

REAL-TIME AUTOMATIC VEHICLE REGISTRATION PLATE
RECOGNITION FOR PAKISTANI TRAFFIC USING YOLOV8S AND OCRMuhammad Jan^{*1}, Hizbullah Mujahid², Hifza Nawaz³, Abdul Basit⁴^{*1,2,3}Lecturer Computer Science, Lasbela University of Agriculture, Water, and Marine Sciences, LUAWMS, Wadh,
Khuzdar, Balochistan⁴Visiting Lecturer Computer Science, Balochistan University of Engineering and Technology, Khuzdar, Balochistan^{*1}muhammadjan@wadh.luawms.edu.pk, ²hizbullah@wadh.luawms.edu.pk, ³hifzanawaz35@gmail.com,
⁴raisjan9@gmail.comDOI: <https://doi.org/10.5281/zenodo.22007204>**Keywords**

Automatic vehicle registration
plate recognition, YOLOv8s,
object detection, optical character
recognition, intelligent
transportation systems, license
plate recognition.

Article History**Received:** 17 May 2026**Accepted:** 13 June 2026**Published:** 16 August 2026**Copyright @Author****Corresponding Author: *****Muhammad Jan****Abstract**

Automatic license plate recognition, or ALPR, plays a key role in intelligent transport and traffic law enforcement. Most conventional ALPR techniques struggle when lighting is poor, plates are tilted, or when dealing with the varied formats used on Pakistani vehicles. Given these challenges, we developed an end-to-end framework combining YOLOv8s for plate detection with OCR for character recognition. We trained and tested the model on a Roboflow dataset of 3,000 Pakistani vehicle images collected under different lighting, occlusion, and angle conditions. On the test set, the system reached 94% mAP@0.5, 94% precision, and 94% recall, while running at 28 FPS on a GPU, which makes real-time use feasible. The results show the model handles low-contrast scenes well. However, it still faces challenges with badly damaged plates, extreme angles, and very fast-moving vehicles. Overall, YOLOv8s proves to be an efficient and accurate choice for traffic monitoring and automated tolling in complex, real-world traffic.

INTRODUCTION

In the era of rapid urbanization and increasing global population, efficient management and monitoring of vehicular traffic have become critical concerns for governments and law enforcement agencies worldwide. Among the various challenges in intelligent transportation systems, accurate vehicle identification remains a key requirement with significant implications for road safety, security, and traffic management. The exponential growth in the number of vehicles has placed immense pressure on existing transportation infrastructure. Traffic congestion, violations, accidents, and criminal activities

involving vehicles are rising at an unprecedented rate, particularly in developing regions. Traditional manual methods of vehicle monitoring are no longer sufficient to handle this scale, leading to delays, human error, and inefficiencies. This has created an urgent need for automated, reliable, and real-time systems that can identify and track vehicles with minimal human intervention. Recent advances in computer vision and deep learning have opened new possibilities for intelligent vehicle identification. Techniques such as Automatic License Plate Recognition (ALPR), vehicle make and model classification, and object detection have shown promising results

in controlled environments. However, real-world deployment faces significant challenges including varying lighting conditions, occlusions, diverse vehicle types, low-resolution images, and complex urban backgrounds. These factors often degrade the performance of existing models and limit their practical applicability. Therefore, there is strong motivation to develop robust vehicle identification frameworks that can operate effectively in dynamic and unconstrained environments. Such systems not only support law enforcement in tasks like stolen vehicle recovery and traffic rule enforcement but also contribute to smart city initiatives by enabling data-driven traffic management and planning. This research aims to address these gaps by proposing an approach that enhances accuracy and reliability of vehicle identification under real-world conditions, thereby contributing to safer and more efficient transportation systems. [1].

