Optimized YOLOv8 for Automatic License Plate Recognition on Resource Constrained Devices

Barka Satya
Faculty of Information Technology, Satya Wacana Christian University, Salatiga, Indonesia | Faculty of Computer Science, Universitas Amikom Yogyakarta, Sleman, Indonesia
barka.satya@amikom.ac.id (corresponding author)

Danny Manongga
Faculty of Information Technology, Satya Wacana Christian University, Salatiga, Indonesia
danny.manongga@uksw.edu

Hendry
Faculty of Information Technology, Satya Wacana Christian University, Salatiga, Indonesia
hendry@uksw.edu

Afrig Aminuddin
Faculty of Computer Science, Universitas Amikom Yogyakarta, Sleman, Indonesia
afrig@amikom.ac.id

Received: 19 December 2024 | Revised: 25 January 2025 and 11 February 2025 | Accepted: 14 February 2025

 DOI: https://doi.org/10.48084/etasr.9983

ABSTRACT

This paper presents an optimized Automatic License Plate Recognition (ALPR) system designed for resource-constrained devices, leveraging YOLOv8 for real-time object detection and Optical Character Recognition (OCR) to extract license plate information under challenging conditions such as low-light, motion blur, and occlusions. Unlike traditional ALPR systems that rely on high computational resources, our approach balances detection accuracy, processing speed, and efficiency. The system is evaluated on three benchmark datasets: the Chinese City Parking Dataset (CCPD) with 250,000 images under diverse conditions, the UFPR-ALPR Dataset (Universidade Federal do Paraná, Brazil) containing 4,500 real-world traffic images, and the RodoSol-ALPR Dataset with 20,000 highway surveillance images for high-speed license plate recognition. Among various YOLOv8 variants tested, YOLOv8-s achieved the best performance, with a mean Average Precision (mAP) of 99.3% while sustaining over 30 Frames Per Second (FPS), making it suitable for real-time ALPR applications. Furthermore, image sharpening and contour segmentation techniques improved OCR recognition accuracy by 5.1% under low-light conditions, improving robustness. Comparative analysis against state-of-the-art OCR-based ALPR methods (EasyOCR, FastOCR, and CR-NET) demonstrated that our approach surpasses existing models in both recognition rate and computational efficiency. These findings establish YOLOv8 as a highly effective and deployable solution for intelligent transportation, surveillance, and law enforcement applications requiring real-time license plate recognition with minimal computational overhead.

Keywords-computer vision; image processing; license plate recognition; object detection

I. INTRODUCTION

Deep learning in computer vision has revolutionized the field of object detection and recognition, particularly for tasks such as Automatic License Plate Recognition (ALPR) [1]. ALPR plays a crucial role in traffic management, toll collection, and law enforcement [2]. The advancements in deep learning architectures, specifically those using the You Only Look Once (YOLO) models, have significantly improved detection speed and accuracy [3]. YOLO models, known for real-time detection, are preferred for high-throughput and low-latency applications [4].

Authors in [5] presented a multinational ALPR system using Tiny YOLOv3 for license plate detection and YOLOv3-SPP for character recognition, demonstrating high accuracy and speed across datasets from South Korea, Taiwan, Greece, the USA, and Croatia. The system effectively classified single and double-line plates and introduced KarPlate, a new Korean car plate dataset. However, it required extensive annotated data for training, making it difficult to adapt to new license plate layouts not included in the dataset. The approach performed well in structured environments but struggled with plate variations from countries not represented in the training data. Authors in [6] presented an end-to-end YOLOv5-based ALPR system incorporating a channel attention mechanism to enhance feature extraction and GRU for Optical Character Recognition (OCR) recognition without requiring segmentation. The model achieved 98.98% precision on the Chinese City Parking Dataset (CCPD) dataset, demonstrating high accuracy and robustness in complex environments. While integrating attention mechanisms improved accuracy, it slightly reduced detection speed due to increased computational complexity. Additionally, the model's performance was influenced by the quality and diversity of training data, limiting its generalizability across datasets with significantly different license plate structures. Authors in [7] presented a vehicle pose estimation system that utilized YOLOv5 trained on RGB images of license plates and wheels to assist Automated Guided Vehicles (AGVs) in parking scenarios. The system aimed to improve maneuverability by reducing the need for multiple sensors while ensuring precise vehicle positioning. The model demonstrated high precision and recall in detecting license plates and wheels, allowing AGVs to estimate vehicle positions accurately. However, the detection performance for wheels was lower than for license plates, making pose estimation less reliable in certain conditions. Furthermore, the model's reliance on specific datasets limited its applicability to environments with different vehicle structures.

