



ORIGINAL RESEARCH ARTICLE

PERFORMANCE EVALUATION OF YOLO OBJECT DETECTION MODELS FOR
AUTOMATED LICENSE PLATE RECOGNITION

O. F. Odeyinka^{1*}, W. A. Raheem¹, C. O. Folorunso¹ and A. Fajinmi¹

¹Department of Systems Engineering, University of Lagos, Akoka, Nigeria

*Corresponding author's email: oodeyinka@unilag.edu.ng

ARTICLE INFORMATION

Received: 25th February 2025

Revised: 27th April 2025

Accepted: 28th April 2025

Keywords:

Automated license plate
recognition system
Deep learning
Object detection
YOLO

ABSTRACT

This paper outlines an evaluation and comparison of three different You Only Look Once (YOLO) object detection models - YOLOv3, YOLOv8, and YOLOv10 for use in automated license plate recognition (ALPR) systems. To analyze these models, a total of 261 license plate images were collected from the car park of an auditorium inside the University of Lagos, Nigeria. Metrics of each model (accuracy, precision, recall, F1 score, and training efficiency) were used to measure the performance of the models. Results indicates that YOLOv8 (accuracy of 86.9%, precision score 100%, recall of 87%, and an F1 score of 0.93) significantly outperformed the other models, demonstrating its potential as a robust solution for object detection. In contrast, YOLOv3 had an accuracy of 62.1%, precision of 75%, a recall of 78.3%, and an F1 score of 0.766, reflecting balanced performance but slower training times. YOLOv10, despite being the latest version, showed mixed results, achieving an accuracy of 43.2%, a precision of 47.5%, a recall of 82.6%, and an F1 score of 0.603. This study highlights the critical importance of model selection based on specific application needs and suggests that further optimization may enhance the capabilities of YOLOv10 for future developments in ALPR systems.

© 2025 Faculty of Engineering, University of Maiduguri, Nigeria. All rights reserved.

1.0 Introduction

Automated License Plate Recognition (ALPR) is an intelligent system that interprets car license plates using optical character recognition on snapshots (Musa, 2023). This technology works remotely, removing the need for physical contact or vehicle engagement. Law enforcement agencies have adopted it for comparing vehicle license plates with lists of stolen, wanted, and other vehicle interests. The important steps in ALPR are license plate localization and detection, which affect the recognition process's overall accuracy, efficiency, and adaptability (Bakshi et al. 2023). However, some factors can affect license plate localization and detection accuracy (Jenitta et al. 2023). These include blurred or occluded plates, skewed camera angles, bad weather conditions, poor lighting conditions, non-standard plates, and irregular image acquisition. Despite the challenges posed by these factors, ongoing research and technological advancements exist (Jiang et al. 2022). Detecting instances of objects from one or more classes in a picture is the main idea behind object detection. Classes of objects in an image include people, cars, or faces (Amit et al. 2021).

ALPR systems integrate three primary techniques - object detection, image processing, and pattern recognition (Batra et al., 2022). Object detection employs various models which are categorized into traditional computer vision techniques and deep learning-based methods, where the You Only Look Once (YOLO) models are classified. Despite advances in technology, traditional computer vision techniques are becoming obsolete and have many limitations (Jamtsho et al. 2021; Solunke & Gengaje, 2023). The rise of deep learning has made object detection more efficient, quick, and accurate, hence the integration of models like the YOLO versions (O'Mahony et al. 2020). These models address challenges such as the inability of pixel connectivity to capture all characters in joined sequences, illumination issues, and poor recognition rates on ambiguous characters (e.g., B-8, O-0, A-4, C-G, D-O, and K-X) (Wang et al. 2024). Other concerns addressed include broken characters and issues of multi-style plates (plates with different national fonts and syntax) (Solunke & Gengaje, 2023). For example, Bakshi et al. (2023) enhanced ALPR systems for diverse conditions by using a Spatial

Transformer Network to align distorted license plates and a CNN to recognize characters with various fonts and sizes, achieving superior accuracy in challenging scenarios compared to other methods.

