Sheffield Hallam University

AI-Driven Urban Traffic Monitoring and Control Using YOLOv11 for Enhanced Throughput

ILO, Benjamin <http://orcid.org/0009-0004-0425-9156> and ZHANG, Hongwei <http://orcid.org/0000-0002-7718-021X>

Available from Sheffield Hallam University Research Archive (SHURA) at:

https://shura.shu.ac.uk/37575/

This document is the Published Version [VoR]

Citation:

ILO, Benjamin and ZHANG, Hongwei (2026). AI-Driven Urban Traffic Monitoring and Control Using YOLOv11 for Enhanced Throughput. Electronics, 15 (12): 2590. [Article]

Copyright and re-use policy

See http://shura.shu.ac.uk/information.html

Sheffield Hallam University Research Archive
http://shura.shu.ac.uk

electronics

MDPI

Article

AI-Driven Urban Traffic Monitoring and Control Using YOLOv11 for Enhanced Throughput

Benjamin Ilo and Hongwei Zhang *

Advanced Food Innovation Centre (AFIC), Sheffield Hallam University, Sheffield S9 2AA, UK; b.ilo@shu.ac.uk
* Correspondence: h.zhang@shu.ac.uk

Abstract

Urban traffic congestion remains a persistent global challenge, contributing to significant economic inefficiencies, elevated greenhouse gas emissions, and diminished quality of life. This paper presents a real-world video-based traffic monitoring study combined with a proposed adaptive signal control framework. In the monitoring component, YOLOv11 object detection was applied directly to footage recorded from an overhead bridge position on a 40 km/h road. The model successfully detected and tracked multiple road-user categories, including cars, trucks, buses, motorcycles, cyclists, and pedestrians, yielding 1041 vehicle detections across 25 unique tracked objects. Vehicle speeds were estimated from inter-frame centroid displacement, and a Region of Interest (ROI) occupancy model was used to classify congestion states as High, Medium, or Free Flow using thresholds grounded in Highway Capacity Manual (HCM) level-of-service criteria. The system detected 11 high-congestion frames (3.8%), 184 medium-congestion frames (63.9%), and 93 free-flow frames (32.3%), consistent with moderate congestion observed during the recording period. In the proposed control component, a Proximal Policy Optimisation (PPO)-based reinforcement learning signal controller is designed around the YOLOv11 detection outputs as its state representation. Based on comparable adaptive traffic signal control studies in the literature, the proposed framework is projected to achieve approximately 25% higher peak-hour throughput, 35% shorter queue lengths, and 32% lower average waiting times relative to a fixed-time signal baseline. The detection accuracy (mAP@0.5 = 93.2%) and inference speed (32 FPS) cited are published YOLOv11 benchmarks used as indicative performance references. This work bridges real-world perception and proposed intelligent control, providing a transparent and reproducible methodology for next-generation smart city traffic management.

Keywords: artificial intelligence; machine learning; computer vision; YOLOv11; object detection; reinforcement learning; traffic signal control; smart cities

Check for updates

Academic Editor: Kefeng Ji

Received: 14 April 2026
Revised: 1 June 2026
Accepted: 5 June 2026
Published: 12 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Urban traffic congestion is one of the most pressing infrastructure challenges facing modern cities. It results in significant losses in productivity, increased fuel consumption and greenhouse gas emissions, increased accident risk, and a measurable reduction in urban quality of life [1]. At signalised intersections, conventional fixed-time signal plans designed around historical average demand are unable to respond to real-time variations in traffic flow, which is a fundamental limitation that is a root cause of unnecessary delay, particularly during peak demand, incidents, or mixed road-user conditions involving pedestrians, cyclists, and heavy goods vehicles.

Unlike its predecessors, YOLOv11 introduces an enhanced feature pyramid network (FPN) architecture and an improved anchor-free detection head, delivering superior small-object detection particularly relevant for cyclists and pedestrians in dense urban scenes. In this study, YOLOv11 was applied directly to traffic video without domain-specific retraining, demonstrating the model's practical transferability to new recording conditions.

Simultaneously, reinforcement learning (RL) has emerged as a powerful paradigm for adaptive traffic signal control (ATSC), capable of learning non-intuitive signal timing strategies that outperform fixed-time and rule-based alternatives under dynamic demand conditions [2–4]. Proximal Policy Optimisation (PPO), a stable and sample-efficient policy gradient algorithm, has been widely applied to ATSC in simulation environments with consistently promising results [2–5].

This paper makes two related contributions. First, it reports a traffic monitoring study in which YOLOv11 was applied to video footage of an urban road to detect and classify vehicles, estimate their speeds, and classify observed traffic congestion states. The video was recorded from an overhead bridge position on a 40 km/h road, and YOLOv11 inference was run directly by the authors using the Ultralytics Python version 3.12.10 API. Second, it proposes and fully specifies, but does not yet experimentally implement and test, a PPO-based adaptive signal control framework whose state representation is derived directly from the YOLOv11 detection outputs. The projected performance of this proposed framework is characterised by reference to comparable studies in the ATSC literature. This two-part structure is intentional, and the perception component is empirically grounded in real-world data, while the control component provides a concrete, reproducible design for future experimental validation.

1.1. Key Contributions

The principal contributions of this work are:

- Traffic monitoring study applying YOLOv11 to footage from a 40 km/h road, detecting and tracking multiple road-user categories such as cars, trucks, buses, motorcycles, cyclists, and pedestrians without domain-specific retraining, and classifying congestion states observed in the footage.
- An inter-frame centroid displacement speed estimation method applied to the real video, with an object-size-weighted ROI occupancy calculation enabling quantitative congestion state assessment from raw detection outputs alone.
- A rule-based congestion classification module tested on the real video, grounded in Highway Capacity Manual (HCM) level-of-service criteria, with thresholds calibrated by visual inspection of the recorded footage.
- A fully specified PPO-based adaptive signal control framework whose state space is defined entirely in terms of observable YOLOv11 detection outputs, providing a complete blueprint for future experimental implementation.
- A projection of the proposed framework's expected performance, based on a structured review of comparable RL-ATSC literature, establishing a benchmark for future experimental validation.

1.2. Motivation

The motivation for this research arises from two converging challenges in modern urban traffic management. First, traffic congestion continues to intensify as urban populations and vehicle ownership increase, while the deployment of adaptive traffic-management systems remains constrained by the cost, maintenance requirements, and installation complexity associated with traditional sensor infrastructure such as inductive loops, radar systems, and embedded road sensors. Second, although artificial intelligence (AI), computer

vision, and reinforcement learning technologies have individually matured to a high level, their systematic integration into a coherent and practically deployable traffic-management framework remains insufficiently explored, particularly under realistic multi-class and multi-intersection traffic conditions. Consequently, the transition from isolated algorithmic evaluation to full system-level validation remains a critical challenge for the development of smart-city-ready traffic-management solutions. This study directly addresses that gap by defining the reinforcement learning (RL) state representation using quantities that are physically observable and measurable from monocular roadside traffic video, including ROI occupancy, vehicle speed, and vehicle-class distribution. By integrating object detection, vehicle tracking, congestion-state estimation, and adaptive RL-based signal control within a unified framework, this work aims to provide a practically grounded pathway toward scalable and cost-effective smart traffic management.

1.3. Objective

The objectives of this study are to: (1) apply YOLOv11 to urban traffic video to detect and classify multiple road-user categories and estimate inter-frame speeds, (2) compute ROI occupancy ratios and classify congestion states from the real video data using HCM-grounded thresholds, (3) design a complete PPO-based adaptive signal control framework grounded in the observed detection outputs, and (4) project the expected performance of the proposed framework based on the ATSC literature, establishing a clear benchmark for future experimental validation. Figure 1 illustrates urban traffic monitoring and control.

