



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 15, Issue 3, March 2026

AI-Powered Vehicle Number Plate Detection System

Pranita V. Shitole¹, Parth Gaikwad², Prathamesh Gore³, Om Lokhande⁴, Aanas Shaikh⁵

Mentor, Department of Artificial Intelligence & Machine Learning, AISSMS's Polytechnic, Pune, Maharashtra, India¹

Student, Diploma in Artificial Intelligence & Machine Learning, AISSMS's Polytechnic, Pune,

Maharashtra, India²⁻⁵

ABSTRACT: This project presents an AI-powered Vehicle Number Plate Detection and Recognition System capable of automatically detecting, extracting, and processing license plates from images, video files, and live camera feeds. It uses the YOLOv11 deep learning model for fast and accurate plate detection, supporting multiple vehicles and plates per frame under varying lighting and environmental conditions.

For text recognition, a multi-tier OCR pipeline combining EasyOCR, PaddleOCR, and Google Gemini Vision API is employed, achieving high accuracy (85–95%) even with low-resolution or challenging images. Additionally, the system includes an intelligent vehicle classification module based on color, automatically identifying Electric, Commercial, Private, and Temporary vehicles according to registration standards.

To enhance security and surveillance applications, a real-time watchlist matching feature flags stolen, blacklisted, or restricted vehicles, while all detected data is stored in an SQLite database for historical analysis. The system is deployed through an intuitive Streamlit web interface, providing a scalable solution for traffic monitoring, smart parking management, and automated vehicle identification.

I. INTRODUCTION

Vehicle Number Plate Detection and Recognition (VNPR), also known as Automatic Number Plate Recognition (ANPR), is a key technology in modern transportation, urban surveillance, and intelligent traffic management systems. As cities transition into smart ecosystems, the demand for automated vehicle identification solutions has increased significantly. Traditional methods, relying on manual observation, checkpoints, and offline record keeping, are limited by high human dependency, slow response times, and low scalability in high-traffic environments.

With the advancement of Artificial Intelligence (AI), Machine Learning (ML), and Computer Vision, automated VNPR systems have emerged as efficient and real-time alternatives. Deep learning-based models like YOLOv11 allow rapid and accurate detection of license plates in dynamic environments, while Optical Character Recognition (OCR) tools such as EasyOCR, PaddleOCR, and Google Gemini enable automated extraction of plate information. By integrating these technologies, modern systems provide end-to-end automation—from detecting vehicles to extracting, classifying, and storing their registration details.

Despite these advancements, current solutions face challenges such as environmental variability, motion blur, occlusion in crowded scenes, low-resolution inputs, and limitations in OCR performance. Most traditional systems also lack intelligent vehicle classification based on plate colors (Electric, Commercial, Private, Temporary), and fail to provide seamless integration with backend databases or real-time watchlist alerts. This project addresses these gaps by developing an AI-powered VNPR framework capable of multi-plate detection, high-accuracy text extraction, automatic vehicle classification, security alerts, and centralized data storage for analytics. The system is scalable, supports multiple input formats, and offers an intuitive Streamlit interface for live monitoring, media uploads, and historical data management, making it suitable for toll booths, parking systems, border security, highway enforcement, and smart city applications.

II. PROBLEM STATEMENT

The automated identification and tracking of vehicles is crucial for traffic control, urban surveillance, parking management, and security enforcement. However, existing systems face multiple challenges that affect their performance. Environmental conditions such as poor lighting, bright sunlight, rain, fog, dust, and motion blur can reduce detection accuracy. Detecting multiple vehicles in crowded scenes is also difficult due to occlusion, overlapping vehicles, and plates appearing at different angles or partially visible. Additionally, OCR systems often struggle with distorted, dirty, or low-resolution plates, incorrect character segmentation, and confusion between similar characters like O/0, I/1, or B/8.

