IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

AI-Based Detection of Two-Wheeler Safety Violations at Toll Plazas

Nileshwari Prakash Girase[1], Disha Rajendra Nikumbh[2], Gaurav Bhojraj Borse[3], Tushar Sanjay Patil[4], Dr. Priti Shivaji Sanjekar[5]

Department of Computer Science & Engineering, R. C. Patel Institute of Technology (R.C.P.I.T.), Shirpur, Maharashtra, India[1]

Department of Computer Science & Engineering, R. C. Patel Institute of Technology (R.C.P.I.T.), Shirpur, Maharashtra, India[2]

Department of Computer Science & Engineering, R. C. Patel Institute of Technology (R.C.P.I.T.), Shirpur, Maharashtra, India[3]

Department of Computer Science & Engineering, R. C. Patel Institute of Technology (R.C.P.I.T.), Shirpur, Maharashtra, India[4]

Guide, Department of Computer Science & Engineering (Data Science), R. C. Patel Institute of Technology Shirpur, Maharashtra, India[5]

Abstract: Walk onto any bustling street in India, and you will immediately see that two-wheelers are the lifeblood of daily transit. From college students on scooters to families commuting on motorcycles, these vehicles offer an unbeatable mix of affordability and the agility needed to navigate dense, unpredictable traffic. However, this convenience is often shadowed by a stark reality: unsafe riding habits are a daily norm. Practices like riding without a helmet or cramming three or more people onto a single bike drastically increase the likelihood of severe, life-altering injuries on the road.

Traffic authorities are well aware of this, and toll plazas naturally serve as strategic checkpoints. Since all vehicles must slow down to pass through these structured gates, they provide a rare, controlled window to monitor compliance. But in reality, relying on human operators to catch these violations is an impossible task. During rush hours, bikes zip through lanes in a chaotic blur. Toll booth staff are already overwhelmed with managing payments and directing traffic; asking them to simultaneously spot a missing helmet, count passengers on a moving bike, and note down a license plate is simply beyond human limits.

To bridge this gap, we are proposing an AI-powered computer vision system designed to act as a tireless, automated set of eyes. Rather than replacing human oversight, this system works in the background to monitor two-wheeler safety in real time, filtering out the noise and only flagging actual violations. The technology operates through a clear, multi-step pipeline that continuously scans live video feeds to instantly identify incoming two-wheelers as they approach the plaza. Once a bike is detected, the system calculates the exact number of people on board, immediately flagging dangerous "triple riding" scenarios. Simultaneously, the AI assesses each detected rider to confirm whether they are wearing a helmet. When a violation is spotted, the system captures a high-quality snapshot of the incident, serving as undeniable visual proof for human reviewers to verify before any ticket is issued. Additionally, if the lighting is good and the license plate is clearly visible, an optional smart recognition module automatically reads and logs the vehicle's registration details, streamlining the enforcement process.

While current tech literature is full of deep learning models that boast near-perfect accuracy in traffic monitoring, we know that lab results rarely survive the chaos of the real world. A model trained on clear, sunny highways in Europe will not necessarily work for a crowded toll booth in monsoon-season India. Therefore, we are not taking existing performance claims at face value. Our system's true capabilities will be rigorously benchmarked against a highly specific, locally sourced dataset to ensure it actually works where it matters most.

Finally, deploying this technology outside of a controlled lab introduces a host of gritty, real-world challenges that this project aims to solve. Fixed toll cameras are often positioned for reading car license plates, not checking motorcycle helmets, meaning riders can easily obscure the view of the passengers sitting directly behind them. Furthermore, environmental hurdles like blinding sun glare, heavy rain, and poor nighttime lighting can instantly blind a standard camera, requiring smart image-enhancement techniques. The AI must also strike a delicate balance with its confidence thresholds. If the threshold is too high, it misses real violations, but if it is too low, it floods operators with false alarms,

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

such as mistaking a rider's cap or turban for a helmet. Underlying all of this is the crucial issue of data privacy. Because the system captures images of the general public, it must be designed with strict data-handling protocols to ensure that the faces and information of law-abiding citizens are protected and not needlessly stored. By addressing these practical hurdles, this system aims to move beyond theoretical AI research and deliver a functional, scalable tool that genuinely makes roads safer.

Keywords: Computer vision, deep learning, two-wheeler safety, helmet detection, triple riding, toll plaza, YOLO, number-plate recognition.

I. INTRODUCTION

For millions of daily commuters, two-wheelers are the practical backbone of urban transit, offering an affordable and agile way to beat gridlocked traffic. However, this convenience often comes with a steep cost to public safety. Routine shortcuts—like riding without a helmet or squeezing three or four passengers onto a single motorcycle—drastically increase the risk of severe, preventable road injuries. Policing these habits is a logistical nightmare for traffic authorities, who are already stretched thin by the sheer volume of vehicles on the road.

To systematically address this, we need to utilize the natural bottlenecks in our transit networks: toll plazas. With their structured lanes, predictable traffic flow, and existing fixed-camera infrastructure, toll gates provide an ideal, controlled environment for monitoring rider safety.

Yet, expecting human toll staff to catch every violation is an impossible ask. During peak hours, motorcycles zip through the lanes in a blur. Riders easily obscure one another from view, and environmental factors like harsh sun glare, heavy rain, or poor nighttime lighting make it incredibly difficult for a human eye to quickly verify helmet use. Manual oversight is easily exhausted, prone to error, and simply cannot scale to meet modern traffic demands.

This is where computer vision offers a highly practical lifeline. By analyzing live camera feeds in real time, AI can instantly spot potential violations and flag them for human review. This technology is not designed to replace authorized personnel or automate punishment. Instead, it serves as a tireless, intelligent support tool that filters out the noise, allowing human operators to focus on verifying the flagged evidence rather than staring blankly at a screen.

This paper introduces a comprehensive, AI-driven computer vision pipeline engineered to detect helmet violations and excess passengers specifically at toll plazas. Our proposed architecture handles the entire workflow seamlessly: isolating the vehicle, counting the individuals on board, assessing helmet compliance, and automatically logging the visual evidence—including number-plate recognition when camera clarity permits.

Rather than resting on unverified, theoretical performance claims, this work is deeply grounded in practical application. We focus on building a realistic operational workflow and a structured evaluation framework, demonstrating exactly how this AI system can be deployed in the real world to improve compliance, assist enforcement teams, and ultimately save lives on the road.

II. RELATED WORK AND LITERATURE REVIEW

The idea of using computer vision to monitor motorcyclist safety is not entirely new; in fact, researchers have spent years developing algorithms specifically for this purpose. Previous studies, such as the YOLO-based transfer-learning framework of Srusti *et al.* [10] and the YOLOv5 ensemble of Agorku *et al.* [4], have successfully built frameworks designed to quickly and accurately spot helmets in video feeds. Other teams have tested similar custom setups, carefully mapping out just how well these AI models can perform [8]. While these early experiments provide excellent building blocks and establish strong benchmarks, their impressive accuracy rates were usually achieved under carefully controlled test conditions. Because of this, we cannot simply assume those same algorithms will perform identically within the unique, complex framework we are proposing.

Looking beyond just spotting helmets, the focus of recent research has naturally shifted toward all-in-one systems that can juggle multiple enforcement tasks simultaneously. Innovators in the field have proposed unified designs capable of checking for helmets, catching triple-riding violations, and reading license plates all in one sweep [1], [3]. Recognizing that installing high-end computers at every street corner is impractical, researchers have also proposed ensemble deep-learning approaches for edge-based traffic video analytics [15]. Others have taken the concept a step further by building complete, automated pipelines that not only flag a violation but immediately fire off an automated alert to the offender, such as embedded surveillance with GSM alerts [7] and helmet and number-plate detection with SMS notifications [9].

