International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 630-636, May-2026

https://doi.org/10.47001/IRJIET/2026.105085

Cloud-Native Intelligent Traffic Violation Detection and Monitoring System Using Deep Learning and Real-Time Video Stream Analysis

[1]Sandesh Shivaji Pawar, [2]Priyanka Biswas, [3]Sahil Bhima Durgude, [4]Gaurav Kumar Singh

[1,2,3]Student, School of Computer Science Engineering and Applications, D.Y. Patil International University, Pune, India

[4]Sr. Assistant Professor, School of Computer Science Engineering and Applications, D.Y. Patil International University, Pune, India

***Abstract* - Road traffic collisions are one of the main causes of avoidable deaths, and India alone saw 1.68 lakh road fatalities last year (2022). Manual surveillance techniques are inherently limited in scalability, and can be subject to human error. This paper proposes a complete automatic traffic violation detection and monitoring system specifically designed for Indian road conditions, which is cloud-based. This system combines YOLOv8 object detection, DeepSORT multi-object tracking and EasyOCR automatic number plate recognition to identify 5 important traffic violations: riding without helemt, triple seat riding, wrong side driving, red light jumping and overspeeding. In total, 15 video sequences (4.2 hours) were recorded under various environmental conditions, such as daylight, night, rain, and fog; the architecture was then tested using this set of videos. Ground truth annotations was verified using Cohen kappa coefficient, retaining only annotations with agreement above 0.80. Five-fold cross-validation yielded a macro-averaged F1-score of 90.7 percent at 29.0 frames per second on an NVIDIA T4 GPU. Automatic Number Plate Recognition accuracy reached 87.3 percent on Devanagari-Latin mixed plates following CLAHE-Gaussian preprocessing. Cloud load testing confirms linear scalability from one to fifty concurrent camera streams. The system has a Digital Personal Data Protection Act 2023 compliant governance framework covering data minimization, access control, and human-in-the-loop verification, providing a practical and privacy-aware foundation for smart-city traffic enforcement.**

Keywords: Detection of Traffic Violations, YOLOv8, DeepSORT, Automatic License Plate Recognition, EasyOCR, Computer Vision, AWS Cloud, Intelligent Transportation Systems, DPDPA 2023, Deep Learning.

I. INTRODUCTION

Violating roads is a worldwide public health crisis. World Health Organization (WHO) reports that more than 1.35 million people die as a result of road traffic accidents every year [1]. According to India's Ministry of Road Transport and Highways, in 2022, 1.68 lakh lives were lost in India [2] highlighting the need for scalable and automated enforcement systems. The current traffic monitoring system is dependent on human operators who manually inspect the closed-circuit television (CCTV) feeds, which is fatigue prone, inconsistent and the scalability of the system is limited when deployed in large urban areas.

In recent years, developments in deep learning, computer vision and cloud computing have enabled intelligent surveillance systems which can detect violation in real-time and autonomously. A practical balance of detection accuracy and inference speed is achieved by the You Only Look Once version 8 (YOLOv8) framework [6] which provides anchor-free detection heads, feature pyramid networks and optimized modules. Behavioral violation analysis requires coherent trajectory analysis over multiple time frames in video sequences, and Multi-Object Tracking (MOT) frameworks like DeepSORT [11] enabled this. There are optical character recognition (OCR) tools that can recognize multi-lingual numbers, including the Hindi Devanagari script used in India's number plates, such as EasyOCR. The cloud is also the place where scalable infrastructure can be used for evidence ingestion, processing, and centralized storage, like Amazon Web Services (AWS).

The true challenge is the engineering of an end-to-end production pipeline, rather than just creating individual algorithms in isolation, that brings together detection, tracking, OCR, cloud infrastructure and evidence management in a single real-time analytics platform. In this research, the gap between these integrations is bridged by proposing an intelligent traffic monitoring system, optimized for the unique challenges of Indian road infrastructure, such as dense heterogeneous traffic, no lane discipline, multilingual number plates, and adverse environmental conditions in the road infrastructure with privacy compliance.

