STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **x**, Month 202x, pp 0–15.
Published online in International Academic Press (www.IAPress.org)

Detection and Recognition for Iraqi Modern License Plate Using Deep Learning Approach

Mayaf Zakawa Hasan [1,2,*], Salim Ganim Saeed Al-Ali [3]

[1] *Information Technology Management Dept., Technical College of Administration, Duhok Polytechnic University, Iraq*
[2] *Computer Information System Dept., Technical College of Zakho, Duhok Polytechnic University, Iraq*
[3] *Information Technology Management Dept., Technical College of Administration, Duhok Polytechnic University, Iraq*

Abstract In recent years, as the accuracy of deep learning techniques has become spectacular, object identification and recognition has recently become an increasingly popular target of computer vision applications. Among them, automatic number plate recognition (ANPR) has attracted much attention and is already a popular subject of research, but is still difficult in cases where there are few public datasets and differing plate formats. In this paper, the pipeline and a combination of YOLOv11n for detection and LPRNet as optical character recognition is proposed as deep learning-based. The experiments are performed on a dataset collected to perform the identification and detection of new Iraqi license plates. Our method had good detection performance on our dataset, in spite of applying cross-dataset pre-trained detection weights based on a CCPD dataset showing that it had good cross-domain generalization. In the recognition stage, an LPRNet model, the training and evaluation of which was conducted on our collected data only. The accuracy of detector and recognizer of this system is 96.0% and 99.8% respectively. The findings emphasize the power of deep learning models to perform cross domain ALPR tasks, and are an indication that future expansion of datasets will result in an increase in robustness on a variety of real-world circumstances.

Keywords Automatic License Plate Recognition(ALPR), License Plate(LP), License Plate Detection(LPD), License Plate Recognition(LPR)

DOI: 10.19139/soic-2310-5070-2801

1. Introduction

*A*utomatic License Plate Recognition (ALPR) is required in a modern intelligent transportation system especially where traffic management is involved and safety of people is necessary in resource-limited environment. However, the majority of existing models struggle to adapt to local difference, such as in Iraq, where modern LP exhibit structural and visual modifications that cannot be captured in widely used benchmark datasets, even with impressive developments enabled by deep learning (DL). Globally, ALPR systems are being included into smart surveillance platforms due to the growing demand for automation and real-time vehicle recognition [1]. Addressing traffic safety concerns [2], which continue to be a major issue in many more-income nations, is where this technology is more important. It has drawn a lot of attention because of its critical function in assisting traffic law enforcement and improving public safety through its capacity to precisely identify and recognize LPs. Even though DL has significantly improved ALPR, a number of issues still cause ALPR systems to operate poorly. For instance, variations in illumination limit a number of methods. Even while some strategies concentrate on creating systems that function both during the day and at night, their performance varies [3]. Furthermore, every country has different requirements for LPs (size, material, use of alphabetic characters, language of characters, etc.). Because of these

*Correspondence to: Mayaf Zakawa Hasan (Email: mayaf.zakawa@dpu.edu.krd). Computer Information System Dept., Technical College of Zakho, Duhok Polytechnic University, Iraq

ISSN 2310-5070 (online) ISSN 2311-004X (print)

intra-class differences, it is challenging to create a single ALPR system that is flawless for all LPs in all nation [4, 5]. Due to these problems, computer vision scientists will always be trying to develop trustful and efficient algorithms to identify a vehicle and its number plates. In the past years, DL strategies have exhibited promising results in this area with a resulting growth in the accuracy and stability of these systems [6]. Many nations across the world use LPR technology, especially in smart cities. Each nation's government designs its LP, which usually consists of a mix of letters, numbers, and words in the national tongue [7, 8]. For example, the unified model, which debuted in 2022 [9], The Arabic language and numbers were changed from LP to English language and numbers that differ from the old Iraqi LP. Derived from the German LP model [10]. Most publicly available data and state-of-the-art models are tailored to specific countries and LP types, and therefore, they cannot be used in such regions as Iraq where there exist limited locally available data, and where local LP designs differ and may diverge significantly over time. The new Iraqi LPs test the current LPDR technology. This challenge is caused by the fact that some of the LP characters have had their designs modified in ways that break away with conventional design used in the standard dataset. This paper fills this gap by building a novel pipeline specifically to be used on new Iraqi LPs and analyzing whether DL models trained on other data can be transferred. This is aimed at evaluating the viability of implementing a high-accuracy ALPR solution in Iraq by adapting existing models on transfer learning, with little preprocessing on the one hand, and being independent of any kind of large-scale labeled data on the other.

This rest of the paper is structured as follows: the research from the literature is described in the following section. The suggested methodology and its operation are covered in Section 3. The empirical findings are discussed in Section 4. The discussion in Section 5, while Section 6 offers closing thoughts.

2. Related Work

Several methods have been developed to improve the accuracy of both detection and recognition. Conventional approaches frequently depend on traditional image processing methods including template matching, morphological filtering, and edge identification. DL models have gained popularity recently because of their exceptional performance in challenging real-world situations. Convolutional neural networks (CNNs) have been the mainstay of recent developments in ALPR for both detection and recognition applications. While sequence recognition models like LPRNet [11] have increased the accuracy of character identification on LPs, techniques like YOLOv11 [12] have shown strong object detection performance.

On the potential to develop an intelligent LPDR system in a real-life traffic situation, Sutikno et al. (2025) [12], proposed that there is a need to combine YOLOv11 and CLAHE. They obtained an accuracy level of 99.5% in detection and 93.1% in identification. This paper backup our inclination on its execution on a personalized level of a working circumstance of Iraq as proven the practicality of the YOLOv11 in the actual world LPR situation. which aligns with our motivation to adopt model in the context of Iraqi LPs under uncontrolled conditions. Similarly, as for the research by Chan et al. (2020)[13], introduced one of the largest EU LP datasets (TLPD) containing over 18,000 tagged plates acquired in diverse environmental surroundings. Their models, using unseen video frames yielded over 92.0% precision and recall, trained by Fast-YOLO and Tiny-YOLO v3. Their research reveals the effectiveness of YOLO trains-based detectors although they are focused on the GDPR in terms of anonymization and detection. Shi and Zhao (2023)[14], on the other hand, their research provides the complete LPR system, the model has a channel attention optimization mechanism which enhances the process of feature extraction with the overall concept of the sampling process in their model. The researchers used a CCPD dataset of 12,500 disparate pictures and obtains the mean accuracy of 98.98% on recognition and 97.1% on detection with the support of the YOLOv5-LSE model. Likewise, Jasem and Mohammed (2020)[15], developed an Iraqi LPR system in a two-stage procedure. In the initial stage, it identifies and identifies LPs in photos. The second step is focused on the identification of these plates with the use of the SIFT and SURF algorithms and the K-Nearest Neighbor (KNN) method to match the plate characters. The complete accuracy of the system after training on 1300 photos of Iraqi vehicles was 98.17%, 99.2% of the detections, and 97.14% of recognitions. In a more globally oriented study, Salemdeeb and Erturk (2020)[16], provide an innovative idea of detecting and classifying LPs of various countries, languages, and schemes in their work. The LPDC2020 dataset contains the images of Brazil, the United States,

