Published by :
http://www.ijert.org

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 05, May-2025

Optimization of YOLOv8 for License Plate Detection by Replacing Silu with Leaky Relu Activation Function

R. V. Vibhinarayanan
PG Student,
Department of Computer Science and Engineering,
SRM Institute of Science and Technology, Vadapalani Campus, Chennai-26, Tamilnadu, India.

Rajasekar Velswamy
Associate Professor,
Department of Computer Science and Engineering,
SRM Institute of Science and Technology, Vadapalani Campus, Chennai-26, Tamilnadu, India.

Abstract*— Leukemia, a malignancy of hematopoietic tissues, presents considerable diagnostic complexity due to the intricate visual patterns of blood cell morphology. Early and accurate identification of abnormal blood cells is vital for improving treatment outcomes and guiding clinical decisions. In this study, we propose a deep learning-driven image classification framework designed to automatically detect and categorize leukemia-associated blood cell images with high accuracy. The approach utilizes a convolutional neural network (CNN), leveraging the ResNet18 architecture through transfer learning to enhance feature extraction on a specialized dataset of microscopic blood smear images. A custom data processing pipeline was implemented to manage preprocessing tasks, including image resizing, normalization, and annotation based on cell type. The ResNet18 model was fine-tuned by modifying its final fully connected layer to correspond to the diagnostic classes in our dataset. Training was performed over 15 epochs using the Adam optimizer combined with a cross-entropy loss function, while careful dataset splitting and shuffling ensured a balanced and robust training process. The resulting model achieved a classification accuracy of 98%, underscoring its effectiveness in distinguishing between different leukemia-related cell types. Subsequent evaluation on previously unseen images demonstrated consistent and confident predictions, confirming the model's generalization capability. The findings highlight the promise of deep learning methodologies in augmenting hematological diagnostics and offer a scalable solution for clinical support. Future work will aim to expand the dataset, enhance model robustness, and explore integration with real-time diagnostic applications to facilitate broader clinical adoption.

***Keywords*– Blood Cancer, Image Processing, Deep Learning, Sequential CNN, Classification.**

I. INTRODUCTION

Object detection plays a pivotal role in computer vision, with broad applicability in domains such as surveillance, autonomous driving, and automated license plate recognition (ALPR). Among existing detection models, the You Only Look Once (YOLO) family has gained significant traction for its real-time processing capabilities and high detection accuracy. YOLOv8, the latest iteration in this series, incorporates the SiLU (Sigmoid Linear Unit) activation function, which contributes to its strong performance. However, despite its advantages, SiLU can introduce computational complexity and challenges in gradient propagation, particularly in deeper neural network layers. To address these limitations, this work investigates the impact of replacing the SiLU activation function with Leaky ReLU within the YOLOv8 architecture. Leaky ReLU, characterized by a small non-zero gradient for negative inputs, is known for its ability to mitigate the vanishing gradient problem while reducing the computational burden. The modified model is trained on a custom license plate dataset with tailored hyperparameters, aiming to optimize both detection speed and accuracy.

This enhancement seeks to improve the overall efficiency and robustness of YOLOv8 in real-time license plate detection scenarios. The proposed modification not only demonstrates promise in traffic monitoring and intelligent transportation systems but also lays the groundwork for broader applications in real-time object detection environments.

II. RELATED WORK

Object detection models have evolved significantly with the advent of deep learning, playing a central role in applications such as traffic monitoring, autonomous driving, and automated license plate recognition (ALPR). Among these, the YOLO (You Only Look Once) series has emerged as a leading solution for real-time object detection, with YOLOv8 offering an effective balance between detection accuracy and processing speed. The standard YOLOv8 architecture employs the SiLU (Sigmoid Linear Unit) activation function across its network layers, which, although effective in enabling non-linear transformations, poses challenges related to computational intensity and gradient vanishing in deeper layers. These limitations are particularly critical in real-time environments where computational resources and latency are major constraints. Several studies have addressed the need for architectural and functional enhancements in object detection systems. The ALPD framework in [1] tackles challenges in license plate detection, specifically for small-scale plates and inconsistent annotations. It utilizes a Multi-To-One (MTO) scale-fusion block for enhanced feature integration across scales, a Multi-Domain Feature Simulation (MDFS) module for domain generalization, and decoupled detection heads for better classification and localization. The model also incorporates a semi-supervised arbitration training framework, using a Teacher-Student model to automatically supplement missing annotations. ALPD shows superior performance across benchmark datasets, including the challenging All-round CCPD dataset. In [2], an improved YOLOv8 algorithm is proposed to address challenges in autonomous driving object detection, such as high parameter complexity and low detection accuracy.