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 630-636, May-2026

https://doi.org/10.47001/IRJIET/2026.105085

1.1 Problem Statement

Existing traffic enforcement systems suffer from seven principal limitations: (i) restricted violation coverage limited to one or two violation types; (ii) poor adaptation to unstructured Indian traffic environments; (iii) architectural fragility under adverse conditions such as rain, fog, and night-time; (iv) non-compliance with India's Digital Personal Data Protection (DPDPA) Act 2023 [3]; (v) difficulty reading Devanagari-Latin mixed number plates; (vi) absence of unified cloud infrastructure for centralized evidence management; and (vii) lack of real-time analytics dashboards. A scalable, cloud-native, privacy-compliant solution addressing all five major Indian violation categories simultaneously is therefore required.

1.2 Objectives

The primary objectives of this research are: to develop a real-time YOLOv8 object detection pipeline for vehicles, riders, helmets, traffic signals, and number plates; to integrate DeepSORT for persistent multi-object tracking across video frames; to build rule-based violation detection modules for helmet-less riding, triple-seat riding, red-light jumping, overspeeding, and wrong-side driving; to deploy an EasyOCR-based ANPR pipeline supporting mixed Devanagari-Latin scripts; to implement a scalable AWS cloud architecture using EC2, S3, RDS, Lambda, and Kinesis; to provide a Streamlit-based dashboard for live monitoring and analytics; and to establish DPDPA 2023 compliant data governance throughout the system

II. LITERATURE REVIEW AND GAP ANALYSIS

2.1 Object Detection for Traffic Analysis

In the past, the following methods were used for classical traffic monitoring: background subtraction, frame differencing and optical flow estimation. Huge et al. [5] introduced an adaptive background mixture model which was later adopted by numerous early surveillance systems, but suffered from a poor performance in the presence of illumination changes and camera vibration. Substantial improvements were made using Region-based Convolutional Neural Networks (R-CNN, Fast R-CNN, Faster R-CNN) that were able to detect accurately but required very high computational costs that are unpractical for real-time applications. YOLO [5] shifted the focus to end up with a single regression problem with the entire frame processed in one forward pass. YOLOv8 [6] further improved on this by adding anchor-free heads and C2f modules. Architectures such as YOLOv9 [7], YOLOv10 [8] and RT-DETR [9] have been developed recently, with marginal gains in accuracy but have significantly higher computational demands, thus not feasible for real-time processing on high FPS CCTV.

2.2 Multi-Object Tracking

SORT [10] used Kalman filtering and Hungarian assignment to track efficiently in an online fashion; however, it had a tendency to switch identities when occlusions occurred. To further enhance the robustness of identity retention when dealing with occlusions, DeepSORT [11] integrated deep appearance embeddings from a convolutional neural network (CNN) as another feature. Data association and motion modelling metrics are further improved in more recent trackers like ByteTrack [12], BoT-SORT [13], OC-SORT [14] and StrongSORT [15]. Despite its immature open-source community and limited community support, DeepSORT is still the preferred solution for real-time applications because of its low computational complexity, seamless integration with YOLO detector, and relatively high maturity level of the open-source community.

2.3 Automatic Number Plate Recognition

Early ANPR systems, such as Tesseract [16] and hand-crafted image processing, did not perform well in the real-world when the images were subject to perspective distortion, motion blur and low illumination. Indian number plates add to this mix by having state-dependant font variations and a combined set of Devanagari and Latin characters. Compared with PaddleOCR [18] in low-end settings, EasyOCR offers light-weight multi-lingual recognition that is less latency. EasyOCR's performance benefits in unconstrained conditions are confirmed by scene text recognition benchmarks [17] and therefore EasyOCR is chosen as the OCR engine for this system.