What all of this collective research tells us is that the foundational technology for detecting helmets, counting passengers, and recognizing license plates is already quite mature. However, taking these algorithms out of a controlled lab and dropping them into a bustling toll plaza changes everything. Out in the field, the system has to contend with high-speed movement, passengers physically blocking the camera's view of one another, constantly shifting sunlight

DOI: 10.17148/IJARCCE.2026.1511131

and shadows, and the strict, unchangeable angles of fixed toll booth cameras—all while processing the video in fractions of a second. Because of these harsh operational realities, we cannot rely on idealized laboratory performance. Our proposed workflow must be rigorously tested against real, messy footage captured directly from the toll lanes.

To give a clearer picture of where the technology currently stands, the following section features a comparative analysis of fifteen different research papers dedicated to two-wheeler safety. Summarized in Table I, this comparison breaks down the various methodologies, key findings, and specific violations each study tackled, mapping out exactly how they relate to our proposed system. This breakdown not only highlights what current AI is already capable of achieving, but also brings into sharp focus the specific, practical hurdles that still need to be overcome to make automated safety enforcement a reliable reality at toll plazas.

TABLE I SUMMARY OF REVIEWED LITERATURE ON TWO WHEELER TRAFFIC VIOLATION DETECTION

Sr. No.	Title, Authors, Year and Publication Venue	Methodology	Key Findings	Violations Detected	Relevance
1	**Two-Wheeler Vehicle Traffic Violations Detection and Automated Ticketing for Indian Road Scenario.** Charran, R.S., & Dubey, R.K. [1], 2022. *IEEE Transactions on Intelligent Transportation Systems, 23(11).*	Deep learning CNN-based detector, ANPR	High detection accuracy across 5 violation categories on real Indian CCTV footage	Helmet, Phone, Triple riding, Wheeling, Parking	Multi-violation detection with ANPR integration for Indian roads
2	**DashCop: Automated E-ticket Generation for Two-Wheeler Traffic Violations Using Dashcam Videos.** Rawat, D., et al. [2], 2025. *IEEE/CVF WACV 2025 (also arXiv:2503.00428).*	Segmentation and Cross-Association (SAC), multi-object tracking, RideSafe-400 dataset	Significant improvement in violation detection accuracy over baseline tracking methods	Triple riding, Helmet violations	Indian dataset creation and violation tracking methodology
3	**Computer-vision based automatic rider helmet violation detection and vehicle identification in Indian smart city scenarios using NVIDIA TAO toolkit and YOLOv8.** Deshpande, U.U., et al. [3], 2025. *Frontiers in Artificial Intelligence, 8:1582257.*	ResNet18-based DetectNet_v2, YOLOv8n for helmet, integrated OCR for ANPR	91.42% mAP for triple-rider detection under challenging real-world conditions	Helmet violations, Triple-riding, Number plate recognition	YOLOv8 + OCR integration achieving benchmark accuracy levels
4	**Real-time Helmet Violation Detection via Ensemble Learning (YOLOv5).** Agorku, G., et al. [4], 2023. *arXiv preprint 2304.09246.*	YOLOv5 with ensemble of five models, varied hyperparameters	mAP of 0.5267; robust performance with diverse training data	Helmet violations	Ensemble learning approach for improved detection robustness
5	**Detection of the Triple Riding and Speed Violation on Two-Wheelers Using Deep Learning Algorithms.** Mallela, N.C., Volety, R., & RK, N. [5], 2021. *Multimedia Tools and Applications, 80(6):8175-8187.*	Deep learning object detection and counting algorithm, frame-based speed estimation	Reliable detection of triple-riding instances and speed violations in test videos	Triple riding, Over-speeding	Triple-riding specific detection and counting methodology
6	**Detection of License Plate Numbers and Identification of Non-Helmet Riders Using YOLOv2 and OCR Method.** Srilekha, B., Kiran, K., & Padyala, V.V.P. [6], 2022. *2022 International Conference on Electronics and Renewable Systems (ICEARS), IEEE.*	YOLOv2 for rider/helmet detection, OCR pipeline for automatic number plate recognition	Effective helmet/no-helmet classification with successful plate-number extraction	Helmet violations, Number plate recognition	YOLOv2 + OCR integration for helmet and plate detection

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

7	**Advanced Computer Vision-Based Surveillance System for Helmet Detection and Triple Rider Identification on Motorcycles using Machine Learning.** Yenugonda, M., & Hemalatha, M. [7], 2023. *International Journal of Scientific Research in Science and Technology, 10(5):163-168.*	CNN-based object detection, Arduino Uno, RFID, GSM notification system	70% accuracy with CNN-based categorization; SMS alerts via GSM	Helmet violations, Triple riding	Embedded systems and IoT integration for real-time notification
8	**Helmet & Number Plate Detection using Deep learning and Its Comparative Analysis.** Singh, A., Singh, D., Singh, J., Singh, P., & Kaur, A. [8], 2023. *Guru Tegh Bahadur Institute of Technology, SSRN.*	CNN + SVM + Decision Tree comparative analysis on 3500 images dataset	CNN: 99.1%, SVM: 97.27%, KNN: 93.45%, Decision Tree: 85.27% accuracy	Helmet violations, Number plate detection	Deep learning model comparison showing CNN superiority for helmet detection
9	**Automatic Number Plate and Helmet Detection Using Yolo.** Sundar, S.S., Pradhumna, S., Sai, T.S., & Navya, K. [9], 2025. *International Journal of Engineering Research and Science & Technology (IJERST), 21(2):2263-2270.*	YOLOv3, DNN classifier, ALPR (Automatic License Plate Recognition), KNN-based SMS API	Real-time helmet and plate detection with SMS notifications for enforcement	Helmet violations, Number plate recognition	Integrated YOLOv3 + OCR + SMS notification system for enforcement
10	**Helmet Detection Using Machine Learning.** Srusti, C., Deo, V., & Jaiswal, R.C. [10], 2022. *Journal of Emerging Technologies and Innovative Research (JETIR), 9(10):d10-d14.*	YOLO algorithm with DarkFlow framework, transfer learning, OpenCV for real-time detection	96% accuracy with mAP@0.50 = 96.17%, detection time 1.35 seconds per frame	Helmet violations (with/without helmet detection)	YOLO transfer learning approach with real-time performance metrics
11	**Real-Time Multi-class Helmet Violation Detection Using YOLOv8 with License Plate Recognition.** Charef, A., Jarir, Z., & Quafafou, M. [11], 2024. *Springer Conference Proceedings, 10.1007/978-3-031-68675-7_26.*	YOLOv8 multi-class detection, integrated license plate recognition module	Real-time multi-violation detection with integrated plate recognition	Helmet violations, Number plate recognition, Multi-class violations	Latest YOLOv8 + LPR integration for complete violation detection

Sr. No.	Title, Authors, Year and Publication Venue	Methodology	Key Findings	Violations Detected	Relevance
12	**Helmet Violation Detection System Based on YOLO11 and EfficientNet for Real-Time Traffic Monitoring.** Ashuri, R.F., Pramesti, D., & Fakhrurroja, H. [12], 2025. *2025 International Conference on Information and Communication Technology (ICoICT), IEEE.*	YOLO11 architecture, EfficientNet backbone, optimized for real-time inference	Enhanced detection accuracy and speed using latest YOLO11 version	Helmet violations, Real-time traffic monitoring	Cutting-edge YOLO11 architecture for toll plaza deployment
13	**An Intelligent YOLOv8 and OCR-Based System for Real-Time Detection of Wrong-Way Vehicles and License Plate Recognition.** Hamddi, S.A., et al. [13], 2026. *ICAIN 2025, Lecture Notes in Networks and Systems, vol. 1931, 10.1007/978-3-032-23314-1_27.*	YOLOv8 for vehicle detection, integrated OCR for plate recognition and wrong-way detection	Complete integrated system for multi-violation detection and enforcement	Helmet, Wrong-way driving, Number plate recognition	Comprehensive YOLOv8 + OCR system for intelligent traffic enforcement