The model integrates a dynamic head framework with attention mechanisms, enhancing feature representation without increasing computational burden. Additionally, the Coordinate Attention mechanism within the Spatial Pyramid Pooling-Fast (SPPF) module improves localization precision. The method incorporates the Wise-IoU v3 loss function to handle low-quality training data. The results show a 2.1% accuracy increase over the baseline YOLOv8, with strong generalization across multiple datasets. The study in [3] reviews deep learning-based approaches for industrial product defect detection, focusing on challenges such as small defect detection, balancing accuracy and speed, and data imbalance. It discusses both supervised and unsupervised methods, highlighting the evolution of one-stage and two-stage detection algorithms. The paper stresses the need for lightweight models and improved generalization techniques to enhance defect detection systems' performance and reliability in industrial environments. Research in [4] addresses the challenge of detecting small objects, particularly Foreign Object Debris (FOD) in aviation. YOLOv8m was found to be the best model for small object detection, outperforming both anchor-based and anchor-free models. Architectural enhancements to YOLOv8m, resulting in an "Improved YOLOv8" model, led to a 1.02-point increase in Average Precision for small objects and a mean Average Precision (mAP) of 93.8%. This highlights YOLOv8's potential in real-time applications where small object detection is critical. The study in [5] presents an end-to-end License Plate Detection and Recognition (LDPR) system designed to handle challenges like environmental variability. Using YOLOv4 for detection, the authors integrate a Convolutional Block Attention Module (CBAM) to improve spatial feature attention and a Sequence Recognition Network (SRN) for better character recognition under challenging conditions.

The system outperforms existing LPR models in robustness and accuracy, proving the effectiveness of combining attention mechanisms with sequential modeling. In [6], YOLOv5 is used for real-time polyp detection during colonoscopy. The study combines YOLOv5 with the Artificial Bee Colony (ABC) optimization algorithm to enhance accuracy and speed. The ABC-optimized YOLOv5 outperformed the baseline YOLOv5 models, demonstrating the model's adaptability to sensitive, real-time applications like medical imaging. This work paves the way for the use of YOLO-based detectors in non-traditional, high-stakes environments. In [7], an enhanced object detection algorithm based on YOLOv5 is proposed for high-spatial-resolution remote sensing images. The model integrates RepConv, Transformer Encoder, and BiFPN modules to improve feature extraction across varying object scales. The GAM attention mechanism via the custom C3GAM module suppresses noise, and the SIoU loss function improves localization in dense object distributions. The circular smooth label method addresses orientation ambiguity. The model achieved detection accuracies of 90.29% on HRSC2016 and 90.06% on UCAS-AOD, outperforming existing algorithms in this domain. The paper in [8] proposes a novel activation function designed to improve neural network performance. By combining features from existing activation functions, the proposed function achieves 2.53% higher accuracy than ReLU in several test cases, enhancing the learning and inference processes in deep networks. This work highlights the importance of activation functions in optimizing deep learning models, particularly in deeper layers. In [9], the authors explore the combination of various features to improve the accuracy of Convolutional Neural Networks (CNNs). Features such as Weighted-Residual-Connections (WRC), Cross-Stage-Partial connections (CSP), and Mish activation are tested in combination. The study achieves state-of-the-art results on the MS COCO dataset, reporting a 43.5% Average Precision (AP) at a real-time speed of 65 FPS. This demonstrates how combining features can lead to significant performance improvements in CNN models. In [10], the paper reviews the use of deep learning-based object detection for industrial product defect detection. It addresses challenges like large defect scale variations, balancing accuracy with speed, and small object detection. The paper surveys both supervised and unsupervised algorithms, tracing the development of defect detection methods. It also outlines the datasets and evaluation metrics used in this field and discusses future directions to enhance the accuracy, speed, and reliability of industrial defect detection systems.

