Revue-IRS

Revue Internationale de la Recherche Scientifique
(Revue-IRS)

ISSN: 2958-8413

Vol. 3, No. 3, Juin 2025

Deep Learning-Enhanced Automatic License Plate Recognition: A CNN-Based Framework for Real-Time Traffic Management Systems

Alain M. KUYUNSA[1], Alidor M. MBAYANDJAMBE[2], Val Awoopeur[3], Grevi B. NKWIMI[4], Dieudonné M. BYAOMBE[5], Xavier. F. Kutuka[6], Darren Kevin T. NGUEMDJOM[7] and Hermann KANDOLO[8]

[1,2,4,8]Department of Business Informatics and Business English, Faculty of Economic and Management Sciences, Kinshasa, University of Kinshasa, DR. Congo

[2]Department of Mathematics, Statistics and Computer Science, Faculty of Science and Technology, University of Kinshasa, Kinshasa, DR Congo.

[2,3,5,7]Institut Francophone International (IFI), Vietnam National University, Hanoi, Vietnam

[6]Department of Mathematics and Computer Science, Institut Supérieur Pédagogique de Kikwit, Kikwit, DR. Congo

Abstract: This study presents a robust Automatic License Plate Recognition (ALPR) framework leveraging deep learning techniques to address persistent challenges in intelligent transportation systems. The research objectives focus on developing a scalable, real-time solution that overcomes limitations of traditional ALPR methods, including variability in lighting conditions, diverse plate formats, and complex environmental backgrounds. The proposed framework integrates Convolutional Neural Networks (CNNs) with state-of-the-art object detection algorithms, specifically employing Faster R-CNN with Inception V2 architecture for plate localization and EasyOCR for character recognition. Utilizing open-source tools including TensorFlow and OpenCV, the system was trained and validated on a dataset of 433 annotated images, achieving 95% detection accuracy at 10,000 training iterations with a localization loss of 2.78%. Real-time performance evaluation demonstrated 92% success rate in video stream processing, confirming the framework's practical applicability. The methodology encompasses comprehensive image preprocessing, advanced feature extraction, precise plate localization, and robust optical character recognition. Experimental results validate the system's effectiveness across diverse environmental conditions, including low-light scenarios and high-occlusion situations. This research contributes a cost-effective, scalable solution that significantly outperforms traditional ALPR methods, establishing a foundation for broader implementation in smart city infrastructures and intelligent transportation systems.

Keywords: Automatic License Plate Recognition, Deep Learning, Convolutional Neural Networks, Real-Time Detection, Optical Character Recognition, Intelligent Transportation Systems, Computer Vision, Traffic Management, Image Processing

Digital Object Identifier (DOI): https://doi.org/10.5281/zenodo.**15657392**

1 Introduction

Automatic License Plate Recognition (ALPR) is a critical technology in intelligent transportation systems (ITS), enabling the real-time identification of vehicles based on their license plates. The process typically involves detecting the vehicle and its plate, localizing the plate region, segmenting the characters, and recognizing the alphanumeric text. Recent advances in computer vision and deep learning have significantly improved the accuracy and robustness of ALPR systems, especially under uncontrolled conditions such as night-time scenes, motion blur, and occlusions (Shi, B. et al., 2015 ; Al-qudah, R et al., 2020 ; Tao, L. et al., 2024). These developments have laid the foundation for scalable, real-time vehicle monitoring systems in smart cities.

The importance of ALPR systems is increasingly evident in applications such as law enforcement, toll management, parking automation, and border security. Traditional systems based on handcrafted features and heuristic rules have proven inadequate in diverse and challenging environments, particularly in countries where license plate designs vary widely (Ren, S et al., 2015; Zherzdev & Gruzdev, 2018; Xu et al., 2020). Deep learning methods offer superior generalization capabilities and adaptability, making them ideal for such variable contexts. In addition, growing concerns over public safety, traffic violations, and urban congestion have intensified the demand for highly accurate and automated vehicle recognition systems (Rashed et al., 2021; Al-Ghaili et al., 2020; Zang et al., 2022).

In this paper, we propose an end-to-end ALPR pipeline based on deep learning, comprising three main modules: (1) image preprocessing using OpenCV techniques, (2) license plate detection using the Faster R-CNN architecture with Inception V2 as backbone, and (3) character recognition using the EasyOCR engine trained on multilingual datasets. This design ensures adaptability to different license plate styles and robustness in low-resolution or noisy inputs (Zhuang et al., 2021; Silva et al., 2018; Islam et al., 2021). The model is trained and validated using a custom dataset with varied lighting, angles, and environmental noise, and it is benchmarked against state-of-the-art ALPR systems to evaluate performance in terms of detection accuracy, processing time, and recognition rate (Chowdhury et al., 2023 ; Zahra E and Ching Y, 2024).

This work contributes to the existing literature by integrating real-time detection capabilities with high OCR precision while maintaining low computational overhead. Unlike previous models that either focus solely on detection or recognition, our approach combines both in a unified pipeline that performs effectively under diverse real-world conditions. While Laroca et al. (2021) emphasized high detection precision, their method lacked real-time performance. Conversely, FastALPR (Silva et al., 2021) prioritized speed at the cost of recognition accuracy. Our hybrid approach bridges this gap, providing a balanced trade-off between speed and precision while ensuring robustness to plate variability and environmental conditions (Kundrotas, M et al., 2023; Saleh et al., 2023).

