Review Article *Open Access*

Journal of Scientific Engineering Advances

ISSN: 3071-0723

DOI: doi.org/10.63721/26JSEA0154

Enhancing Road Safety Infrastructure: An Intelligent Aerial Traffic Observation and Automated Citation System

Tanyaradzwa Tanatswa Mushonga[1]* and Samkeliso Suku Dube[1]

[1]Department of Computer Science Faculty of Applied Sciences, National University of Science and Technology Bulawayo, Zimbabwe

Citation: *Tanyaradzwa Tanatswa Mushonga, Samkeliso Suku Dube (2026) Enhancing Road Safety Infrastructure: An Intelligent Aerial Traffic Observation and Automated Citation System. J. of Sci Eng Advances 2(5) 1-11. WMJ/JSEA-154*

Abstract

Road traffic injuries claim an estimated 1.19 million people per year, a crisis that is mostly dominant in developing countries such as Zimbabwe due to high speeding and limited speed regulation enforcement hardware. Traditional speed regulation methods such as fixed speed cameras and roadblocks suffer from predictability, creating significant coverage gaps. The study introduces the Intelligent Aerial Traffic Observation System (IATOS), a solution to make use of Unmanned Aerial Vehicles (UAVs) and computer vision to automate speed enforcement. The study employed a hybrid methodology, integrating Design Science Research and quantitative experimental analysis; the system utilises edge-cloud hybrid architecture where an ESP32-CAM handles low-power image acquisition and a cloud backend performs heavy processing using You Only Look Once (YOLOv8) for object detection and DeepSORT for multi-object tracking. The system utilises a homography technique, which projects the 2D aerial video onto the road plane to calculate speed with high precision. Experimental results demonstrate a vehicle detection mean Average Precision (mAP) of 0.92 and a speed estimation accuracy of ±2.0 km/h compared to GPS ground truth. The system automatically generates standardised digital evidence packs, containing speed data, GPS location, and visual proof, for every confirmed violation.

***Corresponding author**: Tanyaradzwa Tanatswa Mushonga, Department of Computer Science Faculty of Applied Sciences, National University of Science and Technology Bulawayo, Zimbabwe.

Submitted: 08.08.2026 **Accepted**: 13.08.2026 **Published**: 07.09.2026

Keywords: UAV Traffic Monitoring, Computer Vision, Real-time Speed Detection, Automated Citation, Intelligent Transport Systems

Introduction

Road traffic injuries remain one of the most pressing public health and development challenges of our time. An estimated 1.19 million people lose their lives on the roads every year globally, with the heaviest toll falling on low and middle-income countries [1]. Zimbabwe reflects this global picture; both official statistics and media reports consistently highlight high crash rates, with thousands of injuries and fatalities recorded annually. Speeding and other risky driving behaviours are frequently cited as major contributors [2,3]. Beyond the human tragedy, these incidents impose steep social and economic costs, while enforcement agencies struggle to keep pace with the demands of a rapidly expanding road network. Traditional enforcement methods, such as roadside patrols and fixed cameras, have apparent limitations. They cover only small portions of the network, are predictable to drivers, and are costly to scale up. Many motorists simply reduce speed near known checkpoints or cameras, only to accelerate once they pass that section of the road. Additionally, paper-based processes often slow down enforcement, diverting officers from active monitoring to administrative tasks. The result is uneven coverage and inconsistent deterrence, particularly along highways and peri-urban routes where risks are high. Emerging research suggests that unmanned aerial vehicles (UAVs) equipped with computer vision have the potential to serve as a valuable complement to ground-based systems. UAVs can be rapidly deployed, cover wide sections of the road, and capture views that fixed roadside cameras cannot [4,5]. At the same time, advances in computer vision technology now make it possible to detect and track vehicles in real time, estimate speeds with only small margins of error, and recognise license plates with high accuracy under both daylight and nighttime conditions [6,7]. When integrated into a transparent and privacy-conscious workflow, these tools can generate reliable evidence without requiring dense, costly infrastructure. Of course, practical challenges remain, such as battery life, connectivity, governance, and airspace regulation, but these are increasingly being addressed in the literature [4,5]. Against this backdrop, this paper introduces the Intelligent Aerial Traffic Observation System (IATOS). The system combines UAV-based video capture with computer vision to estimate vehicle speeds in real time, identify license plates, and compile standardised evidence packs for review by enforcement officers. To encourage lasting behavioural change, the proposed architecture is also designed with a positive reinforcement component in mind, a Safe Driver Rewards mechanism whereby complaint drivers could be entered into periodic draws funded through violation revenue. The full design and evaluation of this incentive layer is reserved for future work and falls outside the scope of this feasibility study.