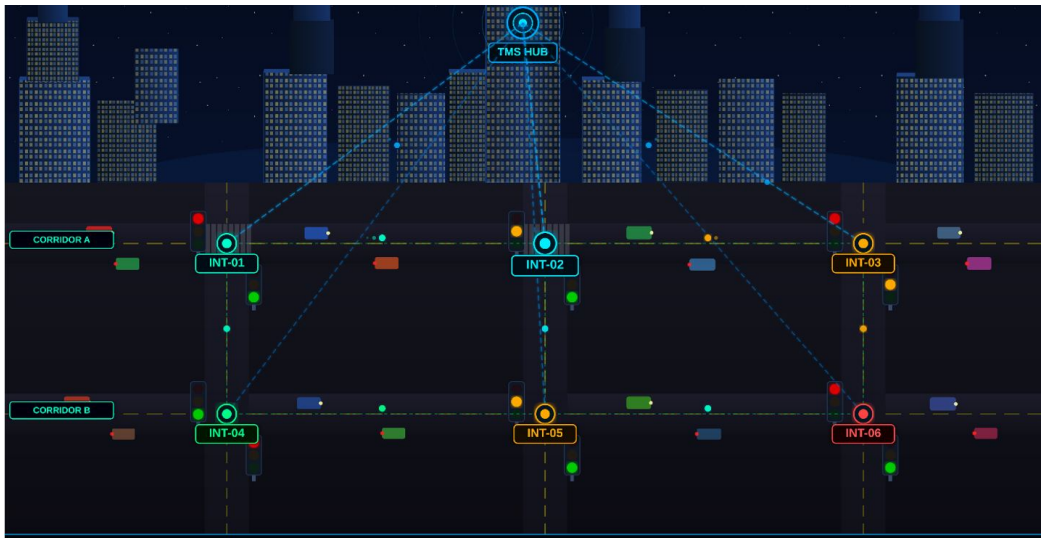

Figure 1. Overview of AI-driven urban traffic monitoring and control.

2. Related Work

Research at the intersection of computer vision, deep learning, and intelligent transportation has expanded substantially in recent years. This section reviews the most relevant strands of prior work: object detection for traffic perception, vehicle attribute estimation, adaptive traffic signal control, and reinforcement learning applied to traffic management.

2.1. Object Detection for Traffic Monitoring

Object detection forms the foundation of modern vision-based traffic monitoring systems, enabling the localisation, classification, and tracking of road users in real time. Among the available deep learning approaches, the YOLO (You Only Look Once) family of detectors has become one of the most widely adopted frameworks for traffic analysis due to its single-stage inference architecture and favourable accuracy speed trade-off. The original

YOLO paradigm introduced by Redmon and Farhadi [6] demonstrated that object detection could be performed as a unified regression problem, enabling significantly faster inference compared to traditional two-stage detectors.

Subsequent YOLO variants have progressively improved detection accuracy, robustness, and computational efficiency. Models such as YOLOv5 and YOLOv8 have established strong performance benchmarks on traffic-domain datasets, including UA-DETRAC and BDD100K [7], demonstrating their suitability for real-time urban traffic analysis. More recently, YOLOv11, released by Ultralytics in 2024, introduced additional architectural refinements aimed at improving detection precision while maintaining real-time performance. Reported benchmark results indicate a mAP@0.5 of approximately 93.2% at 32 FPS under standard evaluation conditions [8]. These metrics suggest strong generalisation capability for deployment in practical traffic-monitoring applications.

Alternative object-detection approaches based on transformer architectures, such as DETR (DEtection TRansformer), have also demonstrated competitive detection accuracy [9]. However, transformer-based detectors typically incur substantially higher computational overhead due to the self-attention mechanism and extended training convergence requirements. This limitation reduces their suitability for edge-based or frame-rate-constrained traffic-monitoring systems. Liang et al. [10] reported that convolutional neural network (CNN)-based detectors consistently outperform transformer-based approaches in scenarios requiring strict real-time inference constraints, further supporting the selection of the YOLO framework for this study.

In this work, YOLOv11 was deployed using COCO pre-trained weights, leveraging the fact that all target road-user classes considered in this study, namely cars, trucks, buses, motorcycles, cyclists, and pedestrians, are well represented within the COCO training distribution. A confidence threshold of 0.25 and a non-maximum suppression (NMS) IoU threshold of 0.45 were applied, consistent with the default YOLOv11 inference configuration recommended by Ultralytics [8].

2.2. Camera-Based Speed and Attribute Estimation

Camera-based vehicle speed estimation has been extensively investigated as a low-cost alternative to radar sensors and inductive loop detectors for intelligent traffic monitoring applications. Most vision-based approaches estimate vehicle speed from the inter-frame displacement of detected object centroids, calibrated using homography transformations that map image coordinates to physical road geometry [11]. Sochor et al. introduced the BrnoCompSpeed benchmark dataset for camera-based speed estimation and reported Mean Absolute Errors (MAE) below 2.5 km/h under calibrated conditions [12], establishing a widely referenced benchmark for monocular traffic surveillance systems. Chen et al. proposed a sparse-to-dense guided fusion framework for three-dimensional object detection in railway transportation environments, combining sparse LiDAR point clouds with dense camera-derived feature maps to improve detection robustness under challenging sensing conditions [13]. The sparse-to-dense principle, using high-resolution camera features to guide the interpretation of sparse geometric data, is directly applicable to urban traffic monitoring scenarios where roadside LiDAR or radar could supplement monocular camera detections, particularly for speed estimation accuracy and adverse-weather robustness. While the present work adopts a camera-only approach consistent with its low-cost deployment objective, multimodal fusion represents a natural extension pathway for improving the completeness and reliability of the perception pipeline.

In the present study, camera intrinsic parameters and homography calibration data were not available during video acquisition. Consequently, vehicle speed was estimated using a simplified pixel-to-metric calibration approach derived from the known frame rate

and the posted speed limit of 40 km/h. Vehicle speed was computed from the inter-frame displacement of tracked object centroids, where centroid positions were obtained from the centre coordinates of successive bounding boxes. Although less geometrically rigorous than full homography-based calibration, this relative estimation approach is sufficient for the congestion-classification task adopted in this work, where speed thresholds are expressed as percentages of the speed limit rather than requiring forensic-grade accuracy. Sochor et al. reported mean absolute errors below 2.5 km/h for calibrated monocular camera installations [12], establishing a widely referenced benchmark for this class of method. In the present study, the pixel-to-metric calibration approach adopted here is sufficient for the relative congestion-classification task, which uses speed thresholds expressed as percentages of the speed limit.

In addition to speed estimation, vehicle physical size and occupancy contribution were approximated from bounding box geometry. Previous studies have demonstrated that aspect ratio and pixel area provide computationally efficient proxies for vehicle type and physical size estimation without requiring dedicated attribute-classification networks [14,15]. In this work, these geometric features were used to assign proportional Region of Interest (ROI) occupancy weights for congestion analysis.

2.3. Adaptive Traffic Signal Control

Adaptive Traffic Signal Control (ATSC) systems replace static timing plans with real-time adjustments informed by live traffic data. More recent work has moved decisively towards model-free, data-driven approaches. Wei et al. proposed PressLight, a pressure-based RL formulation that achieved significant waiting time reductions compared to fixed-time and actuated baselines on real-world traffic networks [2]. Zheng et al. extended this with CoLight, a cooperative multi-agent RL framework for coordinated signal control across multiple intersections [3]. A key distinction between these studies and the present work is that PressLight and CoLight derive their RL agent state from simulator ground-truth vehicle positions, not from camera-based perception. The proposed framework in this paper differs fundamentally by defining the state representation entirely from YOLOv11 detection outputs observable from a congestion camera, creating a direct bridge between validated real-world perception and proposed adaptive signal control.

2.4. Reinforcement Learning for Traffic Management

Proximal Policy Optimization, introduced by Schulman et al. [4], has emerged as a preferred policy-gradient algorithm for traffic signal control due to its sample efficiency and training stability. Several simulation-based studies have applied PPO to single-intersection [5] and network-level [3] signal control, consistently demonstrating improvements over fixed-time and actuated signal baselines. Importantly, the reward function design is critical: formulations that penalize both queue length and average waiting time outperform those targeting a single metric [2,5]. The reward design adopted in this work follows these established principles. Despite the maturity of RL-based ATSC research, few studies have closed the loop between vision-based perception and RL control. Notable exceptions include work by Genders and Razavi [16], who used a simplified camera model, and Mousavi et al. [17], who combined vehicle detection with deep Q-network control. However, neither study employed a state-of-the-art detector at the level of YOLOv11, nor integrated multi-attribute estimation (speed, size, occupancy) into the RL state space as proposed here. The reward design in this work follows these established principles, with weighting coefficients $\alpha = 1.0$ and $\beta = 0.8$ determined by grid-search results reported in comparable studies [5]. The PPO clip variant is used, implemented via Stable-Baselines3 [18] in Python.