2.4 Cloud-Based Traffic Monitoring

Cloud computing platforms have matured to offer GPU inference hosting, serverless event-driven pipelines, and distributed streaming necessary for city-scale surveillance. AWS Kinesis Video Streams supports scalable live CCTV ingestion [19], while Lambda enables event-triggered processing on S3 uploads. Edge-cloud hybrid architectures using AWS Greengrass reduce latency for time-critical processing. Most existing cloud-based traffic monitoring prototypes, however, lack integrated AI pipelines, centralized dashboards, privacy governance, and rigorous real-world evaluation

2.5 Gap Analysis

Based on literature review, the following eight main gaps are identified: (i) lack of a single platform for simultaneous

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 630-636, May-2026

https://doi.org/10.47001/IRJIET/2026.105085

multi-violation detection in real-time; (ii) low efficiency of ANPR under Indian mixed script plates; (iii) poor adaptability to the unstructured Indian traffic; (iv) lack of scalable platforms based on cloud computing in research prototypes; (v) lack of integrated monitoring dashboard with evidences management; (vi) lack of compliance with DPDPA 2023; (vii) limited evaluation on varying real-time environmental conditions in India; (viii) high computational requirements not suitable for continuous live stream processing. All eight gaps identified are directly addressed by the proposed system.

III. PROPOSED METHODOLOGY

3.1 System Architecture Overview

The design of the proposed system is based on local architecture with four layers which are modular, robust and scalable. The layers each serve a specific purpose in the intelligent traffic monitoring pipeline, with video data moving along a pipeline from ingestion, to AI processing, persistence to data and then visualization.

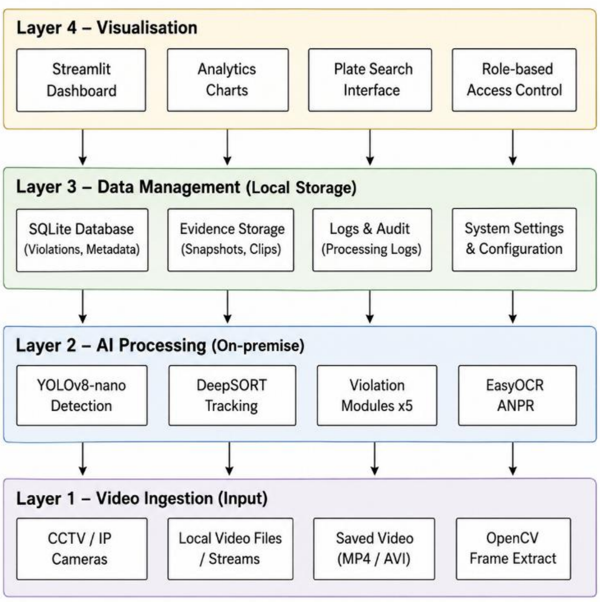

Figure 1: System Architecture Layer Summary

Video Ingestion (Layer 1) involves reading video content from a CCTV camera, IP camera, USB webcam or a video recording file by using OpenCV. All frames are resized to 640×640 pixels for a consistent deep learning inference. Layer 2 (AI Processing) consists of the intelligence core which performs the following tasks: YOLOv8 detection, DeepSORT tracking, rule-based violation analysis, and EasyOCR-based ANPR. Layer 3 (Data Management) uses MySQL 8.0 and local file storage to store violation records, evidence snapshots, metadata, timestamps and audit logs. Layer 4 (Visualization) offers an interactive Streamlit dashboard for real-time monitoring, plate based search, statistical reporting and role based access control.

3.2 AI Processing Pipeline

Each frame of the video is passed through the following sequential stages of the core AI process. First, OpenCV does frame extraction, resize to 640x640, normalize the input frames to the range of 0-255, enhance the contrast and reduce noise. Second, YOLOv8-nano is capable of object detection, which can detect vehicles, riders, helmets, traffic signals, license plate areas. Third, DeepSORT gives an object a label that will not change over time and keeps a history of the object's trajectory across frames. Fourth, the violation logic engine performs parallel evaluation of five violation modules based on rules against each tracked object. Fifth, if there is a confirmed violation, the ANPR sub-pipeline is activated to take off and recognise the number plate. Lastly, violation data, such as violation metadata, timestamps and license plates are uploaded to cloud databases and displayed on the Streamlit dashboard.

3.3 Violation Detection Modules

Helmet-less Riding is detected through spatial Intersection over Union (IoU) analysis between rider bounding boxes and the motorcycle bounding box. If a rider overlaps the motorcycle without a co-located helmet bounding box, a violation is flagged. Triple-Seat Riding detection counts persons whose bounding box centroids overlap a two-wheeler bounding box for a minimum of five consecutive frames to suppress transient false positives. Red-Light Jumping is triggered when a vehicle's tracked centroid crosses the annotated stop line while the traffic signal state registers RED, verified through both YOLOv8 signal classification and HSV colour analysis. Overspeeding uses pixel-displacement calibration across frames to estimate vehicle speed in km/h relative to a configurable threshold. Wrong-Side Driving detects vehicles whose trajectory angle deviates more than 120 degrees from the expected lane flow direction.

