

Advanced Vehicle Surveillance System: License Plate Recognition and Tracking

Prof. Pallavi Pathare¹, Mayur Sonawane², Prasad Raktate³, Yash Sonawane⁴, Aditya Ugale⁵

¹Professor, Dept. of Computer engineering, Sir Visvesvaraya Institute of Technology, Nashik, Maharashtra, India.

^{2,3,4,5}UG Students, Dept. of Computer engineering, Sir Visvesvaraya Institute of Technology, Nashik, Maharashtra, India.

Abstract:

This paper presents a complete framework for a vehicle surveillance system designed for automatic license plate recognition and tracking. The proposed system includes a processing pipeline that consists of video processing, object detection, optical character recognition (OCR), and data visualization. For plate detection, we use the YOLOv5 model. We achieve text recognition through a combination of Easy OCR and Tesseract. Additionally, we integrate post-processing algorithms to improve recognition accuracy by enhancing image quality, normalizing text, and validating plate formats. The system includes a web-based interface that allows users to upload surveillance footage, search for detected license plates, and visualize results through an interactive map and analytical dashboards. Experimental results show the system's performance under real-world conditions with different lighting, camera angles, and vehicle speeds. The proposed approach has strong potential for use in traffic management, law enforcement, and smart city development.

Keywords: License plate recognition, YOLOv5, optical character recognition, vehicle surveillance, computer vision, image processing.

1. INTRODUCTION

Automated license plate recognition (ALPR) is becoming more popular in areas like traffic management, law enforcement, and security. These systems automatically detect and recognize vehicle license plates from images or video footage, providing useful data for traffic monitoring, parking management, toll collection, and investigations.

ALPR systems face challenges such as changes in lighting, different plate formats, obstructions in the scene, and motion blur in images. However, recent advancements in computer vision have significantly improved the accuracy and reliability of these systems. Still, creating a truly complete solution that combines detection, recognition, and data management remains difficult.

In this paper, we introduce a comprehensive vehicle surveillance system designed to overcome these challenges through a multi-stage process.

Our system uses cutting-edge models for object detection and character recognition, along with effective post-processing strategies to achieve high accuracy in real-world scenarios. The main contributions of this work include:

- An end-to-end framework that covers the entire process of vehicle license plate recognition from video footage.
- A hybrid type of optical character recognition (OCR) that employs multiple recognition methods to improve results.
- Advanced post-processing techniques for plate validation and removing duplicates.
- A web-based GUI for data visualization and analysis.
- An experimental evaluation that tests the system under various conditions.

The rest of the paper is organized as follows: Section II surveys existing work related to license plate recognition. Section III discusses the architecture and methodology proposed in this paper. Section IV addresses implementation issues. Section V reviews experimental results and performance evaluation. Finally, Section VI concludes the paper and suggests directions for future work.

2. RELATED WORK

An Automatic License Plate Recognition (ALPR) system detects and interprets vehicle registration plates from images or video streams. It serves various purposes like traffic management, toll collection, parking control, and security analysis. However, ALPR faces several practical challenges, such as changing lighting conditions, different plate designs, motion blur, partial obstructions, and variations in viewing angles. Recent advancements in convolutional neural networks (CNNs) and new object detection models have greatly improved detection and recognition. Still, to achieve a dependable end-to-end ALPR solution, it is essential to address OCR errors, confirm plate formats, and manage and analyze captured data effectively. This paper presents a full-stack ALPR framework that combines a YOLOv5-based detector with a hybrid OCR system and a post-processing pipeline. This setup aims to improve recognition reliability in real-world situations.