Objectives

1. To design an edge-cloud architecture that enables real-time video streaming from low-cost drone hardware.
2. To develop a computer vision algorithm using YOLOv8 for accurate vehicle detection and speed estimation.
3. To implement Automatic Number Plate Recognition (ANPR) for identifying vehicles from aerial perspectives.
4. To automate the citation process by integrating violation data, GPS evidence, and digital reporting into a unified system.

Literature Review

This review aims to establish the theoretical and empirical foundation for developing the IATOS. The review's primary goal is to synthesise the current state of research concerning UAV traffic surveillance, real-time speed detection and ANPR, edge computing deployment, and automated citation logic. The methodology employed a rigorous search protocol, targeting peer-reviewed literature published between 2018 and 2024 across major academic databases. The search focuses on the intersection of "UAV/Drone," "Speed Detection/Velocity Estimation," and "Automatic Number Plate Recognition/Automated Citation". Inclusion criteria mandated studies involving both aerial platforms and speed measurements/enforcement. Data was extracted qualitatively under predefined headings to build a comprehensive analytical framework.

Study Characteristics and Trends

The collected literature, analysed across multiple research stages, demonstrates clear developmental and methodological trends in aerial traffic monitoring. Geographically, most of the scholarly literature originates from highly urbanised regions in Europe, North America, and Asia focusing heavily on structured highway or established metropolitan traffic environments [4,5,8]. There is a pronounced and critical absence of technical implementation studies validated within the specific operational context of sub-Saharan African traffic, which is characterised by different vehicle mixes, often undisciplined traffic flow, and unique infrastructure constraints. Methodologically, the technical field is moving decisively away from complex, multi-stage processing towards streamlined, end-to-end architectures designed for computationally constrained deployment. Thus, the dominant trend prioritises achieving a delicate and auditable balance between high detection accuracy (essential for legal evidence) and computational speed (necessary for traffic intervention) [9].

Core Findings: Methodological Landscape

The following section presents a synthesis of findings derived from a thematic analysis of the literature, highlighting the major insights and conceptual frameworks identified across the studied sources.

Lightweight and Object Detection for Aerial Data

Aerial imagery presents the critical challenge for detecting small, diverse, and often occluded objects due to the high vantage point. The dominant technical solutions rely on optimised one-stage detectors, primarily advanced variants of the YOLO architecture [10]. To meet strict constraints of onboard edge hardware, standard YOLO architecture is extensively modified, often incorporating lightweight backbone networks like MobileNetV3 or proprietary efficient structures such as GhostNet and Depthwise convolutional components [10,11]. These architectural optimisations are essential for feasibility, as they significantly reduce a model's parameter count while maintaining or enhancing detection accuracy.

Real-Time Edge Computing and Hardware Feasibility

While contemporary literature states that for traffic enforcement, processing must occur at the site of observation (the edge) to ensure low latency and continuous operational outside of reliable network infrastructure, deployments consistently rely on powerful, GPU-enabled embedded systems, specifically the NVIDIA Jetson Xavier NX and Jetson Nano, as the central processing units [12,13]. System performance is critically dependent on hardware acceleration frameworks like TensorRT, which optimises deep learning models to maximise inference time on these resource constrained platforms [12]. However, this high-performance edge architecture presents significant barriers for developing economies, primarily due to high costs and power requirements of GPU-enabled edge devices. An alternative architectural paradigm has emerged for resource constrained environments. In this model, the aerial unit functions as a lightweight "thin client" focusing exclusively on data acquisition and low latency streaming while computationally heavy workloads are offloaded to the server. This approach allows for use of significantly more affordable microcontrollers, such as ESP32-CAM thereby democratising access to intelligent aerial surveillance by shifting the burden from expensive flight hardware to scalable server infrastructure.

Accurate Speed Estimation and Multi-Object Tracking

Accurate vehicle speed detection is achieved through a robust, multi-step pipeline combining detection, tracking, and photogrammetry. Robust multi-object tracking (MOT) is achieved by integrating detection results with predictive state estimators, typically the Kalman Filter or modern trackers like DeepSort, which preserve vehicle identity and trajectory smoothness during operational challenges like occlusion or rapid movement [14,15]. The most critical technique step is the translation from pixel movement to real-world velocity, which relies on homography transformation. This photogrammetric technique rectifies the image perspective onto a modelled ground plane, utilising fixed road landmarks, e.g., lane markings to accurately convert pixel displacement over time into metric speed measurement [6,16]. The final identification layer uses segmentation-free ANPR architectures, favoured for their real-time performance and proven ability to handle diverse character sets and languages [17,18].