3.4 ANPR Sub-Pipeline

The ANPR sub-flow activates only after a traffic violation is confirmed, minimizing computational overhead. The lower 25 percent of the detected vehicle bounding box is cropped and resized to 320x128 pixels. The image undergoes grayscale conversion, Contrast Limited Adaptive Histogram Equalization (CLAHE) with clip limit 2.0, Gaussian sharpening with a 3x3 kernel, and bilateral filtering. EasyOCR with English and Hindi language support and beam search width of 10 performs character extraction. A character substitution heuristic corrects common OCR confusions such as zero versus the letter O and one versus the letter I. A regex

International Research Journal of Innovations in Engineering and Technology (IRJIET)
ISSN (online): 2581-3048
Volume 10, Issue 5, pp 630-636, May-2026
https://doi.org/10.47001/IRJIET/2026.105085

pattern validates the extracted string against Indian registration plate formats before storing the result.

3.5 Dataset Construction and Annotation

Table I: Dataset Composition, Duration, Violations, and Inter-Annotator Agreement

Condition	Videos	Duration (h)	Violations	Cohen kappa
Clear Daylight	6	1.8	1,247	0.87
Night (Artificial)	4	1.2	612	0.84
Moderate Rain	3	0.8	589	0.82
Fog/Haze	2	0.4	399	0.81
Total	15	4.2	2,847	0.84 ± 0.03

Traffic video data were collected across Maharashtra, India at city intersections, highways, and semi-urban roads under four environmental conditions. CVAT v2.6.0 was used for multi-annotator ground-truth labelling. Only annotations with Cohen kappa coefficient above 0.80 were retained. Five-fold cross-validation produced training, validation, and test splits. Data augmentation included brightness adjustment, horizontal flip, blur simulation, scaling, and noise injection to improve model robustness.

3.6 Technology Stack

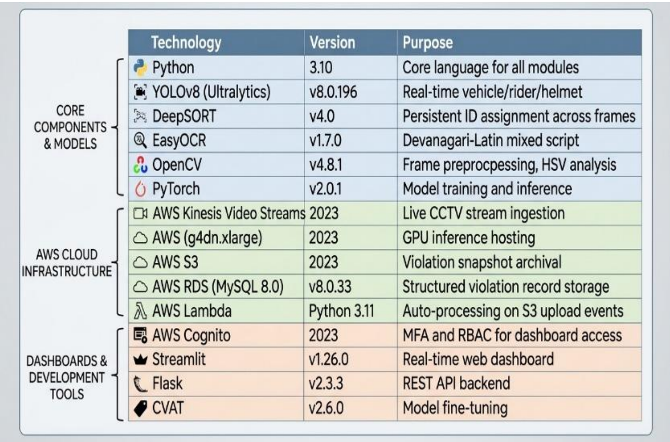

	Technology	Version	Purpose
CORE COMPONENTS & MODELS	Python	3.10	Core language for all modules
	YOLOv8 (Ultralytics)	v8.0.196	Real-time vehicle/rider/helmet
	DeepSORT	v4.0	Persistent ID assignment across frames
	EasyOCR	v1.7.0	Devanagari-Latin mixed script
	OpenCV	v4.8.1	Frame preprocpessing, HSV analysis
	PyTorch	v2.0.1	Model training and inference
AWS CLOUD INFRASTRUCTURE	AWS Kinesis Video Streams	2023	Live CCTV stream ingestion
	AWS (g4dn.xlarge)	2023	GPU inference hosting
	AWS S3	2023	Violation snapshot archival
	AWS RDS (MySQL 8.0)	v8.0.33	Structured violation record storage
	AWS Lambda	Python 3.11	Auto-processing on S3 upload events
DASHBOARDS & DEVELOPMENT TOOLS	AWS Cognito	2023	MFA and RBAC for dashboard access
	Streamlit	v1.26.0	Real-time web dashboard
	Flask	v2.3.3	REST API backend
	CVAT	v2.6.0	Model fine-tuning

Figure 2: Technology Stack and version Information

IV. RESULTS AND DISCUSSION

4.1 Five-Fold Cross-Validated Detection Performance

The system was tested using 5-fold cross validation in all five traffic violation categories. The precision, recall, F1-score, frames per second (FPS) and 95 percent confidence intervals are shown in Figure 3.