3. PROPOSED ALGORITHM

A. Video Upload and Session Management:

The user interface serves as the entry point to the system, where operators can upload surveillance footage along with associated metadata such as location, time, and other relevant details. Each upload is registered as a new session with a unique identifier that tracks the processing status and organizes the corresponding results. The session manager is responsible for storing video files and metadata, as well as initiating the processing pipeline

B. Frame Processing and Plate Detection:

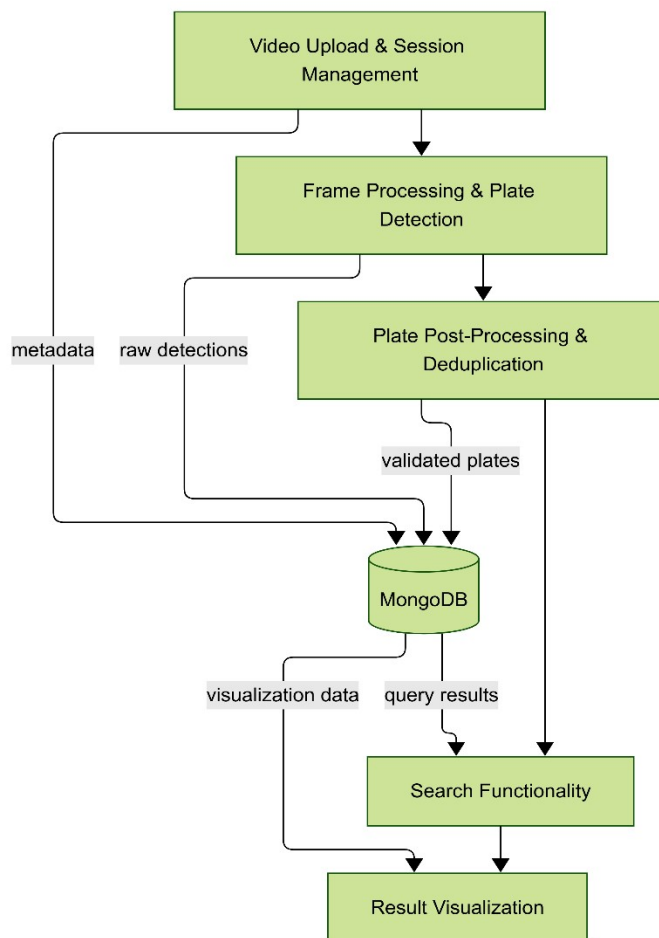


Figure 1. System architecture of the vehicle surveillance system showing the main components and their interactions.

This module processes the uploaded videos frame by frame to detect and recognize vehicle license plates. It uses a YOLOv5 model specifically trained for license plate detection. For each detected plate area, the corresponding image is cropped and sent to the OCR engine for text extraction. The

entire process runs asynchronously in a separate thread. This keeps the user interface responsive during the operation.

C. Plate Post-Processing and Deduplication:

Raw Optical Character Recognition (OCR) outputs often have errors and duplicate entries. The post-processing techniques described in this section are used in the frame processing and plate detection pipeline to improve recognition accuracy and remove duplicate reports. These techniques include image preprocessing, such as converting to grayscale, enhancing contrast, and reducing noise. They also involve text normalization, which corrects visually similar characters, format validation using regular expressions, and clustering of similar detections based on the Levenstein distance.