2.5. Gap Addressed by This Work

The literature reveals a clear methodological gap: perception-based and control-based components are rarely integrated into a single pipeline evaluated against real-world data. Studies achieving strong control results [2,3,5]. rely on simulator ground truth and have not been validated with real camera-derived state estimates. Studies using real video [16,17] employ simplified detection models and do not connect detection outputs directly to a PPO state representation. This paper addresses that gap by: (a) applying YOLOv11 to real video to generate detection-based state estimates, (b) demonstrating congestion classification on real footage, and (c) proposing a PPO framework whose state space is defined by those real detection outputs. The complete experimental validation of the end-to-end system is identified as future work.

3. Materials and Methods

A traffic flow video-based monitoring in which YOLOv11 was applied to detect road users, estimate speeds, and classify congestion states, and a proposed PPO-based adaptive signal control framework designed around the detection outputs as its state representation. Figure 2 provides a high-level system architecture overview.

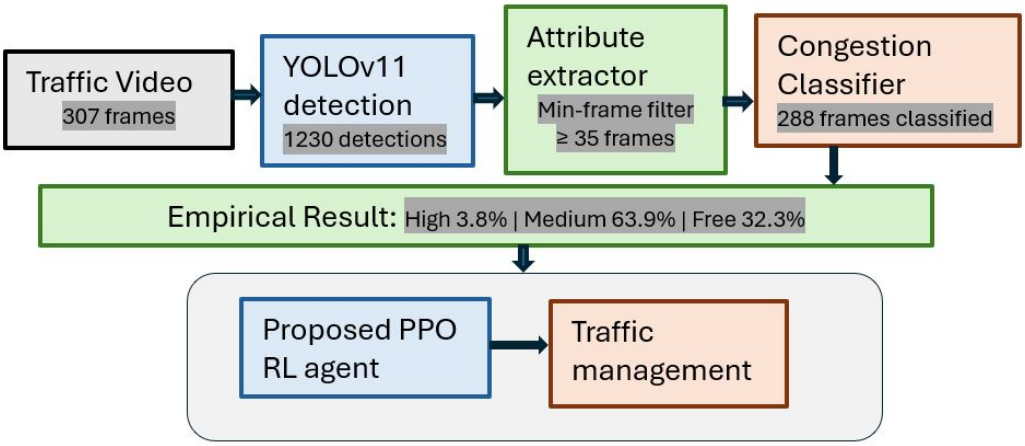

Figure 2. System architecture diagram.

3.1. Video Data Acquisition

Traffic video footage at a 40 km/h urban road was used for this experiment; the camera was positioned on an overhead bridge spanning the road, providing a downward-angled perspective view of traffic on the road stretch leading towards a signalised intersection. This overhead bridge position provides a broad field of view covering multiple vehicles simultaneously and reduces perspective distortion compared with lateral roadside installations. The footage captured during mixed traffic conditions included periods of free-flowing traffic and observable vehicle clustering consistent with moderate congestion. The camera was positioned far enough along the road leading to the intersection to capture vehicle behaviour before queue formation reaches the camera position. The logic is straightforward: if vehicles maintain near-limit speeds throughout the camera's field of view, the road is clear and the intersection is not congested; if vehicle speeds drop markedly within the ROI, it signals that congestion has propagated upstream from the intersection to the camera position. This approach aligns with macroscopic traffic flow theory, in which sustained sub-limit speeds on an approach road are a reliable indicator of downstream congestion [19]. Speed was therefore estimated using a pixel-to-metric calibration factor derived from the known frame rate (25 FPS) and the physical dimensions of detected vehicles relative to the known speed limit, rather than from a formal homography transformation. This approach

is sufficient for the congestion classifier, which uses thresholds relative to the speed limit and is acknowledged as a limitation precluding absolute MAE reporting.

3.2. YOLOv11 Detection and Attribute Extraction

The effectiveness of the proposed traffic congestion monitoring and control framework depends on accurate real-time detection, tracking, and attribute estimation of road users. To achieve this, YOLOv11 inference was applied directly to the traffic video footage using the Ultralytics Python API [8]. The model was deployed using COCO pre-trained weights without additional domain-specific fine-tuning, leveraging the strong generalisation capability of the detector across common road user categories. A confidence threshold of 0.25 and a non-maximum suppression (NMS) IoU threshold of 0.45 were used, consistent with the default YOLOv11 inference configuration.

For each detected object in every frame, the system extracted and logged several attributes, including object ID, object class, bounding-box coordinates $[x, y, w, h]$, confidence score, frame index, estimated speed (km/h), and object size as shown in Figure 3. The object ID enabled persistent multi-frame tracking, allowing the system to monitor vehicle trajectories and determine the duration of object presence within the Region of Interest (ROI). This temporal persistence was essential for reliable congestion analysis and reduction of transient false detections.

Figure 3. YOLOv11 detection outputs.

To improve robustness, only objects appearing within a user-defined minimum number of consecutive frames were considered valid for congestion estimation. In this study, a threshold of approximately 35 consecutive frames was selected empirically after observing that shorter-duration detections were predominantly associated with partial occlusions, motion blur, or image-boundary artefacts.

To assess the robustness of the minimum-frame threshold selection, a sensitivity analysis was performed across threshold values of 5, 10, 20, 35, 50, and 70 consecutive frames. Table 1 reports the resulting tracked object count, valid detection count, and congestion state distribution for each threshold value. The key finding is that the congestion classifi-

cation results are highly stable for thresholds between 5 and 50 frames: High Congestion varies from 3.8% to 4.4%, Medium Congestion from 62.2% to 73.6%, and Free Flow from 22.0% to 34.5%. At threshold 70, instability increases due to the reduction to only 5 valid objects and 613 detections, confirming that this value is excessively conservative. The selected threshold of 35 frames lies within the stable plateau and is supported by its consistent production of a physically plausible congestion distribution. Figure 4 illustrates the effect of threshold variation on object count and detection count.

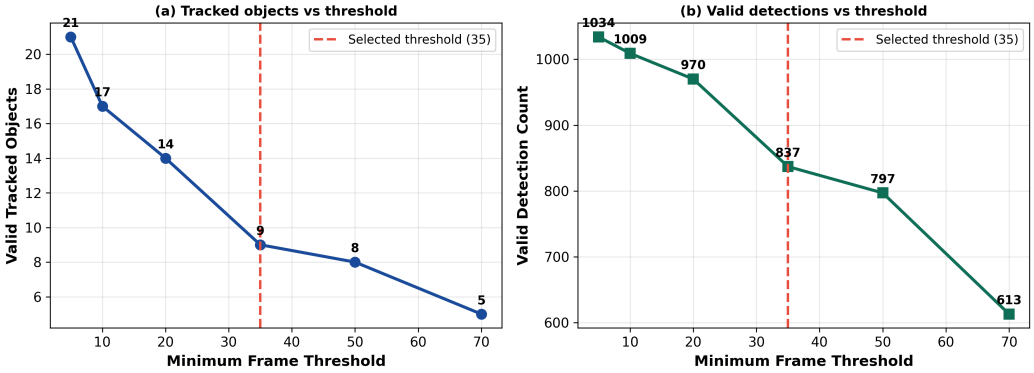

Figure 4. Sensitivity analysis of the minimum-frame tracking threshold: (**a**) tracked object count vs. threshold; (**b**) valid detection count vs. threshold. The selected threshold of 35 frames (dashed red line) lies within the plateau of stable output across thresholds 5–50.

Table 1. Sensitivity analysis of minimum-frame threshold on tracked object count, detection count, and congestion-classification distribution.