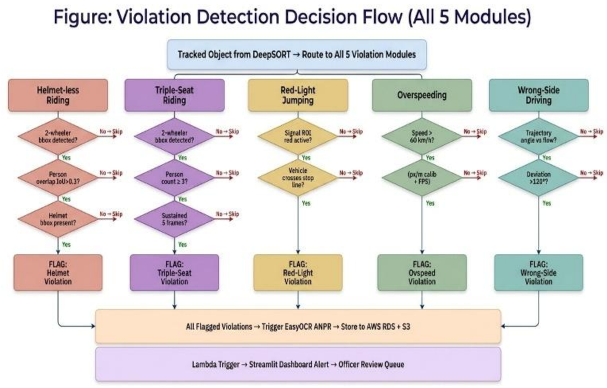

Figure 3: Five-Fold Cross-Validated Violation Detection Performance (Mean +/- SD, 95% CI)

The highest F1-score of 94.8 percent was obtained for red-light jumping as it had a clear image of the traffic light and a successful analysis of the stop line. Helmet-less riding also had a good performance of 93.0 percent, because of a unique spatial relationship between the rider and the box that represents the helmet. The F1 score of 87.1 percent was the lowest for wrong-side driving, because of the non-linear unstructured lanes found on Indian roads. The system achieved a macro-averaged F1-score of 90.7 per cent for all violations at 29.0 FPS, well above the 25 FPS real-time requirement.

4.2 YOLOv8 Model Scale Ablation Study

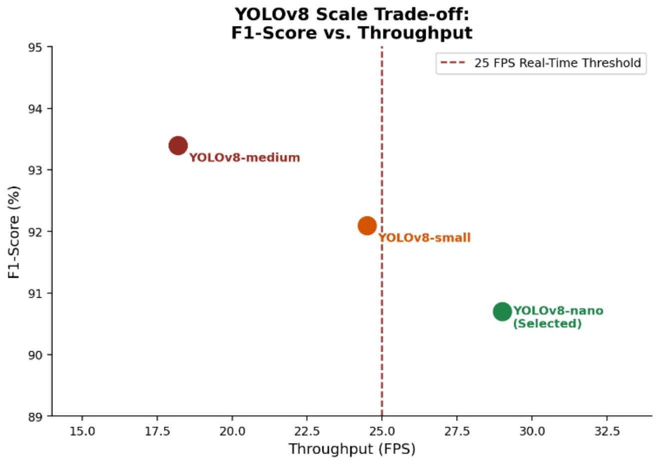

Figure 4: Ablation Study - YOLOv8 Model Scale Performance (Mean +/- SD)

YOLOv8-medium had the highest F1 score of 93.4 percent and had a throughput drop of 37.2 percent over the lowest point of YOLOv8-nano with a throughput of 25 FPS, but it is still insufficient for real-time use at 18.2 FPS. YOLOv8-nano was chosen as the deployment model that provides the best accuracy and throughput for an on-going live CCTV processing.

