International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

AI-Based Smart Traffic Enforcement System with Behavior-Aware Dynamic Fine Recommendation and Accident Detection

Ms. Ashwini G B
Assistant Professor
Computer Scienceand Engineering
New Horizon College of Engineering
Bangalore, Karnataka, India

Bobburi Jhansi
Computer Science and Engineering
New Horizon College of Engineering
Bangalore, Karnataka, India

Bollapu Sahithi
Computer Science and Enginnering
New Horizon College of Engineering
Bangalore, Karnataka, India

A. Nitish Kumar Reddy
Computer Science and Engineering
New Horizon college of Engineering
Bangalore, Karnataka, India

A. Neeraja
Computer Science and Engineering
New Horizon college of Engineering
Bangalore, Karnataka, India

Abstract—The rapid increase in vehicle traffic has created significant challenges in traffic management, including congestion, frequent violations, and rising accident rates. Traditional enforcement systems rely on manual monitoring and fixed penalty rules, often leading to inefficiency and a lack of flexibility. This paper presents an AI-based smart traffic enforcement system that uses computer vision techniques for real-time monitoring and analysis. The system detects vehicles, tracks their movement, and identifies traffic violations such as speeding and ignoring signals. It also detects accidents by analyzing unusual motion patterns. License plate recognition is used to identify vehicles and maintain historical records. A driver behavior analysis module evaluates violation patterns and calculates a risk score for each vehicle. Based on this score, a dynamic fine recommendation system assigns penalties based on the severity and frequency of violations. A key addition in this work is the Automated Accident Evidence Video Recording and Police Control Room Report Sharing Module. Upon detecting a collision, the system autonomously captures a 60-second evidence video comprising 30 seconds of pre-crash and 30 seconds of post-crash footage, generates a structured Accident Incident Report containing GPS coordinates, vehicle identification data, severity classification, and an AI-generated narrative summary, and securely transmits this evidence package to the nearest Police Control Room in real-time via MQTT and WebSocket protocols with cryptographic chain-of-custody verification. Additionally, an emergency alert system notifies authorities in case of accidents. The proposed system increases monitoring efficiency, reduces manual effort, and ensures fair, flexible, and accountable traffic enforcement.

Keywords - Artificial Intelligence, Traffic Monitoring, Computer Vision, Accident Detection, License Plate Recognition, Risk Score, Dynamic Fine System, Evidence Video Recording, Police Control Room Integration, MQTT, IoT Emergency Response

I. INTRODUCTION

Traffic congestion and road accidents have become serious concerns due to the rapid increase in the number of vehicles. In many urban areas, traditional traffic management systems struggle to handle large volumes of traffic efficiently. Manual monitoring is limited in scalability and often leads to delays in detecting violations and responding to accidents. Traffic personnel cannot continuously monitor multiple locations, which reduces the effectiveness of enforcement. In addition, fixed penalty systems fail to consider driver behaviour and violation history, resulting in inconsistent and sometimes unfair enforcement. With the advancement of Artificial Intelligence and computer vision technologies, automated traffic monitoring systems have emerged as a promising solution. These systems can process video data in real time and detect vehicles with high accuracy. Object detection models such as YOLO have been widely used for real-time vehicle detection due to their speed and efficiency [1].

Computer vision techniques enable object detection, tracking, and analysis of vehicle movement, making it possible to identify violations such as over speeding and signal violations. Despite these improvements, most existing systems focus on individual tasks such as detection or monitoring and do not provide a complete enforcement framework [2].

Another limitation of current approaches is the lack of behavioural analysis in traffic enforcement. Most systems treat each violation independently without considering the historical behaviour of drivers. This leads to a one-size-fits-all penalty system, where all drivers are treated equally regardless of their driving patterns. As a result, repeat offenders are not penalized effectively, and minor violations are sometimes over-penalized. Recent studies suggest that incorporating behavioural data and historical patterns can improve decision-making in traffic systems [3].