Most current solutions also suffer from fragmented architectures, lacking seamless integration between detection models and backend storage, fallback OCR mechanisms, and synchronization between edge devices and centralized databases. Furthermore, traditional systems rarely incorporate intelligent vehicle classification, which is essential for differentiating Electric (green plates), Commercial (yellow plates), Private (white plates), and Temporary Registration (red plates) vehicles. These limitations highlight the need for a scalable, AI-driven system capable of real-time multi-plate detection, high-accuracy text extraction, automatic vehicle classification, watchlist-based alerts, and centralized historical data storage. This project proposes an end-to-end AI-powered vehicle identification framework to address these challenges effectively.

III. OBJECTIVES OF THE PROJECT

Primary Objectives:

- To develop a high-accuracy AI-powered vehicle number plate detection system.
- To implement YOLOv11 for real-time, multi-plate detection across images, videos, and live streams.
- To design a multi-layer OCR architecture integrating Google Gemini, EasyOCR, and PaddleOCR.
- To implement fallback mechanisms to ensure maximum recognition reliability and achieve 85–95% accuracy with Gemini Vision API.
- To implement a real-time detection system.
- To develop an algorithm for automatic vehicle classification based on license plate color standards.
- To compare detected plates against a predefined database for security alerts, law enforcement tracking, and access control.
- To provide a Streamlit interface for media uploads, live feed monitoring, detection visualization, and historical data search.
- To store detection results, timestamps, vehicle classifications, and alert status in SQLite for reliable historical tracking and reporting.

Secondary Objectives:

- To support multiple input formats including images, pre-recorded videos, and real-time camera feeds.
- To enable concurrent detection of multiple plates in a single frame.
- To provide detailed logs and analytics dashboards.
- To ensure a modular and scalable system architecture.

IV. LITERATURE REVIEW

Object detection is a fundamental task in computer vision, involving the identification and localization of objects within images using bounding boxes. Over the past decade, the field has evolved significantly, largely driven by advancements in Deep Learning and Convolutional Neural Networks (CNNs). Traditional detection methods, such as Haar Cascade classifiers, Histogram of Oriented Gradients (HOG), and Support Vector Machines (SVM), performed adequately in controlled environments but struggled with complex backgrounds, real-time processing, variations in lighting, and multiple-object detection, making them unsuitable for high-speed vehicle monitoring. The advent of CNN-based architectures like R-CNN, Fast R-CNN, and Faster R-CNN improved detection accuracy but were computationally intensive and often unsuitable for real-time applications due to multi-stage processing pipelines. The introduction of YOLO (You Only Look Once) revolutionized object detection by framing it as a single regression

problem, enabling high-speed, real-time detection with a unified architecture and multi-object detection capability. YOLOv11, with its enhanced backbone network, optimized feature pyramid, and improved inference pipeline, is particularly suitable for license plate detection, offering high inference speed, better small-object detection, and robust performance under diverse environmental conditions.

Optical Character Recognition (OCR) has also evolved significantly, moving from template matching and rule-based segmentation to deep learning-based models that combine CNNs for feature extraction, RNNs or LSTMs for sequence modeling, and attention mechanisms for contextual recognition. Traditional OCR methods struggled with distorted, rotated, or low-resolution plates, while modern engines like EasyOCR and PaddleOCR provide high accuracy, multi-language support, robust segmentation, and handling of rotated or noisy text. Furthermore, cloud-based AI services such as the Google Gemini Vision API enhance OCR performance by leveraging contextual patterns, correcting ambiguous characters, and utilizing large-scale language models for post-processing. In this project, a multi-tier OCR pipeline is implemented with Gemini API as the primary engine, PaddleOCR as the secondary engine, and EasyOCR as a fallback mechanism to ensure maximum robustness and recognition accuracy.