Research Problem

Traditional Vehicle Registration Plate Recognition (VRPR) systems often suffer from limited accuracy due to their reliance on manual entry or human observation, leading to higher error rates and slower processing times. Adverse weather conditions, poor lighting, and inconsistent plate positioning further degrade the performance of these systems, compromising their reliability and efficiency in real-world traffic monitoring applications [2]. To address these limitations, this thesis proposes an accurate, real-time, and efficient AVRPR system based on the YOLO (You Only Look Once) object detection model. A single-stage object detector that processes the entire image in one pass. We use it here to detect license plates in real time without separate region-proposal steps. The proposed system aims to automate vehicle plate detection and recognition, thereby reducing human intervention and improving response time in law enforcement, toll collection, and traffic surveillance applications [3].

Objectives

The primary objective of this research is to design and evaluate an efficient Automatic Vehicle Registration Plate Recognition (AVRPR) system for real-time vehicle surveillance. The specific objectives are as follows:

- To develop an end-to-end AVRPR framework using the YOLO (You Only Look Once) object detection model for accurate plate localization and recognition.
- To evaluate the performance of the proposed system on a Roboflow dataset under diverse real-world conditions including varying illumination, occlusion, and plate orientations.
- To achieve high recognition accuracy while maintaining low computational complexity
- for real-time deployment in traffic monitoring and law enforcement applications.

Despite extensive research on AVRPR for standardized plate formats in countries such as China, South Korea, and Saudi Arabia, limited work addresses the challenges posed by diverse and non-uniform Pakistani license plates. Existing deep learning models often require manual cropping and struggle with variations in font, color, and plate dimensions prevalent in Pakistan. To bridge this gap, this paper presents one of the first comprehensive evaluations of YOLOv8s for end-to-end AVRPR on a publicly available Roboflow dataset of Pakistani vehicles. The proposed single-stage framework eliminates separate segmentation steps, thereby reducing computational complexity while achieving real-time inference suitable for deployment in resource-constrained traffic surveillance systems.

2. LITERATURE REVIEW

Early research on vehicle identification primarily relied on handcrafted features and traditional image processing techniques. Methods such as edge detection, color histograms, and SIFT/HOG features were commonly used for vehicle detection and classification. While these approaches achieved reasonable performance in constrained environments, they suffered from poor generalization under varying illumination,

weather conditions, and complex backgrounds. Moreover, their reliance on manual feature engineering made it difficult to scale for large and diverse datasets. With the advent of deep learning, Convolutional Neural Networks (CNNs) have become the dominant approach for vehicle identification tasks. Architectures such as AlexNet, VGG, ResNet, and more recently YOLO and Faster R-CNN have been successfully adapted for vehicle detection, license plate recognition, and make-model classification. Studies have shown that deep models can automatically learn hierarchical features and achieve significantly higher accuracy compared to traditional methods. However, most existing deep learning models are trained and evaluated on benchmark datasets captured under ideal conditions. Their performance often degrades in real-world scenarios involving motion blur, low resolution, nighttime images, and occlusions caused by traffic density. Despite these advancements, several research gaps persist in the field of intelligent vehicle identification. First, there is limited work focusing on region-specific challenges such as diverse vehicle types, non-standard license plates, and unstructured traffic common in developing countries. Second, many current systems operate in isolation and do not integrate identification with broader traffic management or security applications. Third, computational complexity and real-time requirements remain a bottleneck for deployment on edge devices used by law enforcement. Addressing these gaps requires the development of lightweight, robust, and context-aware models that can perform reliably in dynamic environments. This motivates the need for the proposed work to enhance both accuracy and practical applicability of vehicle identification systems.[4].

AVRPR Techniques for Pakistani Number Plates

Several studies have focused on AVRPR for Pakistani license plates, which pose unique challenges due to diverse plate designs, fonts, and lighting conditions. An efficient approach was proposed using image processing and machine

learning on a dataset of 900 vehicle images captured under varying illumination. The method demonstrated a high identification rate, though scalability to larger, more diverse datasets remained a limitation [5].