Object localization is the process of establishing a specific “location” or spatial coordinates of an object within an image or video. Localizing license plates from car images is a complex operation due to considerable variance in size, shape, color, texture, and spatial orientations of license plate regions in these pictures (Saha *et al.* 2009). In Gnanaprakash *et al.* (2021), localization is highlighted as an objective of an ALPR system that interprets these images using optical character recognition (OCR) technology to get the license number of the vehicle.

YOLO is an approach to object detection that was developed by Redmon *et al.*, (2016). (Redmon *et al.* 2016). The YOLO approach treats object classification as a regression problem by multi-dimensionally separating boxes and applying class probability (Lee & Kim, 2020). The YOLO algorithm processes an input image through a straightforward deep convolutional neural network for object detection. Different YOLO algorithms have been used for objection detection (Sung *et al.* 2020; Zhao and You, 2020; Laroca *et al.* 2021; Jenitta *et al.* 2023, Wadare *et al.* 2023). However, the YOLO algorithm has undergone different forms of improvement. In YOLOv2, a batch normalization layer was added to enhance accuracy and speed, while the convolutional anchor size was set to establish initial training, thus addressing the gaps in YOLOv1. Similarly, logistic regression was used in YOLOv3 to predict the object score of a bounding box, making it significantly faster than other object detection algorithms at equivalent performance (Zhao and You, 2020). In Lin *et al.* (2022), it was noted that YOLOv3 exhibited certain limitations, leading to the introduction of YOLOv4. When compared with YOLOv3, YOLOv4 effectively prevents information loss, a weakness observed in YOLOv3 (Sozzi *et al.* 2022). In real-time images, each frame of YOLOv4 demonstrates greater stability compared to YOLOv3, resulting in recognition frames that are not affected by flickering. Also in exploring the different YOLO versions, Batra *et al.* (2022) highlighted the additions to YOLOv5 by focusing on its lightweight version, YOLOv5s (small). The authors noted its remarkable capabilities, its generating time and memory-efficient model - a substantial improvement over YOLOv4. YOLOv5 demonstrated high time performance, yielding relatively smaller-sized object detection models. Interestingly, Li *et al.*, (2022) developed a suite of deployment-ready networks at various scales, apt for diverse use cases. With the generous permission from YOLO authors, they named it YOLOv6. YOLOv6-N was documented to achieve a 35.9% average precision on the COCO dataset at 1234 Frames per Second (FPS). YOLOv6-S attains a 43.5% AP at 495 FPS, outperforming other mainstream detectors at the same scale, including YOLOv5-S. It particularly stands out for achieving high accuracy at an impressive frame rate. YOLOv7 takes the spotlight as highlighted by Zhao *et al.* (2023) in its application to detect maize seedlings in images. The model exhibits commendable performance in real-time object detection. However, issues surfaced during their analysis, including redundancy in feature extraction, loss of location information, and slow convergence in loss calculation. To address these shortcomings, they proposed a lightweight model derived from YOLOv7. This proposed model boasts the highest mean Average Precision (mAP) value, a smaller model size, and as well as detection speed that meets the requirements for real-time detection. Tamang *et al.* (2023) explored YOLOv8 for detecting face mask-wearing conditions. Their analysis led them to the conclusion that YOLOv8 demonstrates enhanced efficiency compared to its predecessors. This improvement is attributed to the utilization of a larger feature map and a more efficient convolutional network. Wadare *et al.*, (2023) are the first to merge YOLOv8 and EasyOCR into an ALPR system, and they investigated their combined capabilities in advanced license verification, as well as accuracy and efficiency verification.

As outlined by Hussain (2023), each of the previous YOLO models has recorded modest precision scores as shown in **Table I**. The latest release of YOLOv10 has proved to be a masterstroke in improvement, recording well over 70% precision (Wang *et al.* 2024). However, this is not a confirmation of the supremacy of YOLOv10 over the other versions. In fact, Hussain (2023) had suggested a comparative analysis of various YOLO versions to ascertain their efficiency and performance for ALPR systems. This is the focus of this research.