Europe, Turkey, Saudi Arabia, and the United Arab Emirates. The YOLOv2 detector gives an average precision of 99.57% using the dataset as training set. Furthermore, the research proposes a simple architecture of the CNN to classify the LPs which achieved an accuracy of 99.33% which is better when compared to the VGG16 network. The first step to making an ALPR system Global is to come up with a standard system to be able to detect and categorize the LP in a multi-national, multi-language, and multi-layout scheme. Focusing specifically on Iraq, Abbas and Mohammed (2024)[17], the study proposes an innovative ANPR system, which has been developed specifically to resolve the challenges posed by inconsistent and irregular LPs in Iraq. utilizes transfer learning and YOLOv8 models in order to identify and classify LPs. They obtain precision 98.7% rate of detecting Iraqi LPs in three processes sequence, namely; detection, segmentation, and identification.
Luo and Liu (2022)[18], The study presents a LPR method using the YOLOv5m algorithm and the LPRNet model. The YOLOv5m algorithm was enhanced with the K-meansCC algorithm, DIOU loss function, and feature map removal. The lightweight LPRNet network was used to recognize LP characters without segmentation, in frontal, oblique, night scenes, and strong light interference. Their system achieved recognition accuracy 99.49%, recall 98.79%, and mAP 98.56%, in frontal, oblique, night scenes, and strong light interference. The research emphasized the model's robustness and adaptability, addressing shortcomings in accuracy and real-time performance. In a similar line, Huang et al. (2021)[11] suggested that ALPRNet accomplishes the two tasks in tandem with the assistance of two demonstration fully convolutional object detectors- a single object detector operates on characters and the other on LPs. They network was tested on various data sets such as the HZM multi-style dataset and AOLP where the LPD rates were 100%, and recognition was 98.21%. The successful recognition of at least one LP in an image presents a high passing rate which proves its high usability in real life conditions. Multi-lines and variably structured plates are also supported in ALPRNet and therefore it can be used in heterogeneous plate scenarios that occur in the wild.

3. Methodology

The methodology of the study involves DL, which has combined steps of detection and recognition. The limits of the available data at local locations and peculiarities of Kurdistan Region of Iraq plates are supposed to be resolved with the help of the proposed methodology. Under the best designs already established in object detection and optical character recognition, we can design the system to work as a sequential pipeline in the following sense: detect LPs in the vehicle images, then process the identified areas to extract the textual information. Data preparation, model architectures, training strategies, and evaluation protocols of each component of the pipeline are provided in the sections below.

3.1. Dataset

Iraqi LPs included the name of province as written words as shown in Figure 1, until 2022, the normal display of the name of the governorate has been replaced with symbols that are used within the vehicle number as demonstrated in Figure 2. Nevertheless, according to the instructions of the Iraqi General Traffic Directorate, licenses are replaced by pairs of symbols indicating regions on the plates since 2022. This divergence with the recognized standard complicates LPR systems since it demands the use of symbolic representation instead of the common name [10]. The LPs in new Iraq are made up of (two digits, letter, numbers), as shown in Figure 3. Because dataset is not available for Iraqi LPs in the new style, a distinctive pool of LPs of Kurdistan Region of Iraq in new form is chosen. In total, the sample contains 306 photos of LPs, and all of them are real photos of the cars acquired in the Kurdistan Region in the usual working day and nighttime and made it to be gathered randomly at various angles and distances. It is currently limited to the 3 northern governorates: Erbil, Dohuk, and Sulaymaniyah. Due to the small size of the dataset no further preparation or data augmentation is performed. The dataset contains no environmental conditions that are not indicative of weather differences (rain, fog, intense sunlight), lighting differences (intense shadows, dark lighting), and distressed LPs. The changes in vehicles and LPs are restricted to the private cars and do not incorporate the trucks or motorcycles.

(a) (b) (c)

Figure 1. Sample of old LP[19].

Figure 2. The unified coding for all governorates and the new Iraqi painting codes [20].

Figure 3. An example of a private car LP from Iraq [9].

3.2. Experiment Platform

Training and testing of the modules in this study is made by the use of an NVIDIA GPU P100 graphics card. The one-stage detector is the official Ultralytics YOLO PyTorch. And initialization of pre-trained weights is used on dataset CCPD.

3.3. Justification of Model Selection and Transfer Learning

Selection of models in this research is based on a tradeoff between precision, computing performance and applicability to small datasets. The YOLOv11n is chosen as the detector backbone, it has been chosen over the previous versions like YOLOv5 or YOLOv8 because it has a better trade-off of speed, parameter efficiency, and detection accuracy. Although YOLOv8 has shown good results, it uses more computational resources and is inefficiently trained on domain-specific small datasets. Instead, it is able to compete or outperform the accuracy of the original model using fewer parameters, thus better fitting resource-constrained systems (like smart transportation systems) with real-time uses. Furthermore, Design refinements included in YOLOv11n to improve the accuracy of localization and the stability of the system with limited data conditions are a crucial need in the Iraqi LP of that time. These improvements positively make it a moderate option in terms of computational efficiency, robustness, and transfer learning adaptation over its predecessors.
And LPRNet is chosen among other OCR architectures like CRNN or EasyOCR due to a number of reasons. To begin with, it is also an effective end-to-end model that directly predicts the sequence of characters, without making explicit character segmentation, which is especially beneficial in the case of LPs of different layout or space. Second, the architecture it is much lighter than that of CRNN, enabling it to perform inference much faster without compromising the recognition accuracy, which is essential when deploying LPRNet in intelligent transportation systems. Third, it is fine-tuned to LP recognition and unlike EasyOCR, which is a general-purpose OCR tool, it demonstrates high robustness in difficult circumstances (e.g., oblique angles, different illumination, and partial occlusions) in the literature. These characteristics render it a somewhat better fit to our application, particularly due to the small size of the dataset and the necessity of having a high generalization level on the contemporary LPs of Iraqi origin. And also CCPD dataset is selected as a baseline because of its large range of lighting conditions, angles, and realistic scene, which is appropriate to reliably train and test LPD models.
In addition, to confirm the benefit of transfer learning, we performed an ablation study by comparing YOLOv11n and YOLOv8n trained on scratch with YOLOv11n started with pre-trained weights. As will be demonstrated in the Results section, the pre-trained version performed much better than the scratch model that illustrates the usefulness of transfer learning and cross-domain generalization to Iraqi LPs.