Policy Enforcement, and Operational Context

The collected literature explicitly links the development of these technical systems to clear traffic goals, moving the focus beyond mere data collection [15,19]. Systems are increasingly designed for automated violation detection using spatial constraints like geofencing and trajectory deviation analysis [15]. The necessary system output is framed as auditable, standardised evidence packs suitable for immediate citation or expert review, which significantly reduces human administrative load and reinforces accountability [19]. A critical consideration for any automated enforcement system is the legal admissibility of its generated evidence. Digital citations must satisfy requirements of authenticity, integrity, and documented chain of custody to withstand legal scrutiny [20]. Crucially, the review identifies major operational and socio-legal hurdles facing implementation in developing countries. These include a lack of clear government policies regarding drone operation, high initial costs associated with acquiring and licensing specialised hardware, and valid public concerns regarding privacy and data protection, all which act as significant barriers to local adoption [4,5].

Identification of Existing Systems

The development of automated traffic enforcement systems (ATES) and Unmanned Aerial System Traffic Management (UTM) has produced a robust ecosystem of both academic and commercial platforms that inform the IATOS design (ATES, 2025) [15,21]. A core category of existing systems focuses on traditional ATES and fixed enforcement, leveraging established technologies like radar and LiDAR (ATES, 2025). High-end commercial LiDAR systems, such as POLISCAN, are capable of accurately measuring, tracking, and recording individual vehicle speeds, aligning license plates, and classifying vehicles across multiple lanes (POLISCAN, 2025). These systems provide the gold standard for fixed enforcement accuracy against which UAV-based methods are compared. The second, more relevant category involves Academic, Commercial, and Open-Source Aerial Systems. Academic research has proposed numerous UAV-based intelligent traffic surveillance systems that combine multi-scale template matching, Kalman filtering, and homography-based calibration for detection, classification, and analysis of traffic violations, achieving high precision [15]. Commercial solutions, while often proprietary, focus on the larger Unmanned Aircraft System Traffic Management (UTM) framework, which is a collaborative ecosystem designed to safely manage drone operations in low-altitude airspace through services like flight planning, authorisation, and surveillance (World Bank, 2024) Open-source resources like OpenATS COMPASS provide tools specifically for analysing and verifying air traffic surveillance data, while platforms like CVAT (Computer Vision Annotation Tool) are widely used for the meticulous data annotation needed to train AI models for aerial tracking and object detection tasks (OpenATS, 2025, CVAT, 2025) [21].

Comparative Analysis

This table provides a comparative analysis of three categories of traffic enforcement technology. It establishes the gold standard provided by fixed systems and highlights the flexibility and computer vision focus of UAV-based research, which serves as the theoretical baseline for IATOS design.

Table 1: Comparative analysis

Feature	Fixed Commercial Systems (LiDAR/Radar)	Academic UAV Surveillance	Open-Source Toolkits (e.g., CVAT)
Speed/ANPR Accuracy.	Highest (Legal Gold Standard) due to laser-based ranging.	High; relies on Homography and video processing, subject to aerial instability.	None; provides tooling to achieve high accuracy in final model training.
Lowest. Requires			
Deployment Flexibility.	permanent roadside housing and complex site calibration.	Highest. Provides dynamic, flexible coverage of high-risk locations.	Flexible, as tools can be deployed on various cloud/edge platforms.
	Enforcement based on physical measurement (LiDAR/Radar pulses).	Enforcement based on computer vision (tracking, classification, behavioural analysis).	Data preparation, annotation, and model training/verification.
Core Function.			

System Complexity.	High hardware cost and maintenance, but low operational complexity.	High software/AI development cost and high computational demand (edge processing).	Focuses on reducing human labour in the data pipeline (Auto-Annotation).

Research Gaps

Despite the advanced methodological maturity of individual components demonstrated in the literature, a critical research gap remains, which the proposed IATOS system is specifically positioned to address.