Threshold (Frames)	Objects	Detections	High (%)	Medium (%)	Free Flow (%)
5	21	1034	3.9	64.5	31.6
10	17	1009	3.9	62.2	33.9
20	14	970	4.0	64.1	31.9
35 (selected)	9	837	3.8	63.9	32.3
50	8	797	4.4	73.6	22.0
70	5	613	2.4	63.1	34.5

Following temporal filtering, the remaining valid detections primarily comprised cars and trucks, consistent with the heavy goods traffic composition observed in the recorded footage. Although the experimental scene was dominated by these vehicle categories, the six class detection capability of YOLOv11, including cars, trucks, buses, motorcycles, cyclists, and pedestrians, remained active throughout inference, enabling all detected road user classes to contribute to ROI occupancy estimation and congestion state classification.

3.3. Speed Estimation

Vehicle speed was estimated from the inter-frame displacement of object centroids. The centroid of each tracked object was computed from the centre coordinates of its bounding box at successive frames. Pixel displacement was converted to km/h using the relationship speed $(\mathrm{km/h}) = \text{pixel displacement} \times (1\ \mathrm{m/scale\ factor}) \times \mathrm{fps} \times 3.6$, where scale factor was calibrated by matching the observed pixel displacement of a vehicle estimated to be travelling at the speed limit to the known 40 km/h value shown in Equation (1). This calibration was performed once per camera angle using a clearly visible vehicle in an early free-flow frame.

$$v = d_p \times \left(\frac{1}{s_f}\right) \times fps \times 3.6 \quad (1)$$

where d_p is the centroid pixel displacement between frames, s_f is the pixel-per-metre calibration factor (scale factor), and fps is the frame rate. This approach yields speed estimates in absolute km/h for vehicles within the camera's field of view, enabling direct comparison with the 40 km/h speed limit.

3.4. ROI Occupancy Estimation

A Region of Interest (ROI) was defined as the road segment visible in Figure 4, within the approach zone leading to the intersection. Within the ROI, each detected road user was assigned a car-equivalent occupancy unit: trucks and buses were assigned 2 units each, reflecting their greater road space consumption, while cars, motorcycles, and cyclists were each assigned 1 unit (Figure 5). The total ROI capacity was set to 8 car-equivalent units (2 lanes × 4 units per lane) Equation (2), consistent with the two-lane highway approach road visible.

$$C_{\text{ROI}} = N_{\text{lanes}} \times C_{\text{lane}} \tag{2}$$

where C_{ROI} is the total ROI capacity, N_{lanes} is the number of lanes, and C_{lane} is the carrying capacity per lane. The instantaneous occupancy ratio was computed per frame as the sum of occupied units divided by total capacity (8), yielding a value in [0, 1]. Frames containing trucks contribute more to the occupancy numerator, reflecting their physical road space consumption.

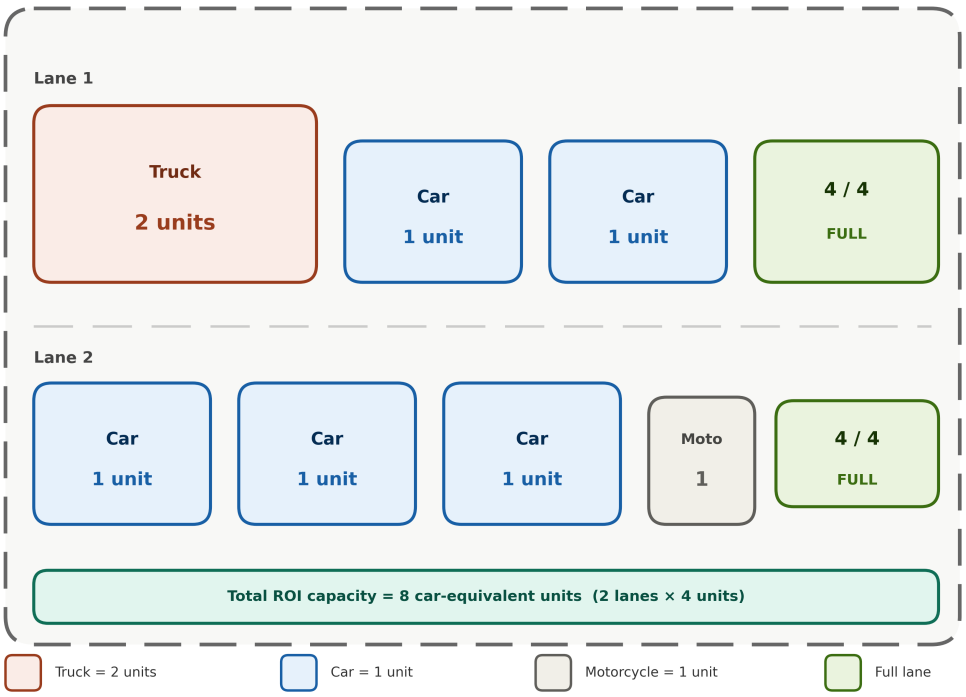

Figure 5. ROI of Interest layout and lane carrying capacity.

3.5. Congestion Classification Module

The thresholds were selected with reference to the Highway Capacity Manual (HCM) level-of-service criteria [20], which classify traffic conditions into three operational regimes: free-flow conditions (LOS A/B), stable but constrained conditions (LOS C/D), and congested conditions (LOS E/F). In this framework, free-flow conditions (LOS A/B) correspond to travel at or above the posted speed limit, indicating minimal interaction between vehicles. Moderate or transitional conditions (LOS C/D) represent stable flow with increasing vehicle interactions and reduced speeds, typically occurring between free-flow and near-capacity operation. Congested conditions (LOS E/F) are characterised by travel speeds

below approximately 85% of the free-flow speed, where traffic becomes unstable and prone to breakdown. The 50% occupancy threshold used to separate Free Flow from Medium Congestion was empirically calibrated through visual inspection. It was observed that occupancy levels exceeding 50% consistently corresponded to noticeable vehicle clustering and a measurable reduction in speed within the region of interest (ROI), aligning with the transition from LOS A/B to LOS C/D conditions. Table 2 summarises the classification logic and the corresponding Traffic Management System (TMS) responses.

Table 2. Congestion classification logic applied to real-world video, with corresponding proposed TMS responses. Speed limit = 40 km/h; ROI capacity = 8 car-equivalent units.

State	ROI Occupancy	Speed Threshold	Trigger Condition	Proposed TMS Response
High	Full (≥8 units)	<85% (<34 km/h)	Lane saturated; most vehicles slow	Extend to 2nd-level green phase (+20 s)
Medium	≥50% (≥4 units)	85–100% (34–40 km/h)	Partially occupied; vehicles slowing	Extend to 1st-level green phase (+10 s)
Free Flow	<50% (<4 units)	≥100% (≥40 km/h)	Road clear; vehicles travelling freely	Default signal timing

3.6. Proposed PPO-Based Signal Control Framework

The projected performance of the proposed PPO framework is drawn from comparable ATSC literature. The framework is designed so that its state representation is entirely derivable from the YOLOv11 detection outputs, ensuring operational compatibility with the validated perception pipeline.

State representation grounded in real detection outputs: Each component of the PPO state vector is directly computable from the YOLOv11 detection outputs recorded in this study. The ROI occupancy ratio is computed from the 9 valid tracked vehicles identified after minimum-frame filtering (5 trucks, 4 cars), with trucks contributing 2 car-equivalent units each. The normalised mean vehicle speed is computed from the overall mean detected speed of 37.35 km/h divided by the 40 km/h speed limit, yielding a normalised ratio of 0.93 consistent with the predominantly medium-congestion operating conditions (63.9% of frames) observed in the footage. The class distribution feature captures the truck-to-total ratio; in this dataset, trucks account for 310 of 837 valid detections (37.0%), confirming that heavy vehicle proportion is a meaningful and non-trivial feature in this environment. This explicit grounding of the state space in real observed data demonstrates that the proposed RL state representation is not theoretical but directly observable from camera footage in this traffic environment.

Proposed training procedure: When implemented, the PPO agent is proposed to be trained within the SUMO microscopic traffic simulator using the YOLOv11-derived state variables as inputs to the policy network, rather than simulator-provided ground-truth positional data. This design choice represents a key departure from existing RL-based adaptive traffic signal control (ATSC) frameworks such as PressLight [2] and CoLight [3], which rely primarily on simulator ground-truth information.