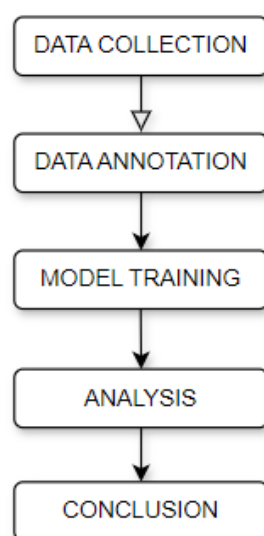
The growing demand for excellent object detection, particularly when recognizing license plates emphasizes the importance of evaluating automated license plate system models (YOLO versions). This is the motivation behind this research. Hence, the primary aim of this research is to evaluate and compare the performance of three different YOLO models - YOLOv3, YOLOv8, and YOLOv10 for their effectiveness in recognizing license plates. These algorithms are also used to demonstrate advances in deep learning and machine learning techniques. Similarly, investigating these models will help to understand their potential and further the adoption of advances in ALPR systems.

Table 1: Progression of Yolo Versions Hussain (2023)

Yolo Version	Year Released	Framework	Backbone	Average Precision (%)
YOLOv1	2015	Darknet	DarkNet24	63.4
YOLOv2	2016	Darknet 19	DarkNet24	63.4
YOLOv3	2018	Darknet-53	DarkNet53	36.2
YOLOv4	2020	Darknet-53	CSPResNext-50, CSPDarknet-53, & EfficientNet-B3	43.5
YOLOv5	2020	PyTorch	CSPDarknet53	55.8
YOLOv6	2022	PyTorch	EfficientRep	52.5
YOLOv7	2022	PyTorch	RepConvN	56.8
YOLOv8	2023	PyTorch	GhostblockV2	53.9

2. Materials and Method

This research is divided into six phases as described in **Figure 1**.

**Figure 1:** Flowchart for Methodology

- A. **Data Collection:** The dataset for this research work was obtained by capturing the plate number of vehicles parked at the J. F. Ajayi auditorium in the University of Lagos, Nigeria. In order to ensure clarity and consistency of the images taken, a single camera was used to take the pictures over a period of five days. The camera produced full high definition pictures of 16 megapixels with 261 images collected to ensure a diverse and comprehensive dataset. These images include a variety of plate numbers of various types of vehicles and in different conditions. The different pictures also helped to ensure diversity in the dataset, thus providing the foundation for the subsequent analysis and model training.
- B. **Data Annotation:** For this part, the dataset compiled was carefully labeled, outlining the license plate regions within each image. This task was carried out using Labellmg, a Python tool designed for image annotation and labeling, which ensured accuracy and consistency during the dataset creation process. After each labeling, the resulting annotations were saved in a standard “.txt” format. This careful approach makes sure that the labeled data is easily useable and optimized for seamless incorporation into the model’s training pipeline. A typical image annotation is outlined in Plate 1.



Plate 1: Image Annotation with LabelImg

When this was done, the dataset was then split into **Train, Val, and Test**, for both images and labeled images.




Figure 2: Code flow for YOLOv3, YOLOv8, and YOLOv10

- C. **Model Training:** Model training is the process of teaching the model to recognize patterns and make predictions by exposing it to a labeled dataset.
- D. **Analysis:** Statistical tests are run to see if any reported performance differences between the YOLO versions selected are statistically significant, strengthening our conclusions regarding their relative effectiveness.

This project uses four models for object detection – YOLOv3, YOLOv8, and YOLOv10. YOLOV3 has three closely related object detection models - YOLOv3, YOLOv3-Ultralytics, and YOLOv3u. YOLOv3 was chosen for this analysis because it is an older version of the YOLO series. Comparing its performance metrics with those of the newly released versions will help to understand the improvements and differences. YOLOv8 is the latest iteration in the YOLO series of real-time object detectors. This model is selected for its flexibility and fast training time. It has a lightweight network architecture, efficient feature fusion algorithms, and more accurate detection outcomes (Lou *et al.*, 2023). The YOLOv8 network solves challenges in categorization, object detection, and image segmentation. All of these methods recognize things in photos or videos in various ways. The YOLOv10 model is a new addition to the YOLO family, though it wasn't developed by the original creators. Instead, researchers at Tsinghua University built it using the Ultralytics Python package. This model was chosen for its ability to accurately predict object categories and positions in images with low latency, meaning it has minimal delay between input and output. Additionally, it was selected because it is the latest addition to the YOLO series.