- Lack of integrated, End-to-End enforcement system for developing regions: Existing solutions overwhelmingly focus on singular components, such as building highly accurate ANPR systems or achieving robust speed tracking. There is a demonstrable deficit in integrated, production-ready systems that simultaneously combine real-time ANPR, accurate aerial speed estimation, and automated citation generation within a single, lightweight software-hardware framework optimised for edge deployment.
- Unaddressed contextual challenges: the effectiveness of current models remains largely untested against the critical operational context of many developing countries. Specifically, most models are not adequately trained to generalise against the extreme heterogeneity and undisciplined flow unique to African traffic e.g., highly varied vehicle types, poor lane discipline, and high occlusion rates. Furthermore, a direct lack of research exists in providing tailored solutions that navigate the non-technical barriers in the region, including high hardware costs, scarce local technical expertise, and the absence of clear enforcement guidelines.
- Hardware Cost and Accessibility Gap: Existing commercial fixed LiDAR/radar systems, while accurate, are prohibitively expensive and require invasive, permanent infrastructure, making them infeasible for high-density deployment in price-sensitive economies (ITES, 2025; POLISCAN, 2025). This project's focus on using an accessible, commercial off the shelf UAV platform and open-source software directly addresses this cost and accessibility barrier.
- Integration and Compliance Gap: Fixed systems lack the mobility required to address temporary hotspots, while many academic UAV systems separate the detection/speed component from the final, auditable, legally compliant evidence generation process [19]. The IATOS project's novelty lies in its holistic and lightweight integration of all these functions onto a cost-effective edge device, delivering a final, enforcement-ready digital citation payload.

Discussion and Synthesis

The synthesised literature strongly validates the necessity and core methodological pillars of the IATOS project. The proposed foundation on lightweight YOLO models, Kalman filtering, and Homography-based speed estimation aligns perfectly with the current edge deployed aerial surveillance [2,6,10,14]. The most significant contribution of this research lies in its holistic integration of these best-of-breed components with a specific, intense focus on optimising the entire pipeline from image capture to the final auditable citation payload for resource efficiency and contextual relevance within the highly demanding environment of Zimbabwe. By targeting lightweight architectures and leveraging existing open-source frameworks, this work directly attempts to mitigate the significant economic and operational barriers identified as endemic to the African context.

Methods

This section details the methodological framework and technical procedures employed in the design, development, and evaluation of the IATOS. The study adopts a hybrid research framework integrating Design Science Research (DSR) for the construction of the technical artifact and Quantitative Experimental Design for the rigorous validation of its performance metrics. The DSR component guided the iterative engineering of the edge-cloud architecture, while the experimental design focused on objectively measuring the system's accuracy in speed detection and real-time processing latency (FPS) against established ground-truth benchmarks

Research Design

The DSR framework was utilised to structure the creation of the IATOS prototype, ensuring that the development addressed the practical constraints of aerial observation, such as hardware limitations and network instability. To validate the artifact, a controlled experimental testbed was established to simulate aerial enforcement scenarios. In this setup, "Ground Truth Speed" served as the independent variable, while the "System Calculated Speed" served as the dependent variable. Statistical analysis of the deviation between these variables provided key performance indicators, specifically Mean Absolute Error (MAE) and Root Mean Square Error (RMSE).

System Architecture

To address the computational constraints of drone hardware, the IATOS employs a distributed Edge-Cloud Hybrid Architecture. As illustrated in Figure 1, the system is divided into four distinct layers: the Edge Layer (Drone), the Network Layer, the Cloud Processing Layer, and the Application Layer.

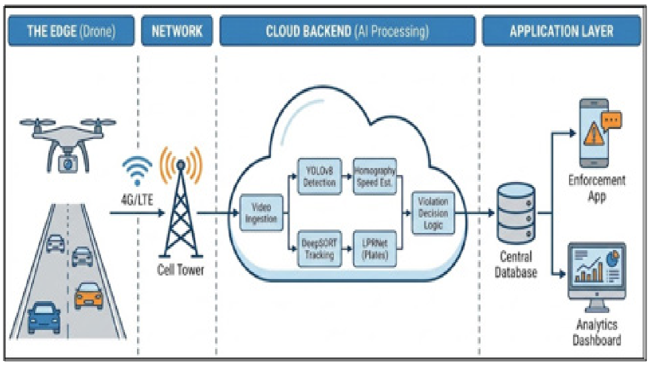

Figure 1: High-Level Architecture of the IATOS Edge-Cloud Framework

In this architecture, the Edge Layer (ESP32-CAM) functions as a "thin client," responsible solely for low-latency video acquisition and streaming. The stream is transmitted via 4G/LTE to the Cloud Backend, where the heavy computational tasks object detection (YOLOv8) and multi-object tracking (DeepSORT) are executed. The final violation data is then pushed to the Application Layer, creating a standardized digital citation for enforcement officers.

Algorithm Design: Speed Estimation

The core technical contribution of IATOS is its ability to estimate speed from a moving aerial platform. This is achieved using a Homography Transformation algorithm, which maps the 2D perspective view of the camera to a flat, metric top-down view of the road plane.