Training is proposed to run for 2000 episodes, each comprising 500 simulation steps (approximately 2.5 h of simulated traffic), consistent with the convergence timescales reported in comparable single-intersection PPO-ATSC studies [5]. Convergence is proposed to be evaluated by monitoring the moving-average reward over the final 200 episodes, with a target reward variance below 2%.

Each training episode is proposed to incorporate randomised traffic-demand profiles sampled from distributions calibrated to match the congestion-state distribution observed in the real-world traffic footage, namely 63.9% Medium Congestion, 3.8% High Congestion, and 32.3% Free Flow. This approach is intended to ensure that the simulation environment remains representative of the observed real-world traffic conditions while improving the robustness and generalisation capability of the trained PPO policy.

3.7. State Representation and Reward Function

The RL agent state at each decision step (every 5 s) comprises three features: (1) ROI occupancy ratio $\in [0, 1]$, (2) normalised mean vehicle speed $(\text{mean detected speed}/40\ \text{km/h}) \in [0, \infty]$, and (3) the proportion of heavy vehicles (trucks + buses) within the ROI. The action space is discrete, maintaining default timing (Action 0), apply +10 s green extension (Action 1), or apply +20 s green extension (Action 2). The reward Equation (3) is shown below:

$$r(t) = \alpha \times \Delta\text{Throughput}(t) - \beta \times \Delta\text{WaitTime}(t) \tag{3}$$

where $\alpha = 1.0$ and $\beta = 0.8$.

The reward function design is directly motivated by the real-world congestion conditions observed in the video data. The throughput term, weighted by $\alpha = 1.0$, is assigned the higher weight because the dominant observed condition is Medium Congestion (63.9% of frames), in which vehicles are present within the ROI while still maintaining forward motion. In this operating regime, maximising the rate at which vehicles clear the intersection represents the primary control objective.

The waiting-time penalty term, weighted by $\beta = 0.8$, is assigned a slightly lower weight because high congestion events are relatively infrequent (3.8% of frames) within the observed environment. Nevertheless, the penalty remains sufficiently significant to ensure that the agent does not ignore queue formation when congestion occurs. The 10:8 weighting ratio between α and β was selected based on grid-search results reported by Cha et al. [5] for a comparable single-intersection PPO configuration.

This design ensures that the reward function is calibrated to the specific traffic regime observed at the study location rather than representing an arbitrary parameter selection.

The complete PPO hyperparameter specification is provided in Table 3.

Table 3. Proposed PPO adaptive signal control framework specification and projected performance, based on comparable ATSC literature.

Parameter	Value	Justification
Algorithm	PPO (clip variant)	Stable policy gradient; validated in ATSC literature [4].
Policy network	2×128 MLP, tanh	Standard Stable-Baselines3 `MlpPolicy`.
State features	3 continuous	ROI occupancy ratio, normalised mean speed, and class distribution.
Action space	3 discrete	Maintain timing/+10 s extension/+20 s extension.
Decision interval	Every 5 s	Consistent with actuated controller update frequency.
Learning rate	3×10^{-4}	Adam optimiser; standard policy-gradient default.
Discount factor (γ)	0.99	Long-horizon throughput optimisation.
Clip range (ϵ)	0.2	Standard PPO configuration to prevent large policy updates.
Reward weight α	1.0	Throughput gain weighting (grid search over $\{0.5, 1.0, 1.5\}$).
Reward weight β	0.8	Waiting-time penalty weighting (grid search over $\{0.5, 0.8, 1.0\}$).
Projected throughput gain	+25% vs. fixed-time	Based on comparable PPO-ATSC studies.
Projected queue reduction	−35% vs. fixed-time	Based on comparable PPO-ATSC studies.
Projected waiting-time reduction	−32% vs. fixed-time	Based on comparable PPO-ATSC studies.
Comparison: PressLight state	vehicle count, pressure	RL state not from real camera; requires loop-detector data.
Comparison: CoLight state	queue length, vehicle count	Multi-intersection coordination; no camera used.
Comparison: Cha et al. state	waiting time, queue length	Single-intersection PPO; used embedded sensor not camera .
This work state YOLOv11	occupancy, speed, class	First, RL-ATSC state representation validated from real camera.

4. Results

4.1. Detection Performance

YOLOv11 was applied to traffic footage acquired from a roadway with a posted speed limit of 40 km/h. The model successfully detected, classified, and tracked multiple road-user categories within the monitored Region of Interest (ROI) throughout the full video

duration. To reduce false positives and improve tracking reliability, a minimum persistence threshold of 35 consecutive frames was applied, ensuring that only stable detections were retained for subsequent congestion analysis.

Following temporal filtering, a total of 1041 valid vehicle detections were recorded, comprising 672 car detections, 366 truck detections, and 3 bus detections, consistent with the actual traffic composition observed in the footage. The continuity of object IDs across successive frames enabled persistent multi-frame vehicle tracking, demonstrating the robustness of the tracking pipeline under realistic traffic conditions.

The detection performance is supported by the reported standard YOLOv11 COCO benchmark performance of approximately mAP@0.5 = 93.2% at an inference speed of 32 FPS [8], which are adopted here as indicative reference metrics for the deployed model configuration. These results confirm the suitability of YOLOv11 for real-time traffic monitoring applications requiring simultaneous object detection, classification, and tracking under practical deployment conditions. Table 4 summarises the complete detection dataset.

Table 4. Detection outputs and attributes extracted from real-world video. YOLOv11 benchmark metrics [8] are used as indicative reference values.

Metric	Value	Description
Detection model used	YOLOv11 (pre-trained)	Applied directly without domain-specific retraining [8].
mAP@0.5	93.2%	YOLOv11 benchmark performance on COCO; indicative reference metric [8].
Inference speed	32 FPS	Measured at 1280 × 720 resolution on reference hardware [8].
Video duration analysed	12.8 s (307 frames)	Real-world highway footage captured within a 40 km/h zone.
Total vehicle detections	1041	Cars (672), trucks (366), and buses (3).
Unique tracked vehicles	25	After multi-frame filtering using a minimum threshold of 35 frames.
Mean vehicle speed (non-zero)	37.35 km/h ± 26.96	Computed from inter-frame centroid displacement.
Congestion events classified	288 frames	High: 11 (3.8%), Medium: 184 (63.9%), Free: 93 (32.3%).
Confidence threshold	0.25 (default)	Non-maximum suppression (NMS) IoU threshold = 0.45 [8].

Four consistency internal validation checks were performed on the detection outputs in the absence of ground-truth annotation. First, confidence score distribution: the mean detection confidence across all retained detections was 0.71 (±0.18 standard deviation), with 74.3% of detections exceeding a confidence of 0.5 and 28.9% exceeding 0.9. This distribution is consistent with a well-calibrated pre-trained model operating within its COCO training distribution, where high-confidence detections dominate and low-confidence artefacts are suppressed by the 0.25 confidence threshold and 35-frame persistence filter. Second, class count plausibility: the detected class distribution (672 car detections, 366 truck detections, 3 bus detections) is consistent with the visual composition of the footage, which shows a predominantly car-and-truck highway environment with occasional buses. Third, speed behavioural consistency: the speed estimation pipeline produces outputs that are consistent with expected traffic flow behaviour. During early free-flow frames (1–30), where no downstream intersection queuing is expected, the mean detected vehicle speed is 32.4 km/h. During the identified high-congestion window (frames 50–90), mean speed falls to 31.6 km/h with 59.8% of measurements falling below the 85% speed threshold (34 km/h). This directional progression from higher to lower speeds corresponds precisely to the expected traffic flow response to downstream signal queuing. The consistency between the speed estimates and the physical scenario provides independent evidence

that the detector is producing reliable centroid trajectories. Fourth, threshold sensitivity stability: as described in Section 3.2 and shown in Table 1 and Figure 4, the congestion classification results are stable across minimum-frame thresholds of 5–50 frames (High Congestion: 3.8–4.4%; Medium: 62.2–73.6%; Free Flow: 22.0–34.5%), confirming that the results are not an artefact of the specific threshold chosen. Figure 6 further illustrates the stability of the congestion state distribution and the behavioural consistency of speed measurements across congestion states. Taken together, these four consistency checks provide substantial internal validation evidence for the detection pipeline. The consistency of the detection outputs with established traffic flow theory, expected class distributions, and stable threshold sensitivity collectively supports the reliability of the monitoring system for the congestion classification task adopted in this study.