DOI: 10.17148/IJARCCE.2026.1511131

14	**Multi-Object Vehicle Detection and Tracking Algorithm Based on Improved YOLOv8 and ByteTrack.** You, L., et al. [14], 2024. *Electronics, 13(15):3033.*	Improved YOLOv8 with ByteTrack multi-object tracking algorithm	Enhanced tracking performance for multiple vehicles and violation detection	Triple riding detection, Multi-vehicle tracking, Violation association	Advanced tracking methodology for triple-riding and multi-violation detection
15	**Ensemble Deep Learning Approach for Traffic Video Analytics in Edge Computing.** Sathyamoorthy, M., et al. [15], 2026. *Scientific Reports, 16:274.*	Ensemble deep learning, edge computing deployment, real-time video analytics	Optimized edge-based deployment for toll plaza real-time monitoring	Multi-violation detection with edge computing optimization	Edge AI deployment strategy for toll plaza systems

The reviewed studies fall into several groups. Multi-violation systems for Indian roads combine detection with ANPR and automated ticketing across categories such as helmet use, triple riding, and parking [1], and dashcam-based tracking has been used to generate e-tickets for helmet and triple-riding violations on the RideSafe-400 dataset [2]. YOLOv8 pipelines with OCR have reported high mAP for triple-rider detection in smart-city settings [3], and multi-class YOLOv8 with licence-plate recognition [11] and YOLO11 with an EfficientNet backbone [12] target real-time helmet violation monitoring. Ensemble and comparative approaches include a YOLOv5 ensemble for helmet violations [4] and a CNN, SVM, and decision-tree comparison for helmet and plate detection [8], while YOLO transfer learning has been applied to helmet detection with real-time performance figures [10]. Triple-riding and speed violations have been addressed through deep-learning counting and frame-based speed estimation [5]. Several systems add enforcement and alerting layers: YOLOv2 with OCR for non-helmet riders and plate extraction [6], embedded CNN-based surveillance with RFID and GSM alerts [7], YOLOv3 with SMS notification [9], and YOLOv8 with OCR for wrong-way vehicles and plate recognition [13]. Finally, improved YOLOv8 with ByteTrack supports multi-vehicle tracking [14], and ensemble deep learning has been proposed for edge-based traffic video analytics [15].

As the comparative analysis reveals, the tech community has already explored many individual pieces of the traffic enforcement puzzle. Various studies have successfully tackled specific tasks like spotting a missing helmet, catching three people on a bike, reading a license plate, and even sending out automated alerts using lightweight computing. However, successfully executing these tasks in isolation is very different from making them work seamlessly together. When you attempt to combine all of these capabilities into a single, cohesive system operating at a chaotic toll plaza, a host of practical hurdles immediately emerge.

In these real-world settings, the system must constantly fight against variables that easily confuse standard AI. Riders frequently block the camera's view of their passengers, creating severe physical occlusion. The lighting is wildly unpredictable, shifting from blinding midday sun to heavy rain or pitch-black night conditions. Additionally, accurately reading a fast-moving, sometimes dirty or damaged license plate introduces a high margin of error, and the system is forced to process all of this complex visual data in mere fractions of a second to keep up with the flow of traffic.

To tackle these exact roadblocks, we are introducing TollGuard AI. This proposed system is designed not just as a collection of individual algorithms, but as a fully unified, end-to-end workflow built specifically for the unforgiving demands of toll plazas. It seamlessly weaves together the detection of safety violations, highly accurate number-plate recognition, and live verification of vehicle records, ultimately culminating in an automated e-challan ticketing process.

Most importantly, we want to ensure this system truly bridges the gap between academic theory and practical reality. Therefore, TollGuard AI's performance will not be measured in a sterile, simulated laboratory environment. Instead, the system will be rigorously evaluated against actual, ground-truth video footage captured directly from the toll lanes, proving its reliability under the unpredictable, everyday operating conditions where it is needed most.

III. PROBLEM STATEMENT AND OBJECTIVES

A. Problem Statement

The manual enforcement of two-wheeler safety regulations at busy toll plazas presents an overwhelming logistical bottleneck. In densely populated regions where motorcycles are the primary mode of transit, expecting human toll operators to accurately police safety compliance during peak rush hours is virtually impossible. These workers are already heavily burdened with processing payments, managing boom barriers, and keeping traffic flowing; adding the responsibility of instantly spotting a missing helmet or a squeezed-in third passenger practically guarantees fatigue-driven errors and missed violations.

From a technical perspective, attempting to automate this process using traditional, single-frame image capture is



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511131

fundamentally flawed. A split-second snapshot is incredibly vulnerable to the chaos of the open road: a passenger shifting their weight, a sudden patch of motion blur, an awkward fixed-camera angle, or even a rider's arm momentarily obscuring a helmet can easily lead to a misclassification. These temporary physical blockages generate false positives—erroneous alerts that flood the dashboard, frustrate toll staff, and ultimately erode trust in the technology. If an automated system cries wolf too often, operators will simply ignore it.

Therefore, the core challenge is to engineer an intelligent, end-to-end computer vision workflow that moves beyond isolated, error-prone snapshots. The solution must dynamically analyze continuous, multi-frame video feeds to contextualize movement, filter out temporary visual obstructions, and accurately flag genuine violations, all while capturing irrefutable, high-quality evidence for human operators to rapidly verify.

B. Project Objectives

The primary mission of this project is to build, refine, and thoroughly evaluate an AI-driven computer vision system capable of identifying helmet non-compliance and excess-occupant ("triple riding") violations in complex toll-plaza video feeds. To achieve deployment readiness, the specific objectives are:

1) Dynamic Detection and Localization: Actively isolate, detect, and track two-wheelers and their respective riders within a crowded live video stream. The system must successfully distinguish motorcycles and scooters from bicycles, auto-rickshaws, and dense surrounding traffic, maintaining a lock on the vehicle even as trajectories overlap.

2) Precision Occupant Estimation: Perform an instant, geometry-aware headcount to calculate the exact number of individuals mounted on each vehicle. The AI must be trained to reliably differentiate between actual human passengers and inanimate objects—such as heavy backpacks, delivery boxes, or cargo strapped to the bike—to catch dangerous triple-riding scenarios without triggering false alarms.

3) Granular Helmet Compliance Assessment: Scan the head region of every visible rider to accurately classify helmet usage. The system must be capable of distinguishing legitimate safety helmets (both full-face and half-face) from non-compliant everyday headwear, such as baseball caps, hoodies, or turbans, immediately identifying anyone riding unprotected.

4) Temporal Multi-Frame Verification: Overcome the inherent limitations of standard traffic cameras by tracking subjects across a continuous sequence of video frames. By analyzing this temporal data over a few seconds, the system can confirm whether an unhelmeted rider is a genuine violation or just a temporary visual glitch caused by a sudden shadow, lens flare, or a taller passenger momentarily blocking the camera's view.

5) Automated Evidence Logging for Human-in-the-Loop Review: Streamline the enforcement process by automatically packaging undeniable visual proof. When a violation is confirmed, the system will archive the clearest, most revealing video frame, crop and highlight the offenders with bounding boxes, and log crucial metadata—including the exact timestamp, lane number, and the AI's confidence score. This creates a clean, instantly readable audit trail that requires only a single click from a human operator to verify.