3.4. Detection and Recognition License Plate

This has been proposed by two major stages of the ALPR pipeline namely detection and recognition.

3.4.1. License Plate Detection (LPD)

It is well known that YOLO's real-time detection capabilities contributed to its popularity. This algorithm, is trained independently on the CCPD[21] dataset. In this work YOLOv11n model used to train on the CCPD dataset. Organized images and labels files in both directory structures. To make sure the model is thoroughly evaluated, the dataset is divided into subsets using K-fold cross-validation, 10% of each set is set aside for internal validation, 10% for testing, and 80% for training. Using hyperparameters that optimize this dataset, the model is trained on 30 training epochs with 16 models in a batch with a learning rate of 0.0001 and 640x640 pixels. AdamW optimizer to enhance performance and generalization on different situations is applied. A set of data augmentation methods are used to help in generalization and reduce the risk of overfitting, and these included color transformations of the HSV color space. Most of the proposed detection pipeline, in which YOLOv11n detects LPs, is illustrated in Figure 4. Then saved the weight to test our dataset, and prepared the data as follows: Once we had our dataset, the labelImg program is used to label the images. After that, the data are restructured once again to match the YOLOv11n hierarchy, and pictures and annotations files are sorted into specific directories. The dataset path is specified in an automatically created YAML configuration file, and the pre-trained data is then tested on the test set by use of a fixed image size of 640x640 pixels and a batch size of 16 pixels. The model also presented standard object detection measures, such as Average Precision (mAP@0.5 and mAP@[0.5:0.95]) and Recall, and Precision, which are a summary of the performance on object detection. And there is no preprocessing done.

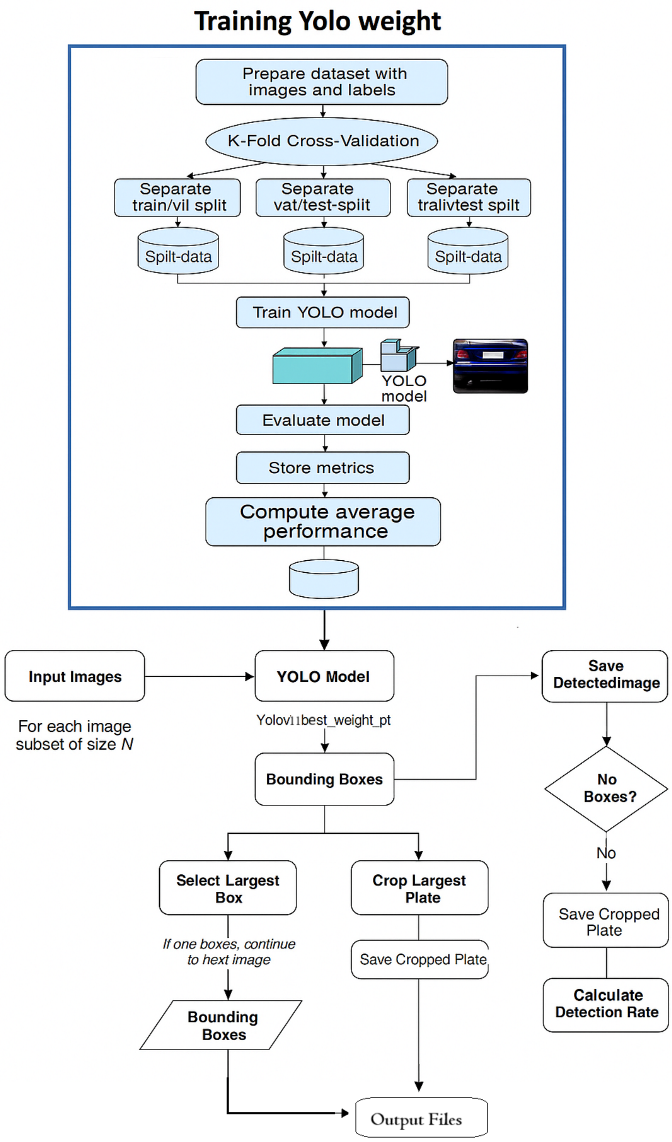

Figure 4. Structure of our detection model.

Along with testing with fixed weights, performed an experiment in which YOLOv11n and YOLOv8n are trained independently on the modern Iraqi LP dataset in a fresh run to give a baseline to compare against. In this setup, to ensure a fair comparison, the model is initialized and optimized on the same data splits of 80% training, 10% validation and 10% test as in the transfer learning experiment. The training is done using 30 epochs consisting of a batch size of 16, input resolution of 640x640 pixels, and AdamW optimizer with a learning rate of 0.0001. Generalization is enhanced by standard data augmentation methods like color conversion using HSV.

This is a scratch-trained model that is a simple yet significant point of reference to measure the advantage of transfer learning. As reported in the Results section, the scratch model had significantly weakest performance than the pre-trained version, especially in Recall and mAP, which is an important indicator of the critical importance of pre-trained weights when using limited and domain-specific data to form a robust detection.

3.4.2. License Plate Recognition (LPR)

The LPRNet network has been very effective at end-to-end LPR, without any explicit character segmentation. Its recurrent layers and CNNs have also been effective in its use to identify different conditions. After the detection phase, once we had cropped the detected area LP and preconditioned the cropped LP images, the LP is read with the help of the LPRNet model. This is an effective neural net based on the convolutional layers, then preceded by two Greek LSTM networks. In this experiment, the model took in clipped LP data to train the model.