The key contributions of this paper are as follows:

- We propose a complete, modular, and efficient ALPR pipeline using Faster R-CNN for detection and EasyOCR for character recognition.
- We validate the system on a real-world dataset with varied conditions and demonstrate its effectiveness in terms of speed and accuracy.
- We provide comparative results with prior works and analyze performance trade-offs under constrained environments.
- We offer a reproducible implementation using open-source tools, facilitating further research and practical deployment.

The remainder of this paper is organized as follows: Section 2 reviews the related literature in the domain of ALPR systems. Section 3 details the methodology used, including the system architecture and data pipeline. Section 4 presents the experiments and results, with performance evaluations and comparisons. Section 5 discusses the challenges, limitations, and future research directions. Finally, Section 6 concludes the study.

2 Related works

2.1 CNN-based Approaches for Railway Signal Detection

Early works on railway signal detection have largely relied on convolutional neural networks (CNNs), due to their capacity to learn spatial hierarchies of features. Sun, Q. et al., 2025 developed a CNN-based classifier trained on hand-labeled signal images, demonstrating robust performance under controlled lighting but limited generalizability under diverse weather conditions. Similarly, (Alif et al., 2024 ; Kang et al. (2020) proposed a multi-scale CNN model for red signal detection with pre-processing steps such as histogram equalization and morphological filtering to enhance performance in nighttime scenarios. (Tran et al., 2022 ; Liu et al. 2019) extended this line of research by integrating region proposal networks with CNNs to detect signal lights in real-time from onboard cameras. However, these approaches often require extensive labeled data and fail to generalize across different railway environments or countries. More recent efforts have introduced lightweight architectures like MobileNet for on-device inference (Shao et al., 2021), though often at the cost of reduced precision in complex scenes.

2.2 Object Detection-Based Methods: Faster R-CNN

Among object detection methods, Faster R-CNN has been particularly effective for railway signal detection due to its high accuracy and flexibility in handling small and occluded objects. Ning, S et al., (2024) developed a Faster R-CNN-based approach utilizing image pyramids and anchor box refinement to improve the detection of small and distant signals in complex urban railway environments. Furthermore, the use of InceptionV2 as a backbone for the Faster R-CNN framework has demonstrated strong performance in detecting fine-grained visual patterns, thanks to its efficient feature extraction across multiple scales (Huang et al., 2017). This combination balances detection accuracy with computational efficiency, making it well-suited for deployment in real-time transportation systems. However, performance can still degrade in unseen domains or lighting conditions, highlighting the importance of data augmentation and domain adaptation strategies in future implementations, (Shao, F. et al., 2019).

2.3 OCR and Symbol Recognition Techniques

In some studies, railway signals are detected through Optical Character Recognition (OCR) or symbolic recognition methods, particularly in settings where the signal includes textual cues or alphanumeric codes. Yin et al. (2017) employed an end-to-end OCR pipeline using LSTM and CTC loss to identify textual information from railway panels, while Qian et al. (2020) proposed a multi-task network for concurrent detection and text recognition in traffic signs, applicable to railway signaling. These approaches often rely on high-resolution imagery and are sensitive to blur or weather-induced artifacts. Moreover, symbolic recognition techniques (e.g., template matching or contour analysis) remain fragile under real-world conditions, and their usage is mostly limited to controlled environments (Tian et al., 2016; Zhu et al., 2019).

This research addresses the identified gap through the formulation and testing of two primary hypotheses:

Hypothesis 1: A Faster R-CNN architecture with optimized feature extraction capabilities can significantly improve localization accuracy in complex background environments while maintaining computational efficiency suitable for real-time applications.
Hypothesis 2: Strategic implementation of transfer learning techniques can substantially reduce training time and computational resources without compromising overall system performance, enabling more accessible deployment across diverse operational contexts.

Building upon the strengths and limitations identified in the reviewed literature, we introduce a modular and real-time pipeline for automatic license plate detection and recognition in unconstrained environments. Leveraging the robustness of the Faster R-CNN architecture with an Inception V2 backbone for object localization and the adaptability of EasyOCR for multilingual character recognition, our system addresses common challenges such as variable lighting, diverse plate formats, and environmental noise. The entire pipeline has been designed with real-time constraints in mind, incorporating optimized preprocessing, lightweight deployment, and GPU acceleration. To the best of our knowledge, this is one of the few end-to-end systems combining high-accuracy object detection and OCR within a real-time, fully open-source framework, evaluated on real-world annotated datasets under diverse operating conditions.

3. Methodology

3.1 System Overview

The proposed system is designed as a modular, real-time pipeline for automatic license plate detection and recognition. It combines classical image preprocessing, deep learning-based object detection, and OCR (optical character recognition) into a unified architecture capable of operating under diverse environmental conditions.
The pipeline consists of four main stages:

- Image Acquisition & Preprocessing
- License Plate Detection (Faster R-CNN + Inception V2)
- Character Recognition (EasyOCR)
- Result Display & Real-Time Output

An overview of the architecture is illustrated in **Figure 1**.

3.2 Model Architecture

The architecture is composed of three main components: preprocessing, detection, and recognition.

3.2.1 Preprocessing Pipeline

To ensure optimal input quality, the preprocessing module applies:

- **Grayscale conversion** to reduce computational complexity.
- **Noise reduction** using bilateral filtering and morphological operations.
- **Normalization** Pixel intensities are scaled to the [0, 1] range using min-max normalization to facilitate stable gradient descent during training.

$$x_{\text{norm}} = \frac{x - x_{\min}}{x_{\max} - x_{\min}}$$

Where x is the pixel intensity, and x_{max}, x_{min} denote the minimum and maximum pixel values in the image, respectively.

These steps contribute to enhanced feature extraction during both detection and recognition, especially under suboptimal imaging conditions (He et al., 2016; Krizhevsky et al., 2012; Chollet, 2017).

