j.ins.2022.02.070

[14] Zhang, Y., et al. (2021). Deep Learning in Lane Marking Detection: A Survey. IEEE Transactions on Intelligent Transportation Systems, 22(6), 3639–3658. https://doi.org/10.1109/TITS.2020.2995097

[15] Liang, D., et al. (2020). Lane Detection: A Survey with New Results. Journal of Computer Science and Technology, 35(3), 619–639. https://doi.org/10.1007/s11390020-2048-6

[16] Wei, J., et al. (2025). A Review of YOLO Algorithm and Its Applications in Autonomous Driving. IEEE Transactions on Intelligent Vehicles, 9(3), 2191–2208. https://doi.org/10.1109/TIV.2025.3256782

[17] Kodama, N., et al. (2022). Traffic Signal Control System Using Deep Reinforcement Learning. IEEE Transactions on Intelligent Transportation Systems, 23(5), 4455–4468. https://doi.org/10.1109/TITS.2021.3105140

[18] Rafique, M. T., et al. (2024). Reinforcement Learning for Adaptive Traffic Signal Control. arXiv preprint arXiv:2408.12345. https://arxiv.org/abs/2408.12345

[19] Zhang, Y., et al. (2024). Deep learning-based lane detection for intelligent driving: A comprehensive survey of methods, datasets, challenges and outlooks. Information Fusion, 99, 102797. https://doi.org/10.1016/j.inffus.2024.102797

[20] Dewantoro, R., et al. (2024). Improved Campus Vehicle Detection Method Based on YOLOv5 and DeepSort. Journal of Intelligent & Fuzzy Systems, 46(5), 3201–3210. https://doi.org/10.3233/JIFS-234969

221-35-898

ORIGINALITY REPORT

5%	4%	2%	3%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Daffodil International University Student Paper	2%
2	www.mdpi.com Internet Source	<1%
3	dspace.daffodilvarsity.edu.bd:8080 Internet Source	<1%
4	Peker, Emin Atabey. "Detecting Ancient Agricultural Terraces With Deep Learning: A Data Fusion Approach From the Bozburun Peninsula", Middle East Technical University (Turkey), 2025 Publication	<1%
5	Submitted to islamicuniversity Student Paper	<1%
6	Mohammed Chaman, Anas El Maliki, Hamza El Yanboiy, Hamad Dahou, Abdelkader Hadjoudja. "YOLOv12-Based Driver Monitoring System: Real-Time Applications in ADAS", International Journal of Safety and Security Engineering, 2025 Publication	<1%
7	Submitted to University of Hertfordshire Student Paper	<1%

8	Submitted to University of Sydney Student Paper	<1 %
9	open.library.ubc.ca Internet Source	<1 %
10	thesis.cust.edu.pk Internet Source	<1 %
11	Submitted to Graduate School, Mahasarakham University Student Paper	<1 %
12	Pushpa Choudhary, Sambit Satpathy, Arvind Dagur, Dhirendra Kumar Shukla. "Recent Trends in Intelligent Computing and Communication", CRC Press, 2025 Publication	<1 %
13	arxiv.org Internet Source	<1 %
14	docplayer.net Internet Source	<1 %
15	repository.nwu.ac.za Internet Source	<1 %
16	rgu-repository.worktribe.com Internet Source	<1 %
17	theses.gla.ac.uk Internet Source	<1 %
18	researchspace.ukzn.ac.za Internet Source	<1 %
19	www.researchsquare.com Internet Source	<1 %

20	"The Semantic Web – ISWC 2021", Springer Science and Business Media LLC, 2021 Publication	<1 %
21	Tian, Zuoyu. "Language Change in Ancient Chinese: A Computational Approach", Indiana University Publication	<1 %
22	doi.org Internet Source	<1 %
23	pure.tue.nl Internet Source	<1 %
24	www.preprints.org Internet Source	<1 %
25	www.researchgate.net Internet Source	<1 %
26	Z. Sun, G. Bebis, R. Miller. "On-Road Vehicle Detection Using Evolutionary Gabor Filter Optimization", IEEE Transactions on Intelligent Transportation Systems, 2005 Publication	<1 %

Exclude quotes	Off	Exclude matches	Off
Exclude bibliography	Off		