Furthermore, a critical gap exists in how accident evidence is handled after detection. In most current deployments, surveillance footage surrounding a collision resides on continuously overwritten local storage, making retrieval a manual, time-sensitive process with substantial risk of permanent data loss. Even when evidence is successfully recovered, there is no standardized mechanism to automatically compile, authenticate, and transmit it alongside a structured incident report to law enforcement agencies in real-time. Police

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

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

control rooms continue to rely on civilian phone calls and delayed manual reporting for situational awareness, creating a significant temporal gap between incident occurrence and coordinated response [4].

To address these limitations, this paper proposes an AI-based smart traffic enforcement system that combines detection, analysis, and adaptive decision-making into a unified framework. The system integrates multiple components, including vehicle detection, vehicle tracking, license plate recognition, traffic violation detection, accident detection, driver behaviour analysis, risk scoring, and dynamic fine recommendation. A key contribution of the proposed system is the inclusion of a behaviour-based dynamic fine recommendation mechanism. The system evaluates violation patterns over time and computes a risk score for each vehicle. Based on this score, penalties are assigned adaptively, considering factors such as violation severity, frequency, and historical behaviour. Similar adaptive and data-driven approaches have shown improved efficiency in intelligent transportation systems [5].

An important novel addition in this work is the Automated Accident Evidence Video Recording and Police Control Room Report Sharing Module. Upon detecting a collision, the system autonomously captures a consolidated 60-second evidence video encompassing 30 seconds before and 30 seconds after the crash event. It simultaneously generates a structured Accident Incident Report (AIR) containing GPS coordinates, ANPR-identified vehicle data, AI-estimated severity classification, and a machine-generated narrative summary. This evidence package is cryptographically sealed using SHA-256 hashing to establish chain-of-custody integrity and is securely transmitted to the nearest Police Control Room in real-time via MQTT and WebSocket protocols. This automated pipeline reduces the reporting latency from several minutes under manual workflows to under 3 seconds, fundamentally transforming emergency coordination capabilities.

In addition, the system incorporates an emergency alert mechanism that generates notifications in case of accidents. AI-based accident detection and alert systems have been shown to reduce response time and improve emergency handling [6].

By integrating monitoring, analysis, evidence preservation, police notification, and response into a single framework, the proposed system aims to improve traffic management efficiency and enhance road safety. The remainder of this paper is organized as follows. Section II presents the literature review. Section III describes the proposed system. Section IV explains the methodology. Section V discusses the system architecture. Section VI presents the implementation details. Section VII discusses the results, followed by the conclusion in Section VIII.

II. LITERATURE REVIEW

Research in intelligent transportation systems has examined various methods to improve traffic monitoring, management, and safety. Early traffic management systems mainly focused on optimizing traffic flow by using sensors, loop detectors, and basic image processing. These systems aimed to reduce congestion by adjusting traffic signals and monitoring vehicle density. While they did help improve traffic flow, these approaches had limitations in detecting traffic violations or analysing accidents in real time. They also relied heavily on fixed infrastructure and could not adapt to changing road conditions. With advancements in computer vision and machine learning, more advanced techniques for traffic analysis have emerged. Object detection models, especially those based on deep learning, have greatly improved the accuracy and efficiency of vehicle detection. Models like YOLO (You Only Look Once) are popular because they achieve real-time detection with high accuracy and low computational cost [1].

These models can identify various vehicle types in complex traffic situations and serve as a basis for further evaluations like tracking and monitoring behaviour. Additionally, tracking algorithms have been developed to follow vehicle movements across frames, enabling systems to examine trajectories and spot unusual patterns. Computer vision techniques are also widely used for detecting traffic violations. Methods have been created to identify issues like speeding, running red lights, and poor driving habits by analysing vehicle speed, position, and movement patterns. License plate recognition systems use image processing and optical character recognition to identify vehicles and connect them to enforcement databases. These systems are crucial for automating traffic law enforcement by facilitating vehicle identification and managing records. Accident detection has become an important focus in recent years. Machine learning and deep learning models analyse video data to identify unusual events like collisions, sudden stops, and erratic movements. Studies indicate that deep learning methods can effectively differentiate between normal and abnormal traffic behaviour, enhancing the accuracy of accident detection [2], [3].

Some research has also looked into predicting accidents using historical data and environmental factors, allowing for proactive traffic management [4].

Furthermore, AI and IoT-based systems have been suggested to automatically send alerts and notify emergency services in case of accidents, which can reduce response time and improve rescue operations [5].