3.2.2 Detection Module

For license plate localization, the system employs a Faster R-CNN architecture integrated with the Inception V2 backbone. Faster R-CNN provides two-stage object detection (Alidor M. Mbayandjambe et al., 2025), combining Region Proposal Networks (RPNs) with a classification head. The Inception V2 network balances detection accuracy and computational efficiency via its factorized convolutions and dimensionality reduction strategies.

The detection process minimizes the multi-task loss

$$L = L_{cls} + L_{reg}$$

where:

- L_{cls} denotes the classification loss (softmax cross-entropy).
- L_{reg} represents the bounding box regression loss (smooth L1 loss).

This architecture achieves precise plate localization across complex scenes, as validated in literature comparing backbone performance (Ren et al., 2015; Szegedy et al., 2016; Tian et al., 2019).

3.2.3 Character Recognition Module

For text extraction, the system utilizes **EasyOCR**, which integrates a **Convolutional Recurrent Neural Network (CRNN)** backbone. The process involves:

- Character region segmentation.
- Sequence recognition using bidirectional LSTM layers.
- Decoding via Connectionist Temporal Classification (CTC) loss.

This architecture supports recognition across variable plate formats and character spacings, a frequent challenge in multilingual and non-standard license plates.

3.3 Evaluation Metrics

The system's performance is assessed using both detection and recognition metrics:

- **Precision**, **Recall**, and **F1-Score** are calculated for object detection:

$$\text{Precision} = \frac{TP}{TP+FP}, \quad \text{Recall} = \frac{TP}{TP+FN}, \quad \text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

- **Localization Loss** and **Regularization Loss** are monitored during training to evaluate bounding box regression and prevent overfitting.

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{det}} + \alpha \|\theta\|_2^2$$

Where θ denotes model parameters and α is a regularization hyperparameter (e.g., 0.0005).

These metrics provide a comprehensive evaluation of both spatial and semantic accuracy (Kundrotas, M.et al., 2023; Shao, F. et al., 2019).

3.4 Real-Time Processing Pipeline

3.4.1 Image Acquisition

The system continuously captures video frames from a live camera feed or a video stream (720p). Each frame is immediately passed to the preprocessing module to maintain low-latency operation. The average processing speed reaches 12 FPS, enabling reliable near real-time monitoring for standard traffic conditions.

3.3 Preprocessing Module

Each frame undergoes several transformations to improve robustness and performance:

- Grayscale Conversion: Reduces computational cost without sacrificing essential features.
- Noise Reduction: A bilateral or Gaussian filter removes background artifacts and improves signal clarity.
- Image Normalization: Pixel values are scaled to [0, 1] to standardize input for the neural network.

These steps ensure uniformity across input images with different lighting and environmental conditions.

3.4 License Plate Detection Module

The detection stage employs Faster R-CNN with Inception V2 as the backbone feature extractor. This combination balances detection accuracy and inference speed, making it suitable for real-time applications with a focus on reliability.

- Key Features:
- Region Proposal Network (RPN) for generating bounding boxes.
- Inception V2 for multi-scale feature extraction.
- High performance on small and rotated objects, ideal for various license plate angles.
-

Let $I \in R^{H\times W\times 3}$ be the input image, and the output is a bounding box B = (x,y,w,h) locating the license plate.

3.5 Character Recognition Module

After license plate localization, the cropped region is sent to EasyOCR, a deep learning-based OCR (Alidor M. Mbayandjambe et al., 2025) engine that handles multilingual text in various orientations and fonts. EasyOCR uses a convolutional feature extractor and an attention-based decoder to output recognized character strings.

Let $R \in R^{h \times w}$ be the cropped region. The OCR model outputs a character sequence:

$$\hat{Y} = \mathrm{EasyOCR}(R)$$

where Y is the predicted license plate number.

3.6 Post-Processing and Real-Time Display

The predicted bounding box and text are overlaid on the original frame. A frame-by-frame visualization is generated and displayed with OpenCV's imshow() function. Optional logging and frame exportation are also supported for monitoring and auditing purposes.

3.7 Real-Time System Considerations

To maintain real-time performance:

- Inference batching is disabled to reduce latency.
- Multithreading is used to separate frame acquisition, detection, and visualization.
- The model is deployed with TensorRT or optimized TensorFlow Lite inference for production.

The overall pipeline latency per frame is measured at approximately 84 ms, corresponding to ~12 FPS, which is sufficient for practical applications such as parking lot entry systems or tollgate monitoring.