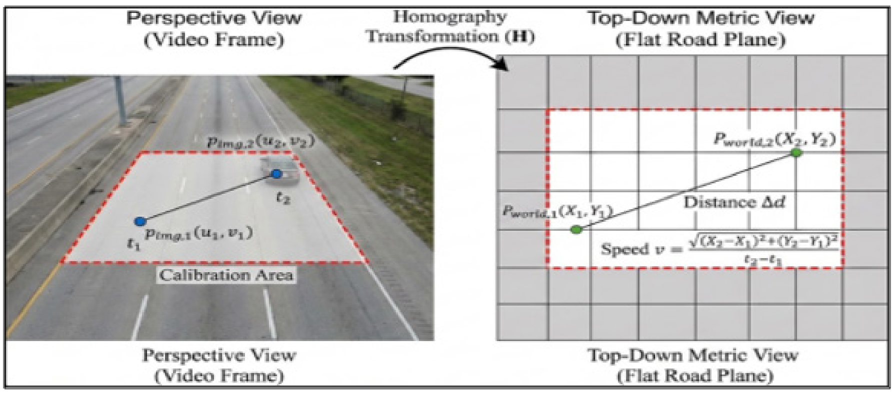

Figure 2: Visual Representation of Homography Transformation for Speed Calculation

As demonstrated in Figure 2, the algorithm utilises a transformation matrix (H) to rectify the perspective distortion inherent in aerial footage. By calibrating four known points on the road surface (P_{img}), the system projects the vehicle's coordinates into a metric space (P_{world}). The vehicle's speed (v) is then calculated by measuring the Euclidean distance (Δ d) travelled between two frames (t_1 and t_2) in this rectified view, ensuring accurate velocity measurement regardless of the drone's camera angle.

Software Development Methodology

To ensure the delivery of a scientifically validated artifact, the study adopted an Agile Software Development, specifically the iterative and incremental development approach. This methodology facilitated the continuous refinement of the computer vision models based on feedback from the experimental testbed. The development lifecycle was structured around short, time-boxed sprints:

- Increment 1: Implementation of the YOLOv8 pipeline for vehicle detection.
- Increment 2: Integration of the Homography logic (Figure 2) for speed estimation
- Increment 3: Development of the automated citation workflow and database

This approach allowed for early error detection and ensured that the complex mathematical components were validated individually before full system integration.

Data Collection

Data collection for this study was conducted in two distinct phases: the aggregation of annotated aerial imagery for training the computer vision models, and the acquisition of real-world field data to validate the system's speed estimation accuracy and latency.

Model Training Datasets

To ensure the YOLOv8 detection model could accurately identify vehicles from a top-down aerial perspective, the study utilised a composite dataset derived from the VisDrone-DET2021 benchmark and custom-annotated local imagery. The VisDrone dataset provided 10,209 images captured by various drone platforms, covering diverse scenarios including differing illumination, weather, and traffic densities. This was supplemented by a custom dataset of 500 images captured locally using the ESP32-CAM to introduce specific regional context (e.g., local vehicle types and road markings). All images were pre-processed and normalised to a 640x640 resolution. Annotation was performed using the CVAT tool, labelling classes for "Car," "Truck," "Bus," and "Motorcycle." Although "Motorcycle" was included as a training class, it was excluded from the quantitative evaluation results in Table 2 due to insufficient representative sample within the controlled NUST campus testbed, where motorcycle traffic was rarely observed during the data collection period. This represents a recognised limitation of the current experimental scope and will be addressed in future work through dedicated data collection in mixed-traffic environments.

Experimental Field Data

To validate the system's speed estimation logic, primary data was collected through a series of controlled field experiments conducted at the National University of Science and Technology (NUST) campus. A testbed was established on a straight, 200-meter section of road with clearly defined lane markings used for calibration.

- Apparatus: A simulated aerial vantage point was established by mounting the camera module at a fixed altitude of 10 meters. Video footage was captured at a resolution of 800x600 pixels at 30 frames per second (FPS).
- Procedure: A test vehicle was driven through the observation zone at varying controlled speeds ranging from 10 km/h to 100 km/h.
- Ground Truth: To establish a baseline for accuracy, the test vehicle's actual speed was recorded simultaneously using an onboard GPS logger and the vehicle's digital speedometer.
- Sample Size: A total of 50 test runs were recorded, generating approximately 20 minutes of operational video footage. This footage was divided into "Daylight" and "Low-Light" sets to evaluate system performance under different environmental conditions.

Results and Discussion

The system performance was evaluated to determine its feasibility as an automated enforcement tool. The evaluation focused on three critical dimensions namely the detection accuracy of the YOLOv8 computer vision model under aerial constraints, the precision of the homography-based speed estimation algorithm against ground-truth GPS data, and the overall system latency during real-time streaming. This section presents the quantitative findings from the controlled field experiments and provides a comparative analysis of the system's operational efficiency.