Authors in [8] presented a real-time ALPR system that combined YOLOx, YOLOv4-tiny, Paddle OCR, and SVTR-tiny to perform vehicle make/model identification and license plate recognition. The system achieved 97.5% accuracy, demonstrating robustness in adverse conditions like fog, rain, and low-light environments. The study also implemented GradCam technology to visualize the model's decision-making process and identify areas for improvement. However, the system struggled with recognizing rare vehicle models due to the limited availability of training images, and its performance was affected by environmental factors such as varying lighting conditions and camera angles. These limitations made the system less adaptable to new and uncommon vehicle types. Authors in [9] presented an Iranian ALPR system using dual YOLOv3 networks to perform both License Plate Detection (LPD) and Character Recognition (CR). The model achieved 95.05% accuracy on a dataset of 5,719 images, processing license plates in an average of 119.73 milliseconds, demonstrating its real-time applicability. The system was designed to function without preprocessing or calibration, making it efficient for large-scale traffic monitoring and toll collection applications. However, challenges remained in recognizing small license plates and Persian characters with similar structures, which sometimes led to misclassification. Additionally, the system struggled with low-resolution images, necessitating improvements in data augmentation techniques to enhance recognition accuracy. Authors in [10] presented a license plate detection method in the compressed domain using High-Efficiency Video Coding (HEVC) attributes, eliminating the need to fully decompress encoded data and reducing computational costs. The method, implemented with YOLOv3 Tiny Object Detector, demonstrated comparable precision and recall to state-of-the-art models while achieving over 30% savings in inference time. The study also introduced a new compressed domain LP database that is publicly available for future research. Despite its efficiency, the method focused only on license plate detection and did not address character recognition, limiting its usefulness for end-to-end ALPR systems. Furthermore, its performance in diverse real-world conditions was not extensively tested, leaving room for generalizability and OCR integration improvements.

Despite advancements in ALPR, several challenges remain. State-of-the-art YOLO-based models achieve high accuracy but rely heavily on large, annotated datasets, limiting their adaptability to unseen license plate variations [11]. Feature extraction enhancements, such as attention mechanisms, often increase computational complexity, reducing real-time efficiency. While some models perform well in adverse weather, their accuracy declines under poor lighting, low resolution, and noisy environments. This underscores the need for efficient, adaptable ALPR systems that maintain high accuracy and speed while minimizing training data dependency.

This paper presents an optimized YOLOv8-based ALPR system that balances real-time performance and accuracy for deployment in low-power environments. Unlike previous studies that rely on large-scale computing resources, our system focuses on lightweight, efficient ALPR deployment without sacrificing accuracy. Our main contributions are:

Dataset-Specific Performance Optimization: We evaluate three ALPR datasets—CCPD, UFPR-ALPR, and RodoSol-ALPR, testing multiple YOLOv8 variants to determine the most efficient configuration.

- Efficient Model Deployment: We optimize YOLOv8 for mobile and embedded applications by reducing computational complexity while maintaining a high mean Average Precision (mAP) (99.3%) and real-time processing speeds (30+ FPS).
- Enhanced Preprocessing for Robust ALPR: We implement image sharpening and contour segmentation to improve character recognition, especially in low-light and noisy conditions.
- Benchmarking Against State-of-the-Art Methods: Our approach is compared to existing OCR-based ALPR techniques, demonstrating superior recognition rates across diverse datasets.