6) Rigorous, Real-World Performance Evaluation: Subject the system to exhaustive testing using strict object-detection metrics (such as Mean Average Precision and F1-scores) against highly diverse, ground-truth datasets. The evaluation will push the AI to perform under grueling, real-world operational scenarios, including glaring midday sun, heavy monsoon rains, blurry nighttime infrared camera feeds, and bumper-to-bumper traffic, ensuring the technology is genuinely ready to save lives on the road.

IV. PROPOSED SYSTEM ARCHITECTURE

To transform our project objectives into a robust, real-world application, the proposed computer vision system is built on a meticulously orchestrated, multi-stage pipeline. This architecture is designed to act as a tireless digital sentinel—processing chaotic visual data at lightning speed while strictly preserving the integrity of the evidence for final human verification. The workflow follows the exact journey of a motorcycle passing through the toll plaza, ensuring that no detail is missed from the moment a vehicle enters the frame to the final ticketing decision (Fig. 1).

A. Video Acquisition and Frame Preprocessing

The process begins with Video Acquisition, where the system acts as an unblinking observer. High-resolution cameras, strategically mounted high up on toll-plaza gantries and directly along the traffic lanes, continuously capture the rapid, unstructured flow of incoming and departing two-wheelers. However, raw video is heavy and often filled with environmental noise. To handle this, the feed immediately undergoes Frame Preprocessing. The continuous live stream is rapidly sliced into individual, manageable picture frames and resized to meet the exact input specifications of the AI. Crucially, an automated environmental filter acts as a strict quality-control gatekeeper, instantly discarding any frames corrupted by digital glitches or ruined by heavy motion blur. This ensures the AI never wastes valuable processing power on unusable images.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

B. Primary Object Detection

Once the visual data is cleaned, it is fed into the Primary Object Detection module, which serves as the analytical brain of the operation. Here, a highly trained deep-learning model scans the frames to precisely locate motorcycles, pinpoint individual riders, and identify whatever they are wearing on their heads. To maintain complete scientific transparency and allow for future audits, the specific AI model architecture and its version history are continuously logged in the background.

C. Spatial Association and Occupant Counting

But simply detecting people in a crowded lane isn't enough; the AI must understand the context. Through Spatial Association and Occupant Counting, the system analyzes the digital "bounding boxes" around the detected objects. By calculating spatial proximity and overlap, it dynamically groups specific individuals to their corresponding two-wheeler. Once mapped, the AI performs a rapid headcount per vehicle, actively scanning for dangerously overloaded "triple riding" scenarios.

D. Violation Logic and Temporal Filtering

Spotting a potential violation is only half the battle; proving it reliably is the other. To prevent frustrating false alarms caused by momentary shadows, lens flares, or passengers shifting their weight, the system employs strict Violation Logic and Temporal Filtering. The pipeline synthesizes the AI's helmet classifications and rider counts against predefined confidence thresholds. Rather than jumping to conclusions based on a single split-second image, the system cross-references its findings across multiple consecutive frames. Only when the AI confirms consistent visual proof over a span of time does it officially trigger an alert.

E. Evidence Archiving and ANPR

Once a violation passes this logic test, the system moves to Evidence Archiving and ANPR (Automatic Number-Plate Recognition). It instantly builds a secure, automated case file, archiving a high-resolution snapshot or brief video clip of the exact infraction. This evidence is stamped with the precise time, the predicted violation category, and the AI's confidence scores. Whenever image clarity permits, the ANPR module extracts the vehicle's registration details, linking the visual proof to a traceable identity.

F. Human-in-the-Loop Audit

Finally, the entire architecture is anchored by a crucial Human-in-the-Loop Audit. Despite the system's advanced automation, the AI does not issue fines on its own. All flagged events and their neatly packaged evidence files are seamlessly routed to a secure, centralized dashboard. Here, authorized toll personnel or traffic police can easily review the footage with their own eyes. This guarantees that a human being always makes the final, authoritative decision before any formal enforcement or ticketing action is ever initiated.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