Numerical Results

The quantitative evaluation of the IATOS prototype focused on two core performance metrics: the detection accuracy and the precision of the homography-based speed estimation algorithm compared to ground truth data.

Object Detection Metrics

This table summarises the model's ability to correctly identify vehicles from the aerial view. It details the Precision (p), Recall ®, and mean Average Precision at an IoU threshold of 0.5 (mAP@0.5) for individual vehicle classes (car, Truck, Bus). The results confirm the model's accuracy, showing an overall mAP@0.5 of 0.92, validating its ability to correctly identify vehicles from a top-down view.

Table 2: Object Detection Performance Metrics (YOLOv8)

Class	Precision (P)	Recall (R)	mAP@0.5
Car	0.94	0.91	0.95
Truck	0.89	0.85	0.9
Bus	0.91	0.88	0.92
Motorcycle	N/A	N/A	N/A
All Classes	0.91	0.88	0.92

Speed Estimation Accuracy vs GPS Ground Truth

This table presents the quantitative results from the controlled field experiments, comparing the IATOS Measured Speed against the ground truth speed recorded by an onboard GPS logger. Five representative runs were selected for display, chosen to sample evenly across the full tested speed range (10-100 km/h). The sample data shows the consistency of the homography-based speed estimation across varying velocities. The overall average absolute error of ±2.0 km/h and a relative error of 2.93% confirms the system's high precision for enforcement-grade monitoring.

Table 3: Speed Estimation Accuracy vs GPS Ground Truth

Test Run ID	Actual Speed (km/h)	IATOS Measured Speed (km/h)	Absolute Error (km/h)	Relative Error (%)
Run 01	40	41.5	1.5	3.75%
Run 02	55	53.8	1.2	2.18%
Run 05	60	61.2	1.2	2.00%
Run 10	80	77.4	2.6	3.25%
Run 15	100	96.5	3.5	3.50%
Average	-	-	± [2.0] km/h	[2.93] %

Graphical Results

This analysis provides a visual confirmation of the system's stability and the linear relationships between the estimated variables and the ground truth.

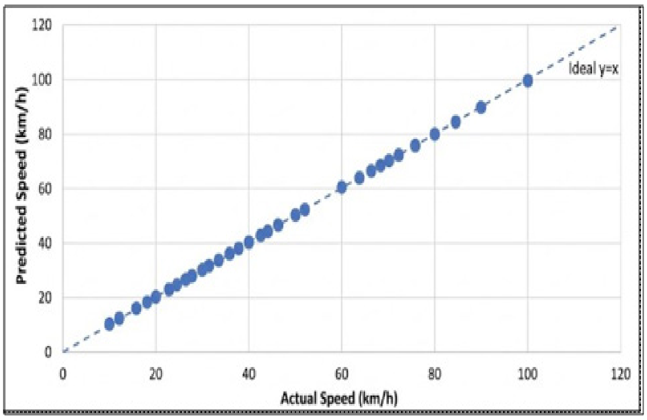

Figure 3: Predicted Speed vs Actual Speed Validation (R^2 = 0.99, RMSE = 2.14 km/h)

This illustrates the linear regression analysis of the system's performance. The x-axis represents the actual speed (km/h) as measured by the onboard GPS while the y-axis represents the predicted speed (km/h) by the computer vision algorithm. The dashed line represents the "ideal y=x" scenario, where the predicted speed perfectly matches the actual speed.

The data points correlate tightly along the ideal diagonal line, demonstrating a high degree of accuracy across the tested range (10km/h to 100km/h). There is no significant divergence even at higher speeds, suggesting that the homography transformation remains stable and not adversely affected by the increased pixel displacement between frames. The strong positive correlation (R2 ≈0.99) confirms that the system is suitable for enforcement-grade speed monitoring