These contributions establish a scalable and efficient ALPR framework that can be deployed in real-world intelligent transportation and surveillance systems.

II. PROPOSED METHODS

We propose a deep learning-based system for ALPR that integrates YOLO object detection and OCR to detect and extract vehicle and license plate information. The proposed ALPR method is presented in Figure 1.

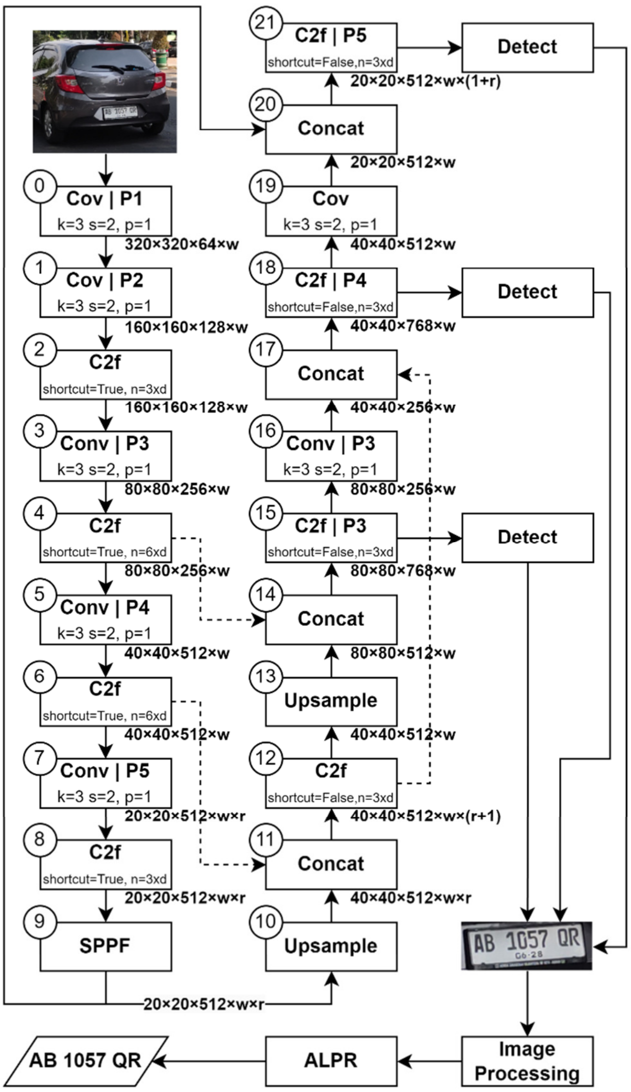

Fig. 1. The architecture of the proposed ALPR.

The method is designed to handle large-scale datasets and efficiently process real-time images for ALPR tasks in diverse environmental conditions [12]. The following sections outline our proposed method's key components and processes.

A. Data Preparation and Preprocessing

During data preprocessing, bounding boxes for vehicles and license plates are converted to the YOLO format [13]. The original bounding box is typically described by its top-left corner coordinates (x_{min}, y_{min}) with width w and height h, along with the image dimension (img_width, img_height). The conversion to YOLO format involves transforming these values into normalized center coordinates (x_{center}, y_{center}) and the normalized width and height of the bounding box. The following equations are used for this transformation:

$$x_{center} = \frac{x_{min} + \frac{w}{2}}{img\_width} \quad (1)$$

$$y_{center} = \frac{y_{min} + \frac{h}{2}}{img\_height} \quad (2)$$

$$w_{normalized} = \frac{w}{img\_width} \quad (3)$$

$$h_{normalized} = \frac{h}{img\_height} \quad (4)$$

These normalized values ensure that the bounding box coordinates are scaled between 0 and 1 relative to the image dimensions, which is the format required for YOLO training.