License Plate Recognition (LPR) systems combine object detection and OCR to automatically detect, segment, and recognize license plates. Conventional LPR systems follow a sequential pipeline of detection, segmentation, recognition, and validation, which works in controlled conditions but faces challenges in real-world scenarios such as camera angle variations, occlusion, crowded traffic, and limited scalability. Modern LPR systems address these limitations through end-to-end deep learning models, multi-scale feature extraction, attention-based recognition, transformer-based OCR, and integration with cloud AI services. These systems support multi-plate detection, real-time video streams, watchlist matching, and vehicle classification based on metadata such as color, region, and plate format, making them more robust, scalable, and suitable for intelligent traffic monitoring and security applications.

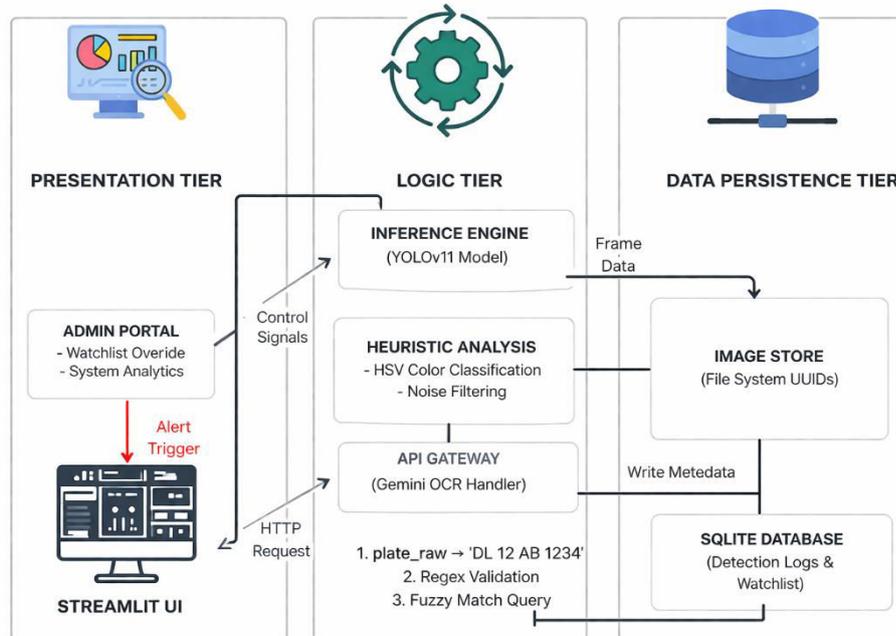
V. PROPOSED SYSTEM

The proposed system is an AI-powered Vehicle Number Plate Detection and Recognition framework designed to operate across multiple input sources, including static images, video files, and live camera feeds from webcams, IP cameras, or CCTV systems. The system employs YOLOv11 for real-time multi-plate detection, generating bounding boxes even for angled, partially visible, or distant plates. Detected regions are preprocessed—cropped, normalized, and enhanced—before being passed to a robust multi-tier OCR pipeline. The OCR architecture integrates Google Gemini as the primary engine, with PaddleOCR and EasyOCR as secondary and fallback engines, ensuring high accuracy (85–95%) and reliability even under challenging conditions.

Beyond detection and recognition, the system includes a vehicle classification module that determines the type of vehicle based on license plate color (green for Electric Vehicles, yellow for Commercial, white for Private, and red for Temporary Registration). Additionally, a watchlist management feature cross-references detected plates against a database of flagged vehicles, generating real-time alerts and logging incidents for security and law enforcement purposes. A user-friendly interface built with Streamlit allows media uploads, live monitoring, visualization of detection results, and historical analytics, while SQLite stores all detection logs, timestamps, vehicle classifications, and alert data.

The system is designed to be modular, scalable, and deployable across cloud servers, edge devices, or distributed camera networks. It emphasizes high performance, reliability, and accuracy through GPU acceleration, parallel processing, robust error handling, and fallback mechanisms for OCR failures. Leveraging Python, OpenCV, NumPy, Ultralytics YOLOv11, multiple OCR engines, Streamlit, and SQLite, the proposed framework provides a comprehensive solution for intelligent traffic monitoring, automated vehicle identification, and proactive security enforcement.