In machine learning, evaluating a model's prediction accuracy is vital for determining its effectiveness. Metrics like precision and recall reveal how well the model identifies positive cases, highlighting areas for improvement and allowing for targeted optimization (Sozzi *et al.*, 2022).

Recall is the percentage of Real Positive instances that are accurately predicted as Positive (Powers, 2020). In other words, it is the measure of a model's ability to identify all important instances in a dataset. When recall is low, it suggests the model failed to detect a substantial number of true positive events. In simpler terms, the model incorrectly classifies many actual positives as negatives. As a result, despite the high accuracy of the model in positive predictions (when anything is predicted positively, it is true), it misses a large number of genuine positives.

It is expressed as the number of true positive predictions divided by the total number of actual positives (true positives plus false negatives). Mathematically:

$$\text{Recall} = \frac{\text{True Positives (TP)}}{\text{True Positives (TP)} + \text{False Negatives (FN)}} \quad 1$$

Precision, on the other hand, measures the ratio of correct positive predictions of correct positive predictions (Ganesan, 2022). It is the ratio of true positive predictions to the total number of positive predictions. It measures how accurate the positive prediction is

$$\text{Precision} = \frac{\text{True Positives (TP)}}{\text{True Positives (TP)} + \text{False Positives (FP)}} \quad 2$$

The F1- score curve graph helps in understanding the performance of a classification model, focusing on the balance between precision and recall across different confidence levels. It helps visualize the relationship between the F1 score and the confidence threshold used for predictions. Where the F1 score is the harmonic mean of precision and recall. It is expressed mathematically:

$$F1 \text{ score} = 2 \times \frac{\text{Precision (P)} \times \text{Recall (R)}}{\text{Precision (P)} + \text{Recall (R)}} \quad 3$$

The confusion matrix is a key technique for evaluating machine learning models. It is used to assess the accuracy of classification algorithms by comparing predicted outcomes against actual results (Karimi, 2021).

$$\text{Accuracy} = \frac{\text{True Positives (TP)} + \text{True Negatives (TN)}}{\text{Total}} \quad 4$$

3. Results and Discussion

The analysis was done on an 11th Gen Intel (R) Core i7-1165G7 @ 2.80GHz 16.0GB Installed RAM 64-bit system.

YOLOv3

The training of this model yielded three distinct batches of detection results. This section details the analysis and outcomes of the training run. The confusion matrix is given in **Figure 3**.




Figure 3: Confusion Matrix for the Training of YOLOv3

From **Figure 3**, the TP indicated by the cell at the top-left corner (18) represents the number of instances where the model correctly predicted “License Plate” when the actual label is “License Plate”. The FN denoted by the cell at the bottom-left corner (5) represents the number of instances where the model predicted “background” when the actual label is “License Plate”. The FP given by the top-right corner (6) represents the number of instances where the model predicted “License Plate” when the actual label is “background”. The

TN (0) represents the number of instances where the model correctly predicted “background” when it is “background”.

Upon evaluation, the accuracy is given as 0.621 (approximately 62.1%). The precision is calculated as 0.75, meaning 75% of predicted “License Plate” instances are correct. The recall is calculated as 0.783 indicates that the model correctly identified about 78.3% of actual “License Plate” instances while the F1 score is approximately 0.766, showing a balance between precision.




Figure 4: Precision-Recall curve for first training YOLOv3

The Average Precision at the 0.5 threshold is 0.832. This indicates that the model can distinguish well between relevant and irrelevant instances, leading to enhanced overall performance. From the **Figure 4** above, the optimal threshold is 0.574, where the model achieves its highest F1 score of 0.82. This high F1 score means the model performs well, making accurate predictions while maintaining a balance between precision and recall. Table 2 is a summary of results for YOLOv3 while **Figure 5** is the F1 score curve.

Table 2: Summary of results for YOLOv3

Training Batch	Dataset Size	mAp@50	Accuracy	Precision	Recall	F1 score
Batch A	261	0.832	62.1%	75%	78.3%	0.766




Figure 5: F1-Score Curve for YOLOv3

The Confusion Matrix for the training of YOLOv8 is given in **Figure 6**.