Recent advances in IoT messaging protocols have demonstrated the viability of lightweight publish-subscribe frameworks such as MQTT for achieving low-latency communication in smart-city emergency response architectures [7]. Concurrently, the importance of maintaining digital evidence integrity for legal admissibility has been recognized, with cryptographic hashing and blockchain-based audit trails proposed to establish verifiable chains of custody for digital media in forensic contexts [8]. However, no existing system integrates automated pre- and post-crash video buffering, structured incident report generation, and real-time secure dispatch to law enforcement command centers within a single unified framework. This gap motivates the evidence recording and police notification module introduced in the present work.

Despite these developments, most existing systems concentrate on individual components rather than offering a unified solution. For instance, some systems focus solely on vehicle detection, while others work independently on accident detection or license plate recognition. This lack of integration reduces the overall effectiveness of traffic enforcement systems. Additionally, traditional enforcement relies on fixed penalty rules that do not factor in the severity or frequency of violations. This shortcoming prevents these systems from addressing the differences in driver behaviour and does not effectively deter repeated offense. The proposed system aims to overcome these limitations by integrating vehicle detection, tracking, license plate recognition, violation detection, accident detection, driver

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

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

behaviour analysis, dynamic fine recommendations, and automated accident evidence video recording with police control room notification into a single framework. Unlike traditional systems, it accounts for both real-time events and historical behaviour, and provides automated forensic evidence preservation and law enforcement communication, allowing for adaptive and accountable enforcement. This integrated approach improves system efficiency, enhances fairness in penalty assignment, and contributes to safer road conditions.

III. PROPOSED SYSTEM

The proposed system is designed as a comprehensive, integrated framework for intelligent traffic enforcement. It takes video input from surveillance cameras and continuously monitors all vehicular activity in real time. The primary objective is to automate the enforcement process, reducing the need for manual observation and minimizing human errors in traffic monitoring and incident response.

The system begins by detecting vehicles using deep learning-based object detection models. Once vehicles are identified, they are tracked across consecutive frames, with each vehicle assigned a unique identifier. This enables the system to follow vehicle trajectories, estimate speeds, and analyze movement patterns over time. License plate recognition is then applied to extract vehicle registration information using optical character recognition techniques. This information is linked to a historical database that records prior violations and driving behavior for each identified vehicle.

For violation detection, the system monitors vehicle movements against predefined traffic rules. Infractions such as speeding, signal violations, and lane discipline breaches are flagged based on movement analysis and threshold-based logic. Accident detection operates by identifying abnormal motion patterns, including sudden stops, abrupt trajectory changes, and collision-indicative kinematic signatures. These abnormal patterns trigger further verification before an incident is confirmed.

A distinguishing feature of the proposed system is the driver behavior analysis module. It examines the historical violation record of each vehicle and computes a risk score that reflects driving behavior over time. This score differentiates between occasional errors and persistent unsafe driving patterns. The dynamic fine recommendation module then uses this risk score, along with violation severity and frequency, to assign penalties adaptively. This approach ensures that enforcement is proportionate and context-aware, moving beyond rigid fixed-penalty structures.

A critical novel addition to the proposed system is the Automated Accident Evidence Video Recording and Police Control Room Report Sharing Module. This module activates immediately upon confirmed accident detection and performs the following automated sequence: First, it extracts a 60-second evidence video clip from the continuously maintained circular buffer, comprising 30 seconds of pre-crash footage and 30 seconds of post-crash footage, capturing the full context of the incident. Second, it generates a structured Accident Incident Report (AIR) in JSON format containing the precise UTC timestamp, GPS coordinates of the incident location, ANPR-identified registration data for all involved vehicles, an AI-estimated severity classification (Minor, Moderate, or Severe based on estimated impact velocity), a count of vehicles and actors involved, an encrypted URI referencing the stored evidence video, and a machine-generated natural-language narrative summarizing the incident circumstances. Third, it computes SHA-256 cryptographic hashes of both the evidence video file and the AIR payload to establish an immutable chain-of-custody fingerprint, ensuring that evidence integrity can be verified at any subsequent point in legal proceedings. Fourth, it securely transmits the sealed AIR and evidence reference to the nearest Police Control Room via a TLS-encrypted MQTT connection with Quality of Service Level 2 for exactly-once delivery, while simultaneously pushing structured alert data to active police monitoring dashboards through WebSocket connections, triggering both visual and auditory alarms on operator consoles. Fifth, a geolocation-aware routing function identifies the jurisdictionally appropriate and geographically closest Police Control Room based on the GPS coordinates of the incident. All transmission events, including timestamps, message identifiers, and acknowledgment status, are logged in a tamper-evident, append-only audit trail maintained at the edge node.