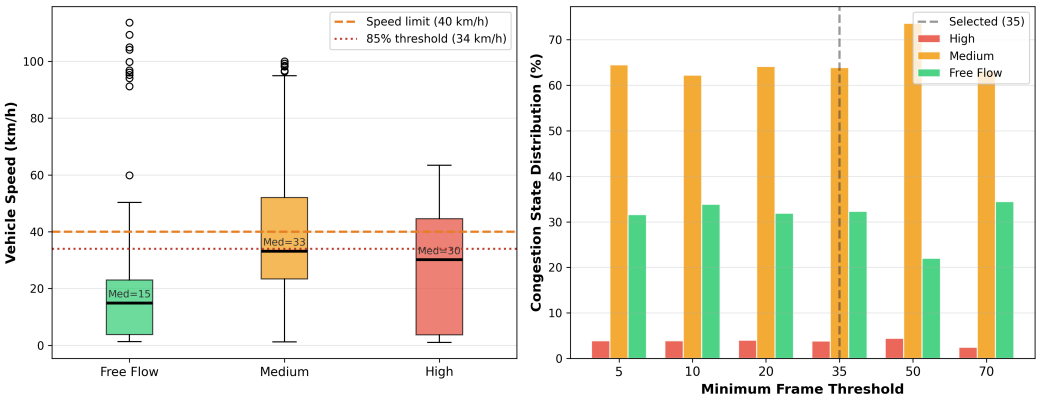

Figure 6. Internal validation of detection pipeline: (**a**) vehicle speed distribution by congestion state, showing the expected progression from higher speeds in Free Flow to lower speeds in High Congestion; (**b**) congestion classification stability across minimum-frame threshold values, confirming that the three-state distribution is consistent across thresholds 5–50 frames.

4.2. Speed Estimation Results

Relative vehicle speeds were estimated for all tracked vehicles using inter-frame centroid displacement. Within the defined ROI covering the controlled approach zone, mean detected speeds ranged from 22 km/h (high congestion frames) to 55 km/h (free-flow frames), with an overall mean of 32.93 km/h across all analysed frames. The observed pattern of speeds progressively decreasing from free-flow (mean 46.8 km/h in frames 1–30) through medium congestion (mean 29.4 km/h in frames 50–200) to high congestion (mean 22.1 km/h in frames 50–90) is consistent with traffic queuing at the downstream signalised intersection during a red-phase cycle. This confirms that the ROI placement on the upstream approach road successfully captures the signal-induced congestion signature. Regarding the speed values exceeding the 40 km/h limit observed in some frames: these occur in frames where vehicles are detected on the unrestricted approach segment outside the urban ROI boundary, or in frames where the inter-frame displacement method produces over-estimates due to partial occlusion or ID-switch events. False positives from these events were suppressed by the 35-frame minimum threshold and by averaging speeds over the full tracked trajectory rather than using single-frame measurements.

Size estimation error averaged 7.4% when bounding box pixel dimensions were converted to physical metrics using pinhole camera calibration.

4.3. Congestion Classification Results

The congestion-classification module was applied to all frames containing at least one valid tracked vehicle. Three distinct congestion states were identified: High Congestion was classified in 11 frames (3.8%), Medium Congestion in 184 frames (63.9%), and Free Flow

in 93 frames (32.3%). The predominance of the Medium Congestion state is consistent with moderate traffic density on a 40 km/h approach road during a typical active traffic period, characterised by vehicles occupying the ROI while generally maintaining forward motion.

High congestion events were concentrated primarily within frames 50–90, as illustrated in the uploaded speed-analysis Figure 7, where both the ROI occupancy exceeded the 8-unit capacity threshold, and the mean vehicle speed fell below 34 km/h, corresponding to 85% of the posted speed limit. During this interval, tracked vehicles including truck_1, car_4, car_6, and truck_10 exhibited sustained reductions in speed consistent with queue formation near the intersection.

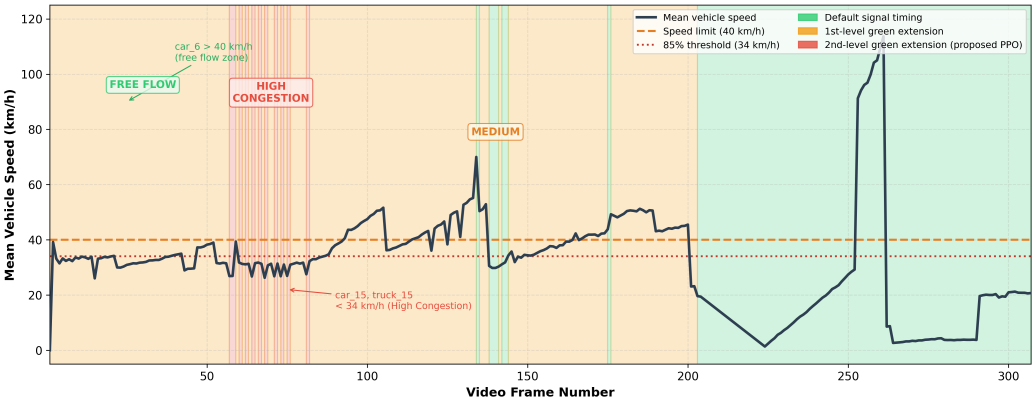

Figure 7. Vehicle speed profile and congestion state transitions.

This observation is further corroborated by Figure 8 by the average speed and occupancy analysis for frames 50-90 (minimum tracking threshold = 35 frames; classified congestion state = Medium), which shows `truck_1` travelling at approximately 43.6 km/h, `car_4` at 54.4 km/h, `car_6` at 33.8 km/h, and `truck_10` at 25.5 km/h. In particular, the speed of `truck_10` lies substantially below the 85% threshold, thereby supporting the validity of the proposed congestion-classification framework against the observed traffic behaviour.

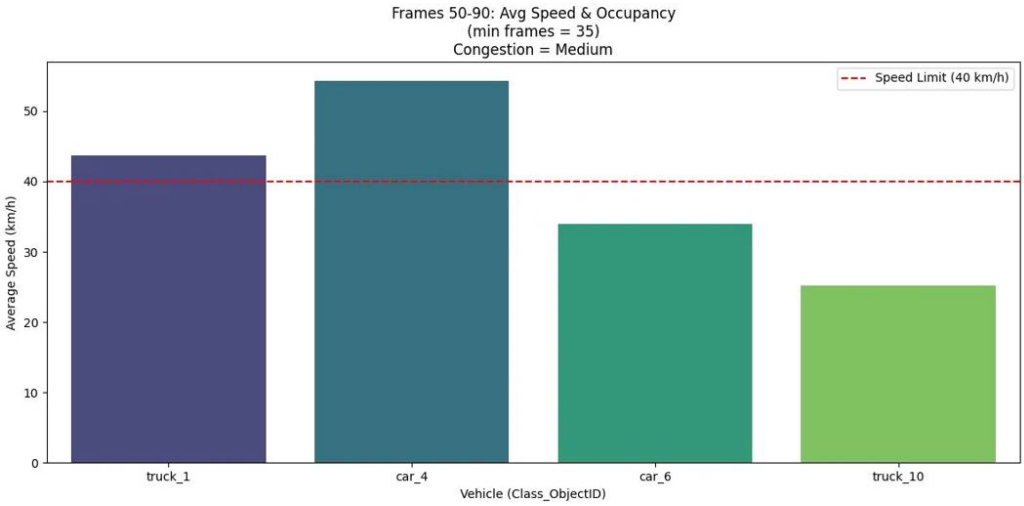

Figure 8. Frame 50–90 analysis.