Figure 6: Confusion Matrix for the third training of YOLOv8

In **Figure 6**, the TP is 20, FN is 3, FP and TN are 0 giving a model accuracy of 86.95%. The precision is 1 meaning that every predicted “License Plate” is correct. Similarly, the recall is 0.87 showing that the model correctly identified about 87% of actual “License Plate” instances while the F1 score is 0.93 giving a good balance between precision and recall.




Figure 7: Precision-Recall Curve for the third training of YOLOv8

From the Precision-Recall Curve of YOLOv8 in **Figure 7**, the mAP@50 significantly increased from 0.506 and 0.4659 in the previous curves to 0.872 in the current batch training curve. This indicates a substantial improvement in the model’s ability to detect the license plates with higher accuracy. A higher F1 score indicates better performance. From **Figure 8**, it is observed that the optimal threshold is 0.667 when the model achieves its highest F1 score of 0.93. This high F1 score means the model performs exceptionally well, making accurate predictions while maintaining a good balance between precision and recall. **Table 3** summarizes the results of the training of the YOLOv8 model.




Figure 8: F1-Confidence Curve for training YOLOv8

Table 3: Summary of results for YOLOv8

Training Batch	Dataset Size	mAp@50	Accuracy	Precision	Recall	F1 score
Batch A	261	0.872	86.9%	100%	87%	0.93

The Confusion Matrix for the training of YOLOv10 is given in **Figure 9**.




Figure 9: Confusion Matrix for the second training of YOLOv10

From **Figure 9**, the TP is 19, FN is 4, FP is 21 and TN is 0. The model accuracy obtained is 43.2% while the precision is 47.5%, indicating that every predicted “License Plate” is 47.5% correct. The recall is approximately 0.826, showing that the model correctly identified about 82.6% of actual “License Plate” instances. The F1 score is approximately 0.603, giving a good balance between precision and recall.

Figure 10 and **12** are the plots of precision-recall and F1 score for YOLOv10 respectively.




Figure 10: Precision-Recall Curve for the second training of YOLOv10




Figure 11: F1-Confidence Curve for second training for YOLOv10

As shown in **Figure 11**, the confidence threshold reaches 0.637 when the F1 score attains its maximum value (0.76). This indicates better precision and recall up to that point. Beyond the 0.637 confidence threshold, the F1 score starts to decline, suggesting that the model becomes overly conservative, possibly leading to higher precision but lower recall, or vice versa. The summary of the training of YOLOv10 is given in **Table 4**.

Table 4: Summary of results for YOLOv10

Training Batch	Dataset Size	mAp@50	Accuracy	Precision	Recall	F1 score
Batch A	261	0.786	43.2%	47.5%	82.6%	0.603

The summary from the training (see **Table 5**) reveals compelling insights into the performance of various YOLO models for image objects. These results were obtained from the best training outcomes using a diverse dataset of over 200 images, including a wide range of challenging scenarios, including blurry images, skewed perspectives, long-distance captures, and images taken at optimal angles and distances. This variety in the dataset ensured a robust evaluation of the models’ capabilities across different real-world conditions.

Table 5: Summary of Model Training

Metrics	YOLOv3	YOLOv8	YOLOv10
Accuracy	62.1%	86.9%	43.2%
Precision	75%	100%	47.5%
Recall	78.3%	87%	82.6%
F1 Score	0.766	0.93	0.603
Average total training time	5.458hrs	About 30 minutes	1hr

In this evaluation, YOLOv3 (which is an earlier version of the YOLO family) demonstrated moderate performance with an accuracy of 62.1%. Its precision and recall scores of 75% and 78.3% respectively indicate a balanced performance in terms of correct detection and false positives. The F1 score of 0.766 further backs up its solid, albeit not exceptional performance. However, YOLOv3's most significant drawback during its training was its extensive training time, requiring an average of 5.458 hours to process the dataset.

YOLOv8 emerged as the standout performer in the evaluation. It achieved an impressive accuracy of 86.9%, significantly outperforming its counterparts. Notably, YOLOv8 achieved a perfect precision score of 100%, indicating that all its positive detections were correct. The recall of 87% suggests that it successfully identified a high proportion of the relevant objects. These metrics culminated in an excellent F1 score of 0.93, reflecting a harmonious balance between precision and recall. Furthermore, YOLOv8's efficiency was remarkable, completing the training process in approximately 30 minutes- a substantial improvement over YOLOv3.