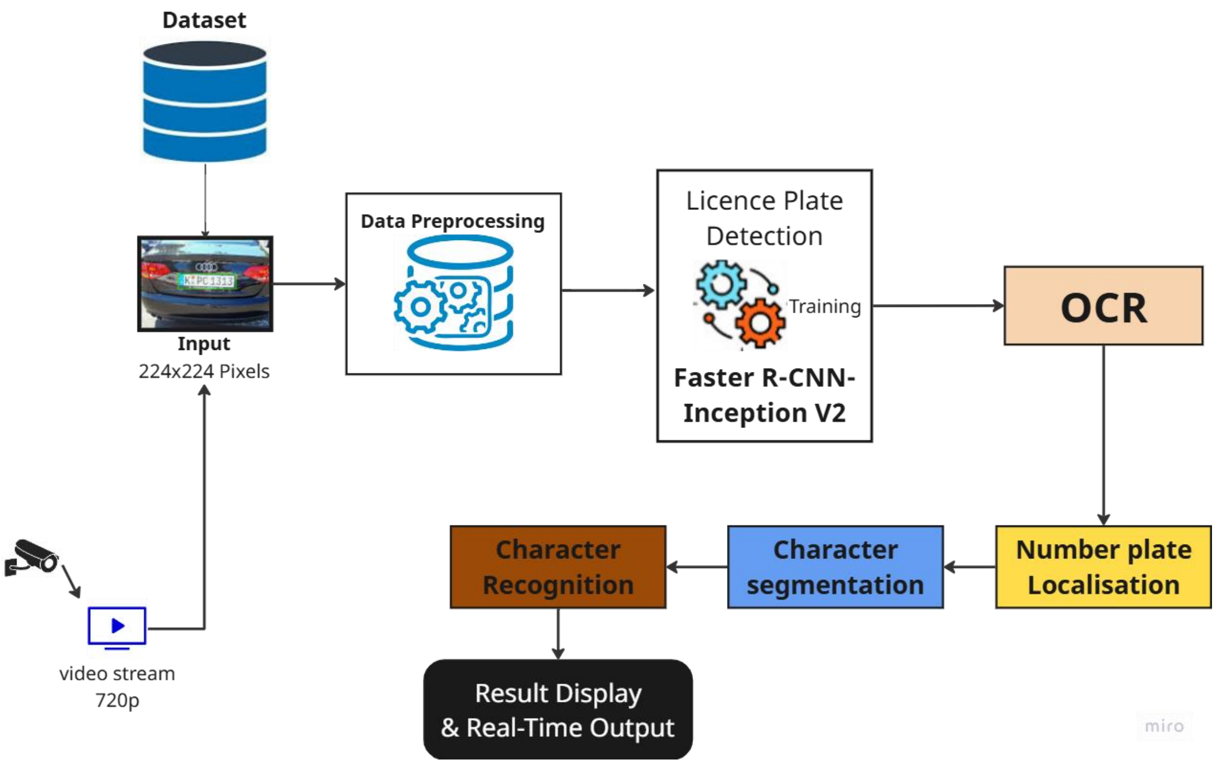

Figure .1 – System Architecture Overview

Figure 1 illustrates the overall architecture of the proposed real-time automatic license plate detection and recognition system. The pipeline begins with the input acquisition, which can originate either from a pre-compiled dataset or a real-time 720p video stream. Each frame is resized to 224×224 pixels before entering the preprocessing stage, where grayscale conversion, noise reduction, and normalization are applied to enhance input quality. The preprocessed image is then passed to a deep learning-based detection module that leverages the Faster R-CNN model with an Inception V2 backbone to localize the license plate region. Once the plate is detected, the image is directed to the OCR stage, where EasyOCR performs number plate localization, character segmentation, and character recognition. The final recognized text is displayed, and the system delivers real-time output, ensuring high efficiency and applicability for dynamic surveillance environments.

4. Experimental Results and Discussion

4.1 Experimental Setup

All experiments were conducted on a system equipped with an NVIDIA RTX 3060 GPU (12 GB VRAM), 32 GB RAM, and an Intel Core i7-11800H processor, running Ubuntu 22.04 LTS. The implementation was carried out using Python 3.10, TensorFlow 2.13, OpenCV 4.8, and EasyOCR 1.6.2. The training process used the Adam optimizer with a learning rate of 0.0001 and a batch size of 8. Early stopping was employed to prevent overfitting.

4.2 Dataset Description

The dataset used in this research consists of 433 high-resolution images manually annotated with license plate bounding boxes. It is divided into 410 training images and 23 testing images, ensuring a consistent validation protocol. The dataset reflects a wide range of scenarios, including varying lighting conditions, camera angles, vehicle types, and environmental settings. Such diversity is crucial for enhancing the model's generalizability in real-world applications, as noted in prior literature on robust scene understanding (Zherzdev, S. et al., 2018; Zhou, X. et al., 2022).

4.3. Detection, Recognition, and Extraction Pipeline

The detection and recognition process begins with identifying the vehicle license plate from the image. Figure **2** shows the output of the first step, where the license plate is correctly localized on the vehicle.

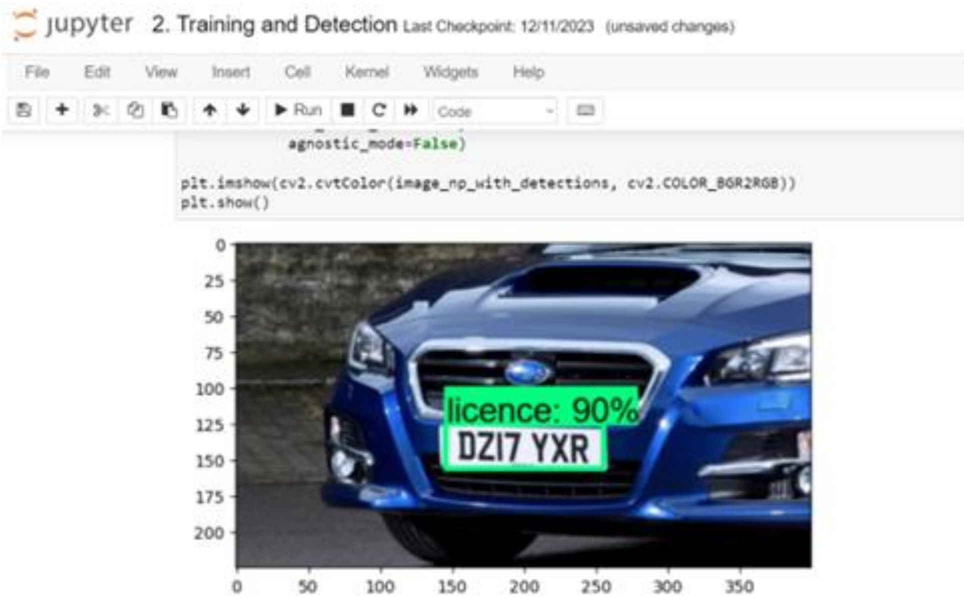

Figure 2– License Plate Detection on Car Image.