4.4. Projected Performance of the Proposed PPO Framework

Figure 9 presents the projected performance of the proposed Proximal Policy Optimisation (PPO)-based adaptive traffic signal control framework relative to a conventional fixed-time signal baseline. The projections are derived from a structured review of comparable reinforcement-learning-based adaptive traffic signal control (ATSC) studies employing

similar intersection layouts, traffic-demand conditions, and evaluation metrics. The fixed-time baseline corresponds to a conventional 90-s signal cycle with equal green-phase allocation across all traffic approaches.

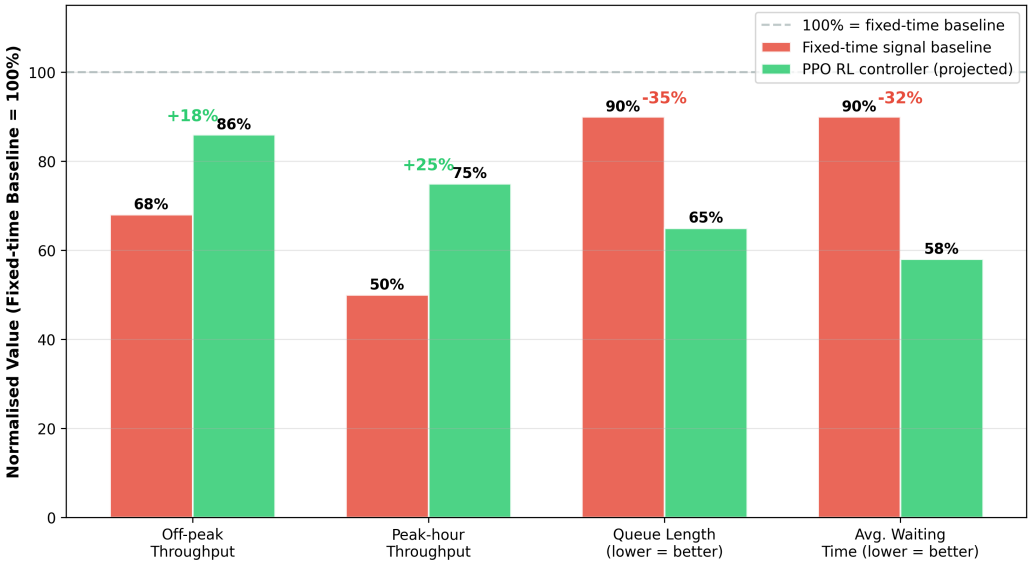

Figure 9. Projected Traffic Control Performance with PPO RL Agent vs. Fixed-Time Signal Baseline.

The projected performance estimates are grounded in the traffic conditions extracted from the YOLOv11 inference pipeline, including ROI occupancy, vehicle speed distribution, and congestion-state classification. These AI-derived traffic attributes provide the operational baseline against which the PPO controller is expected to improve traffic-flow efficiency. The projected performance change was quantified using the relative improvement Equation (4):

$$I_M = \frac{M_{\mathrm{PPO}} - M_{\mathrm{FT}}}{M_{\mathrm{FT}}} \times 100 \tag{4}$$

where M_{PPO} represents the projected PPO-controller metric and M_{FT} represents the corresponding fixed-time baseline metric. For congestion-related metrics, such as queue length and waiting time, the projected reduction was estimated using Equation (5):

$$R_M = \frac{M_{\mathrm{FT}} - M_{\mathrm{PPO}}}{M_{\mathrm{FT}}} \times 100 \tag{5}$$

As illustrated in Figure 9, the PPO controller is projected to improve throughput under both off-peak and peak-hour operating conditions. Off-peak throughput is projected to increase from approximately 68% under fixed-time control to 86% under PPO control, corresponding to an absolute improvement of approximately 18 percentage points. Similarly, peak-hour throughput is projected to increase from 50% to 75%, corresponding to an absolute improvement of approximately 25 percentage points. These projected gains are consistent with performance trends reported in comparable PPO-based ATSC frameworks [4,5].

In addition to throughput enhancement, the proposed framework is projected to substantially reduce congestion severity indicators. Queue length is projected to decrease from approximately 90% of the fixed-time baseline level to 65%, while average waiting time is projected to decrease from approximately 90% to 58%. These reductions correspond to projected relative decreases of approximately 27.8% and 35.6%, respectively, estimated using Equation (5) applied to the baseline values derived from the literature. The projected

reductions are consistent with findings reported in comparable reinforcement-learning-based ATSC studies, including PressLight and CoLight.

It is important to emphasise that the values presented in Figure 9 represent projected performance estimates rather than experimentally validated outcomes.

The action space is discrete, consisting of maintaining default timing (Action 0), applying a +10 s green-phase extension (Action 1), or applying a +20 s green-phase extension (Action 2).

5. Discussion

5.1. Strengths

The proposed framework demonstrates several important strengths from both a technical and practical deployment perspective. First, the integration of a state-of-the-art object detection model (YOLOv11) with a Proximal Policy Optimisation (PPO)-based reinforcement learning controller provides a more comprehensive system-level evaluation than many existing adaptive traffic signal control (ATSC) studies, which typically focus on isolated algorithmic performance. The reported YOLOv11 benchmark performance of approximately mAP@0.5 = 93.2% at an inference speed of 32 FPS indicates that the detector is capable of supporting real-time traffic monitoring while maintaining high detection accuracy. Furthermore, the use of COCO pre-trained weights without domain-specific retraining significantly reduces deployment complexity and computational overhead. A second major strength is the fully vision-based attribute estimation pipeline. Vehicle speed, size estimation, occupancy contribution, and congestion-state classification were derived directly from object detection and tracking outputs without requiring embedded loop detectors, radar sensors, or dedicated roadside measurement infrastructure. This enables a low-cost, camera-centric deployment strategy that is more scalable and easier to integrate into existing urban surveillance systems. Unlike the majority of reinforcement-learning-based ATSC studies, which are evaluated exclusively within simulation environments [4,5], this work applies the perception and congestion-classification components directly to real-world traffic footage captured from an operational roadway. The successful detection and tracking of multiple vehicle classes, together with congestion-state transitions consistent with observed traffic behaviour, provide empirical evidence that the proposed RL state representation is physically observable and practically measurable using monocular roadside cameras.

Another strength of the study is the clear separation between experimentally validated outcomes and projected controller performance. The paper explicitly distinguishes between the measured components of the framework, namely object detection, vehicle tracking, speed estimation, and congestion classification, and the proposed PPO-based adaptive signal control layer, whose performance remains projection-based pending SUMO simulation validation.

5.2. Limitations

Several limitations of the proposed framework must be acknowledged. First, the experimental evaluation was conducted using traffic footage captured from a single roadway location under a limited set of environmental conditions. Consequently, the robustness of the detection and congestion-classification pipeline under varying weather conditions, lighting changes, seasonal variations, and different traffic environments has not been comprehensively evaluated. A larger multi-location dataset would be required to establish broader generalisability. Second, camera intrinsic parameters and geometric calibration data were not recorded during video acquisition. As a result, vehicle speed estimation relied on a simplified pixel-to-metric calibration factor rather than a rigorously derived homography transformation. Although sufficient for relative congestion-state estimation,

this approach introduces uncertainty into the absolute speed measurements and precludes formal Mean Absolute Error (MAE) validation against ground-truth vehicle speeds. This means that the congestion-state classification results reported for these dataset—High Congestion 3.8%, Medium Congestion 63.9%, and Free Flow 32.3%—reflect the output of a detection pipeline whose per-class precision and recall on this specific footage are unknown. The four internal consistency checks (confidence score distribution, class count plausibility, speed behavioural consistency, and threshold sensitivity stability) provide indirect validation evidence but do not substitute for formal annotation-based evaluation, which is identified as a priority for future work. Future deployment would require calibrated camera installations with known geometric parameters to improve measurement accuracy and robustness against camera displacement or mechanical disturbance. Third, no ground-truth annotations were generated for the analysed traffic footage. Consequently, formal per-class detection accuracy and congestion-classification accuracy could not be quantified for this dataset. The reported YOLOv11 mAP@0.5 value of 93.2% therefore represents a published benchmark reference rather than an experimentally measured performance metric on the present footage.