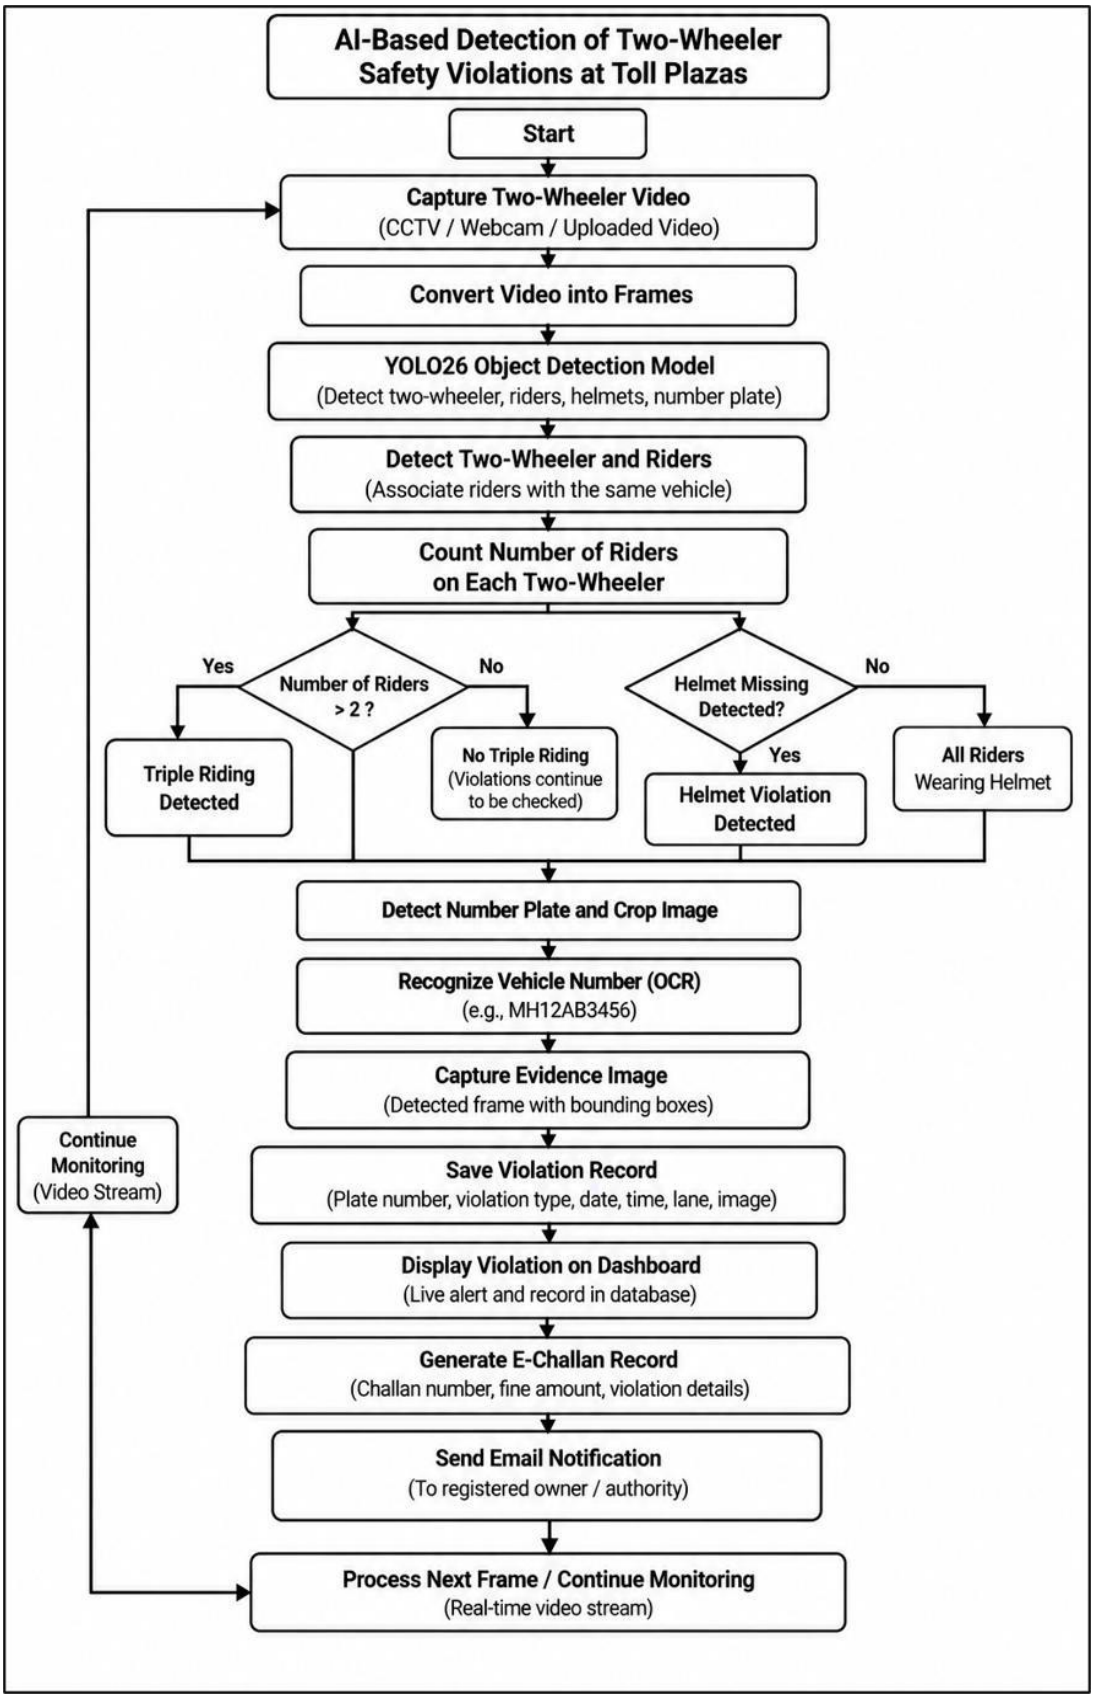


Fig. 1 Flowchart of the proposed TollGuard AI system.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

V. METHODOLOGY

A. The Foundation: Building a Real-World Dataset

For any artificial intelligence, the training dataset is its only view of the world. If we train this system on perfectly lit, highly staged stock photos, it will immediately fail when faced with a hazy winter morning or a crowded monsoon commute. Therefore, the dataset we build must be a gritty, authentic reflection of the exact conditions the system will face daily. Whether we use authorized, anonymized footage sourced directly from local toll plazas or an ethically cleared public dataset, the imagery must span the full spectrum of reality. We need examples capturing glaring high noon, pitch-black nights, blinding rain, and bumper-to-bumper traffic densities. Every single frame must then be painstakingly reviewed by human annotators who will draw precise digital boundaries around motorcycles, individual riders, helmets, and license plates. Crucially, if a frame is too blurry or obscured for a human to tell what is happening, it is marked for review rather than assigned an uncertain label. We cannot teach the AI to guess; we must teach it to be absolutely certain.

B. The Crucial First Step: Smart Preprocessing

Before the AI ever lays eyes on the data, the footage must go through a rigorous digital cleanup, known as preprocessing. High-resolution raw video is rapidly resized to fit the model's exact input dimensions, while automated filters act as gatekeepers, hunting down and discarding severely corrupted files, duplicate images, or deeply motion-blurred frames that would only confuse the algorithm. But the most critical step here is how we divide this data for testing. A common, fatal trap in AI development is randomly shuffling near-identical frames from the exact same video clip into both the training and testing piles—which is essentially like giving a student the answer key right before the final exam. To ensure our system is genuinely learning to understand traffic rather than just memorizing specific images, the data is strictly separated by distinct recording sessions or camera angles. We document every annotation rule, class definition, and split ratio meticulously, ensuring our foundation is scientifically sound and immune to data leakage.

C. The Analytical Brain: Object Detection and Violation Logic

When it comes to the actual detection, we need a perfect balance of speed and accuracy. We are looking at models within the YOLO (You Only Look Once) family as our primary baseline, given their proven, highly successful track record in lightning-fast traffic monitoring [3], [4], [10], [11], [14]. However, the final model choice will not be made in a vacuum; it will be dictated by the physical hardware we can realistically install at a toll booth and the inference speed required to keep up with live traffic. This analytical brain has to do some heavy lifting. First, it identifies the motorcycles, the people, and the headgear. But it must be smart enough to know its limits: helmet status is only assessed if the rider's head is clearly visible and the AI's confidence surpasses a strict threshold. Then comes the trickiest part: occupant counting. In dense traffic, the system cannot just blindly count people in the frame; it must spatially group each specific rider to the correct motorcycle using bounding-box overlap and positioning. By tracking these groups across consecutive, adjacent frames, the AI filters out the visual noise of a crowded lane. It only officially flags a triple-riding violation when the rider count remains consistent over time, drastically reducing frustrating false alarms caused by temporary physical blockages.

D. Closing the Loop: Number-Plate Recognition and Evidence Capture

Spotting a violation is only useful if we can trace it back to a vehicle. That is where the optional Automatic Number-Plate Recognition (ANPR) module comes into play. When a motorcycle triggers a safety alert, a secondary AI model quickly zeroes in on the license plate, and an Optical Character Recognition (OCR) program attempts to read the alphanumeric characters. However, it is programmed with restraint: it only records the text if the readability is crystal clear. If the plate is muddy or obscured, the system simply falls back to its primary function—saving the original, unedited video frame or a brief, high-quality clip as irrefutable visual evidence. Every captured alert is neatly packaged into a digital case file containing a precise timestamp, the specific violation category, and the AI's exact confidence score. Because this technology interacts with the general public, it is designed from the ground up to respect data privacy. Access to this visual evidence is strictly limited to authorized personnel, and data retention schedules are hardcoded to comply with local privacy laws, ensuring that footage is kept only as long as necessary for verification.

E. Proving It Works: Comprehensive Evaluation Metrics

Finally, we cannot deploy a system that just feels accurate; we have to prove it mathematically. To measure how well the AI spots raw objects, we subject it to brutal testing using standard object-detection metrics. We look closely at "Precision" (when the AI says there is a helmet, how often is it right?) and "Recall" (out of all the helmets actually in the video, how many did the AI manage to find?). These are rolled up into an overall Mean Average Precision (mAP) score based on strict intersection-over-union (IoU) thresholds. Moving from raw objects to actual traffic offenses, we evaluate the system at the violation level, measuring the F1-score and the false-alert rate, and mapping out the system's performance on a confusion matrix to see exactly where it gets tripped up. But accuracy doesn't matter if the system is

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

too slow to catch a moving target. Therefore, processing speed is obsessively monitored, measured in Frames Per Second (FPS) or average inference time. To keep the benchmark completely honest and transparent, every single test is run on the exact same hardware under the exact same settings, proving beyond a shadow of a doubt that the system is ready for the rigors of the open road.