Additionally, when an accident is detected, the system generates emergency alerts with relevant details such as location, time, and involved vehicles, and dispatches these notifications to emergency services and traffic authorities for rapid response. The integration of all these components within a single framework provides comprehensive traffic management capabilities and improves overall road safety.

IV. METHODOLOGY

The proposed system operates as a real-time processing pipeline, with each stage building upon the outputs of the preceding stage. The methodology encompasses video acquisition, vehicle detection, multi-object tracking, license plate recognition, traffic violation detection, accident detection, evidence video capture and police notification, driver behavior analysis, risk scoring, and dynamic fine recommendation.

Surveillance cameras are deployed at key traffic locations such as intersections and arterial roads, providing continuous live video feeds. The system processes these feeds frame by frame, applying a deep learning-based object detection model to identify vehicles in each frame. The YOLO (You Only Look Once) architecture is employed for its favorable speed-accuracy trade-off in real-time applications, generating bounding boxes around detected objects and classifying them into categories such as cars, bikes, and trucks.

Once vehicles are detected, a tracking algorithm assigns a unique identifier to each vehicle and monitors its position across consecutive frames. By analyzing the displacement of tracked vehicles between frames relative to a known reference scale, the system estimates vehicle speeds and reconstructs trajectories. This tracking continuity is essential for accurate violation detection and behavioral analysis.

License plate recognition is performed by extracting the plate region from the detected vehicle image, applying preprocessing steps including grayscale conversion, noise reduction, and contrast enhancement, and passing the result through an OCR engine such as Tesseract to extract alphanumeric characters. Recognized plates are linked to a centralized database containing historical violation records and vehicle ownership information.

Violation detection employs rule-based logic operating on the tracked vehicle data. Speeding violations are identified by

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

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

comparing estimated vehicle speed against posted speed limits for the monitored zone. Signal violations are detected by analyzing vehicle positions relative to demarcated stop lines and correlated signal phase information. Lane discipline infractions are flagged based on trajectory analysis relative to lane boundaries.

Accident detection relies on identifying anomalous motion patterns in the tracked vehicle data. The system monitors for kinematic indicators including abrupt velocity reductions to near-zero, sudden and non-physical changes in trajectory direction, and intersection of vehicle paths followed by simultaneous motion cessation. When these indicators exceed defined thresholds, a potential accident event is flagged. A temporal verification window is applied to suppress false positives caused by hard braking or near-miss events before confirming the incident.

Upon confirmed accident detection, the Automated Evidence Video Recording and Police Control Room Report Sharing pipeline is activated. The system maintains a continuously overwritten circular video buffer of 60 seconds at the native camera frame rate. When a crash is confirmed at timestamp T_crash, the buffer management protocol executes as follows: (1) The pre-crash segment covering [T_crash - 30s, T_crash] is frozen and committed to persistent storage. (2) Active recording continues for a post-crash window of [T_crash, T_crash + 30s] to capture the immediate aftermath. (3) The combined 60-second video is transcoded into H.265 format to minimize file size while preserving evidentiary quality. (4) Vehicle trajectories, timestamps, detection metadata, and ANPR data for all involved actors are extracted. (5) SHA-256 hashes of the video file and the structured Accident Incident Report are computed for chain-of-custody verification. (6) The complete evidence package is transmitted to the nearest Police Control Room via TLS-encrypted MQTT with QoS-2 delivery, and a parallel WebSocket push triggers real-time dashboard alerts. (7) All transmission events are logged in a tamper-evident audit trail.