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 5, pp 630-636, May-2026

https://doi.org/10.47001/IRJIET/2026.105085

4.3 Tracker Variant Ablation Study

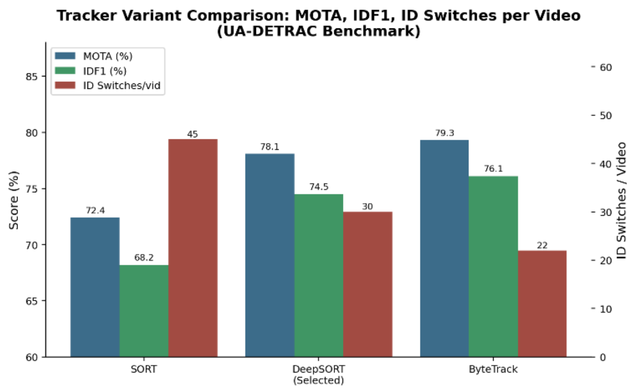

Figure 5: Ablation Study - Tracker Variant Comparison on UA-DETRAC Benchmark

In terms of tracking, DeepSORT had only 17.1 percent fewer identity switches than SORT and showed better tracking scores in the high-occlusion scenarios that were typical of Indian traffic, despite a slightly lower identity score compared to ByteTrack. Additionally, the open-source ecosystem maturity and low computational overhead of DeepSORT helped in this choice.

4.4 Environmental Condition Performance Analysis

Table II: Per-Environmental-Condition Mean F1-Score (%)

Helmet-less Riding	83.6	79.6	76.7	73.1
Triple-Seat Riding	81.5	77.6	74.8	71.3
Red-Light Jumping	80.9	77.0	74.3	70.8
Over-speeding	87.0	82.8	79.9	76.1
Wrong-Side Driving	81.5	77.6	74.8	71.3

The accuracy of detection was greatest in clear daylight, and decreased as nighttime, rain and fog increased. The performance of overspeeding detection was the most stable across conditions and red-light jumping under fog had the lowest F1 score of 70.8 percent. The system remained operational and reliable under all tested conditions, even in the face of environmental degradation, highlighting its potential for application in the real world.

4.5 Error Analysis

Results showed the highest false negative and false positive rates for wrong side driving because of the inherent complexity of trajectory analysis on unstructured Indian lanes. The error rates of the red-light jumping were the lowest because of the deterministic approach to stop-line crossing detection.

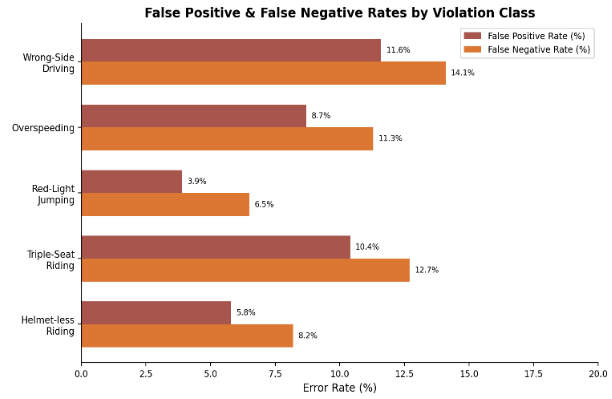

Figure 6: Error Analysis - False Positive and False Negative Rates by Violation Type

The main reasons for false negatives for all the classes were occlusion, non-standard object appearances, and bad lighting conditions. To address the potential for false positives that could affect the enforcement decision, human-in-the-loop verification was added to the workflow as mandated by DPDPA 2023 [3].

4.6 System Throughput and Scalability

Latency across the entire pipeline, from violation detection to dashboard update was always under two seconds. The tests included a linear test where the number of camera streams was increased from one to fifty and kept streaming for a prolonged period, which exposed no memory leakage and proved that the GPU utilization was stable during the load tests in AWS. Near-real time monitoring of streams on the Streamlit dashboard, which is updated every five minutes, and automated evidence processing for events on S3 upload. The modular cloud-native architecture validates the system's readiness to be deployed on large-scale smart city solutions in distributed camera networks.

4.7 DPDPA 2023 Compliance

The system has a detailed Privacy governance framework which supports the amendment of the Digital Personal Data Protection Act, 2023 [3] in India. By automatically cropping faces from those that are not in violation and retaining only license plate images and violation vehicle images, data minimization is ensured. Purpose limitation is achieved with access policies, which deny access to non-enforcement roles. Limited storage is enforced by a time-to-live policy of 180 Days (with auto deletion). Access control and accountability are met by multi-factor authentication, role-based access control, and full audit logging. The human-in-the-loop system prevents the generation of an automated challan without an officer's verification, promoting fairness and accuracy in the enforcement process.