Another study introduced a color-detection-based AVRPR method for Pakistani plates. The approach extracted the number plate using basic image processing and applied OCR for character recognition. Despite being trained on a limited dataset, it showed promising accuracy and provided a foundation for automated plate recognition in Pakistan [6].

Traditional Image Processing and OCR-Based Methods

Traditional AVRPR systems often rely on multi-stage pipelines: preprocessing, plate localization, character segmentation, and OCR. Preprocessing typically involves grayscale conversion, noise removal using Gaussian or median filters, and contrast enhancement. Plate localization is commonly performed using edge detection techniques such as Canny, followed by rectangular contour analysis based on aspect ratio and size. For character segmentation, projection-based histograms and connected component labeling are widely used. Recognition is then performed using OCR engines like Tesseract or custom classifiers such as AdaBoost with features like Histogram of Oriented Gradients (HOG) and Local Binary Patterns (LBP). One such system achieved 97.3% accuracy on a Bangladeshi dataset using shape validation and tilt correction. However, these methods are highly dependent on manual tuning and struggle with real-world variations in plate orientation and image quality.

Furthermore, traditional approaches face significant limitations when deployed in unconstrained environments. The sequential nature of the pipeline means that an error in any early stage, such as inaccurate plate localization or poor segmentation, propagates and severely impacts final recognition accuracy. These systems also assume relatively uniform lighting and frontal plate views, which rarely hold in practice due to shadows, glare, motion blur, and varying camera

angles. Additionally, the handcrafted features used for classification, such as HOG and LBP, are not robust to scale changes and font variations across different regions and vehicle types. As a result, extensive parameter tuning is required for each new dataset or deployment location, making these systems difficult to maintain and scale.

The emergence of deep learning has started to address many of these shortcomings. End-to-end CNN-based models can learn features directly from raw images, reducing dependence on manual preprocessing and feature engineering. Object detectors like YOLO, SSD, and Faster R-CNN have been adapted to perform plate localization in a single forward pass, while recognition networks such as CRNN and attention-based OCR models handle character sequence prediction without explicit segmentation. Despite these advantages, many deep learning approaches still treat detection and recognition as separate tasks, leading to increased computational overhead. This motivates the need for a unified and lightweight framework such as the proposed AVRPR based on YOLOv8, which can perform both plate detection and character recognition efficiently while maintaining high accuracy under real-world conditions. [7].

Deep Learning-Based Approaches

Recent work has shifted toward Convolutional Neural Networks (CNNs) and YOLO-based models for end-to-end AVRPR. CNNs have been used for both plate detection and character recognition, achieving 88.67% accuracy on a 1700-image dataset. While effective, these approaches often require manual cropping of plates and characters during training [8]. YOLO and other single-stage detectors have gained attention for real-time AVRPR due to their ability to detect and localize plates directly from images without separate segmentation steps. This reduces computational complexity and improves speed, making them suitable for traffic surveillance applications [9].

Research Gap

Although significant progress has been made in AVRPR, most state-of-the-art models are trained and optimized for standardized license plate formats with fixed fonts, colors, and layouts, such as those in China, Korea, and Middle Eastern countries. Pakistani license plates lack a unified standard, exhibiting high variability in font styles, background colors, and aspect ratios, which significantly degrades the performance of existing systems. Furthermore, publicly available large-scale datasets for Pakistani plates remain scarce, limiting the training and evaluation of robust deep learning models. Traditional multi-stage OCR pipelines are also highly sensitive to adverse weather, poor illumination, and non-uniform plate orientation, making them impractical for real-time traffic monitoring. These limitations motivate the development of a robust, single-stage AVRPR framework based on YOLOv8s that can efficiently detect and recognize Pakistani license plates under complex, real-world conditions without requiring manual intervention.