YOLOv10, despite being the most recent model, showed mixed results. Its accuracy of 43.2% was lower than expected, failing behind both YOLOv3 and YOLOv8. The precision and recall scores of 70% and 60.9% respectively, while not poor, indicate room for improvement. The F1 score of 0.651 further reflects these moderate performance levels. However, the YOLOv10 did show improvement in training efficiency compared to YOLOv3, completing the process in approximately 1 hour. The stark contrast in performance between YOLOv8 and YOLOv10 is particularly intriguing. While one might expect the newer model to outperform its predecessor, our results suggest that YOLOv10 may still be in a developmental phase, requiring further refinement to reach its full potential. This discrepancy also highlights the importance of thorough testing and optimization in the deployment of new models. In terms of training efficiency, there is a clear progression from YOLOv3 to YOLOv8 and YOLOv10. This trend suggests that advancements in the YOLO architecture have significantly improved computational efficiency, a crucial factor for real-world applications where training time can be a limiting factor.

4. Conclusion

This study evaluated the performance of three YOLO (You Only Look Once) models – YOLOv3, YOLOv8, and YOLOv10 – for object detection tasks. Each model demonstrated distinct strengths and weaknesses in terms of accuracy, speed, and overall performance. The results indicated that the YOLOv8 currently offers the best combination of accuracy, precision, recall, and training efficiency among the tested models. However, the potential of YOLOv10, despite its current limitations, necessitates further investigation and optimization. These findings provide valuable insights for researchers and practitioners in the field of object detection, guiding model selection and highlighting areas for future development and refinement.

Future work could focus on fine-tuning YOLOv10 for better performance, by exploring potential improvements such as experimenting with different hyper parameters (including batch sizes, network architectures, and learning rates). Work on investigating how it performs on larger and more diverse datasets to assess its scalability and performance can be crucial. Introducing diverse datasets like varied lighting conditions, object sizes, and environmental contexts will improve the model's generalization capabilities. Potential modifications to the architecture of YOLOv10 can also enhance its speed without compromising accuracy. It is also recommended that the evaluation of other prominent object detection models alongside YOLO variants. This could provide valuable insights into their relative strengths and weaknesses and offer a more comprehensive view of current object detection capabilities.

References

- Amit, Y., Felzenszwalb, P. and Girshick, R. 2021. Object detection. In *Computer Vision: A Reference Guide*. 875-883, Cham: Springer International Publishing.
- Batra, P., Hussain, I., Ahad, MA., Casalino, G., Alam, MA., Khalique, A. and Hassan, SI. 2022. A novel memory and time-efficient ALPR system based on YOLOv5. *Sensors*, 22(14): 5283.
- Bakshi, A., Gulhane, S., Sawant, T., Sambhe, V. and Udmale, SS. 2023, January. ALPR-An Intelligent Approach Towards Detection and Recognition of License Plates in Uncontrolled Environments. In *International Conference on Distributed Computing and Intelligent Technology* (pp. 253-269). Cham: Springer Nature Switzerland.
- Ganesan, K. 2022. What are Precision & Recall in Machine Learning? Available at: <https://www.opinosis-analytics.com/ai-glossary/precision-and-recall-machine-learning> (Accessed: 30 October 2024).

Gnanaprakash, V., Kanthimathi, N. and Saranya, N. 202. Automatic number plate recognition using deep learning. In IOP Conference series: materials science and engineering, 1084(1): 012027).

Hussain, M. 2023. YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing and Industrial Defect Detection. Machines and Tooling, 11: 677. <https://doi.org/10.3390/machines11070677>.

Jamtsho, Y., Riyamongkol, P. and Waranusast, R., 2021. Real-time license plate detection for non-helmeted motorcyclist using YOLO. Ict Express, 7(1): 104-109.

Jenitta, J., Kusuma, S., Madhura, R. and Prakash, MR. 2023. License Plate Recognition Using Raspberry Pi. 2023 World Conference on Communication & Computing (WCONF), 1–5. <https://doi.org/10.1109/WCONF58270.2023.10235084>.