It reserved,10% of each set to internal validation, 10% to test, and 80% to train. The processed and encoded of the data has turned around by resizing each image with fixed resolution of 160 x 40 pixels. The images have been saved to retain the color attributes in RGB format. By means of using a predefined alphabet, the matching tags are removed out of the file names and converted to a chain of character and numbers indexes. Tags are padded to a maximum of 10 characters, and a special index is allocated for padding. As shown in Figure 5, a lightweight specialized CNN is build where LP OCR is end-to-end using a condensed LPRNet. It consists of three convolutional layers with pooling and ReLU activations, and the other layer is the one of adaptive average pooling. The feature maps in the extracted are flattened, then passed through a fully linked layer in order to generate logits at each position of type.

Model training is with Adam optimizer to 120 epochs using a batch size of 16 and the learning rate of 0.001. The training loss is computed with the help of cross-entropy loss, it achieved a train loss of 0.0011 and validation loss of 0.0012, considering every character location a distinction task. Each epoch is followed by performance estimation on the validation set to monitor the trends of overfitting and the generalization. The trained model saved the parameters and is used and deployed later.

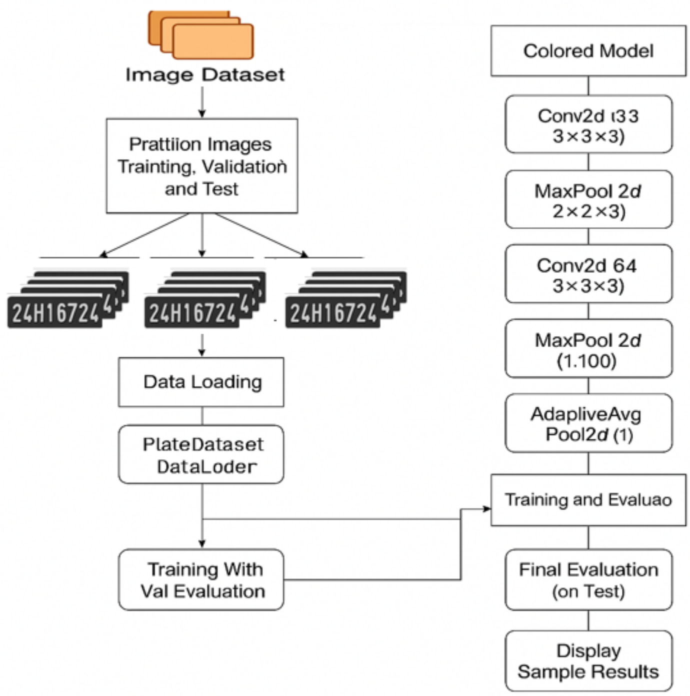

Figure 5. Structure of our recognition model.

3.5. Basis of Measurement

Once the training is completed, a set of typical metrics applied in computer vision studies are applied to determine the efficiency of the proposed LPRD system. Our detection stage involved precision, recall, F1-score and mAP. Recall measures the proportion of detected plates that really were identified, whereas Precision is a measure of the proportion of the detected plates that are correctly identified out of all the observed events. The F1-score provides a balance between harmonic average of precision and recall. The object detection tasked benchmark assessment measure is the mAP measure, which encapsulates the cumulative accuracy of detection and it is calculated at different Intersection over Union (IoU) thresholds. These are defined as follows:

$$Precision = \frac{TP}{TP + FP} \tag{1}$$

$$Recall = \frac{TP}{TP + FN} \tag{2}$$

$$F1 = 2 * \frac{Precision * Recall}{Precision + Recall} \tag{3}$$

$$mAP = \frac{1}{N} \sum_{i=1}^{N} AP_i \tag{4}$$

where TP is the number of true positives, FP is the number of false positives, and FN is the number of false negatives. Where APi is the Average Precision for class i, and N is the number of classes.

Sequence-level accuracy (also known as plate-level accuracy), per-character accuracy, and a confusion matrix are used to evaluate performance for the recognition step. While per-character accuracy assesses recognition ability at the individual character level, offering a more detailed analysis, sequence-level accuracy determines if the full LP string is accurately recognized. Common character-level misclassifications are further highlighted by the confusion matrix, which is helpful for identifying weaknesses in the recognition pipeline.
The combination of these complementary metrics provides a complete and equitable assessment of both detection and recognition phases and provides comparability with existing studies and points to the robustness of our suggested framework.

4. Experiments and Results

4.1. License Plate Detection Performance

LP performance is evaluated and experimented on the pre-trained weights. The technique is performing better knowing that the model is not trained on our data. Table 1 indicates to result of our data. using standard object detection metrics, including precision, recall, F1-score, and mean Average Precision (mAP).

Table 1. Detection performance using standard metrics on our data.

Model	mAP@0.5	mAP@0.5-0.95	Precision	Recall	F1-Score
Pre-Training(YOLOv11n)	96.0%	72.0%	90.1%	93.2%	91.0%

These outcomes show that the model performs reliably in terms of detection. While the lower mAP@[0.5:0.95] 72.0% reveals difficulties at stricter IoU thresholds, which is predicted in short datasets and under difficult conditions, the relatively high mAP@[0.5] 96% suggests great localization accuracy at standard IoU. However, robustness in reducing false positives and false negatives is confirmed by the balance between Precision 90.1% and Recall 93.2%.

4.2. Ablation Study on Transfer Learning

To highlight the benefit of transfer learning, YOLOv11n and YOLOv8n are trained from scratch on the Iraqi dataset and compared it against the model initialized with pre-trained weights. Table 2 shows the results.

Table 2. Comparison of YOLOv11n, YOLOv8n training and pre-trained strategies.

Model	mAP@0.5	mAP@0.5-0.95	Precision	Recall	F1-Score
YOLOv8n	71.4%	58.3%	84.4%	54.2%	65.9%
YOLOv11n	71.9%	54.4%	78.7%	61.7%	61.7%
Pre-Training(YOLOv11n)	96.0%	72.0%	90.1%	93.2%	91.0%

Overall, as Table 2 shows, both YOLOv11n and YOLOv8n underwent poor performance when trained directly on the small Iraqi dataset, reaching Recall under 62.0% and mAP@0.5 of approximately 71.0%. Conversely, the out-of-the-box YOLOv11n reached 93.2% Recall and 96.0% mAP@0.5, very much higher than both scratch-trained models. These findings indicate that using transfer learning is crucial, as it can promote high-performance when using training from scratch cannot, and verify that YOLOv11n with pre-trained weights can offer the best tradeoffs in terms of current Iraqi LP detectors.