3. METHODOLOGY

The AVRPR - Automated Vehicle Registration Plate Recognition framework is designed as a modular pipeline to ensure robustness and efficiency in real-world traffic scenarios. In the data acquisition phase, images and video frames are collected from roadside surveillance cameras and traffic monitoring systems. To address variations in lighting, resolution, and viewing angles, a diverse dataset containing vehicles during daytime, nighttime, and adverse weather conditions is utilized. This ensures that the trained model generalizes well across different environmental settings. In the preprocessing stage, acquired images undergo several enhancement operations to improve detection accuracy. These include noise reduction, contrast adjustment, and resizing to a standard input dimension compatible with YOLOv8. Data augmentation techniques such as rotation, scaling, brightness variation, and mosaic augmentation are also applied during training to increase dataset diversity and reduce overfitting. The preprocessed data is then

annotated with bounding boxes around license plates and segmented characters, which serves as ground truth for model training. For model training and inference, the proposed framework leverages YOLOv8 due to its balance between detection speed and accuracy. YOLOv8 is first fine-tuned on the license plate detection task to localize plate regions within a vehicle image. Once the plate is detected and cropped, a secondary recognition module based on YOLOv8 or a CNN-OCR hybrid is employed for character segmentation and recognition. During real-time inference, the trained model processes live video streams, detects plates, extracts characters, and outputs the complete registration number with minimal latency. This end-to-end design makes the AVRPR system suitable for deployment in intelligent transportation and law enforcement applications where real-time performance is critical.

A-Dataset

The experimental dataset was collected from Roboflow and comprises vehicle images where license plates are labeled. Image acquisition was performed under varied environmental conditions, encompassing changes in illumination, camera-to-vehicle distance, and plate

rotation angles. The dataset is split into training, validation, and testing sets in a 70:15:15 ratio for model development and evaluation [10].

B-Image Preprocessing and Augmentation

To improve model robustness, the input images undergo preprocessing including resizing, normalization, brightness/contrast adjustment, and histogram equalization. Data augmentation techniques such as rotation, flipping, and scaling are applied during training to increase dataset Splitting the Dataset.

C-Proposed YOLOv8-Based AVRPR Model

The system employs YOLOv8, a single-stage detection model, to locate license plates. This version of YOLO partitions the image into a grid, where every cell forecasts bounding boxes along with their associated confidence values and class probabilities. Non-maximum suppression is then applied to eliminate redundant detections. After plate detection, the cropped plate region is passed to an OCR module for character segmentation and recognition. Optical Character Recognition converts the text image into machine-readable format, enabling comparison with authorized plate records.

Figure 1 illustrates the neural network structure of YOLO.




Figure 1. Neural network structure of YOLO

D-Training Setup

The YOLOv8s model was trained using the Ultralytics framework. Training was performed using the training set, while hyperparameters such as learning rate and regularization were tuned using the validation set. The final model performance was evaluated on the unseen testing set. During inference, the model achieved a speed of 28 FPS, confirming real-time capability.

E-System Workflow

The proposed AVRPR system operates in four sequential stages as shown in Fig. 2: Image Acquisition: Vehicle images or video frames are captured using a fixed camera. Plate Detection: YOLOv8 detects and localizes the license plate region. Character Segmentation: Individual characters are segmented from the detected plate using connected component analysis. Character Recognition: OCR recognizes the segmented characters to obtain the final plate number.




Figure 2. System's operation in four sequential stages

4. RESULTS AND DISCUSSION**Performance Evaluation**

The proposed YOLOv8-based AVRPR system was evaluated on video sequences captured under diverse conditions. The overall plate detection and recognition accuracy achieved is 94.7% on static

images and 88.3% on moving vehicles. The drop in accuracy for moving vehicles is attributed to motion blur, low resolution, and high vehicle speed. Table I summarizes the system's performance under different conditions.