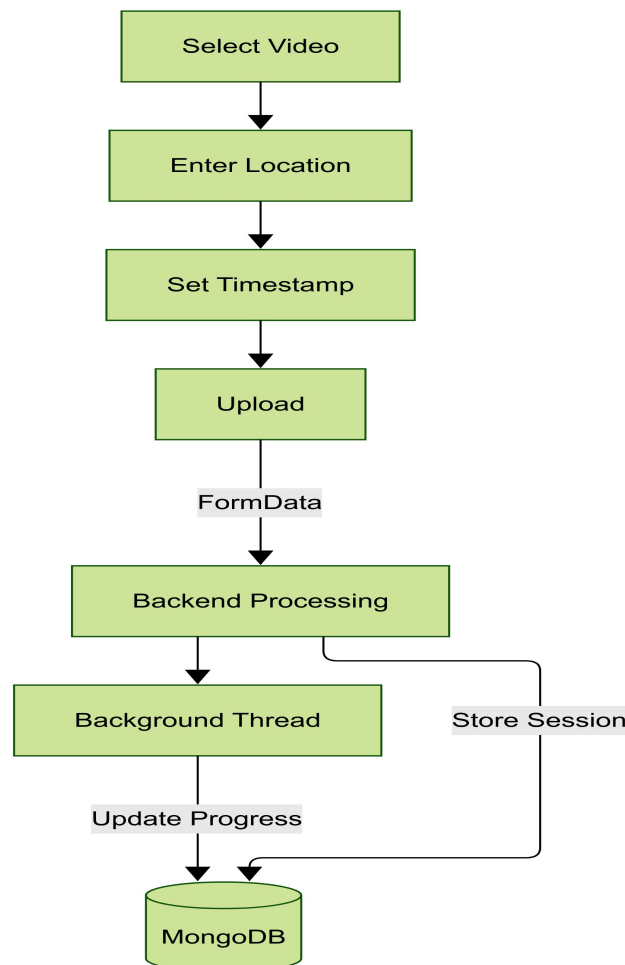


Figure 2. Workflow of the video upload and session management process.

D. SEARCH FUNCTIONALITY

The search component lets users look up specific license plates. It adjusts the input query to ensure correct searching and retrieves all matching detections from the database, along with their timestamps and capture locations. This feature allows users to track a vehicle's movement from one detection point to another.

E.RESULT VISUALIZATION

This component provides several ways to visualize detection results:

- Map View: Shows detection locations on an interactive map with optimized path calculations.
- Statistics View: Displays temporal patterns and spatial distributions using different analytical charts.
- Image View: Offers thumbnails of all detected plate images for visual verification and review.
- PDF Report: Creates detailed reports for documentation and easy sharing.

F. DATA STORAGE

The system uses MongoDB for data storage, organized into four main collections:

- Sessions, which store metadata related to each processing session.
- Plates, which contain raw OCR outputs along with frame-specific details.
- Plate Detections, which maintain cleaned and validated license plate detection records.
- Error Detections, which log failed recognitions for further analysis and system improvement.

○

4. PSEUDO CODE

A. Video Processing Pipeline

1. Initialize YOLOv5 model with license plate detection weights Initialize OCR reader (EasyOCR)

2. Open video file

3. Retrieve video properties (FPS, resolution)

4. Create output video writer

5. for each frame in video do

6. Detect license plates using YOLOv5

7. for each detected plate do

a. Crop license plate region

b. Perform OCR on cropped image

c. if extracted text is valid then

d. Save plate image

e. Store detection metadata in database

f. Draw bounding box on frame

g. Overlay recognized text on frame

h. end if

8. end for

9. Write annotated frame to output video

10. Update processing status in database

11. end for

12. Close video streams and writer

13. Trigger post-processing module

14. end procedure

B. Plate Post-Processing (Pseudocode)

1. procedure PROCESS_PLATES(session_id)

2. Load all plate images for session

3. for each plate image do

a. Apply preprocessing:

- i. Convert to grayscale
- ii. Apply CLAHE
- iii. Apply denoising filter
- iv. Apply adaptive thresholding
- v. Apply morphological dilation

- b. Perform OCR using:
 - i. EasyOCR (greedy search)
 - ii. EasyOCR (beam search)
 - iii. Tesseract (fallback)

- c. Clean extracted text:
 - i. Replace 'O' → '0', 'l' → '1'
 - ii. Remove non-alphanumeric characters

- d. Validate plate format
- e. Compute confidence score
- f. Store processed data in database

4. end for

5. Group similar plates using Levenstein distance
6. Select best candidate from each group
7. Update plate_detections collection

8. end procedure.

The YOLOv5 detection module is set up using the pre-trained Kerem Berke/yolov5m-license-plate model. This allows the system to identify license plates accurately without needing a lot of retraining. A confidence threshold of 0.5 helps balance detection accuracy and reduces false positives. An IoU threshold of 0.45 ensures that overlapping bounding boxes are reliably suppressed. When possible, the detector automatically uses GPU acceleration to improve processing speed.

For the recognition phase, EasyOCR is set up with English language support and GPU optimization. This enhances both the speed and reliability of character extraction. Extracted text is only considered valid if it has at least four characters and does not match any predefined noisy or incorrect OCR patterns. This approach leads to better recognition results.