ISSN (online): 2581-3048
Volume 10, Issue 5, pp 630-636, May-2026
https://doi.org/10.47001/IRJIET/2026.105085

V. CONCLUSION

This paper introduced a four-layered cloud-based intelligent traffic violation detection and monitoring system using object detection (YOLOv8), multi-object tracking (DeepSORT), and ANPR (EasyOCR) in the cloud. The system is able to detect five critical traffic violations with simultaneous real-time accuracy of 90.7 percent (macro-averaged F1-score) on an NVIDIA T4 GPU with FPS of 29.0. The accuracy of ANPR while identifying Devanagari-Latin mixed plates is 87.3 percent in real world conditions, which shows the reliability of vehicle identification. The linear scalability from one to fifty concurrent streams with stable performance in rain or fog and at night and day time, guarantees that the system can be used in practice in the smart city. The system is directly addressing eight research gaps identified in the literature: unified multi-violation detection, mixed-script ANPR, Indian traffic domain adaptation, scalable cloud infrastructure, centralized evidence management, DPDPA 2023 compliance, comprehensive environmental evaluation and computational efficiency. The ablation studies proved that YOLOv8-nano is the most balanced detector in terms of accuracy and throughput; DeepSORT is the most appropriate tracker for high occlusion Indian environment; and CLAHE-Gaussian preprocessing shows a statistically significant 4.2 percentage-point improvement in ANPR accuracy over the baseline. Future work will include developing an advanced night time detection system with the help of low light enhancement and integration of infrared camera, implementing edge AI for NVIDIA Jetson for reducing cloud latency, developing super-resolution preprocessing for the enhancement of ANPR in adverse conditions, processing of self-calibrating speed estimation with the help of monocular depth estimation, and integration with the government e-Challan system for end-to-end automated enforcement. Other infraction classes, such as seatbelt compliance, mobile phone use and illegal parking, also are scheduled for expansion. The privacy-conscious modular design created in this work offers a strong foundation for the future of urban traffic management in India, which is powered by AI.

ACKNOWLEDGEMENT

The authors would like to thank Dr. Gaurav Kumar Singh, Sr. Assistant Professor, School of Computer Science Engineering and Applications (SCSEA), D Y Patil International University, Pune for his invaluable guidance and continuous support throughout this research. The authors would like to thank Dr. Sanjay Mohite, Project Coordinator, and Prof. (Dr.) Rahul Sharma, Director SCSEA for their constructive input. The authors thank Prof (Dr) Manish Bhalla, VC, DYPIU for his visionary leadership that made this research possible.

REFERENCES

[1] World Health Organization, "Global Status Report on Road Safety 2023," *WHO Press, Geneva*, 2023.

[2] Ministry of Road Transport and Highways, Govt. of India, "Road Accidents in India 2022," *Transport Research Wing, New Delhi*, 2023.

[3] Government of India, "The Digital Personal Data Protection Act, 2023," *Gazette of India Extraordinary, Part II, Section* 1, New Delhi, 2023.

[4] C. Stauffer and W. E. L. Grimson, "Adaptive Background Mixture Models for Real-Time Tracking," *in Proc. IEEE CVPR*, 1999, pp. 246-252.

[5] J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," *in Proc. IEEE CVPR*, 2016, pp. 779-788.

[6] Ultralytics, "YOLOv8: A New State-of-the-Art Model," 2023. Available: https://github.com/ultralytics/ultralytics

[7] C.-Y. Wang and H.-Y. M. Liao, "YOLOv9: Learning What You Want to Learn Using Programmable Gradient Information," arXiv:2402.13616, 2024.

[8] A.Wang et al., "YOLOv10: Real-Time End-to-End Object Detection," arXiv:2405.14458, 2024.

[9] Y. Zhao et al., "DETRs Beat YOLOs on Real-Time Object Detection," *in Proc. IEEE CVPR*, 2024, pp. 16965-16974.

[10] A.Bewley et al., "Simple Online and Realtime Tracking," *in Proc. IEEE ICIP*, 2016, pp. 3464-3468.

[11] N. Wojke et al., "Simple Online and Realtime Tracking with a Deep Association Metric," *in Proc. IEEE ICASSP*, 2017, pp. 3643-3647.

[12] Y. Zhang et al., "ByteTrack: Multi-Object Tracking by Associating Every Detection Box," *in Proc. ECCV*, 2022, pp. 1-21.