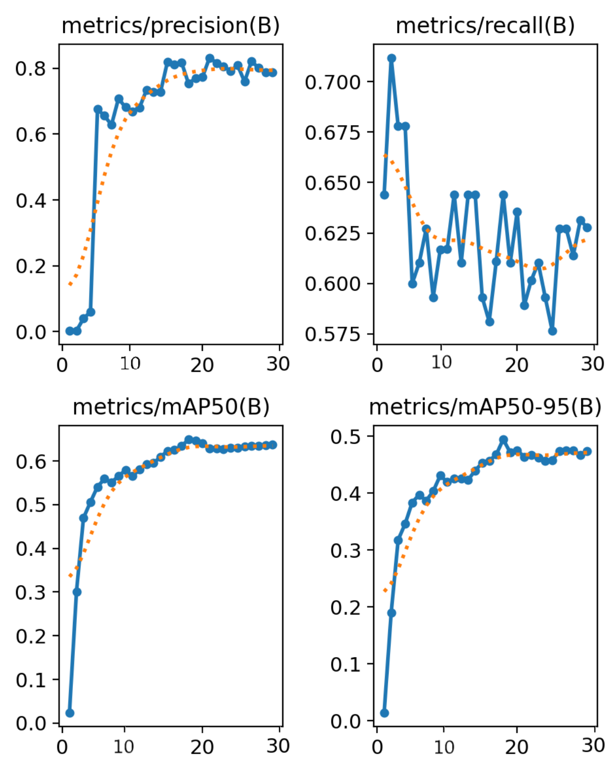

(a) YOLOv8n detection metrics from scratch.

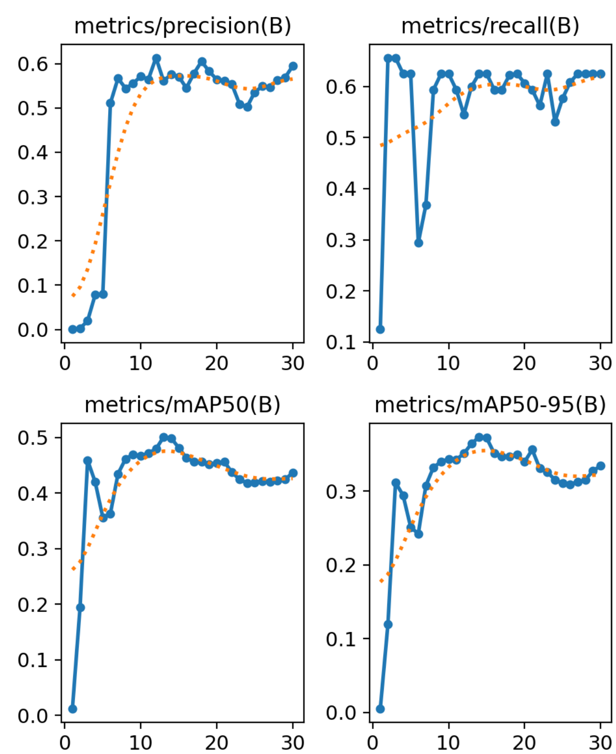

(b) YOLOv11n detection metrics from scratch.

The Figure 6 (a,b), displays the development of measures of detection over 30 epochs that are precision, recall, mAP at 0.5, and mAP at [0.5:0.95]. Both precision and mAP measures continuously increased with the epochs and reached stabilization at 54.2% (recall), 84.4% (precision), 71.4% (mAP at 0.5) and 58.3% (mAP at [0.5:0.95]) in YOLOv8. And reached stabilization at 61.7% (recall), 78.8% (precision), 71.9% (mAP at 0.5) and 54.4% (mAP at [0.5:0.95]) in YOLOv11n. It is important to remember, though, that fluctuated is more erratic, indicating that it is challenging to have high recall on a small data set. These trends support the findings of the ablation study of which transfer learning offered a significant enhancement in comparison to training from scratch.

4.3. License Plate Recognition Performance

The identification step entailed the use of LPRNet model. LPRNet technology records better performance as indicated in Table 3. We measured both per-character accuracy (what fraction of the sample is correctly recognized) and full-sequence accuracy (all the characters in a plate are correctly recognized). There is also a confusion matrix which is made to highlight the most common character-level errors.

Table 3. Recognition performance of our data.

Model	Full-sequence accuracy	Character-level accuracy
LPRNet	99.8%	99.5%

The Figure 7 indicates the validation accuracy (full-sequence and character-level) as well as the training and validation loss during 120 epochs. As can be seen, the accuracies gradually upcurve and approach 100 percent, and the curves of the losses fall gradually to almost zero, which indicates a stable training procedure and good generalization outcomes.

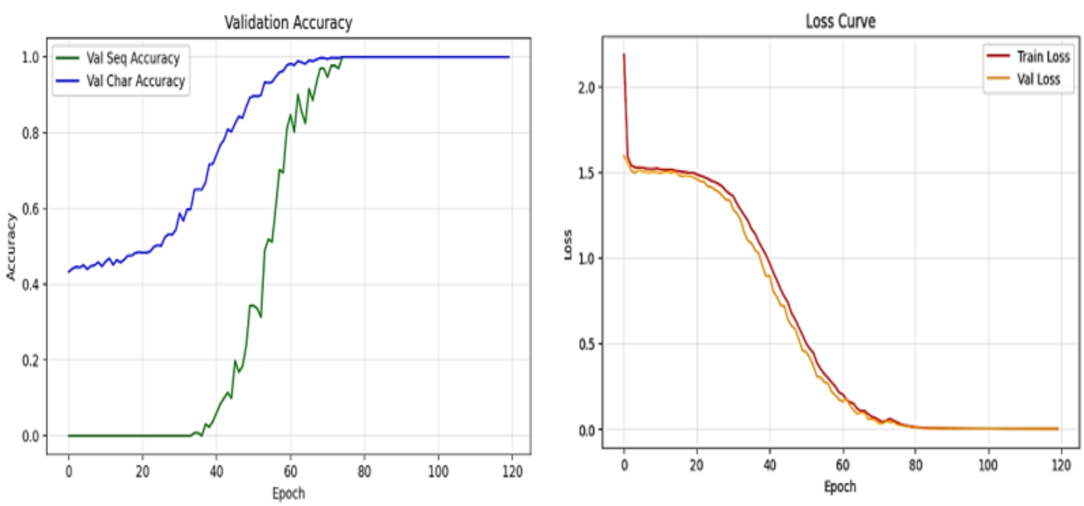

Figure 7. Validation accuracy and loss curves for the LPRNet recognition model.