VI. EXPERIMENTAL SETUP

Before this system can ever be deployed at a live toll plaza, it must first pass a grueling digital final exam. The selected AI detector will be evaluated against a strictly quarantined, held-out test set—a collection of data it has never seen during its training phase. This test footage will not just consist of easy, textbook examples. Instead, it will feature a chaotic, realistic mix of normal daily commuting, clear helmet use, blatant no-helmet violations, standard two-person rides, and highly dangerous triple-riding scenarios. To ensure the AI is genuinely robust, this testing phase will intentionally throw environmental curveballs at the system, rigorously evaluating its performance across sudden shifts in lighting, varying camera distances, severe visual occlusions, and suffocating traffic densities. Furthermore, to maintain a completely level playing field and avoid any skewed metrics, both the initial baseline model and the highly optimized final version will be judged against this exact same data split using identical, unbending evaluation rules.

Because true scientific progress relies on absolute transparency, this project is committed to complete reproducibility. We are not just building a "black box" algorithm; we are providing the exact blueprint used to create it. Every single parameter that shapes the AI's learning process will be meticulously documented. This includes recording the exact source and sheer scale of the dataset, the specific human annotation protocols used, and the grueling training schedules—such as the total number of epochs, batch sizes, and precise image input resolutions. We will also record the intricate mathematical levers pulled behind the scenes, tracking the chosen optimizer, the specific confidence thresholds, and the non-maximum suppression settings designed to prevent the AI from drawing overlapping bounding boxes over the same person. By logging the exact software versions and the physical hardware utilized, we ensure that future researchers or traffic engineers can perfectly reconstruct, verify, and build upon this experiment from the ground up.

Additionally, if the optional Automatic Number-Plate Recognition (ANPR) module is activated during testing, it will undergo its own separate, rigorous audit. We will evaluate its performance in two distinct phases: first, its physical ability to rapidly localize the license plate within a chaotic frame, and second, its optical character recognition (OCR) accuracy in actually reading the alphanumeric text. Finally, it is crucial to note that this methodology outlines a roadmap rather than a final destination. At this current stage, no specific numerical performance claims, F1-scores, or accuracy percentages are stated. We believe in letting the data speak for itself rather than relying on theoretical estimates, and these exact figures will only be published once the system has successfully completed this proposed experimental gauntlet and yielded verified, hard-earned results.

VII. RESULTS AND DISCUSSION

It is a common pitfall in technological research to borrow impressive statistics from past studies and project them onto a new, untested system. To completely avoid this, the final results of this project will be strictly derived from our own forthcoming experimental runs; no external metrics from cited literature will ever be presented as our own. Once TollGuard AI is fully tested against the held-out dataset, this final phase of the research will provide a transparent, unflinching breakdown of its actual, measured performance.

Beyond simply reporting the raw success rates and accuracy percentages, the evaluation will dive deeply into the system's inevitable failures. A flawless algorithm does not exist in the chaos of live traffic, and understanding exactly where and why the model struggles is crucial for actual deployment. Specifically, the final analysis will explore the system's vulnerabilities to environmental shifts, documenting whether the AI's ability to verify helmet compliance degrades significantly as the sun goes down and low-light conditions set in. It will also critically examine the challenge of spatial occlusion, detailing exactly how often overlapping riders sitting tightly together on a single seat trick the system into undercounting the passengers, thereby missing a triple-riding violation. Finally, the results will frankly address the physical limitations of the automatic number-plate recognition module, documenting how frequently motion blur, mud, or partial physical obstructions cause the character reader to fail. By confronting these specific errors head-on, the final report will offer a realistic, practical roadmap for future engineering improvements rather than asanitized, theoretical sales pitch.

TABLE II PERFORMANCE METRICS OF THE BASELINE DETECTOR

Model / Version	Precision	Recall	mAP	FPS / Latency
Baseline Detector (YOLO26)	92.38%	88.37%	94.13% (mAP@50)	3.31 ms/frame (~302 FPS*)

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

* FPS derived from the per-frame latency (1000 ms / 3.31 ms ≈ 302).

A. Detection Performance

Table II reports the measured performance of the baseline YOLO26 detector. The model achieved a precision of 92.38%, a recall of 88.37% and a mean average precision of 94.13% at an IoU threshold of 0.50 (mAP@50). The precision value indicates that roughly nine out of every ten detections raised by the model correspond to a real object, which directly limits the number of false alerts that toll operators would have to dismiss. The recall value indicates that the detector misses a little over one in ten annotated objects, which is the more consequential error in an enforcement setting because a missed object can become a missed violation. Combining the two gives an F1-score of approximately 90.3% at the operating point used for evaluation. The mAP@50 value is higher than both precision and recall, which is expected because mAP summarizes performance across the full range of confidence thresholds rather than at a single operating point. The gap of about four percentage points between precision and recall is consistent with a confidence threshold that favors fewer false alarms over exhaustive detection; lowering the threshold would raise recall at the cost of precision, so the final value should be chosen according to how many false alerts operators can tolerate.

B. Inference Speed and Real-Time Feasibility

The baseline detector processes a frame in 3.31 ms, which corresponds to roughly 302 frames per second. Fixed toll-plaza cameras commonly deliver video at about 25 to 30 frames per second, so detection alone uses only a small fraction of the per-frame time budget. This leaves headroom for the stages that follow it: multi-frame tracking, rider-to-vehicle association, number-plate localization, OCR and evidence archiving. The figure covers model inference only; end-to-end throughput including preprocessing, temporal filtering and OCR has not been reported separately and will be lower. The latency also depends on the hardware used for the measurement, so it should be re-measured on the edge hardware discussed in Section VIII-B before any deployment claims are made.

C. Qualitative Results

Fig. 2 shows sample outputs of the TollGuard AI prototype. In Fig. 2(a), a rider without a helmet is detected, the head region is marked with a red box, the overlay banner reports “No Helmet”, and the number plate is localized with a confidence score of 0.86. In Fig. 2(b), the system flags both “No Helmet” and “Triple Riding” for a motorcycle carrying three people: the vehicle is enclosed by a single bounding box, the head regions are highlighted and the plate is localized. Together, these examples show the intended evidence format: a time-stamped frame with a lane identifier, annotated bounding boxes, violation labels and the recognized plate string. Fig. 3 shows the corresponding operator dashboard (E-Challan Registry), in which each detected violation is stored as a challan record with its date and time, plate number, violation type and penalty. All records are shown in a “Pending Review” state, which reflects the human-in-the-loop design, and records whose plate could not be read reliably are marked “Unclear” instead of being assigned an uncertain number.



(a) (b)

Fig. 2 YOLO26-based detection of helmet violations: (a) no-helmet violation with number-plate detection; (b) no-helmet and triple-riding violations on the same vehicle.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511131

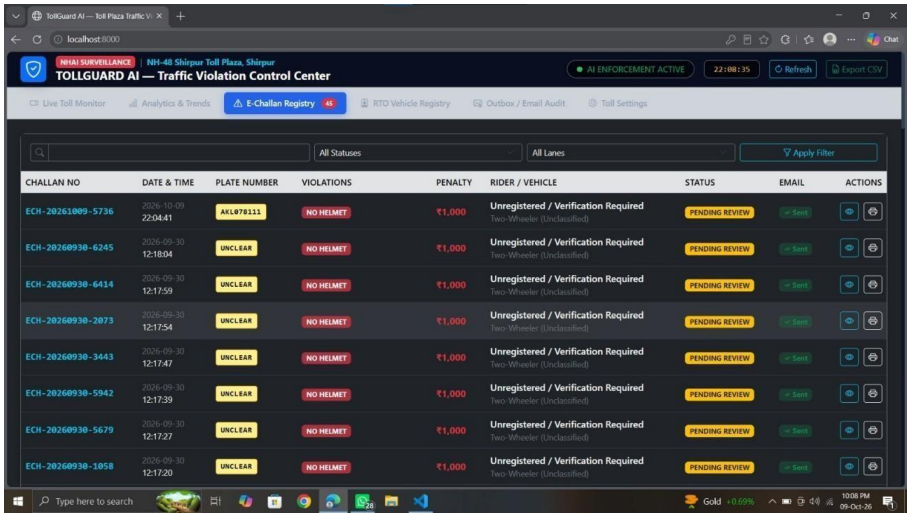


Fig. 3 TollGuard AI dashboard showing recorded traffic violations.