B. Plate Post-Processing
To further improve the robustness of license plate recognition, a multi-stage post-processing pipeline is applied. The process enhances image quality, refines OCR output, and selects the most credible plate text among alternatives.

5. WEB INTERFACE IMPLEMENTATION

The system's frontend is built with React, while the backend services use Flask (Python). The frontend consists of modular components that help with video uploading, searching, and displaying detection results:

- UploadFootage.jsx – Manages the video uploading process through a clear, step-by-step interface.
- SearchDetection.jsx – Lets users search for plate records using text queries or filters based on dates.
- DetectionResults.jsx – Shows detected plates along with related metadata and visual representations.

The backend offers RESTful APIs to manage video processing requests, run database queries over specific time intervals, and retrieve detection metadata. Communication between the frontend and

backend happens through HTTP requests with JSON payloads, making data exchange efficient and scalable.

For spatial visualization, the system includes the Google Maps API, which shows vehicle detection points and helps generate optimized routes between selected locations. Chart.js is also used to create interactive visualizations, such as charts for detection frequency over time and geographical distribution, improving situational awareness for operators.

6. SIMULATION RESULTS

The proposed system is tested with a dataset of traffic surveillance videos gathered from urban roads, highways, and parking facilities. This dataset includes footage of about 500 unique vehicles, taken in different lighting conditions, weather situations, and camera angles. The main goal of this evaluation is to measure the system's recognition accuracy and processing efficiency in real-world conditions.

A. Detection and Recognition Accuracy

Table 1. presents the detection and recognition performance of the proposed system in comparison with manually annotated ground-truth labels.

Table 1 Detection and recognition accuracy

The improvement over the baseline shows how effective the proposed multi-stage processing strategy is. Character-level accuracy goes up from 82.7% to 91.5%. Full-plate recognition accuracy increases from 76.3% to 88.9%. Additionally, the false-positive rate decreases significantly, which means the overall system is more reliable.

B. Processing Performance

To evaluate computational efficiency, we measured execution times for each core module across various hardware setups. Table II summarizes the average processing times per frame. With GPU acceleration, the system reaches speeds of about 6 to 8 FPS on an RTX 2080 and 8 to 10 FPS on an RTX 3090. This performance is suitable for practical applications that involve analyzing pre-recorded video streams.

| Metric | Raw YOLOv5 | Raw OCR | Full Pipeline |
|-----------------------------|------------|---------|---------------|
| Plate Detection Rate | 94.2% | - | 94.2% |
| Character Recognition Rate | - | 82.7% | 91.5% |
| Full Plate Recognition Rate | - | 76.3% | 88.9% |
| False Positive Rate | 3.8% | 8.2% | 2.1% |

TABLE 2. PROCESSING PERFORMANCE (ms/frame)

| Component | CPU Only | GPU (RTX 2080) | GPU (RTX 3090) |
|-----------|----------|----------------|----------------|
| | | | |

| | | | |
|-------------------------------|-------|-------|-------|
| YOLOv5 Detection | 125.3 | 18.7 | 9.2 |
| OCR (per plate) | 342.1 | 78.5 | 42.3 |
| Post-processing | 56.8 | 56.8 | 56.8 |
| Total (avg. 1.2 plates/frame) | 592.5 | 169.7 | 117.5 |

C. Impact of Environmental Factors

To assess how well the system works in real-world situations, we examined its recognition performance across various environmental conditions. The results are shown in Figure (no). The system reaches its highest accuracy in clear daylight, exceeding 95%. Performance drops somewhat in low-light settings, with an accuracy of 78.5%. It further declines in heavy rain or snowfall, where visual distortions cut accuracy down to 72.3%. Recognition struggles most with partial occlusion, identifying only 67.8% of plates correctly. These findings highlight that lighting and unobstructed plate visibility are key factors for achieving the best ALPR performance.




Fig. 3. Recognition accuracy under different environmental conditions.