VI. SYSTEM ARCHITECTURE



Description:

The proposed system is an AI-powered Vehicle Number Plate Detection and Recognition framework designed for real-time monitoring and automated vehicle identification. It employs a three-tier architecture consisting of a Streamlit-based **presentation tier** for user interaction, a **logic tier** for core processing including YOLOv11-based plate detection, heuristic noise filtering, HSV-based vehicle color classification, and OCR integration with Google Gemini, PaddleOCR, and EasyOCR, and a **data persistence tier** using SQLite for storing detection logs, timestamps, vehicle classifications, and watchlist alerts. The system supports multiple input sources, including static images, pre-recorded videos, and live camera feeds, making it suitable for deployment at toll booths, parking gates, traffic intersections, and security checkpoints.

The system ensures robustness through a multi-tier OCR fallback strategy, real-time detection, multi-plate support, and automated vehicle classification. HSV color-based classification enhances accuracy under varying lighting conditions, while watchlist integration enables proactive security enforcement. High-resolution plate crops are stored in structured directories indexed by UUIDs, ensuring rapid retrieval and data integrity. This modular and scalable design allows deployment on edge devices, cloud servers, or distributed camera networks, providing a comprehensive and intelligent solution for traffic monitoring, parking management, and law enforcement applications.

VII. METHODOLOGY

Step 1: Frame Acquisition

- Capture raw frames from static images, video files, or live camera feeds using OpenCV.
- Validate input formats and handle corrupted media gracefully.

Step 2: Object Detection (YOLOv11)

- Resize frames to \$640 \times 640\$ pixels and run YOLOv11 for license plate localization.
- Detect multiple plates per frame, including angled, distant, or partially visible plates.
- Generate bounding boxes around each detected plate.

Step 3: Region of Interest (ROI) Extraction

- Crop detected license plates from the original high-resolution frame.
- Normalize and enhance cropped images (contrast adjustment) for OCR processing.

Step 4: Multi-Tier OCR Processing

- **Primary OCR:** Google Gemini Vision API extracts alphanumeric characters.
- **Secondary OCR:** If Gemini fails (timeout, network issue, low confidence), process ROI using PaddleOCR.
- **Fallback OCR:** If PaddleOCR fails, use EasyOCR as a last resort.
- Validate output using regex patterns matching regional license plate formats to eliminate false positives.

Step 5: Vehicle Classification

- Convert cropped plate images from BGR to HSV color space.
- Apply Gaussian blur to reduce noise.
- Create color masks for predefined categories: Green (EV), Yellow (Commercial), White (Private), Red (Temporary).
- Assign the vehicle category based on the mask with the highest pixel coverage.

Step 6: Watchlist Management & Alerting

- Normalize OCR output (uppercase, remove spaces) and compare with the security watchlist using fuzzy matching.
- Trigger real-time alerts (UI notification, audio alert) if a match is found.
- Log the complete detection record in the SQLite database.