Table 1

System's Performance

Condition	Accuracy (%)	Detection Range (m)
Static vehicle, good illumination	94.7 → 94.0	1-6
Moving vehicle, good illumination	88.3	1-6

Poor illumination / broken plate	76.2	1-4
Distance > 6m	71.5	6-10

Table 2

Yolov8s Performance

Yolov8s Performance Metrics on Test Set

Model	mAP@0.5 (%)	Precision (%)	Recall (%)	FPS
YOLOv8s	94.0	94.0	94.0	28

Effect of Environmental Factors

Experimental results indicate that detection performance degrades with increasing distance and adverse lighting. The system reliably detects plates up to 6 m with a camera angle of 20° relative to the vehicle. Beyond 6 m, accuracy drops to 71.5% due to reduced pixel density. Poor illumination and physical defects such as broken plates or similar characters like 0 vs O, 1 vs I also reduce OCR accuracy.

Real-Time Testing at Khuzdar District

The model was deployed locally in Khuzdar District for real-world validation. Fig. 3 presents sample output images of local vehicles successfully recognized by the system. Real-time inference speed was measured at 28 FPS on GPU using YOLOv8s, confirming suitability for real-time traffic monitoring.



Figure 3. Recognition results locally in Khuzdar

5. CONCLUSION

This paper presented an end-to-end AVRPR framework using YOLOv8s for real-time license plate detection and recognition in Pakistani traffic conditions. The proposed system attained 94.0% mAP@0.5, 94.0% precision, and 94.0% recall on the test dataset, with a processing speed of 28 FPS on GPU, verifying its real-time capability for traffic monitoring and toll management. Compared to traditional multi-stage OCR pipelines, the single-stage YOLOv8s approach reduces computational complexity while maintaining high detection

accuracy. Despite these results, challenges persist in recognizing damaged plates, distinguishing similar characters, and detecting plates from fast-moving vehicles at long distances. Future work will focus on integrating super-resolution techniques and expanding the dataset with more diverse Pakistani plate samples to improve robustness under adverse conditions.

Conflict of Interest

The authors declare that they have no conflict of interest.

REFERENCES

- M. Dahiru, "Advantages of Automated License Plate Recognition Technology," Global Journal of Research in Engineering & Computer Sciences, vol. 3, no. 5, p. 7, 2023.
- S. N. C. Shekhar, "VEHICLE REGISTRATION PLATE DETECTION," UGC CARE Group 1 Journal, vol. 53, p. 6, 2023.
- J. A. Pacaldo, "Utilizing synthetically-generated license plate automatic detection and recognition of motor vehicle plates in Philippines," in in Proc. Int. Conf. Digit. Transf. Appl, Philippines, 2024.
- S. u. Rehman, "An efficient approach for vehicle number plate recognition in Pakistan," Open Artificial Intelligence, vol. 2, no. 3, p. 11, 21 2020.
- S. S. K. M. Javed, "Automatic vehicle number plate recognition,," International Journal for Innovation Science & Technology, vol. 12, pp. 1618-1630, 2022.
- S. Ranjan, "A review paper on automatic vehicle number plate recognition techniques," in in Proc. Int. Conf. Recent Trends Eng. Technol, Nidasoshi, India, 2023.
- J. A. Pacaldo, "Utilizing synthetically-generated license plate automatic detection and recognition of motor vehicle plates in Philippines,," in in Proc. Int. Conf. Digit. Transf. Appl, Philippines, 2021.
- P. Dhar, "A system design for license plate recognition by using edge detection and convolutional neural network," in in Proc. Int. Conf. Comput., Commun., Chem., Mater. Electron. Eng, Bangladesh, 2018.
- M. S. a. S. K. K, "An application of image processing - license plate recognition," International Journal Advance Research Computing Scienec, vol. 8, pp. 5-6, 2017.
- Roboflow, "Roboflow," YOLOv8 License Plate Detection Dataset, [Online]. Available: <https://universe.roboflow.com/kanwal-masroor-gv4jr/yolov8-license-plate-detection/dataset/3>. [Accessed 03 01 2026].