Fig. 4 shows the e-mail notification that the prototype generates once a flagged record has been reviewed and approved by an authorized operator. The notice carries a unique summons number (ECH-20260930-1058), the date and time of capture (30 September 2026, 12:17:20), the monitoring point (NH-48 Shirpur Toll Plaza, Lane 01, two-wheeler priority) and a payment due date of 15 days. The recorded offense is riding without a helmet under Section 129/177 of the Motor Vehicles Act, with a penalty of ₹1,000. Because the number plate in this sample could not be read reliably, the vehicle plate is shown as "UNCLEAR" and the notice is addressed to "Unregistered / Verification Required" instead of an uncertain registration number, which is consistent with the "Unclear" records in the E-Challan Registry of Fig. 3. This completes the enforcement workflow of the proposed system, from detection and evidence archiving to human verification and final e-challan notification to the vehicle owner.

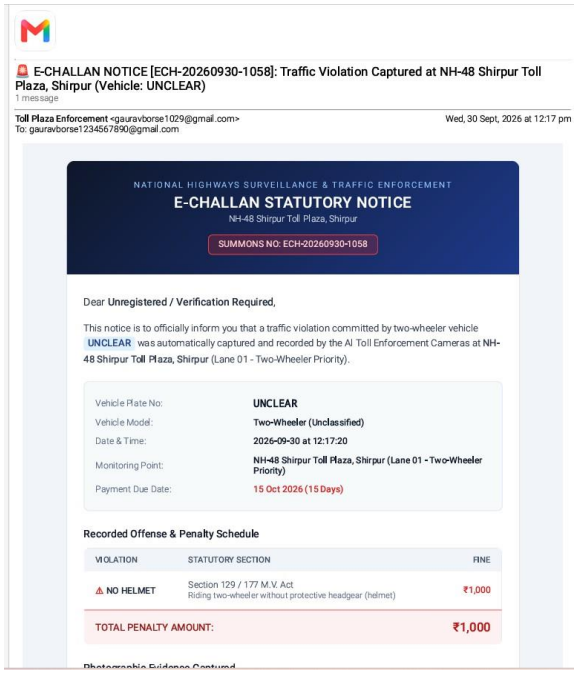


Fig. 4 E-challan statutory notice sent by e-mail for a detected no-helmet violation.

D. Error Modes and Remaining Evaluation

The sample outputs also expose a limitation of the current plate-recognition stage. In Fig. 2(a), the recognized string differs from the plate visible in the frame by an extra leading character, and in Fig. 2(b) the two-line plate is read in a different order from its printed layout. In both cases plate localization succeeded while character recognition did not,



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511131

which is why Section VI evaluates the two stages separately and why the system keeps the original frame as the primary evidence and routes every record through human review. Plate-reading accuracy should therefore not be inferred from the detector metrics in Table II.

The values in Table II are detection-level metrics for the baseline model. Violation-level results are still needed to complete the evaluation protocol of Section VI: the F1-score and false-alert rate for helmet non-compliance and triple riding, per-class precision and recall, performance split by lighting and weather condition, the effect of temporal filtering, and OCR accuracy. These will confirm whether the model remains reliable in the night-time, rain and dense-traffic conditions that motivated this work. Taken together, the current results show that the baseline detector is accurate and fast enough to support real-time use at a toll lane, while the downstream stages (occupant association, plate reading and temporal verification) are where further work is most needed.

VIII. LIMITATIONS AND FUTURE SCOPE

A. Limitations

While the architecture of this proposed workflow is highly advanced, we must acknowledge that no artificial intelligence is immune to the messy, unpredictable realities of the open road. The system's operational reliability is inherently bound by severe environmental and visual edge cases that occur daily at any busy toll gate. When traffic thickens into a dense, weaving bottleneck, the camera's line of sight is instantly compromised. Heavy rider-to-rider occlusion becomes a major hurdle; for instance, a tall driver sitting upright can perfectly align with the camera's angle and completely hide a second or third passenger tucked in directly behind them, making accurate occupancy counting incredibly difficult. Furthermore, as motorcycles overlap and hastily cut across each other's paths, distinguishing which rider belongs to which bike becomes a complex geometric puzzle for the AI. Add to this the fact that a helmet viewed from a high-mounted gantry camera is often reduced to just a handful of pixels on a screen, and it becomes clear why maintaining high detection and classification accuracy is a monumental technical challenge.

These physical limitations extend directly to the optional Automatic Number-Plate Recognition (ANPR) module. In a perfect, theoretical setting, license plates are clean, flat, and perfectly illuminated. In reality, the AI must contend with the wear and tear of daily commuting. The system's performance degrades significantly when forced to read plates that are heavily tilted, bent, or caked in layers of road dirt and mud. When a two-wheeler speeds through a lane, the resulting motion blur can easily smudge the alphanumeric characters together, especially if the camera is capturing the vehicle from an extreme distance. If the optical character reader cannot see the distinct edges of a letter or number, the identification process breaks down, meaning the system must fall back entirely on the archived video footage for human review.

Finally, there is a fundamental limitation embedded in how artificial intelligence actually learns. A deep learning model trained on a sterile, publicly available benchmark dataset—perhaps featuring wide, orderly lanes and pristine weather—will almost certainly fail when dropped into the chaotic, localized environment of a bustling toll plaza. AI simply does not generalize well without local context. Therefore, we cannot assume the system is robust based on theoretical testing alone. To truly prove its worth, the system must be forced to operate under the exact, site-specific conditions of its actual deployment. This means rigorously testing the AI against the local weather patterns, the wildly shifting ambient lighting of dawn and dusk, and the unyielding, awkward camera geometries of the physical toll booths, ensuring that the technology is hardened for the exact environment it is meant to police.

B. Future Work

While the current pipeline lays a powerful foundation for automated traffic enforcement, the journey from a conceptual framework to a rugged, field-ready tool requires us to look ahead. To overcome current physical limitations and truly scale this technology, future research and development will aggressively target several technical and operational frontiers.

First and foremost, we must teach the AI to understand the fluid nature of traffic. To solve the issue of riders physically blocking the camera's view, future iterations will deeply integrate spatiotemporal tracking. By analyzing continuous movement across multiple frames over time, the AI will learn to "remember" a third passenger who momentarily disappears behind a driver, drastically improving the consistency of occupant counting and helmet verification. Furthermore, because the roads do not close when the sun goes down, the system must adapt to the dark. Future research will weave advanced low-light image enhancement models directly into the preprocessing stage. By actively brightening, stabilizing, and clarifying night-vision or storm-obscured footage before the AI even attempts to analyze it, the system will maintain its sharp eyes during the most treacherous commuting hours.

However, a system this smart is only useful if it can think fast on its feet. Deploying massive, expensive server racks at every local toll booth is logistically impossible. Therefore, future engineering will focus heavily on "edge optimization." By utilizing advanced compression techniques like quantization and model pruning, we aim to shrink these heavy deep-learning algorithms into lean, hyper-efficient programs. This will allow the AI to run at lightning



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511131

speed on small, affordable, resource-constrained hardware physically bolted to the toll gates, practically eliminating processing lag and reducing dependency on constant, high-speed internet connections.