B. Object Detection Model

The YOLOv8 model detects objects in an image by dividing it into an $S{\times}S$ grid. Each grid cell predicts bounding boxes, confidence scores, and class probabilities. The model is trained to minimize the difference between the predicted and ground truth values using a loss function composed of three key components: bounding box regression loss, object confidence loss, and classification loss [14]. The types of YOLOv8 models are presented in Table I.

TABLE I. TYPES OF YOLOV8 MODELS.

Models	d (depth multiple)	w (width multiple)	r (ratio)
YOLOv8-n	0.33	0.25	2.0
YOLOv8-s	0.33	0.50	2.0
YOLOv8-m	0.67	0.75	1.5
YOLOv8-l	1.00	1.00	1.0
YOLOv8-x	1.00	1.25	1.0

Table I showcases all the YOLOv8 versions. In this paper, the YOLOv8-n is selected in the experiment as it is the most suitable model for resource-constrained devices due to the small size of the model.

1) Bounding Box Regression Loss

This term minimizes the difference between the predicted and actual bounding box coordinates. It can be expressed as:

$$L_{bbox} = \sum_{i=1}^{S^2} \sum_{j=1}^{B} 1_{ij}^{obj} \left[(x - \hat{x})^2 + (y - \hat{y})^2 + \left(\sqrt{w} - \sqrt{\widehat{w}}\right)^2 + \left(\sqrt{h} - \sqrt{\hat{h}}\right)^2 \right] \quad (5)$$

where x ,y, w, h are the ground truth bounding box center and dimensions $\hat{x}$, $\hat{y}$, $\widehat{w}$, $\hat{h}$, are the predicted bounding box center and dimensions, S is the number of grid cells, B is the number of bounding boxes per cell, 1_{ij}^{obj} is an indicator function that is 1 if the object is present in the i-th grid cell and 0 otherwise.

2) Object Confidence Loss

YOLO predicts a confidence score for each bounding box. The confidence score is the product of the objectless probability (whether the box contains an object) and the Intersection over Union (IoU) between the predicted and ground truth boxes [11]. The confidence loss is given by:

$$\mathcal{L}_{\text{conf}} = \sum_{i=1}^{S^2}\sum_{j=1}^{B}\left[1_{ij}^{\text{obj}}\left(C_i - \widehat{C}_i\right)^2 + 1_{ij}^{\text{no\_obj}}\left(C_i - \widehat{C}_i\right)^2\right] \quad (6)$$

where C_i is the ground truth confidence score (1 if the object is present, 0 otherwise), $\widehat{C}_i$ is the predicted confidence score, $1_{ij}^{\text{no\_obj}}$ is 1 if no object is present in the grid cell and 0 otherwise.

3) Classification Loss

For each bounding box, the model predicts the probability of the object belonging to a particular class (e.g., car, motorcycle, license plate). The classification loss is calculated as the cross-entropy between the true and predicted class:

$$\mathcal{L}_{\text{class}} = \sum_{i=1}^{S^2} 1_i^{\text{obj}} \sum_{c=1}^{C} \left(p_i(c) - \widehat{p}_i(c)\right)^2 \quad (7)$$

where $p_i(c)$ is the ground truth probability for class c, $\widehat{p}_i(c)$ is the predicted class probability for class c, C is the total number of classes. The total loss $\mathcal{L}$ for the YOLO model is the sum of these three components:

$$\mathcal{L} = \lambda_{\text{bbox}}\mathcal{L}_{\text{bbox}} + \lambda_{\text{conf}}\mathcal{L}_{\text{conf}} + \lambda_{\text{class}}\mathcal{L}_{\text{class}} \quad (8)$$

where λ_{bbox}, λ_{conf}, λ_{class} are weighting factors for the respective loss terms.