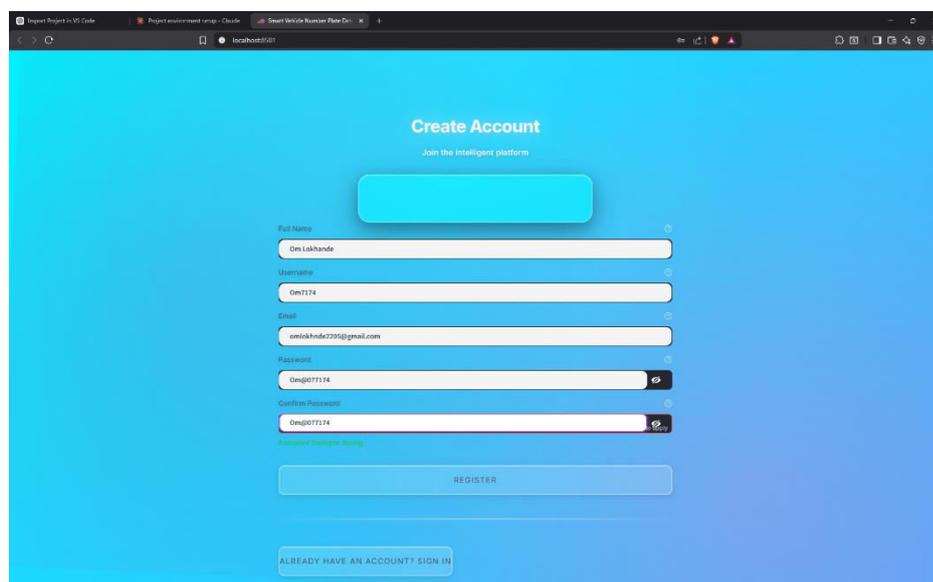
Step 7: Data Storage & Visualization

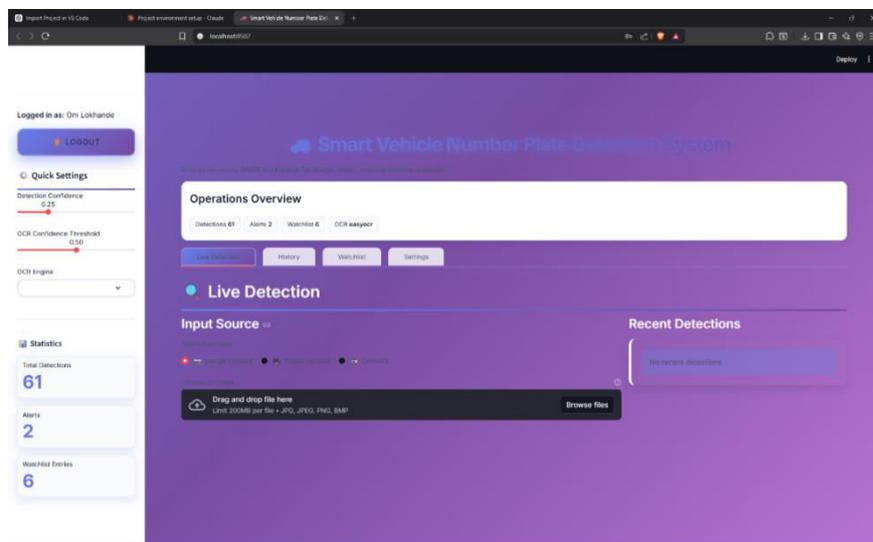
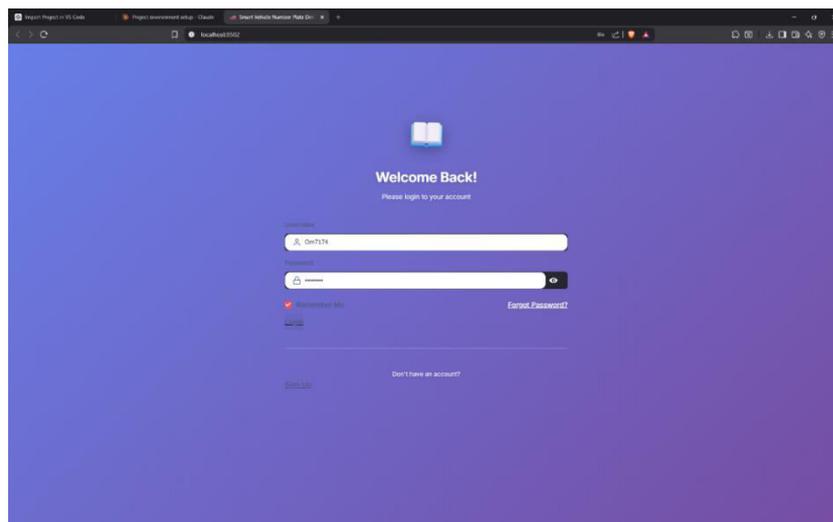
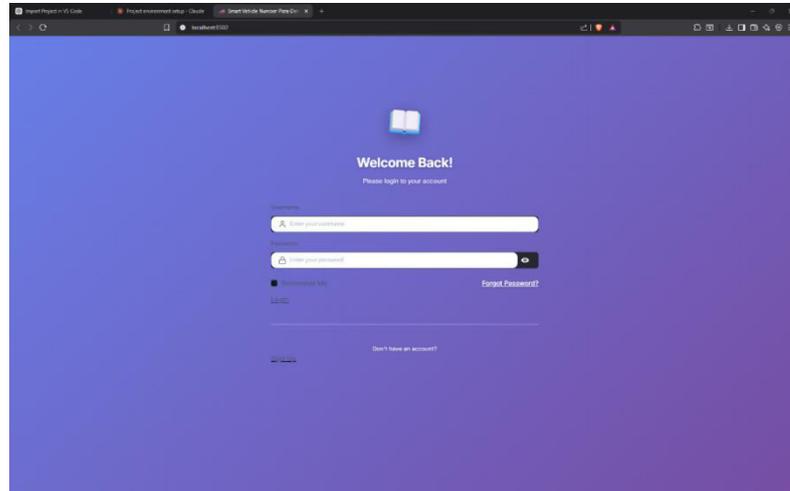
- Store detection metadata (timestamp, OCR text, confidence score, vehicle category, flagged status) in SQLite.
- Store high-resolution cropped images in structured directories indexed by UUID.
- Visualize results, live feeds, and historical data via the Streamlit interface.