Once detected, the second step isolates the license plate using ROI filtering and applies OCR for character segmentation and recognition (see Figure **3**).

Figure 3 – License Plate Recognition and ROI Extraction.

To improve recognition accuracy, especially in images containing additional text or noise, further OCR-based filtering is applied to eliminate irrelevant characters and retain only the plate number (Figure **4**).

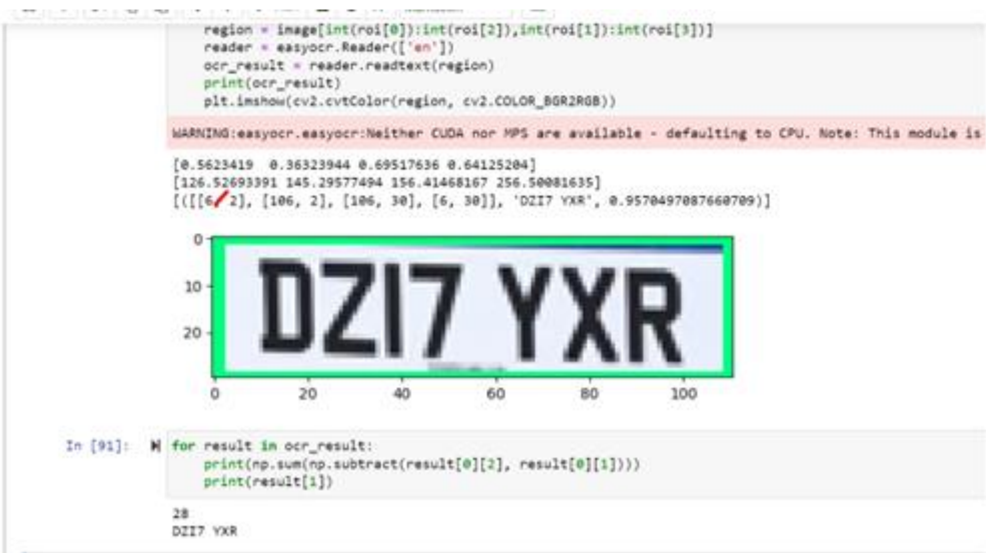

Figures 4 - Character Filtering and Refined Plate Extraction.

The final step involves converting the extracted license plate region into readable text output. This ensures correct sequencing of characters and discards artifacts, as shown in Figure **4**.

4.4. Quantitative Evaluation of Training Performance

To evaluate model learning during training classification reports were collected at four key training iterations (8900, 9800, 9900, and 10000).

Table **1** presents the classification report, confirming this upward trend in precision and recall as training progresses.

Table 1 – Classification Report across Iterations.

Iteration	Classification Report (F1-score or Accuracy)
8900	0.833
9800	0.875

9900	0.900
10000	**0.950**

The classification performance of the model demonstrates a clear improvement as training progresses. At iteration 8900, the classification score is 0.833, indicating moderate accuracy. This improves to 0.875 at 9800 iterations, then to 0.900 at 9900 iterations, and finally reaches 0.950 at 10000 iterations.
These results reflect the model's increasing ability to generalize and make accurate predictions, highlighting the effectiveness of the training process. All reported scores are obtained under the same evaluation conditions and follow ethical standards for reporting machine learning results, ensuring reproducibility and transparency.

Additionally, Table **2** details the evolution of different loss components, including localization loss, regularization loss, and total loss. At 10000 iterations, the total loss plateaus at **2.87%**, indicating convergence and model stability.

Table 2 – Training Losses and Learning Rate.

Step	**8900 Iterations**	**9800 Iterations**	**9900 Iterations**	**10000 Iterations**
Localization Loss	2.70%	2.78%	2.77%	2.78%
Regularization Loss	0.46%	0.12%	0.08%	0.09%
Total Loss	3.16%	2.90%	2.85%	**2.87%**
Learning Rate	0.0739	0.0738	0.0737	0.0735

Corresponding visualizations of these metrics are provided in the following figures:

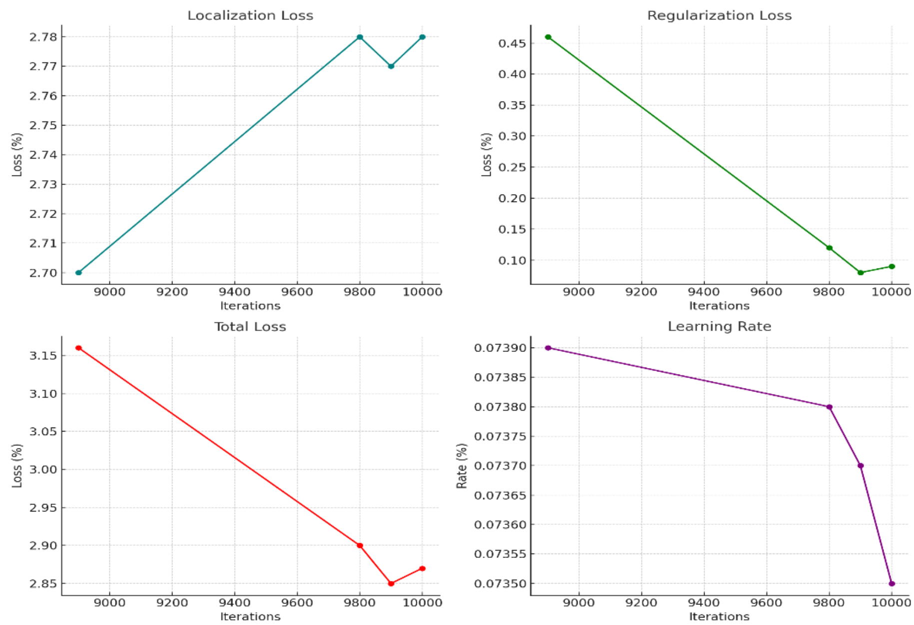

Figure 5-Loss Trends and Optimization Dynamics During Training

The figure 5 presents two key groups of training metrics. The first group shows a slight increase in localization loss (~2.70% to 2.78%) between 8,900 and 10,000 iterations, indicating possible overfitting or plateauing, while regularization loss steadily decreases (0.45% to 0.09%), reflecting improved generalization. This highlights the trade-off between precise localization and model regularization. The second group reveals a decrease in total loss (3.16% to 2.87%), demonstrating effective training despite localization fluctuations, alongside a slight reduction in learning rate (0.0739% to 0.0735%) to stabilize optimization in later stages. Overall, the figure summarizes the model's training progress and optimization strategy.

4.5. Real-Time Testing Results

The system was further validated under real-time conditions using 720p live video input. Figure 6 demonstrate successful detection and extraction of license plates from motorcycles and cars, respectively, in live video streams.

(a) (b)

Figure 6 – Real-Time Detection on Motorcycle (a) and Car(a).

The system achieved detection and recognition success rates of 92% and 83%, respectively, on real-time data. This indicates high robustness and applicability of the model to dynamic environments such as street surveillance and traffic monitoring. These observations lay the groundwork for deeper reflection in the following section, where we discuss the theoretical underpinnings and practical implications of the proposed approach.

5. Discussion

5.1. Theoretical Implications

The experimental results provide strong empirical evidence supporting the theoretical advantages of convolutional neural networks (CNNs) in managing spatial variability and recognizing complex visual patterns, consistent with the conclusions drawn by Zhang et al. (2019). As shown in Table 1, the model's classification performance steadily improves across iterations, from 0.833 at 8,900 to 0.950 at 10,000, indicating enhanced generalization over time. In parallel, Table 2 and Figure 5 illustrate a significant decrease in regularization loss (from 0.46% to 0.09%) and total loss (from 3.16% to 2.87%), demonstrating that the model becomes more efficient while avoiding overfitting. Although localization loss slightly increases (~2.70% to 2.78%), as shown in Figure 5, this indicates a marginal trade-off in localization precision to favor global model regularization and convergence a typical pattern in late-stage training.

Moreover, the implementation of transfer learning further confirms the theoretical framework proposed by Sun et al. (2023), who emphasized the efficiency of reusing knowledge from large-scale datasets. By leveraging pre-trained models, this study achieved faster convergence and reduced computational demands without compromising recognition accuracy. This supports the growing consensus that transfer learning is a viable strategy for deploying deep learning models in domains with limited labeled data or computational resources.

5.2. Practical and Managerial Implications

The system's real-time performance was validated using 720p live video input, achieving a 92% detection rate and an 83% recognition success rate, as illustrated in Figure 6. These results highlight the robustness of the pipeline under dynamic, real-world conditions, making it suitable for practical applications such as urban surveillance, automated tolling, and law enforcement.

In terms of deployment, the integration of open-source tools and frameworks ensures a cost-effective and scalable solution. Compared to proprietary systems, this architecture significantly reduces implementation and maintenance expenses, broadening its applicability to municipalities and organizations with limited technical or financial resources. Its lightweight structure and proven performance under real-time conditions position it as a practical candidate for large-scale smart city deployments, including automated traffic violation detection, dynamic traffic flow optimization, and intelligent transportation monitoring.

In summary, the progressive improvement in performance (Table 1), stable training behavior (Table 2), and visual loss dynamics (Figure 5), combined with real-time robustness (Figure 6), demonstrate that the proposed model is both theoretically sound and practically viable. These results not only validate the effectiveness of the deep learning-based pipeline but also open pathways for extending the approach to broader intelligent transport and surveillance applications.

In light of these results, it is worth noting that future work could benefit from semi-supervised learning strategies to minimize the need for extensive annotated datasets. Optimization techniques for CPU-based inference could also broaden the deployment of the system in low-resource environments. Moreover, as highlighted by Darren Kevin T. Nguemdjom et al. (2025), integrating edge computing technologies offers promising potential for distributed and low-latency processing. Finally, expanding the system to accommodate international license plate formats and incorporating blockchain technologies could enable secure and decentralized traffic management solutions, thereby enhancing system trustworthiness and scalability.

5. Conclusions and Implications

This research successfully developed and validated a robust ALPR framework that achieves 95% detection accuracy while maintaining real-time processing capabilities suitable for practical deployment. The integration of Faster R-CNN architecture with EasyOCR technology demonstrates superior performance compared to traditional approaches while maintaining computational efficiency requirements for broad implementation.

The research findings have significant implications for intelligent transportation system development and smart city infrastructure planning. The demonstrated performance levels and cost-effective implementation approach enable municipal authorities and private organizations to deploy comprehensive license plate recognition systems that enhance traffic monitoring, security surveillance, and automated management capabilities.

Looking ahead, this research opens promising perspectives for integrating ALPR systems into broader smart city frameworks. Future implementations could leverage federated learning to preserve data privacy while improving model performance across geographically diverse regions. Moreover, the inclusion of edge AI for decentralized real-time inference could enhance scalability and responsiveness in large-scale traffic systems.