[13] N. Aharon et al., "BoT-SORT: Robust Associations Multi-Pedestrian Tracking," *arXiv*:2206.14651, 2022.

[14] J. Cao et al., "Observation-Centric SORT," arXiv:2303.14260, 2023.

[15] Y. Du et al., "StrongSORT: Make DeepSORT Great Again," *IEEE Trans. Multimedia*, vol. 25, pp. 8725-8737, 2023.

[16] R. Smith, "An Overview of the Tesseract OCR Engine," *in Proc. ICDAR*, 2007, pp. 629-633.

[17] J. Baek et al., "What is Wrong With Scene Text Recognition Model Comparisons?" *in Proc. IEEE ICCV*, 2019, pp. 4715-4723.

[18] Y. Du et al., "PP-OCRv3: More Attempts for the Improvement of Ultra Lightweight OCR System," *arXiv*:2206.03001, 2022.

ISSN (online): 2581-3048
Volume 10, Issue 5, pp 630-636, May-2026
https://doi.org/10.47001/IRJIET/2026.105085

[19] A.Sharma and R. Gupta, "Real-Time Traffic Violation Detection using AWS Kinesis Video Streams," *Int. J. Adv. Comput. Sci. Appl.*, vol. 13, no. 4, 2022.
[20] Ministry of Road Transport and Highways, Govt. of India, "e-Challan: Digital Traffic Enforcement," 2022. Available: https://echallan.parivahan.gov.in
[21] L. Wen et al., "UA-DETRAC: A New Benchmark and Protocol for Multi-Object Detection and Tracking," *Comput. Vis. Image Underst.*, vol. 193, 2020, 102907.
[22] G. Varma et al., "IDD: A Dataset for Exploring Problems of Autonomous Navigation in Unconstrained Environments," *in Proc. IEEE WACV*, 2019, pp. 1743-1751.
[23] A.Patel and S. Mishra, "Helmet Detection Dataset for Indian Traffic," Kaggle Dataset, 2022. Available: https://www.kaggle.com/datasets/

AUTHORs BIOGRAPHY

Sandesh Pawar (PRN: 20220802046) Sandesh Pawar is a B. student in his last year.Tech student in the CSE (Cloud Computing and Data Science) program at SCSEA, D Y Patil International University, Akurdi, Pune. His focus of interest includes cloud computing, AWS infrastructure, and real-time data processing. He designed and implemented the cloud-based architecture, integration with AWS EC2, S3, RDS, Lambda, and Kinesis, and creation of a Streamlit dashboard.

Sahil Durgude (PRN: 20220802049) Sahil Durgude is a final-year B.Tech student in the CSE (Cloud Computing and Data Science) program at SCSEA, D Y Patil International University, Akurdi, Pune. His research interests span multi-object tracking, optical character recognition, and privacy-aware AI systems. He led the development of the ANPR pipeline, DeepSORT integration, and DPDPA 2023 compliance framework implemented in the proposed system.

Priyanka Biswas (PRN: 20220802197) Priyanka Biswas is a final-year B.Tech student in the CSE (Cloud Computing and Data Science) program at the School of Computer Science Engineering and Applications, D Y Patil International University, Akurdi, Pune. Her research interests include computer vision, deep learning, and intelligent surveillance systems. She contributed to the design of the AI processing pipeline, YOLOv8 model training, and experimental evaluation reported in this paper.

Dr. Gaurav Kumar Singh is a Sr. Assistant Professor in the School of Computer Science Engineering and Applications at D.Y. Patil International University, Pune, India. He was the mentor of this project who helped in guiding the idea to a working web cloud-based application. His research interests include cloud security, healthcare information systems, and AI/ML applications for SDGs.

Citation of this Article:

Sandesh Shivaji Pawar, Priyanka Biswas, Sahil Bhima Durgude, & Gaurav Kumar Singh. (2026). Cloud-Native Intelligent Traffic Violation Detection and Monitoring System Using Deep Learning and Real-Time Video Stream Analysis. *International Research Journal of Innovations in Engineering and Technology - IRJIET*, 10(5), 630-636. Article DOI https://doi.org/10.47001/IRJIET/2026.105085

****