Jiang, P., Ergu, D., Liu, F., Cai, Y. and Ma, B., 2022. A Review of Yolo Algorithm Developments. Procedia Computer Science, 199: 1066–1073. <https://doi.org/10.1016/j.procs.2022.01.135>.

Karimi, F., 2021. Urban Expansion Modeling Using Machine Learning Algorithms. The University of North Carolina at Greensboro.

Laroca, R., Zanolensi, LA., Gonçalves, GR., Todt, E., Schwartz, WR. and Menotti, D., 2021. An efficient and layout-independent automatic license plate recognition system based on the YOLO detector. IET Intelligent Transport Systems, 15(4): 483-503.

Lee, YH. and Kim, Y. 2020. Comparison of CNN and YOLO for Object Detection. Journal of the semiconductor & display technology, 19(1): 85-92.

Lin, CJ., Chuang, CC. and Lin, HY. 2022. Edge-AI-based real-time automated license plate recognition system. Applied Sciences, 12(3): 1445.

Li, C., Li, L., Jiang, H., Weng, K., Geng, Y., Li, L., Ke, Z., Li, Q., Cheng, M., Nie, W. and Li, Y. 2022. YOLOv6: A single-stage object detection framework for industrial applications. arXiv preprint arXiv:2209.02976.

Lou, H., Duan, X., Guo, J., Liu, H., Gu, J., Bi, L. and Chen, H. 2023. DC-YOLOv8: small-size object detection algorithm based on camera sensor. Electronics, 12(10): 2323.

Musa, UF., 2023. Advantages Of Automated License Plate Recognition Technology. TMP Universal Journal of Research and Review Archives, 2(4).

O'Mahony, N., Campbell, S., Carvalho, A., Harapanahalli, S., Hernandez, G.V., Krpalkova, L., Riordan, D. and Walsh, J. 2020. Deep learning vs. traditional computer vision. In Advances in Computer Vision: Proceedings of the 2019 Computer Vision Conference (CVC), 11: 128-144. Springer International Publishing.

Powers, D.M., 2020. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. arXiv preprint arXiv:2010.16061.

Redmon, J., 2016. You only look once: Unified, Real-Time Object Detection. In Proceedings of the IEEE conference on computer vision and pattern recognition.

Saha, S., Basu, S., Nasipuri, M. and Basu, D.K., 2009. License Plate localization from vehicle images: An edge based multi-stage approach. International Journal of Recent Trends in Engineering, 1(1): 284.

Solunke, B. R., & Gengaje, S. R., 2023. A Review on Traditional and Deep Learning based Object Detection Methods. 2023 International Conference on Emerging Smart Computing and Informatics (ESCI), 1–7.

Sozzi, M., Cantalamessa, S., Cogato, A., Kayad, A. and Marinello, F. 2022. Automatic bunch detection in white grape varieties using YOLOv3, YOLOv4, and YOLOv5 deep learning algorithms. Agronomy, 12(2): 319.

Sung, JY., Yu, SB. and Korea, SP. (2020). Real-time Automatic License Plate Recognition System using YOLOv4. 2020 IEEE International Conference on Consumer Electronics - Asia (ICCE-Asia), 1–3.

Tamang, S., Sen, B., Pradhan, A., Sharma, K. and Singh, VK. 2023. Enhancing COVID-19 Safety: Exploring YOLOv8 Object Detection for Accurate Face Mask Classification. *International Journal of Intelligent Systems and Applications in Engineering*, 11(2): Article 2.

Wadare, M., Sharma, M., Patil, Mrs. S., Bewoor, Mrs. M., Pawar, A., Kadam, A., and Patil, S. 2023. Vehicle Number Plate Detection using YoloV8 and EasyOCR.

Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., Han, J. and Ding G. 2024. 'Yolov10: Real-time end-to-end object detection', *Advances in Neural Information Processing Systems*, 37: 107984–108011.

Zhao, K., Zhao, L., Zhao, Y., and Deng, H. 2023. Study on Lightweight Model of Maize Seedling Object Detection Based on YOLOv7. *Applied Sciences*, 13(13): Article 13.

Zhao, S. and You, F. 2020. Vehicle detection based on improved yolov3 algorithm', in *2020 International Conference on Intelligent Transportation, Big Data & Smart City (ICITBS)*. IEEE, 76–79.