As Figure 8 depicts, the confusion matrix shows that all the characters on the LP are correctly identified, and the dark blue cells can be arranged around the main diagonal. This shows that the model has reported very high per-character accuracy, i.e. most of the characters were identified correctly. There are almost no instances of hypothesis errors in the non-diagonal cells and this proves that the trained model is resilient and stable when distinguishing between visually similar characters.

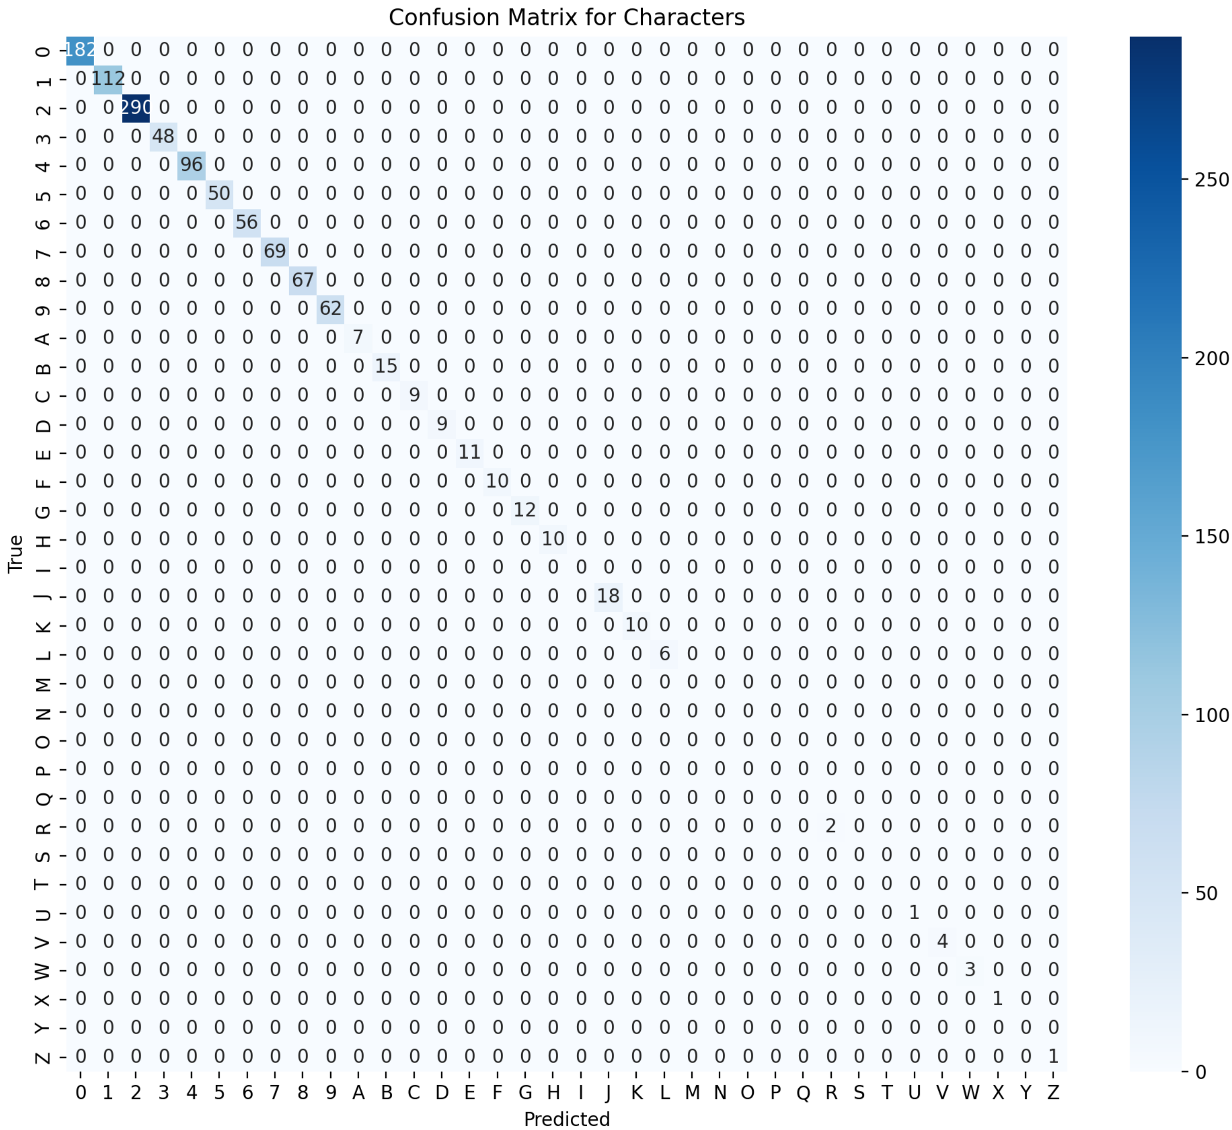

Figure 8. Confusion matrix of each character.

Figure 9 gives qualitative evidence of the effective detection and recognition results. The red bounding boxes show that the pre-trained YOLOv11n detector can detect LPs well regardless of the viewing angle and distance. The text at the top of each bounding box is yellow and contains the plate numbers which in these cases are the same as the ground-truth labels. These tests demonstrate that the capability to do generalization is very high in the proposed pipeline even in real-world scenarios of Iraqi plates in various lighting conditions and orientations. These visualizations, along with the confusion matrix in Figure 8, offer a fair look at the overall system performance, both with its very high accuracy in regular situations and its small number of instances of character-level misclassification.

Figure 9. Detection and Recognition results.

Tests are done on the compiled dataset. Detecting measures, including Precision, Recall, and mAP are calculated to detect, not just basic image-level accuracy. Based on the results, the model had a precision of 90.1%, a recall of 93.2%, a F1-Score of 91.0% a mAP@[0.5] of 96.0% and mAP@[0.5:0.95] of 72.0%. For recognition accuracy is measured by per-character accuracy and percentage of LPs in which the predicted text fully matched the ground truth. The strong detection rate achieved by the system is established by the 99.8% full-sequence accuracy and 99.5% character accuracy of the system.