D. User Experience Evaluation

A usability evaluation involved 15 participants to examine the overall quality of interaction with the web interface. Users rated different aspects of the system on a 5-point Likert scale as follows.

- Ease of uploading videos: 4.7 / 5
- Search functionality: 4.5 / 5
- Result visualization methods: 4.3 / 5
- Overall system experience: 4.5 / 5

The results show that users viewed the interface as intuitive, responsive, and effective for vehicle surveillance operations. Figure (xxxx) shows a visual representation of the collected user experience ratings.




Fig. 4. User experience ratings for different aspects of the system

7. CONCLUSION AND FUTURE WORK

This study presents a detailed vehicle surveillance framework that combines deep learning-based license plate detection and recognition with a strong post-processing pipeline and an easy-to-use web interface. The system aims to keep recognition accuracy high in real-world settings while offering effective querying and visualization tools for operators.

Experimental evaluations show that the proposed post-processing pipeline significantly boosts recognition performance compared to raw OCR outputs. It improves full-plate accuracy and lowers the number of false positives. Although the system works

well in various conditions, its accuracy drops in extreme weather, low light, or when plates are partially blocked. This points out important areas for future research and improvement.

Future improvements will focus on:

- Better image enhancement techniques to improve performance in tough environmental conditions
- Real-time processing support for continuous video streams
- Integration of vehicle make and model recognition for more detailed analysis
- Privacy-preserving methods to ensure the ethical handling of data
- Compatibility with intelligent transportation systems and smart city platforms

Overall, the proposed system has strong potential for use in traffic management, law enforcement, and urban mobility analysis. By streamlining surveillance workflows and generating useful insights from visual data, this approach helps build safer, smarter, and more efficient transportation systems.

8. ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the anonymous reviewers for their valuable feedback and constructive suggestions, which have greatly contributed to improving the quality of this paper.

9. References

1. X. Li, M. Li, W. Jia, C. Guo, and D. Zhang, "YOLO-based license plate detection and recognition in complex scenes," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 5, pp. 5043–5054, May 2023.
2. M. Hussain, "YOLO-v1 to YOLO-v8, the rise of YOLO and its complementary nature toward digital manufacturing and industrial defect detection," *Machines*, vol. 11, no. 7, p. 677, Jun. 2023, doi: 10.3390/MACHINES11070677.
3. C. Zhang, Y. Xu, and Z. Shi, "YOLOv8-transformer: A fast and accurate object detector for embedded systems," *Electronics*, vol. 12, no. 16, p. 3520, Aug. 2023.
4. Y. Feng, H. Sun, R. Liang, Y. Chen, and D. Wu, "A new single-stage detector combining regression and anchor-free for small object detection," *Remote Sens.*, vol. 15, no. 11, p. 2760, May 2023.



5. N. Vishwakarma. (Nov. 1, 2024). Visualizing Model Insights: A Guide to Grad-CAM in Deep Learning. Analytics Vidhya. Accessed: Mar. 2, 2024. [Online]. Available: <https://www.analyticsvidhya.com/blog/2023/12/gradcam-in-deep-learning>
6. K. S. Rathore, R. R. Battu, M. Kukreja, N. Sahni, and A. Singh, “A review on different techniques for vehicle number plates identification,” in Proc. AIP Conf., vol. 2978, Jan. 2024, Art. no. 060010.
7. K. H. K. Khoshnaw, Z. A. A. Shwany, T. Mustafa, and S. K. Ismail, “Mobile recommender system based on smart city graph,” Indonesian J. Electr. Eng. Comput. Sci., vol. 25, no. 3, pp. 1771–1776, Mar. 2022.
8. N. Thakur, E. Bhattacharjee, R. Jain, B. Acharya, and Y.-C. Hu, “Deep learning-based parking occupancy detection framework using ResNet and VGG-16,” Multimedia Tools Appl., vol. 83, no. 1, pp. 1941–1964, Jan. 2024.
9. T. Mustafa and M. Karabatak, “Feature selection for phishing websites by using naive Bayes classifier,” in Proc. 11th Int. Symp.