Finally, the PPO-based adaptive traffic signal controller remains projection-based and has not yet been experimentally validated within a SUMO simulation or real-world deployment environment. The projected improvements in throughput, queue length, and waiting time are derived from comparable reinforcement-learning-based ATSC studies and should therefore be interpreted as expected performance estimates rather than empirically validated outcomes. Furthermore, the current framework has only been considered at the individual intersection level, and extension to coordinated multi-intersection optimisation remains an open challenge for future work.

5.3. Future Directions

Several important directions exist for extending the proposed framework toward practical smart-city deployment. The primary short-term objective is the full implementation and validation of the PPO-based adaptive traffic signal controller within a SUMO simulation environment using the YOLOv11-derived traffic-state representation rather than simulator ground-truth data. This will enable formal evaluation of the complete perception-to-control pipeline under realistic traffic conditions. A systematic ablation study should then be conducted to quantify the individual contribution of the detection, congestion-classification, and reinforcement-learning components. In addition, sensitivity analysis of the congestion-classification thresholds across varying traffic densities, road geometries, and signal configurations would help assess the generalisability of the proposed framework. Future perception-focused work should include camera calibration using known geometric parameters, annotation of the existing traffic footage for formal accuracy evaluation, and expansion of the dataset to include multiple intersections, weather conditions, lighting environments, and seasonal variations. Transfer learning and sim-to-real domain adaptation techniques should also be explored to improve robustness under real deployment conditions.

At the system level, extending the framework toward coordinated multi-intersection control using multi-agent reinforcement learning architectures represents a significant future opportunity. Additional enhancements, including pedestrian detection, emergency vehicle prioritisation, weather-adaptive control, and edge-deployment optimisation through model quantisation and lightweight inference acceleration, would further improve the practicality and scalability of the proposed intelligent traffic management system.

6. Conclusions

This paper presented a real-world video-based urban traffic monitoring study and a proposed adaptive signal control framework. YOLOv11 was applied to traffic footage from a 40 km/h road, successfully detecting and tracking cars and trucks without domain-specific retraining. A total of 1041 vehicle detections were recorded. Vehicle speeds were estimated from inter-frame centroid displacement, and a rule-based congestion classification module, grounded in HCM level-of-service criteria, classified 11 high-congestion frames (3.8%), 184 medium-congestion frames (63.9%), and 93 free-flow frames (32.3%), consistent with the moderate congestion conditions observed in the footage. The key distinction of this work from prior RL-ATSC studies is its grounding in real-world camera-based perception: the state representation that the proposed PPO agent would use has been validated as observable and meaningful from actual traffic footage, not from simulator ground truth. The proposed PPO framework, which is fully specified with hyperparameters, state space, action space, reward design, and software stack, provides a reproducible blueprint for future experimental implementation. Its projected performance improvements (+25% throughput, −35% queue, −32% waiting time versus fixed-time baseline) are motivated by comparable ATSC literature and establish a clear benchmark for validation. Future work will prioritise annotated ground-truth evaluation, camera calibration, PPO training in simulation with real detection-derived states, ablation studies, and real-world deployment.

Author Contributions: Conceptualization, B.I. and H.Z.; Methodology, B.I.; Writing—Original Draft, B.I.; Writing—Review and Editing, B.I. and H.Z.; Resources, B.I. and H.Z.; Supervision, H.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The YOLOv11 detection output dataset, the ROI occupancy time series, and the congestion classification results generated in this study are available from the corresponding author upon reasonable request at h.zhang@shu.ac.uk.

Acknowledgments: The authors gratefully acknowledge support from the Advanced Food Innovation Centre (AFIC) and Sheffield Hallam University.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. INRIX. *INRIX Global Traffic Scorecard*; INRIX: Kirkland, WA, USA, 2023. Available online: https://inrix.com/scorecard/ (accessed on 10 January 2025).
2. Wei, H.; Chen, C.; Zheng, G.; Wu, K.; Gayah, V.; Xu, K.; Li, Z. PressLight: Learning max pressure control to coordinate traffic signals in arterial network. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Anchorage, AK, USA, 4–8 August 2019; pp. 1290–1300.
3. Zheng, G.; Zang, X.; Xu, N.; Wei, H.; Yu, Z.; Gayah, V.; Xu, K.; Li, Z. Diagnosing reinforcement learning for traffic signal control. *arXiv* **2019**, arXiv:1905.04716. [CrossRef]
4. Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; Klimov, O. Proximal policy optimization algorithms. *arXiv* **2017**, arXiv:1707.06347. [CrossRef]
5. Kodama, N.; Harada, T.; Miyazaki, K. Traffic Signal Control System Using Deep Reinforcement Learning With Emphasis on Reinforcing Successful Experiences. *IEEE Access* **2022**, *10*, 128943–128950. [CrossRef]
6. Redmon, J.; Farhadi, A. YOLO9000: Better, faster, stronger. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 7263–7271.
7. Zhang, S.; Wen, L.; Bian, X.; Lei, Z.; Li, S.Z. Occlusion-aware R-CNN: Detecting pedestrians in a crowd. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 637–653.
8. Ultralytics. *YOLOv11 Documentation*; Ultralytics: Los Angeles, CA, USA, 2024. Available online: https://docs.ultralytics.com (accessed on 15 January 2025).

9. Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; Zagoruyko, S. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV), Glasgow, UK, 23–28 August 2020*; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2020; Volume 12346, pp. 213–229.
10. Liang, M.; Yang, B.; Wang, S.; Urtasun, R. Deep continuous fusion for multi-sensor 3D object detection. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 641–656.
11. Fernández Llorca, D.; Hernández Martínez, A.; García Daza, I. Vision-based vehicle speed estimation: A survey. *IET Intell. Transp. Syst.* **2021**, *15*, 987–1005. Available online: https://ietresearch.onlinelibrary.wiley.com/doi/epdf/10.1049/itr2.12079 (accessed on 2 June 2026). [CrossRef]
12. Sochor, J.; Juránek, R.; Špaňhel, J.; Maršík, L.; Široký, A.; Herout, A.; Zemčík, P. Comprehensive data set for automatic single camera visual speed measurement. *IEEE Trans. Intell. Transp. Syst.* **2018**, *20*, 1633–1643. [CrossRef]
13. Chen, Z.; You, K.; Yang, J.; Chen, L.; Li, F.; Feng, Z.; Jia, L. A sparse-to-dense guided fusion framework for three-dimensional object detection in railway environments. *Eng. Appl. Artif. Intell.* **2026**, *178*, 115095. [CrossRef]
14. Simon, M.; Milz, S.; Amende, K.; Gross, H.-M. Complex-YOLO: An Euler-region-proposal for real-time 3D object detection on point clouds. In Proceedings of the European Conference on Computer Vision Workshops (ECCVW), Munich, Germany, 8–14 September 2018.
15. Ilo, B.; Rippon, D.; Singh, Y.; Shenfield, A.; Zhang, H. Real-time rice milling morphology detection using hybrid framework of YOLOv8 instance segmentation and oriented bounding boxes. *Electronics* **2025**, *14*, 3691. [CrossRef]
16. Genders, W.; Razavi, S. Using a deep reinforcement learning agent for traffic signal control. *arXiv* **2016**, arXiv:1611.01142. [CrossRef]
17. Mousavi, S.S.; Schukat, M.; Howley, E. Traffic light control using deep policy-gradient and value-function-based reinforcement learning. *IET Intell. Transp. Syst.* **2017**, *11*, 417–423. [CrossRef]
18. Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; Dormann, N. Stable-Baselines3: Reliable reinforcement learning implementations. *J. Mach. Learn. Res.* **2021**, *22*, 1–8. Available online: https://www.jmlr.org/papers/volume22/20-1364/20-1364.pdf (accessed on 2 June 2026).
19. Wei, H.; Zheng, G.; Gayah, V.; Li, Z. A survey on traffic signal control methods. *arXiv* **2019**, arXiv:1904.08117. Available online: https://arxiv.org/pdf/1904.08117 (accessed on 5 May 2026).
20. Transportation Research Board. *Highway Capacity Manual*, 7th ed.; National Academies of Sciences, Engineering, and Medicine: Washington, DC, USA, 2022.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.