Examples of qualitative findings showing the successful identification and recognition of modern Iraqi LPs are shown in Tables 1 and 2. Despite the small size and limited diversity of our dataset, the strong performance of both detection and recognition models trained initially on plates highlights the usefulness of transfer learning in cross-domain LPR, despite the modest size and restricted diversity of our dataset.

However, A small sample consisting of 306 photos of three governorates in the north of Iraq with minimal differences in weather, lighting or type of vehicles are a significant limitation of the study. This lack of diversity may compromise the generalization capability of the model to other territories and challenging real world scenarios. We will expand the dataset in future studies to include more Iraqi governorates, taken in different weather conditions, and more different types of vehicles. To mimic these variations and enhance the resilience of the system even more, advanced data augmentation methods will be explored as well.

4.4. Comparison

The suggested ALPR system was compared to an earlier study by Ahmed et al. (2023)[10] that also tested new Iraqi LP to evaluate the effectiveness of the proposed system. They applied Tesseract OCR and a CNN-based recognition model. Without an explicit any detection model, paid attention specifically to the stage of recognition accuracy they stated is 95.5%, and no detection metrics are provided. Conversely, LPRNet is applied to character recognition over the whole system and a pre-trained YOLOv11n to detect the LP. Our dataset of 306 images gave the system 99.8% recognition accuracy and 96.0% detection accuracy. This demonstrates the effectiveness of the two-step pipeline, which can achieve excellent performance even with modern Iraqi LPs without pre-processing or per-character segmentation. The benefits of transfer learning are indicated by the high level of performance because the originally trained YOLOv11n model on Chinese plates can generalize to a new area with minimum adaptation. Table 4 provides the two approaches compare in terms of dataset size, detection and recognition model and the general accuracy.

Table 4. Comparison Table.

Study	**Images No.**	**Detect Model**	**Reco. Model**	**Detect Acc.**	**Reco. Acc.**
Ahmed et al. (2023) [10]	180	—-	Tesseract, CNN	—-	35.0% 95.5%
Our	306	Our weight pre-trained	LPRNet	96.0%	99.8%

In general, the findings indicate that the combination of pre-trained detection models and an effective recognition network can be very promising compared to the past methods, especially in end-to-end LP recognition and the generalization across domains.

5. Discussion

The experimental outcomes prove the successfulness of the suggested ALPR system in identifying newly formatted Iraqi LPs in the conditions of limited datasets.

The results indicate that pre-trained models of DL using large, diversified datasets and models are capable of successfully transferring to new domains, such as new Iraqi LPs, with minimal adaptation. The main advantage of such an approach is the ability to achieve high accuracy based on small, tailored datasets, combining with pre-trained models. Nevertheless, the major drawback of this paper is limited data size and scope. The sample of this paper includes only 306 pictures that are taken specifically in three governorates of Iraq, which are located in the north (Erbil, Duhok, and Sulaymaniyah). This is limited by the fact that currently deployed format of the Iraqi LP has not been widely used in other governorates and that the limited samples available are outside the Kurdistan Region and are very limited. Consequently, the dataset fails to represent any significant variations like various weather types (rain, fog, bright sunlight), lighting conditions (night, shadows, glare), image quality problems (motion, focus, and various resolutions), and the overall types of vehicles (commercial vehicles, motorcycles, or plates with dirt and wear and tear). It is geographical and environmental prejudice which restricts the generalizability of the model to other areas. In future research, this study will be important in increasing the sample size by a very large margin once additional plates have been distributed across Iraq in a bid to capture a high number of environmental and vehicle conditions. In the meantime, the state-of-the-art data augmentation algorithms (e.g., synthetic blur, noise, brightness/contrast, rain/fog filters) are the immediate way to artificially improve the robustness until one has access to a wider real-world dataset.

Our proposed system achieved high scores in terms of accuracy which proves that pre-trained weight used in LPD coupled with LPRNet in character recognition may be used with a high degree of success and accuracy even on LPs of a particular country recently formatted like the Iraqi LPs used in our experiment. Remarkably, we do not impose special character segmentation or manual pre-processing, because the system works end-to-end on RGB plate images. In comparison with already existing studies like the one by Ahmed et al. (2023) [10], which assessed the quality of the Tesseract OCR and the quality of a CNN-based categorizer, our approach has a couple of strong points. They only did a study on the recognition stage and no detection accuracy is reported or a specific detection module is used. Conversely, our system incorporates the weight that pre-trained on YOLOv11n as a powerful detecting model, which attained the 96.0% accuracy in the detection of a homemade Iraqi dataset, though being trained with Chinese plates, as Table 1 shown. That is where the power of the transfer learning lies in cross-domain generalization.

Besides, we got 99.8% recognition accuracy using a more efficient pipeline which does not involve per-character annotation through our method compared to 95.5% accuracy using explicit character segmentation through the CNN approach developed by Ahmed et al.[10]. These results indicate that a two-step DL pipeline, with strong detection following sequential recognition may generalize across a variety of plate shapes with minimal domain adaption.

The added assessment measures confirm the quality of our system. In comparison with Ahmed et al. (2023), who merely described the OCR accuracy without detection metrics, we offer a more holistic assessment model because we report the standard object detection metrics (Precision, Recall, F1-score, and mAP). Moreover, the presented methods of per-character accuracy and confusion matrix analysis reveal the practical strengths and weaknesses of the recognition stage that had not been discussed in the work before.

In order to further support the rationale behind our course of action, also examined the performance of YOLOv11n and YOLOv8n, which are trained on the Iraqi dataset, starting with no pretrained weights. The scratch-trained models are significantly very weak (Precision 84.4%, Recall 54.2%, mAP@0.5 71.4%, mAP@[0.5:0.95] 58.3%, F1-score 65.9% in YOLOv8n, and Precision 78.7% , Recall 61.7%, mAP@0.5 71.9%, mAP@[0.5:0.95] 54.4%, F1-score 61.7% in YOLOv11n) when compared to the transfer learning one (Precision 90.0%, Recall 93.0%, mAP@0.5 96.0%, mAP@[0.5:0.95] 72.0%, F1-score 91.0%). This ablation study highlights the crucial role of transfer learning, as the scratch model struggled to generalize with such a small and domain-specific dataset.