C. License Plate Recognition

Once YOLOv8 detects the license plate region, we extract the text. The OCR process can be described as recognizing characters C from the cropped image $\mathrm{I}_{\text{plate}}$. The OCR engine aims to maximize the probability of generating the correct sequence of characters:

$$P\left(C \middle| I_{\text{plate}}\right) = \prod_{t=1}^{T} P\left(c_t \middle| I_{\text{plate}}, \theta\right) \quad (9)$$

where $\mathrm{c_t}$ is the character at time step, θ are the parameters of the OCR model, T is the total number of characters in the recognized sequence. This sequence is generated by selecting the most probable character at each step t, based on the features extracted from the image.

III. RESULTS AND ANALYSIS

This chapter presents the results of the object detection and license plate recognition system using YOLOv8 models for early collision detection. The YOLOv8-based ALPR system was tested across multiple datasets using various performance metrics such as mAP, precision, recall, F1-score, and Frames Per Second (FPS). To assess the performance of the YOLOv8-based ALPR system, we employed three benchmark datasets, each presenting diverse real-world challenges, the CCPD [15], the UFPR-ALPR dataset [16], and the RodoSol-ALPR dataset [17]. The CCPD dataset contains 250,000 images of Chinese vehicle license plates captured from varied angles, lighting conditions, and occlusions, making it ideal for testing detection robustness in urban environments [15]. It includes multiple subsets such as CCPD-Base, CCPD-Weather, and CCPD-Rotate, allowing for a comprehensive evaluation under different conditions. The UFPR-ALPR dataset comprises 4,500 traffic images from Brazilian roads featuring motion blur, occlusions, and illumination variations [16]. Unlike CCPD, it focuses on dynamic traffic scenarios, making it suitable for testing OCR performance on high-speed vehicles and assessing recognition accuracy under real-world distortions. The RodoSol-ALPR dataset comprises 20,000 images collected from highway surveillance cameras to evaluate high-speed license plate recognition [17]. It introduces challenges such as motion blur, adverse weather, and varying font styles, ensuring the system's robustness in real-time enforcement applications.

TABLE II. OBJECT DETECTION SCORES USING YOLO V8-N MODELS.

Dataset	mAP 0.5	mAP 0.5:0.95	Precision (%)	Recall (%)	F1-score (%)	FPS
CCPD	99.3	68	99	99	99.0	30
UFPR-ALPR	94.7	76.95	93	89.73	91.0	36
RodoSol-ALPR	91	81	84	84	84	36

The performance of the YOLOv8-n model for object detection, specifically distinguishing the license plate from the background, is presented in Table II. The results indicate that YOLOv8-n achieved high detection accuracy across the three datasets. Notably, for CCPD, the model attained a mAP of 99.3% at an IoU threshold of 0.5, with a corresponding mAP of 68% at IoU thresholds ranging from 0.5 to 0.95. The precision, recall, and F1-score were exceptionally high at 99%, demonstrating the robustness of YOLOv8-n in detecting license plates with minimal false positives. The FPS performance was 30, ensuring real-time processing capabilities. The performance was similarly strong for UFPR-ALPR, with a mAP of 99% at 0.5 IoU and an improved mAP of 87% at stricter thresholds, with an FPS of 36. However, RodoSol-ALPR exhibited relatively weaker performance with an mAP of 91% at 0.5 IoU and 81% at stricter thresholds, and a slightly reduced precision and recall of 84% each. The confusion matrices for the three datasets are presented in Figure 2.

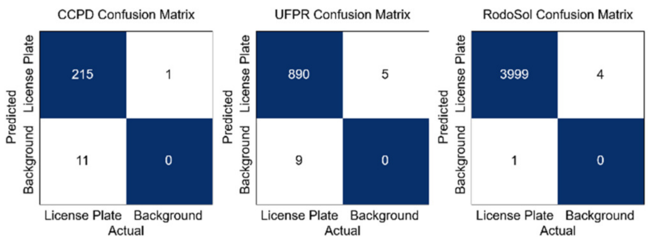

Fig. 2. Confusion matrices of the datasets.