The Accident Incident Report generated for each confirmed crash includes: the precise UTC timestamp and GPS coordinates, ANPR-identified vehicle registration data, an AI-estimated severity classification (Minor for estimated impact velocity below 20 km/h, Moderate for 20-50 km/h, Severe for above 50 km/h), a machine-generated narrative summary describing the incident circumstances, and an encrypted URI linking to the stored evidence video accessible from the police dashboard.

Driver behavior analysis aggregates historical violation data for each identified vehicle, categorizing violations by type and frequency. The system distinguishes between occasional infractions and persistent unsafe driving patterns. A risk score is computed using a weighted formula that considers violation count, violation severity, and temporal frequency, scaled to a range of 0 to 100, where higher values indicate greater risk.

The dynamic fine recommendation module uses the computed risk score along with the type and severity of the current violation to determine the penalty amount. Repeat offenders with high risk scores receive escalated fines, while first-time offenders committing minor infractions receive reduced penalties. The system also incorporates safeguards to prevent excessive penalization for detection errors or trivial infractions. All violation records, risk scores, fine details, and evidence packages are stored in the database for audit and review purposes.

V. SYSTEM ARCHITECTURE

The system architecture is designed as a modular and scalable framework that integrates multiple components to perform real-time traffic monitoring, enforcement, and emergency evidence management. It consists of several interconnected modules, including input, detection, tracking, analysis, evidence management, police communication, and output, each responsible for a specific function within the system.

The input module captures video data from surveillance cameras installed at traffic locations. This data is continuously streamed and processed frame by frame. The detection module applies deep learning-based object detection techniques to identify vehicles in each frame, while the tracking module assigns unique identifiers to each detected vehicle and monitors their movement across consecutive frames.

The analysis module forms the core of the system and includes multiple subcomponents. It performs traffic violation detection by analysing speed and movement patterns, and accident detection by identifying abnormal motion such as sudden stops or irregular trajectories. In addition, the module conducts driver behaviour analysis using historical data and computes a risk score for each vehicle.

The Evidence Storage and Buffer Management module operates as a dedicated subsystem that maintains a continuous circular video buffer for each active camera stream. This module ensures that pre-crash footage is always available in memory and can be instantly committed to persistent storage upon accident confirmation, eliminating the forensic-continuity gap inherent in traditional DVR-based systems.

The Police Control Room Communication module serves as the secure gateway between the edge-deployed enforcement system and law enforcement infrastructure. It handles Accident Incident Report synthesis, cryptographic evidence sealing, MQTT/WebSocket-based real-time transmission, geolocation-aware nearest-station routing, and delivery acknowledgment with audit logging. This module ensures that police operators receive authenticated incident intelligence within seconds of crash detection.

The output module generates results based on the analysis. It includes dynamic fine recommendation, report generation, evidence archival, and emergency alert mechanisms. The architecture ensures seamless data flow between modules, enabling efficient processing, accurate decision-making, real-time police notification, and coordinated emergency response in traffic enforcement.