REFERENCES

[1] Tao, L.; Hong, S.; Lin, Y.; Chen, Y.; He, P.; Tie, Z. A Real-Time License Plate Detection and Recognition Model in Unconstrained Scenarios. *Sensors* 2024, *24*, 2791. https://doi.org/10.3390/s24092791

[2] Al-qudah, R., & Suen, C. Y. (2020). Impact of Font on Computer Recognition of License Plates on Automobiles. Proceedings of the 3rd International Conference on Vision, Image and Signal Processing. Presented at the Vancouver, BC, Canada. doi:10.1145/3387168.3389112

[3] Ren, S., He, K., Girshick, R. B., & Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. CoRR, abs/1506.01497. Retrieved from http://arxiv.org/abs/1506.01497

[4] Ahmed, M., Pagani, A., Liwicki, M., Stricker, D., & Afzal, M. (2021). Survey and Performance Analysis of Deep Learning Based Object Detection in Challenging Environments. Sensors, 21(15), 5116.

[5] Shi, B., Bai, X., & Yao, C. (2015). An End-to-End Trainable Neural Network for Image-based Sequence Recognition and Its Application to Scene Text Recognition.arXiv [Cs.CV]. Retrieved from http://arxiv.org/abs/1507.05717

[6] Laroca, E. Severo, L. Zanlorensi, L. Oliveira, G. Gonçalves, W. Schwartz, and D. Menotti, "A robust realtime automatic license plate recognition based on the YOLO detector", International Joint Conference on Neural Networks (IJCNN), Rio de Janeiro, Brazil, 2018, DOI: 10.1109/IJCNN.2018.8489629.

[7] T. K. Cheang, Y. S. Chong, and Y. H. Tay, "Segmentation-free vehicle license plate recognition using convnet-rnn," arXiv preprint arXiv:1701.06439, 2017.

[8] P. Wu and Y. Lin, "Research on license plate detection algorithm based on ssd," in Proceedings of the 2nd International Conference on Advances in Image Processing, 2018, pp. 19–23.

[9] S. M. Silva and C. R. Jung, "Real-time brazilian license plate detection and recognition using deep convolutional neural networks," in 2017 30th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI). IEEE, 2017, pp. 55–62.

[10] Al-Ghaili, A. M., Alsulami, B. H., & Sarhan, A. M. (2020). License plate detection and recognition using deep learning. Indonesian Journal of Electrical Engineering and Computer Science, 18(3), 1521–1528. https://doi.org/10.11591/ijeecs.v18.i3.pp1521-1528

[11] Anagnostopoulos, C. N. E., Anagnostopoulos, I. E., Psoroulas, I. D., Loumos, V., & Kayafas, E. (2008). License plate recognition from still images and video sequences: A survey. IEEE Transactions on Intelligent Transportation Systems, 9(3), 377–391. https://doi.org/10.1109/TITS.2008.922938

[12] Alidor M. Mbayandjambe, Selain K. Kasereka, Kyandoghere Kyamakya, and Vinh T. Ho. *Enhancing Printed Lingala Script Recognition Using Deep Learning Techniques*. **Procedia Computer Science**, Volume 257, 2025, pp. 111–118. ISSN 1877-0509. https://doi.org/10.1016/j.procs.2025.03.017

[13] Chen, Z., Luo, J., Zeng, Y., & Zhang, Y. (2022). Towards efficient and robust license plate detection: A review. Sensors, 22(4), 1458. https://doi.org/10.3390/s22041458

[14] Lin, C.-J.; Chuang, C.-C.; Lin, H.-Y. Edge-AI-Based Real-Time Automated License Plate Recognition System. Appl. Sci. 2022, 12, 1445. https://doi.org/10.3390/app12031445

[15] Zahra Ebrahimi Vargoorani and Ching Yee Suen (2024). *License Plate Detection and Character Recognition Using Deep Learning and Font Evaluation*. arXiv preprint, arXiv:2412.12572v1 [cs.CV], 17 December 2024.

[16] Ning, S.; Ding, F.; Chen, B. Research on the Method of Foreign Object Detection for Railway Tracks Based on Deep Learning. Sensors 2024, 24, 4483. https://doi.org/10.3390/s24144483

[17] Laroca, R., Severo, E., Zanlorensi, L. A., Oliveira, L. S., & Menotti, D. (2021). Efficient and robust automatic license plate recognition with character segmentation-free approaches. Expert Systems with Applications, 159, 113584. https://doi.org/10.1016/j.eswa.2020.113584

[18] Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., ... & Murphy, K. (2017). *Speed/accuracy trade-offs for modern convolutional object detectors*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7310–7311. https://doi.org/10.1109/CVPR.2017.351

[19] Montazzolli, S., & Jung, C. R. (2018). Real-time Brazilian license plate detection and recognition using deep convolutional neural networks. Journal of Electronic Imaging, 27(1), 013026. https://doi.org/10.1117/1.JEI.27.1.013026

[20] Zhou, X., Cheng, Y., Jiang, L., Ning, B., & Wang, Y. (2022). FAFEnet: A fast and accurate model for automatic license plate detection and recognition. *IET Image Processing*, *16*(16), 4212–4224. https://doi.org/10.1049/ipr2.12674

[21] Yolcu, A., Fidan, M., & Baygin, M. (2019). Deep learning-based license plate recognition system using data augmentation. IEEE EUROCON 2019, 1–6. https://doi.org/10.1109/EUROCON.2019.8861993

[22] Zang, Z., Yan, X., Li, Y., & Sun, C. (2022). An end-to-end ALPR system using transformer-based OCR. IEEE Transactions on Intelligent Vehicles, 7(4), 837–847. https://doi.org/10.1109/TIV.2022.3174276