Image processing is required to sharpen the image for better plate number detection. The differences between the processed and unprocessed plate numbers are presented in Figure 3. It presents a workflow visualization of an ALPR system, illustrating a vehicle's license plate detection, processing, and recognition. A real-world street scene is displayed on the left side, where a white car is parked on the roadside. The proposed ALPR system has successfully detected the vehicle's license plate, highlighted with a green bounding box labeled "License Plate." This demonstrates the initial detection stage of the system in real-world conditions.

Fig. 3. The plate number detection with and without image processing.

Figure 3 also showcases the step-by-step processing of the detected license plate. The sequence begins with the raw detection of the license plate cropped from the vehicle. To ensure privacy protection, the last two characters of the license plate have been obscured. Next, an image processing stage is applied, incorporating enhancement techniques such as sharpening, contrast adjustments, and noise reduction to improve readability. Finally, the proposed ALPR system extracts and refines the license plate characters. Initially, the system incorrectly recognized the plate as "AB J7B5" likely due to noise or distortion in the image. However, after image processing, the final recognition correctly identifies the plate as "AB 1785" demonstrating the impact of preprocessing in improving OCR accuracy. The effectiveness of this enhancement is detailed in Table III.

TABLE III. ALPR SCORES WITH IMAGE PROCESSING.

Dataset	Precision (%)	Recall (%)	F1-score (%)	Recognition Rates (%)
CCPD	81.6	75.8	78.2	64.8
UFPR-ALPR	92.4	88.4	89.9	83.1
RodoSol-ALPR	91.5	77.2	82.4	71.7

The proposed image processing technique contributed to higher recognition performance, particularly improving the precision and F1-score values. For CCPD, the method achieved a precision of 81.6% and an F1-score of 78.2%, though the recognition rate remained relatively low at 64.8%. Meanwhile, UFPR-ALPR displayed a higher recognition rate of 83.1% with a precision of 92.4%, which suggests that the image processing technique contributed significantly to improved character segmentation. RodoSol-ALPR, although performing slightly worse than UFPR-ALPR, still demonstrated a recognition rate of 71.7%, an F1-score of 82.4%, and a precision of 91.5%. To evaluate the performance of the proposed method, the results were compared with several existing techniques, as presented in Table IV.

TABLE IV. RECOGNITION RATES COMPARISON OF STATE-OF-THE-ART OCR

Method	UFPR-ALPR (%)	RodoSol-ALPR (%)
Proposed Method	83.1	71.7
FastOCR [18]	81.6	56.7
CR-NET [19]	82.6	59.0
Rosetta [20]	75.5	89.0
Multi-task [21]	72.3	86.6

The proposed approach outperformed FastOCR [18], which had a recognition rate of 81.6% for UFPR-ALPR and 56.7% for RodoSol-ALPR, demonstrating the superiority of the image processing enhancement in handling character segmentation. Compared to CR-NET [19], which achieved 82.6% on UFPR-ALPR and 59.0% on RodoSol-ALPR, the proposed method performed better on UFPR-ALPR but was outperformed by CR-NET on RodoSol-ALPR. Interestingly, Rosetta [20] achieved the highest recognition rate on RodoSol-ALPR (89.0%), but its performance was significantly lower on UFPR-ALPR (75.5%), indicating a dataset-dependent effectiveness. The multi-task approach [21] also showed competitive results, especially on RodoSol-ALPR, where it achieved an 86.6% recognition rate, but its UFPR-ALPR performance (72.3%) was notably lower than the proposed method.

Overall, these results validate the effectiveness of the proposed optimized YOLOv8-n model for license plate detection and the proposed method with image processing enhancements for character recognition. The comparative analysis with other OCR methods highlights the advantages of the proposed approach in achieving higher recognition rates, particularly on UFPR-ALPR, while indicating areas for further improvements in handling datasets like RodoSol-ALPR. Future work may explore additional image enhancement techniques and fine-tuning strategies to further improve recognition rates, particularly for challenging datasets with complex background conditions.