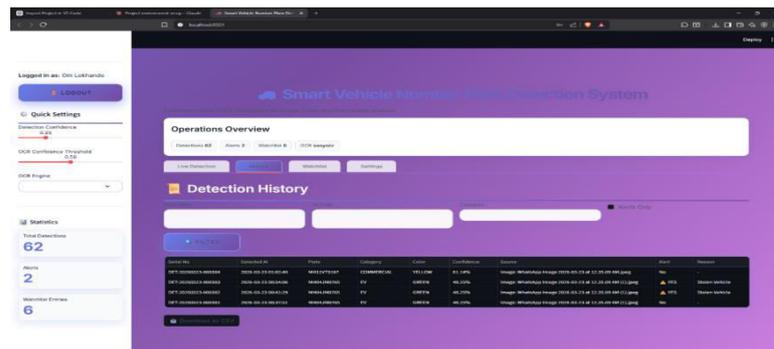
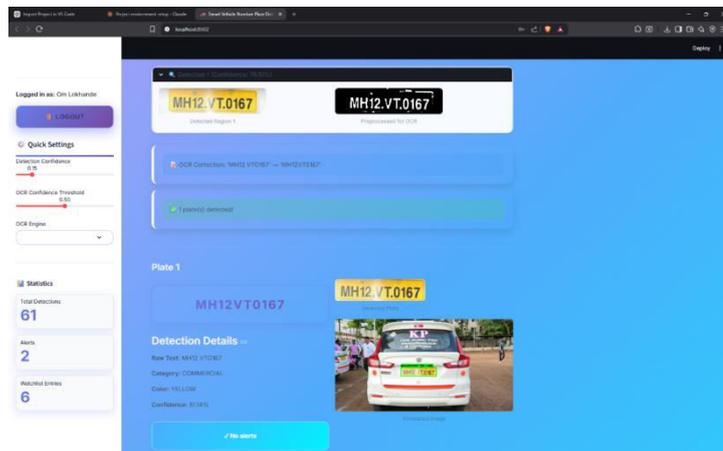
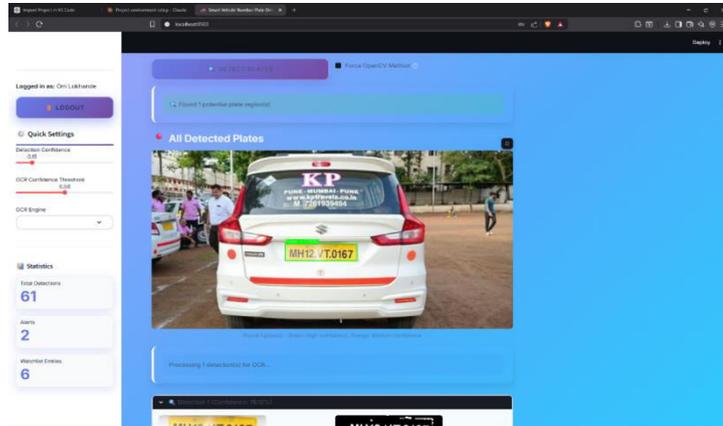
Step 8: System Scalability & Modularity