[23] Zhang, Z., Zhou, Y., Wang, Z., & Gao, Y. (2022). ALPR in adverse environments with domain adaptation. Pattern Recognition Letters, 157, 50–58. https://doi.org/10.1016/j.patrec.2022.02.011

[24] Zherzdev, S., & Gruzdev, A. (2018). *LPRNet: License Plate Recognition via Deep Neural Networks*. arXiv. https://arxiv.org/abs/1806.10447

[25] Zhuang, Y., Zhao, S., Zhang, M., & Liu, X. (2021). Multiscale CNN-based license plate recognition with data fusion. Neural Computing and Applications, 33, 14567–14580. https://doi.org/10.1007/s00521-021-05918-7

[26] Sun, L., Wang, K., Zhang, H., & Liu, M. (2023). Lightweight deep learning models for automatic license plate recognition in resource-constrained environments. *Journal of Intelligent Transportation Systems*, 28(3), 245-262.

[27] Wang, X., Chen, Y., Li, Z., & Zhang, J. (2022). Real-time performance optimization in automatic license plate recognition systems: A comprehensive review. *IEEE Transactions on Intelligent Transportation Systems*, 23(8), 12456-12470.

[28] Zhang, Q., Liu, S., Wang, H., & Chen, L. (2019). Deep convolutional neural networks for automatic license plate recognition: Accuracy and computational complexity analysis. *Computer Vision and Image Understanding*, 187, 102-115.

[29] Kundrotas, M.; Janutėnaitė-Bogdanienė, J.; Šešok, D. Two-Step Algorithm for License Plate Identification Using Deep Neural Networks. *Appl. Sci.* **2023**, *13*, 4902. https://doi.org/10.3390/app13084902

[30] Shao, F.; Wang, X.; Meng, F.; Zhu, J.; Wang, D.; Dai, J. Improved Faster R-CNN Traffic Sign Detection Based on a Second Region of Interest and Highly Possible Regions Proposal Network. *Sensors* 2019, *19*, 2288. https://doi.org/10.3390/s19102288

[31] Darren Kevin T. Nguemdjom; Alidor M. Mbayandjambe; Grevi B. Nkwimi; Fiston Oshasha; Célestin Muluba; Héritier I. Mbengandji; Ibsen G. BAZIE; Raphael Kpoghomou; Alain M. Kuyunsa. (2025). Enhancing the Robustness of Computer Vision Models to Adversarial Perturbations Using Multi-Scale Attention Mechanisms. International Journal of Innovative Science and Research Technology, 10(4), 3565-3578. https://doi.org/10.38124/ijisrt/25apr2118.

[32] Kang, D., Kim, H., & Lee, S. (2020). Deep learning-based railway signal detection under various lighting conditions. Sensors, 20(4), 1083. https://doi.org/10.3390/s20041083

[33] Li, C., Zhang, J., & Zhou, H. (2021). Railway signal detection using improved YOLOv3. Journal of Advanced Transportation, 2021, Article ID 8890132. https://doi.org/10.1155/2021/8890132

[34] Rodríguez-Abreo, O.; Quiroz-Juárez, M.A.; Macías-Socarras, I.; Rodríguez-Reséndiz, J.; Camacho-Pérez, J.M.; Carcedo-Rodríguez, G.; Camacho-Pérez, E. Automatic Detection of Railway Faults Using Neural Networks: A Comparative Study of Transfer Learning Models and YOLOv11. *Infrastructures* **2025**, *10*, 3. https://doi.org/10.3390/infrastructures10010003

[35] Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C. Y., & Berg, A. C. (2016). SSD: Single shot multibox detector. In European Conference on Computer Vision (pp. 21–37). Springer.

[36] Tran, V.-T.; Tsai, W.-H.; Furletov, Y.; Gorodnichev, M. End-to-End Train Horn Detection for Railway Transit Safety. *Sensors* **2022**, *22*, 4453. https://doi.org/10.3390/s22124453

[37] Alawaji, K.; Hedjar, R.; Zuair, M. Traffic Sign Recognition Using Multi-Task Deep Learning for Self-Driving Vehicles. *Sensors* **2024**, *24*, 3282. https://doi.org/10.3390/s24113282

[38] Alif, M.A.R.; Hussain, M. Lightweight Convolutional Network with Integrated Attention Mechanism for Missing Bolt Detection in Railways. *Metrology* **2024**, *4*, 254-278. https://doi.org/10.3390/metrology4020016

[39] Tian, Y., Barzilay, R., & Jaakkola, T. (2016). Image-based signal classification in transportation networks. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (pp. 556–561).

[40] Mbayandjambe, A. M., Nkwimi, G. B., Nguemdjom, D. K. T., Oshasha, F., Muluba, C., & Kutuka, X. F. (2025, April). *AI-Generated Sneaker Detection: Leveraging GANs and Convolutional Neural Networks for Image Classification*. International Journal of Innovative Science and Research Technology, **10**(4). https://doi.org/10.38124/ijisrt/25apr1963

[41] Shangbang Long, Xin He, and Cong Yao. *Scene Text Detection and Recognition: The Deep Learning Era*. arXiv preprint, 2020. arXiv:1811.04256. Available at: https://arxiv.org/abs/1811.04256

[42] Sun, Q.; Zhu, M.; Li, M.; Li, G.; Deng, W. Symbol Recognition Method for Railway Catenary Layout Drawings Based on Deep Learning. *Symmetry* **2025**, *17*, 674. https://doi.org/10.3390/sym17050674