Furthermore, the fact that LPRNet has been chosen in the place of other OCR architectures like CRNN or EasyOCR can be justified by its efficiency as an end-to-end recogniser, which does not require any character-level segmentation and still achieves high recognition accuracy. Combining these results, it is clear that the design options of YOLOv11n with pre-trained weights in detecting and LPRNet in recognizing can be used in a strong and efficient pipeline, especially in the situation of limited resources and data like modern Iraqi LPs.

The most constraint limitation of this research is the very small and geographically limited dataset. There are a total of 306 collected images, and the images are all found in three of the northern governorates (Erbil, Duhok, and Sulaymaniyah). As a result, the dataset is not diverse in such important areas as weather changes (rain, fog), lights (shadows, glare), and image quality (motion blur, focus blur). Additionally, the dataset is narrowed down to only private vehicles and does not cover other types of vehicles (e.g., commercial vehicles, motorcycles) and LPs that are dirty or partially covered or damaged. These considerations present a bias in geography and environment because the limited samples available are outside the Kurdistan Region and are very limited, which limits the applicability of the proposed model to other areas of Iraq and to the conditions of the real world.

6. Conclusion and Future Work

In conclusion, this paper has offered an end-to-end system of detecting and recognizing newly created Iraqi LPs through the use of a YOLOv11n based detection block with pre-trained weights and the LPRNet recognition module. The advantage of transfer learning in the case of adapting the already trained models (trained on the CCPD dataset) to a new domain is seen by the proposed pipeline, which had a high detection precision of 96.0% and a recognition accuracy of 99.8%. A further comparative experiment between (YOLOv11n and YOLOv8n) trained from scratch and training on a small dataset saw transfer learning dramatically performing better in cross domain generalization, and proved its usefulness. However, the limited and geographically limited dataset in this research is the biggest weakness of the research that decreases the capacity of the system to generalize to various conditions including bad weather conditions, low light conditions, motion smears, and various vehicle models. In future study, it is necessary to intensively increase the data volume in order to achieve wider representativeness. Our images will be taken in all Iraqi governorates under varying environmental conditions (rain, fog), at various times, and among a more diverse range of vehicles, such as commercial vehicles and motorcycles. We also wish to have LPs with dirt, damage, and partial occlusion, and other image qualities (e.g., motion blur, variable resolutions). The creation of this scale and variety of data will enable us to enhanced the efficacy and generalizability of the suggested ALPR system to practical deployment conditions in Iraq.

REFERENCES

1. H. Ibrahim, O. M. Fahmy, and M. A. Elattar, "License plate image analysis empowered by generative adversarial neural networks (gans)," *IEEE Access*, vol. 10, pp. 30846–30857, 2022.
2. F. M. Adhari, T. F. Abidin, and R. Ferdhiana, "License plate character recognition using convolutional neural network.," *Journal of Information Systems Engineering & Business Intelligence*, vol. 8, no. 1, 2022.
3. I. Shafi, I. Hussain, J. Ahmad, *et al.*, "License plate identification and recognition in a non-standard environment using neural pattern matching. complex intell syst," Syst, 2021.
4. S. Alghyaline, "Real-time jordanian license plate recognition using deep learning," *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 6, pp. 2601–2609, 2022.
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.
6. A. Ramajo-Ballester, J. M. A. Moreno, and A. de la Escalera Hueso, "Dual license plate recognition and visual features encoding for vehicle identification," *Robotics and Autonomous Systems*, vol. 172, p. 104608, 2024.
7. R. Al-Batat, A. Angelopoulou, S. Premkumar, J. Hemanth, and E. Kapetanios, "An end-to-end automated license plate recognition system using yolo based vehicle and license plate detection with vehicle classification," *Sensors*, vol. 22, no. 23, p. 9477, 2022.
8. G. Y. Abbass and A. F. Marhoon, "Car license plate segmentation and recognition system based on deep learning," *Bulletin of Electrical Engineering and Informatics*, vol. 11, no. 4, pp. 1983–1989, 2022.
9. Y. D. Salman, H. S. Alhadawi, A. S. Mahdi, and F. T. AL-Dhief, "Improved automatic license plate recognition system in iraq for surveillance system using ocr," in *International Conference on Emerging Technologies and Intelligent Systems*, pp. 270–277, Springer, 2022.
10. S. Ahmed, A. Akbas, and E. Naser, "Tesseract opencv versus cnn: A comparative study on the recognition of unified modern iraqi license plates," *Revue d'Intelligence Artificielle*, vol. 37, no. 5, p. 1331, 2023.
11. Q. Huang, Z. Cai, and T. Lan, "A single neural network for mixed style license plate detection and recognition," *Ieee Access*, vol. 9, pp. 21777–21785, 2021.
12. A. Sugiharto, R. Kusumaningrum, *et al.*, "Enhanced automatic license plate detection and recognition using clahe and yolov11 for seat belt compliance detection," *Engineering, Technology & Applied Science Research*, vol. 15, no. 1, pp. 20271–20278, 2025.
13. L. Y. Chan, A. Zimmer, J. L. da Silva, and T. Brandmeier, "European union dataset and annotation tool for real time automatic license plate detection and blurring," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–6, IEEE, 2020.
14. H. Shi and D. Zhao, "License plate recognition system based on improved yolov5 and gru," *Ieee Access*, vol. 11, pp. 10429–10439, 2023.
15. N. H. Jasem and F. Mohammad, "Iraqi license plate recognition system using (yolo) with sift and surf algorithm," *Journal of Mechanics of Continua and Mathematical Sciences*, vol. 15, 2020.
16. 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, 2020.

17. M. H. Abbas and Z. M. Mohammed, "Detection and recognition of iraqi license plates using convolutional neural networks," *Science Journal of University of Zakho*, vol. 13, no. 1, pp. 7–17, 2025.
18. 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.
19. D. Habeeb, F. Noman, A. A. Alkahtani, Y. A. Alsariera, G. Alkawsi, Y. Fazea, and A. M. Al-Jubari, "Deep-learning-based approach for iraqi and malaysian vehicle license plate recognition," *Computational intelligence and neuroscience*, vol. 2021, no. 1, p. 3971834, 2021.
20. iqcar, "iqcar," *https://www.threads.com/@*$iqcar_?hl = en$.
21. Z. Xu, W. Yang, A. Meng, N. Lu, H. Huang, C. Ying, and L. Huang, "Towards end-to-end license plate detection and recognition: A large dataset and baseline," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 255–271, 2018.