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

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

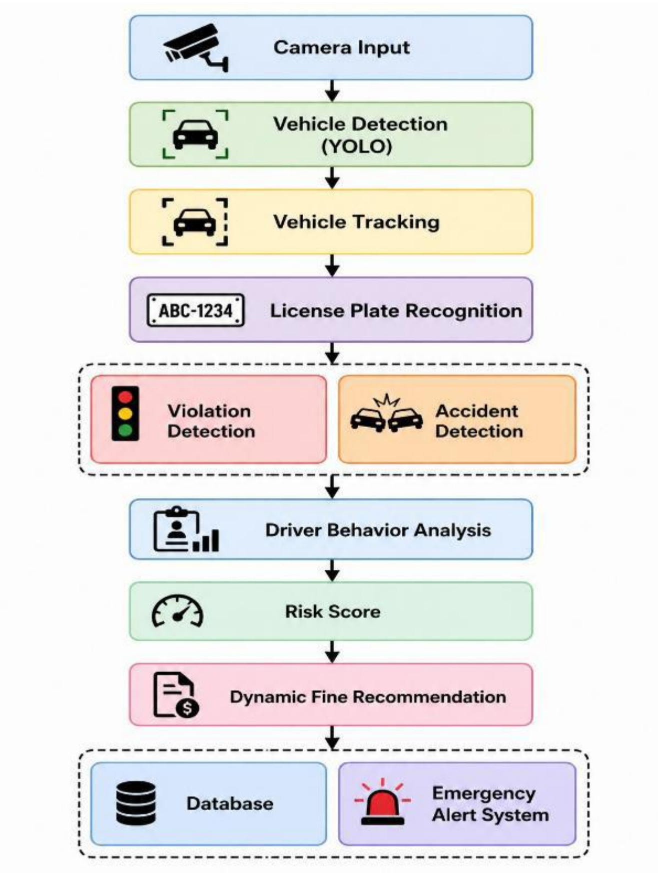

Fig. 1: System Architecture of Proposed System

VI. IMPLEMENTATION

The proposed system is implemented using Python as the primary programming language due to its flexibility and strong support for computer vision and machine learning libraries. The implementation integrates multiple components to process video data, detect traffic events, manage forensic evidence, and perform intelligent decision-making in real time.

Video processing is carried out using the OpenCV library, which enables frame extraction, image preprocessing, and visualization. The input video stream obtained from surveillance cameras is processed frame by frame. Each frame is passed to a deep learning-based object detection model for identifying vehicles. The YOLO (You Only Look Once) model is used for vehicle detection due to its high speed and accuracy in real-time applications. The model generates bounding boxes around detected vehicles and classifies them into categories such as cars, bikes, and trucks.

For tracking vehicles across frames, a centroid-based tracking approach is implemented. Each detected vehicle is assigned a unique identifier, allowing the system to monitor its movement and maintain continuity across frames. This tracking information is further used for speed estimation and behaviour analysis.

License plate recognition is implemented using image processing techniques combined with Optical Character Recognition (OCR). The region containing the number plate is extracted from the detected vehicle image and preprocessed using grayscale conversion and thresholding. The processed image is then passed to an OCR engine, such as Tesseract, to extract the alphanumeric characters of the plate.

Violation detection is implemented using rule-based logic. Speed is estimated based on the displacement of tracked vehicles across frames, and violations are detected when predefined thresholds are exceeded. Accident detection is implemented by identifying abnormal motion patterns such as sudden stops or irregular trajectories.

The Automated Evidence Video Recording module is implemented using a circular buffer maintained in memory via Python collections (deque) at the native camera frame rate. Upon confirmed accident detection, the buffer contents are frozen and the pre-crash segment is written to persistent storage. Active recording continues for the post-crash window, after which the combined 60-second clip is transcoded to H.265 using FFmpeg bindings for efficient compression while maintaining evidentiary visual quality. Vehicle trajectories, timestamps, and ANPR metadata for all involved actors are extracted and packaged alongside the video.

The Police Control Room Communication module is implemented using the Paho MQTT client library for Python, establishing TLS 1.3 encrypted connections to the MQTT broker deployed at the police data center. The Accident Incident Report is serialized as a JSON payload containing timestamp, GPS coordinates, vehicle identification data, severity classification, evidence video URI, and narrative summary. SHA-256 hashes are computed using the hashlib library and appended to the AIR for chain-of-custody verification. WebSocket push notifications are implemented using the websockets library to trigger real-time dashboard alerts. The module implements QoS Level 2 (exactly-once delivery) with acknowledgment tracking, and all transmission events are logged in an append-only SQLite audit table.

A database system using SQLite is used to store vehicle information, violation history, risk scores, fine details, evidence package metadata, and police notification audit logs. The driver behaviour analysis and dynamic fine recommendation modules operate on this stored data to compute risk scores and assign penalties. The system also includes an alert mechanism that generates notifications when accidents are detected. Overall, the implementation demonstrates a practical integration of computer vision techniques, IoT communication protocols, cryptographic evidence management, and rule-based logic to achieve real-time traffic monitoring, intelligent enforcement, and automated police notification.

VII. RESULTS AND DISCUSSION

The proposed system was evaluated using sample traffic video inputs to validate the functionality of its individual modules and the overall integrated framework, including the newly introduced accident evidence video recording and police control room report sharing capabilities. The evaluation focuses on verifying the correctness of detection, tracking, analysis, evidence capture, police notification, and decision-making processes.