IV. CONCLUSION

This study demonstrated the effectiveness of a YOLOv8-based system for Automatic License Plate Recognition (ALPR) on resource-constrained devices, achieving high detection accuracy and real-time performance. Among the YOLOv8 variants, YOLOv8-n balanced precision and speed, achieving a mean Average Precision (mAP) of 99% and maintaining 36 FPS on the UFPR-ALPR dataset, making it suitable for mobile and embedded systems. Image preprocessing techniques, such as sharpening, further improved recognition rates in challenging environments. At the same time, comparisons with state-of-the-art Optical Character Recognition (OCR) methods highlighted the system's robustness and adaptability across diverse datasets. Despite its success, challenges remain in generalizing to uncommon license plate formats and extreme conditions, which can be addressed by expanding training datasets, optimizing lightweight architectures, and exploring advanced transfer learning.

REFERENCES

[1] R. Antar, S. Alghamdi, J. Alotaibi, and M. Alghamdi, "Automatic Number Plate Recognition of Saudi License Car Plates," *Engineering, Technology & Applied Science Research*, vol. 12, no. 2, pp. 8266–8272, Apr. 2022, https://doi.org/10.48084/etasr.4727.

[2] M. S. Sheikh, J. Liang, and W. Wang, "A Survey of Security Services, Attacks, and Applications for Vehicular Ad Hoc Networks (VANETs)," *Sensors*, vol. 19, no. 16, Aug. 2019, Art. No. 3589, https://doi.org/10.3390/s19163589.

[3] Y. Yang, Y. Xiao, Z. Chen, D. Tang, Z. Li, and Z. Li, "FCBTYOLO: A Lightweight and High-Performance Fine Grain Detection Strategy for Rice Pests," *IEEE Access*, vol. 11, pp. 101286–101295, 2023, https://doi.org/10.1109/ACCESS.2023.3314697.

[4] M. Salemdeeb and S. Erturk, "Multi-national and Multi-language License Plate Detection using Convolutional Neural Networks," *Engineering, Technology & Applied Science Research*, vol. 10, no. 4, pp. 5979–5985, Aug. 2020, https://doi.org/10.48084/etasr.3573.

[5] C. Henry, S. Y. Ahn, and S.-W. Lee, "Multinational License Plate Recognition Using Generalized Character Sequence Detection," *IEEE Access*, vol. 8, pp. 35185–35199, 2020, https://doi.org/10.1109/ACCESS.2020.2974973.

[6] H. Shi and D. Zhao, "License Plate Recognition System Based on Improved YOLOv5 and GRU," *IEEE Access*, vol. 11, pp. 10429–10439, 2023, https://doi.org/10.1109/ACCESS.2023.3240439.

[7] X. Li, X. Wang, C. Qu, J. Song, H. Li, and Y. Xi, "Vehicle pose estimation by parking AGV based on RGBD camera," in *2024 9th International Conference on Automation, Control and Robotics Engineering (CACRE)*, Jeju Island, Korea, Republic of, Jul. 2024, pp. 289–294, https://doi.org/10.1109/CACRE62362.2024.10635074.

[8] T. Mustafa and M. Karabatak, "Real Time Car Model and Plate Detection System by Using Deep Learning Architectures," *IEEE Access*, vol. 12, pp. 107616–107630, 2024, https://doi.org/10.1109/ACCESS.2024.3430857.

[9] A. Tourani, A. Shahbahrami, S. Soroori, S. Khazaee, and C. Y. Suen, "A Robust Deep Learning Approach for Automatic Iranian Vehicle License Plate Detection and Recognition for Surveillance Systems," *IEEE Access*, vol. 8, pp. 201317–201330, 2020, https://doi.org/10.1109/ACCESS.2020.3035992.