- Ensure the architecture supports cloud deployment, edge devices, and distributed camera networks.
- Maintain modular design for independent upgrades of detection, OCR, classification, and database modules.

VIII. RESULT







IX. CONCLUSION

The development and successful implementation of the AI-Powered Vehicle Number Plate Detection and Recognition System clearly demonstrate the effectiveness of combining high-speed deep learning-based object detection with generative AI-driven Optical Character Recognition. The project validates that modern Automatic License Plate Recognition (ALPR) systems can significantly outperform traditional rule-based OCR architectures when designed using a multi-layer, intelligent processing pipeline.

By integrating YOLOv11 for real-time and highly precise license plate localization, the system ensures accurate detection even in complex traffic environments containing multiple vehicles, motion blur, and partial occlusions. The transition from rigid segmentation-based OCR to a context-aware generative approach using Google Gemini marks a

major advancement in recognition reliability. Unlike traditional OCR engines that strictly depend on character-level segmentation, the generative AI layer leverages contextual reasoning to infer distorted or partially visible characters, significantly improving accuracy.

Furthermore, the inclusion of local fallback engines such as **EasyOCR** ensures operational continuity during network interruptions or API failures. The HSV-based color classification module enhances the system's functional scope by automatically categorizing vehicles into Electric, Commercial, Private, or Temporary registration types. This elevates the system from a simple recognition tool to a comprehensive smart surveillance and regulatory monitoring solution.

The architecture's modularity allows each subsystem—detection, OCR, classification, database logging, and dashboard visualization—to operate independently while remaining seamlessly interconnected. As a result, the system achieves a balanced blend of accuracy, speed, and reliability, making it suitable for deployment in real-world Indian traffic environments characterized by diverse plate formats and challenging weather conditions

X. FUTURE SCOPE

The AI-powered Vehicle Number Plate Detection and Recognition system has significant potential for expansion and enhancement in the future. One major avenue is **integration with smart city infrastructure**, where the system can be connected to traffic management centers, toll collection systems, and urban surveillance networks to enable automated monitoring, congestion analysis, and real-time traffic control.

Another area of improvement is **enhanced multi-modal vehicle analytics**. Future versions could incorporate additional features such as vehicle make and model recognition, speed estimation, driver identification, and anomaly detection, providing a more comprehensive traffic intelligence system. Integration with cloud-based big data platforms can allow large-scale analytics, pattern detection, and predictive maintenance for transport infrastructure.

The system can also benefit from **AI advancements**, such as transformer-based OCR models, edge AI deployment, and real-time 3D object tracking, improving accuracy, reducing latency, and enabling deployment in resource-constrained environments. Moreover, expanding to **cross-border and multi-country vehicle recognition** with support for multiple regional license plate formats can increase its applicability globally.

Lastly, the system has potential for **IoT and automation integration**, allowing it to trigger automated responses such as barrier control at parking gates, automated fine issuance for traffic violations, or integration with connected vehicles for seamless urban mobility solutions. These enhancements make the system scalable, adaptive, and highly relevant for modern intelligent transportation and security applications.

REFERENCES

1. Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
2. Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767.
3. Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI).
4. Shi, B., Bai, X., & Yao, C. (2016). An End-to-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition. IEEE TPAMI.
5. Silva, S. M., & Jung, C. R. (2018). License Plate Detection and Recognition in Unconstrained Scenarios. European Conference on Computer Vision (ECCV).



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details