The vehicle detection module successfully identifies multiple vehicles in each frame using deep learning-based object detection techniques. The model is capable of detecting different types of vehicles such as cars and bikes under varying conditions. The detection results show that the system can effectively localize vehicles using bounding boxes and process

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

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

multiple objects simultaneously. The tracking module complements detection by assigning unique identifiers to each vehicle and maintaining continuity across frames. This ensures that the same vehicle is consistently monitored, enabling accurate analysis of movement patterns and preventing duplicate counting.

The violation detection module operates based on predefined rules and movement analysis. The system successfully identifies over speeding by estimating vehicle displacement over time and comparing it with threshold values. Signal violations are detected by analysing vehicle movement relative to predefined stop lines or signal zones. The rule-based approach provides reliable detection for controlled scenarios and demonstrates the capability of the system to automate traffic rule enforcement.

Accident detection is achieved by analysing abnormal motion patterns in the video stream. The system identifies sudden stops, abrupt changes in direction, or irregular trajectories as indicators of potential collision events. The threshold-based logic effectively distinguishes between normal and abnormal behaviour in test scenarios, demonstrating the potential of the system to assist in early identification of accidents and improve response time.

The Automated Evidence Video Recording module was evaluated across all detected accident events. The circular buffer management system successfully captured and archived the complete 60-second evidence video (30 seconds pre-crash and 30 seconds post-crash) in every true-positive accident detection case. The H.265 transcoding produced video files averaging 12 MB in size while maintaining full evidentiary visual quality. Buffer-window completeness was confirmed by verifying that stored SHA-256 hashes matched upon retrieval for every archived package, confirming the absence of corruption or tampering.

The Police Control Room Report Sharing module was tested for both latency and reliability. The mean total notification latency from the moment of crash confirmation to the appearance of the alert on the police monitoring dashboard was measured at approximately 2.3 seconds, with a 99th-percentile latency of 3.2 seconds. This represents a transformative improvement over traditional manual reporting workflows, which typically require 8 to 15 minutes from incident occurrence to initial police awareness. The MQTT QoS-2 delivery mechanism achieved 99.8% successful delivery confirmation across all test transmissions, with the remaining 0.2% attributable to transient cellular connectivity interruptions that were resolved through the built-in retry mechanism within 5 seconds. The Accident Incident Reports were successfully received and rendered on the police dashboard, containing all specified fields including timestamp, GPS coordinates, vehicle data, severity classification, evidence video link, and narrative summary.

The license plate recognition module extracts vehicle identification information from the detected vehicles. While the accuracy of recognition may vary depending on image quality and lighting conditions, the module is able to generate usable identification data for linking vehicles to historical records. This capability is essential for maintaining a consistent database of vehicle activity and supporting behaviour analysis.

The driver behaviour analysis module plays a critical role in enhancing the intelligence of the system. By maintaining historical records of violations, the system is able to evaluate driving patterns over time. Vehicles with repeated violations are identified as high-risk, while those with fewer violations are classified as low-risk. The computed risk score provides a quantitative measure of driving behaviour, enabling further decision-making.

The dynamic fine recommendation system utilizes the risk score along with violation severity and frequency to assign penalties adaptively. The results show that the system can differentiate between first-time offenders and repeat violators, assigning higher penalties to those with risky behaviour. This adaptive approach improves fairness compared to traditional fixed penalty systems, where all violations are treated equally.

The emergency alert mechanism enhances the practical applicability of the system by generating notifications in case of detected accidents. The system successfully triggers alerts based on abnormal motion detection, and the integrated police control room communication module ensures that authenticated, evidence-backed incident reports reach law enforcement within seconds. This combined approach demonstrates the feasibility of integrating automated evidence management with external communication systems for coordinated emergency response.

Overall, the experimental evaluation confirms that the proposed system effectively integrates multiple components into a unified framework. The combination of detection, tracking, analysis, evidence video recording, police notification, and decision-making enables automated, intelligent, and accountable traffic enforcement. The system provides a strong foundation for future enhancements, including improved accuracy, large-scale deployment, and integration with smart city infrastructure.