[10] M. S. Beratoğlu and B. U. Töreyin, "Vehicle License Plate Detector in Compressed Domain," *IEEE Access*, vol. 9, pp. 95087–95096, 2021, https://doi.org/10.1109/ACCESS.2021.3092938.

[11] K. Sangsuwan and M. Ekpanyapong, "Video-Based Vehicle Speed Estimation Using Speed Measurement Metrics," *IEEE Access*, vol. 12, pp. 4845–4858, 2024, https://doi.org/10.1109/ACCESS.2024.3350381.

[12] W. Ali, G. Wang, K. Ullah, M. Salman, and S. Ali, "Substation Danger Sign Detection and Recognition using Convolutional Neural Networks," *Engineering, Technology & Applied Science Research*, vol. 13, no. 1, pp. 10051–10059, Feb. 2023, https://doi.org/10.48084/etasr.5476.

[13] D. Habeeb, A. H. Alhassani, L. N. Abdullah, C. S. Der, and L. K. Q. Alasadi, "Advancements and Challenges: A Comprehensive Review of GAN-based Models for the Mitigation of Small Dataset and Texture Sticking Issues in Fake License Plate Recognition," *Engineering, Technology & Applied Science Research*, vol. 14, no. 6, pp. 18401–18408, Dec. 2024, https://doi.org/10.48084/etasr.8870.

[14] S. Luo and J. Liu, "Research on Car License Plate Recognition Based on Improved YOLOv5m and LPRNet," *IEEE Access*, vol. 10, pp. 93692–93700, 2022, https://doi.org/10.1109/ACCESS.2022.3203388.

[15] Z. Xu *et al.*, "Towards End-to-End License Plate Detection and Recognition: A Large Dataset and Baseline," in *Computer Vision – ECCV 2018*, vol. 11217, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds. Cham: Springer International Publishing, 2018, pp. 261–277, https://doi.org/10.1007/978-3-030-01261-8_16.

[16] R. Laroca *et al.*, "A Robust Real-Time Automatic License Plate Recognition Based on the YOLO Detector," in *2018 International Joint Conference on Neural Networks (IJCNN)*, Rio de Janeiro, Brazil, Jul. 2018, pp. 1–10, https://doi.org/10.1109/IJCNN.2018.8489629.

[17] R. Laroca, E. Cardoso, D. Lucio, V. Estevam, and D. Menotti, "On the Cross-dataset Generalization in License Plate Recognition:," in *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, Online Streaming, 2022, pp. 166–178, https://doi.org/10.5220/0010846800003124.

[18] R. Laroca, A. B. Araujo, L. A. Zanlorensi, E. C. De Almeida, and D. Menotti, "Towards Image-Based Automatic Meter Reading in Unconstrained Scenarios: A Robust and Efficient Approach," *IEEE Access*, vol. 9, pp. 67569–67584, 2021, https://doi.org/10.1109/ACCESS.2021.3077415.

[19] S. M. Silva and C. R. Jung, "Real-time license plate detection and recognition using deep convolutional neural networks," *Journal of Visual Communication and Image Representation*, vol. 71, Aug. 2020, Art. no. 102773, https://doi.org/10.1016/j.jvcir.2020.102773.

[20] F. Borisyuk, A. Gordo, and V. Sivakumar, "Rosetta: Large Scale System for Text Detection and Recognition in Images," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, London United Kingdom, Jul. 2018, pp. 71–79, https://doi.org/10.1145/3219819.3219861.

[21] G. R. Gonçalves, M. A. Diniz, R. Laroca, D. Menotti, and W. R. Schwartz, "Multi-task Learning for Low-Resolution License Plate Recognition," in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, vol. 11896, I. Nyström, Y. Hernández Heredia, and V. Milián Núñez, Eds. Cham: Springer International Publishing, 2019, pp. 251–261.