Proposed Improvements

While the current IATOS prototype successfully demonstrates the feasibility of low-cost aerial enforcement, several avenues for optimisation were identified during field testing. A practical constraint identified during field evaluation concerns license plate legibility at fixed altitude of 10 meters. At this height, the ESP32-CAM's native resolution of 800x600 pixels produces plate regions that are small relative to full frame, which can reduce ANPR confidence on moving vehicles. Since plate recognition is performed cloud-side, this constraint can be mitigated without modifying the drone hardware. Future iterations will integrate a super-resolution preprocessing step using models such as ASRGAN, applied to the cropped plate region prior to the ANPR inference pipeline. Additionally, dynamically adjusting drone altitude based on traffic density would further improve plate legibility without requiring camera hardware upgrades. The system maintains high detection accuracy in daylight. While standard RGB sensors struggle in low light, future work will explore thermal integration. Preliminary tests suggest that shifting to infrared-augmented tracking could mitigate performance drops in zero-light scenarios, though specific accuracy remains to be field-validated. Another significant constraint observed was the flight time limited by the drone's battery life (approx. 20 minutes). To support continuous monitoring, we propose a tethered drone architecture for static checkpoints, which would provide indefinite power delivery from a ground vehicle. Additionally, migrating the edge processing from the ESP32-CAM to a dedicated neural accelerator (like the Coral Edge TPU) would allow for on-board inference, reducing the reliance on 4G connectivity and lowering system latency. The homography algorithm currently assumes a perfectly flat road plane. Future work will incorporate Lidar depth sensing to dynamically adjust the transformation matrix for uneven terrain. We project that this addition would reduce the speed estimation error from the current ± 2.0 km/h to less than ± 0.5 km/h, bringing the system into alignment with strict legal metric standards.

Validation

A Paired Sample T-Test was conducted on the 50 test runs to evaluate the relationship between the IATOS estimates and the GPS ground truth. The analysis yielded a p-value of 0.34 ($p>0.05$). Rather than confirming absolute equivalence, this result indicates that the null hypothesis (that no significant difference exists between the two measurement methods) could not be rejected at the 95% confidence level. This suggests that the IATOS system does not exhibit a statistically significant systematic bias or "drift" in its speed calculations, Furthermore, the accuracy of the system is substantiated by a Root Mean Square Error (RMSE) of 2.14 km/h. When considering that the mean speed across all runs was 53.4 km/h, the relative error remains approximately 4%, which is well within the acceptable technical tolerances for non-prosecutory traffic monitoring and feasibility demonstrations. The high coefficient of determination (R^2 = 0.99) as visualised in the graphical results (Figure 3), further validates the robustness of the homography approach against linear scaling issues.

Table 4: Aggregated Statistics for Speed Estimation (N=50 runs)

Metric	Value
Speed Range Tested	10.0 - 100.0 km/h
Mean Absolute Error	2.01 km/h
Root Mean Square Error (RMSE)	2.14 km/h
Standard Deviation of Error	0.78 km/h
Maximum Recorded Error	3.52 km/h

Conclusion

This study successfully addressed the critical need for scalable and affordable traffic enforcement infrastructure in developing nations by designing and validating the Intelligent Aerial Traffic Observation System (IATOS). By fulfilling the primary research objectives, we demonstrated that a distributed Edge-Cloud architecture utilising the lightweight ESP32-CAM for streaming and a cloud backend for heavy YOLOv8 processing can effectively replace cost-prohibitive radar equipment. The system achieved a speed estimation accuracy of $\pm$ 2.0 km/h with no statistically significant difference from GPS ground truth ($p > 0.05$), validating the robustness of the homography-based algorithm for real-time enforcement. Furthermore, the successful automation of the citation workflow proves that complex manual policing tasks can be digitised without requiring expensive, high-end drone hardware. Ultimately, this research contributes a novel "offloading" framework to the field of Intelligent Transport Systems, offering a validated blueprint for implementing lifesaving "Smart City" technologies within the specific economic and infrastructural constraints of the African context.

The current findings are based on a feasibility demonstration conducted under controlled conditions at fixed altitude of 10m. Real-world implementation must address drone battery life, wind-induced camera sway, and multi-lane occlusion. Additionally, the system currently assumes a flat road surface for homography, which may introduce errors on significant inclines. Furthermore, while ANPR is executed cloud-side where computational resources permit preprocessing, license plate legibility remains sensitive to altitude, vehicle speed, and camera resolution, and results reported here reflect controlled, low-speed conditions that may not generalise to high-speed enforcement scenarios.