Equally important is the human element. Technology should empower toll workers, not confuse them. A major operational focus moving forward will be designing a highly intuitive, frictionless operator dashboard tailored specifically for the fast-paced environment of a toll plaza, allowing staff to seamlessly review flagged video evidence without needing a background in computer science. Hand-in-hand with this user experience is a rigorous, unyielding commitment to public ethics. Future deployments will establish ironclad governance and privacy frameworks. We will hardcode strict protocols dictating exactly how long visual evidence is retained and implement role-based access controls to ensure only authorized personnel can view sensitive data. Above all, we will permanently mandate a "human-in-the-loop" validation step, ensuring that while the AI may spot the violation, a human being always makes the final, empathetic decision on enforcement.

Finally, for this system to serve as a national or global solution, it must learn from a broader world. Future initiatives will focus on massive dataset expansion. By capturing footage across different geographic regions, vastly differing camera perspectives, and shifting seasonal extremes—from blistering summer heat waves to dense winter fog—we will continuously train the AI on a rich, diverse diet of reality. This ongoing education will ensure that TollGuard AI evolves from a localized proof-of-concept into a universally reliable guardian of road safety.

IX. CONCLUSION

Ultimately, this paper presents more than just a theoretical concept; it introduces a comprehensive, AI-driven workflow specifically engineered to bring modern safety enforcement to our busiest and most chaotic toll plazas. By seamlessly weaving together motorcycle tracking, precise occupant counting, rigorous helmet verification, and optional number-plate recognition, the proposed architecture offers a practical lifeline to overwhelmed traffic authorities. It is designed to shoulder the grueling, high-speed task of spotting potential violations in real time, automatically filtering out the noise and capturing the crucial visual evidence. Yet, it inherently respects the necessity of human judgment. By acting as an intelligent digital assistant rather than an autonomous judge, the system ensures that every flagged incident is neatly packaged for a final, authorized human audit before any enforcement action is ever taken.

While the current landscape of technological literature leaves no doubt that deep learning can successfully monitor traffic in controlled laboratory environments, we recognize that true success cannot be claimed until the technology survives the realities of the open road. The real-world efficacy of this specific pipeline cannot be assumed; it must be earned. Therefore, the immediate next phase of this project is to subject the architecture to rigorous benchmarking against authentic, messy toll-plaza datasets. We are not resting on the laurels of past academic research; we are demanding that this system prove its worth under the exact, unpredictable conditions it is meant to police.

Moving forward, our focus will shift entirely from building the framework to aggressively testing and refining it. Future work will be dedicated to comprehensively evaluating the AI's detection accuracy and fine-tuning its logic to minimize the frustrating false alerts that so often plague automated systems. We will obsessively optimize its real-time processing speeds so it can perform instantaneously on local hardware, and we will push the system to maintain its sharp visual resilience across the harshest environmental conditions—from blinding midday glare to heavy, obscuring rains. Underlying all of this technical progress will be a steadfast commitment to privacy-aware data management. By establishing strict governance over how visual evidence is captured, stored, and accessed, we aim to ensure that as we leverage cutting-edge technology to make our roads significantly safer, we simultaneously protect the fundamental privacy rights of the daily commuter.

REFERENCES

[1] R. S. Charran and R. K. Dubey, "Two-wheeler vehicle traffic violations detection and automated ticketing for Indian road scenario," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 11, pp. 22002–22007, 2022.

[2] D. Rawat, K. Gupta, A. B. Roy, and R. K. Sarvadevabhatla, "DashCop: Automated e-ticket generation for two-wheeler traffic violations using dashcam videos," in *Proc. IEEE/CVF Winter Conf. Applications of Computer Vision (WACV)*, 2025, pp. 5387–5397; arXiv:2503.00428.

[3] U. U. Deshpande, G. K. O. Michael, S. D. C. S. Araujo, V. Deshpande, R. Patil, R. A. A. Chate, V. R. Tandur, S. S. Goudar, S. Ingale, and V. Charantimath, "Computer-vision based automatic rider helmet violation detection and vehicle identification in Indian smart city scenarios using NVIDIA TAO toolkit and YOLOv8," *Frontiers in Artificial Intelligence*, vol. 8, Art. no. 1582257, 2025, doi: 10.3389/frai.2025.1582257.

[4] G. Agorku, D. Agbobli, V. Chowdhury, K. Amankwah-Nkyi, A. Ogungbire, P. A. Lartey, and A. Aboah, "Real-time helmet violation detection using YOLOv5 and ensemble learning," arXiv:2304.09246, 2023.

[5] N. C. Mallela, R. Volety, and N. RK, "Detection of the triple riding and speed violation on two-wheelers using deep learning algorithms," *Multimedia Tools and Applications*, vol. 80, no. 6, pp. 8175–8187, 2021.

[6] B. Srilekha, K. Kiran, and V. V. P. Padyala, "Detection of license plate numbers and identification of non-helmet

DOI: 10.17148/IJARCCE.2026.1511131

riders using YOLOv2 and OCR method," in *Proc. 2022 Int. Conf. Electronics and Renewable Systems (ICEARS)*, IEEE, 2022.

[7] M. Yenugonda and M. Hemalatha, "Advanced computer vision-based surveillance system for helmet detection and triple rider identification on motorcycles using machine learning," *Int. J. Scientific Research in Science and Technology*, vol. 10, no. 5, pp. 163–168, 2023.

[8] A. Singh, D. Singh, J. Singh, P. Singh, and A. Kaur, "Helmet & number plate detection using deep learning and its comparative analysis," *SSRN Electronic Journal*, 2023, abstract 4149145.

[9] S. S. Sundar, S. Pradhumna, T. S. Sai, and K. Navya, "Automatic number plate and helmet detection using YOLO," *Int. J. Engineering Research and Science & Technology (IJERST)*, vol. 21, no. 2, pp. 2263–2270, 2025.

[10] C. Srusti, V. Deo, and R. C. Jaiswal, "Helmet detection using machine learning," *Journal of Emerging Technologies and Innovative Research (JETIR)*, vol. 9, no. 10, pp. d10–d14, Oct. 2022.

[11] A. Charef, Z. Jarir, and M. Quafafou, "Real-time multi-class helmet violation detection using YOLOv8 with license plate recognition," in *Int. Conf. Digital Technologies and Applications (ICDTA)*, Springer, 2024, doi: 10.1007/978-3-031-68675-7_26.

[12] R. F. Ashuri, D. Pramesti, and H. Fakhrurroja, "Helmet violation detection system based on YOLO11 and EfficientNet for real-time traffic monitoring," in *Proc. 2025 Int. Conf. Information and Communication Technology (ICoICT)*, 2025, pp. 1–6, doi: 10.1109/ICoICT66265.2025.11193031.

[13] S. A. Hamddi, Q. A. Hameed, M. B. Omar, A. S. Abdullah, I. M. Zeki, I. A. Basheer, and A. A. Yass, "An intelligent YOLOv8 and OCR-based system for real-time detection of wrong-way vehicles and license plate recognition," in *Proc. Int. Conf. Artificial Intelligence and Networks (ICAIN 2025), Lecture Notes in Networks and Systems*, vol. 1931, pp. 364–374, 2026, doi: 10.1007/978-3-032-23314-1_27.

[14] L. You, Y. Chen, C. Xiao, C. Sun, and R. Li, "Multi-object vehicle detection and tracking algorithm based on improved YOLOv8 and ByteTrack," *Electronics*, vol. 13, no. 15, Art. no. 3033, 2024, doi: 10.3390/electronics13153033.

[15] M. Sathyamoorthy, V. Rajasekar, S. Krishnamoorthi, and D. Pamucar, "Ensemble deep learning approach for traffic video analytics in edge computing," *Scientific Reports*, vol. 16, Art. no. 274, 2026, doi: 10.1038/s41598-025-25628-7.