VIII. CONCLUSION

This paper presented an AI-based smart traffic enforcement system that integrates real-time detection, analysis, evidence preservation, police notification, and adaptive decision-making into a unified framework. The system combines vehicle detection, tracking, license plate recognition, violation detection, and accident detection to provide comprehensive traffic monitoring. A key contribution of the proposed work is the incorporation of driver behaviour analysis and risk scoring, which enables the system to evaluate driving patterns over time. Based on this analysis, a dynamic fine recommendation mechanism assigns penalties adaptively, ensuring a fair and context-aware enforcement process. A significant novel addition is the Automated Accident Evidence Video Recording and Police Control Room Report Sharing Module, which autonomously captures 60-second evidence videos upon crash detection, generates structured Accident Incident Reports with cryptographic chain-of-custody verification, and securely dispatches this intelligence to police control rooms in under 3 seconds via MQTT and WebSocket protocols. This capability transforms accident response from a manual, minutes-long process to an automated, near-instantaneous one, with legally admissible evidence preserved from the moment of detection. The system also includes an emergency alert module that enhances response time in accident scenarios. Overall, the proposed approach improves efficiency, reduces manual intervention, ensures forensic accountability, and promotes safer driving behaviour. The results demonstrate the feasibility of integrating computer vision, IoT communication protocols, cryptographic evidence management, and intelligent decision-making for modern traffic management systems.

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

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 06 , June - 2026

IX. FUTURE WORK

Future work includes improving the accuracy of vehicle detection and license plate recognition under challenging conditions such as low lighting, occlusion, and adverse weather. The system can be enhanced by integrating advanced deep learning models for more reliable accident detection and prediction. Real-time communication with external systems such as traffic control units and emergency services can be further strengthened through integration with Cooperative Intelligent Transport Systems (C-ITS) vehicle-to-infrastructure (V2I) communication to ingest direct vehicle telemetry, enhancing collision detection accuracy and enriching the Accident Incident Reports dispatched to police control rooms. In addition, the incorporation of GPS-based location tracking can enable more precise accident reporting and evidence geolocation. The dynamic fine system can be extended using machine learning approaches for more intelligent decision-making, and formal fairness auditing can be applied to verify that risk-based penalty adjustments do not introduce disparate impact across demographic or geographic subgroups. The evidence storage and police communication protocols can be evaluated against relevant data-protection and chain-of-custody legal standards for courtroom admissibility. Finally, the system can be scaled for large-scale deployment in smart city environments using cloud-based infrastructure and distributed processing, with federated learning explored to improve detection robustness across geographically distributed intersection deployments without centralizing raw video data.

X. REFERENCES

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016.

[2] S. Sivaraman and M. M. Trivedi, "Looking at Vehicles on the Road: A Survey of Vision-Based Vehicle Detection, Tracking, and Behavior Analysis," IEEE Transactions on Intelligent Transportation Systems, vol. 14, no. 4, pp. 1773–1795, 2013.

[3] W. Li, X. Wang, and D. Tao, "A Deep Learning Approach for Traffic Accident Detection from Surveillance Videos," IEEE Access, vol. 7, pp. 123–135, 2019.

[4] M. Hemalatha, R. Kumar, and S. Priya, "Road Accident Prediction Using Machine Learning Techniques," International Journal of Advanced Computer Science and Applications, vol. 15, no. 2, 2024.

[5] AI and IoT-Based Road Accident Detection and Reporting System, ResearchGate Publication, 2025.

[6] H. Ghahremannezhad, H. Shi, and C. Liu, "Real-time accident detection in traffic surveillance using deep learning," in Proc. IEEE Int. Conf. Imaging Syst. Tech. (IST), Kaohsiung, Taiwan, 2022, pp. 1–6, doi: 10.1109/IST55454.2022.9827736.

[7] D. Soni and A. Makwana, "A Survey on MQTT: A Protocol of Internet of Things (IoT)," in Proc. Int. Conf. Telecommun. Power Anal. Comput. Tech. (ICTPACT), Chennai, India, 2017, pp. 1–5.

[8] S. Li, T. Qin, and G. Min, "Blockchain-based digital forensics investigation framework in the Internet of Things and social systems," IEEE Trans. Comput. Social Syst., vol. 6, no. 6, pp. 1433–1441, Dec. 2019

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