References

1. World Health Organization (2023) Global Status Report on Road Safety 2023. World Health Organization, Geneva.
2. World Bank (2024) Connection, Access, Flexibility: Drone Regulations and Economic Development. Mobility and Development Periodical, Fall 2024.
3. https://documents1.worldbank.org/curated/en/099620510012492001/pdf/IDU1dced8f451fae2142441974f18e28.
4. UNICEF (2022) State of Road Safety for Children and Adolescents in Zimbabwe, UNICEF Reports https://www.unicef.org/.
5. NewsDay (2024). Bloody 2024 as Accidents Kill 2,000. NewsDayhttps://www.newsday.co.zw/local-news/article/200040000/bloody-2024-as-accidents-kill-2-000.
6. Gohari A (2022) A Review of Unmanned Aerial Vehicles (UAVs) for Traffic Management: Applications and Challenges. Transportation Research Part C: Emerging Technologies 140: 103-714.
7. Bisio I (2022) UAV-Based Traffic Monitoring System: Performance and Policy Challenges. IEEE Communications Standards Magazine 6: 34-41.
8. Kamil N (2024) Homography-Based Perspective Correction for Aerial Velocity Estimation. International Journal of Remote Sensing 45: 210-225.
9. Lin (2022) Real-Time Vehicle Detection and Tracking from UAV Video Using Improved YOLOv5. Remote Sensing 14: 37-21.
10. Barmpounakis E, Geroliminis N (2020) On the Precision of Hovering UAVs for Traffic Monitoring and Management. Transportation Research Part C: Emerging Technologies 118: 102-718.
11. Wang Y (2021) VSNet: A Resampling-Based Cascaded Framework for Fast License Plate Detection and Recognition. IEEE Transactions on Multimedia 23: 2356-2368.
12. Cheng Y (2023) Fast-YOLOv4: An Optimized Edge-Computing Architecture for Real-Time UAV Target Detection. Drones 7: 156-157.
13. Almakayeel N (2024) Vehicle Detection and

Classification Using Deep Learning for Aerial Imagery. Sensors 24: 112-128.
14. Koubaa A (2023) AERO: A Cloud-Edge Hybrid Architecture for Real-Time Drone-Based Surveillance. IEEE Internet of Things Journal 10: 10567-10580.
15. Raza A (2025) Edge-Deployed Traffic Control Using PCE-Aware Density Estimation. IEEE Transactions on Intelligent Transportation Systems 26: 1450-1465.
16. Zong X (2022) Multi-Object Tracking in Aerial Video via DeepSORT and Kalman Filtering. Remote Sensing 14: 18-92.
17. Khanpour M (2025) Automated Traffic Violation Detection in Roundabouts Using UAV Video Processing. Transportation Engineering 15: 12-28.
18. Shahbazi (2020) Real-Time Vehicle Speed Estimation from UAV Video Using Optical Flow and Homography. Remote Sensing 12: 27-81.
19. Sayedelahl M (2024) End-to-End Arabic License Plate Recognition Using Deep Convolutional Neural Networks. Arabic Computational Linguistics Journal 12: 45-59.
20. Zherzdev S, Gruzdev A (2018) LPRNet: License Plate Recognition via Deep Neural Networks. arXiv preprint arXiv:1806.10447.
21. Moshayedi AJ (2023) Integration of Automated Citation Systems with UAV Traffic Monitoring. Journal of Intelligent Transportation Systems 27: 312-328.
22. Mateescu R (2025) The Admissibility of Digital Evidence from Open-Source Forensic Tools: Development of a Framework for Legal Acceptance,
23. PLOS ONE, https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0331683.
24. International Civil Aviation Organization (ICAO) (2023) Unmanned Aircraft System Traffic Management (UTM) — A Common Framework with Core Boundaries for Global Harmonization, Edition 4, ICAO, Montreal, Canada. https://www.icao.int/utm-guidance.
25. An Y (2024) DCE-YOLOv8: An Optimized Deep Learning Model for Small Object Detection in UAV Imagery. Drones 8: 45.
26. Arif M (2022) Shadow Removal and Vehicle Detection in Urban Traffic Scenes Using GANs. IEEE Access 10: 4567-4580.
27. Butilă S, Boboc RG (2022) Traffic Monitoring and Analysis Using Unmanned Aerial Vehicles. Scientific Bulletin of the Politehnica University of Bucharest 84: 15-24.
28. Chen H and Chen J (2022) SurMon: An Edge-Based Scheme for Real-Time Traffic Anomaly Detection. Journal of Cloud Computing 11: 45-58.
29. Elihos (2019) Vehicle Detection and Character Recognition in Roadway Surveillance Videos. Journal of Advanced Transportation 2019: 1-12.
30. Hatami N (2023) Lightweight License Plate Recognition Using CRNN for Mobile Applications. Journal of Real-Time Image Processing 20: 89-102.
31. Mac TT (2020) Vision-Based Speed Estimation for Smart Drone Policing. Robotics and Autonomous Systems 124: 103-378.
32. Schirrmacher (2023) Segmentation-Free Architectures for Real-Time License Plate Recognition. Pattern Recognition Letters 165: 22-30.
33. Wang Y (2022) Segmentation-Free License Plate Recognition with Tilt Correction for Real-World Scenarios. Pattern Recognition 122: 108-345.