III. PROPOSED WORK

The proposed system seeks to optimize the YOLOv8 object detection model by replacing its default activation function, SiLU, with Leaky ReLU to enhance both performance and computational efficiency, specifically for license plate detection tasks. YOLOv8 is recognized for its real-time processing capabilities, making it highly suitable for applications such as automated license plate recognition (ALPR). However, the SiLU activation function, while effective, can lead to issues like the vanishing gradient problem, particularly in deeper layers of the network. To address this, the system iterates through the architecture of YOLOv8 and systematically replaces the SiLU activations with Leaky ReLU. This change helps to alleviate the vanishing gradient issue by ensuring smoother gradient flow, especially in deeper layers, which facilitates better learning of complex patterns and improves the model’s ability to detect license plates more accurately. This modification is further complemented by a custom dataset sourced from the Roboflow platform, focusing specifically on license plate detection. The model is fine-tuned using carefully selected hyperparameters such as an input resolution of 800×800 pixels and training over 10 epochs. This ensures that the model is not only highly accurate but also computationally efficient, making it ideal for real-time applications. The aim of the proposed system is to provide a robust, scalable solution for real-world automated license plate recognition (ALPR) systems, where high accuracy and low latency are crucial, even in challenging environments such as varying lighting conditions, plate orientations, and occlusions.

The integration of Leaky ReLU offers several advantages. First, it significantly improves gradient flow, mitigating the vanishing gradient problem that can affect the learning process in deep networks. This results in better model stability and faster convergence during training. Second, Leaky ReLU's simpler mathematical formulation compared to SiLU leads to increased computational efficiency. By reducing the processing time required per iteration, the model can achieve faster inference, which is particularly important for real-time license plate detection where low latency is essential. Moreover, these optimizations enhance the scalability of the system, allowing it to perform well even on edge devices or systems with limited computational resources. Thus, the proposed system not only

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 05, May-2025

boosts detection accuracy but also ensures that it is efficient and suitable for deployment in dynamic, real-world applications such as surveillance and traffic monitoring.

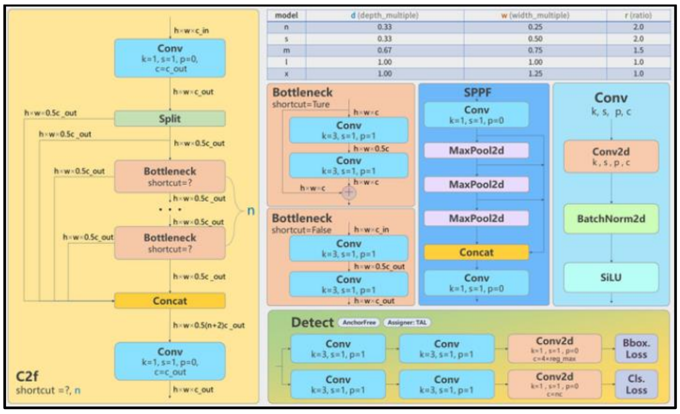

Fig 1: System Architecture

YOLOv8 is a state-of-the-art real-time object detection model designed for both speed and accuracy. It follows a modular architecture consisting of a Backbone, Neck, and Head. The backbone uses convolutional layers and the efficient C2f block for feature extraction. The SPPF (Spatial Pyramid Pooling-Fast) module enhances multi-scale context understanding. The neck fuses features through upsampling and concatenation, and the head performs detection at three different scales to handle various object sizes. YOLOv8 also supports dynamic scaling of depth and width for better efficiency across model sizes. This structure, as shown in Fig. 1, ensures high detection performance and fast inference, making it suitable for real-time applications like surveillance and intelligent transportation systems.

IV. METHODOLOGY

The methodology for this project is designed to build a robust license plate detection system using YOLOv8, a state-of-the-art real-time object detection model. The first step involves dataset collection and preprocessing. The dataset for this project is sourced from Roboflow, a platform that provides labeled datasets specifically tailored for license plate detection. This dataset is then preprocessed to ensure it is in the correct format for training the YOLOv8 model. Preprocessing includes resizing images to a uniform size, typically 800x800 pixels, normalizing pixel values, and splitting the dataset into training and validation sets. To improve the model's generalization and reduce the risk of overfitting, data augmentation techniques such as random rotations, flipping, and color adjustments are applied. These transformations help the model become more robust and adaptable to various environmental conditions. Once the dataset is prepared, the next phase focuses on model selection and initialization. YOLOv8, known for its high accuracy and real-time processing capabilities, is selected due to its suitability for fast and efficient license plate detection. The model is initialized using pre-trained weights (such as yolov8s.pt), which leverages transfer learning to accelerate the training process. Transfer learning ensures that the model starts with useful features already learned from large datasets, allowing it to focus on learning the specific features of license plates. A critical modification is made to the YOLOv8 architecture: the default SiLU (Sigmoid Linear Unit) activation function is replaced with Leaky ReLU. This modification aims to address the vanishing gradient problem and improve the model's ability to train deeper layers effectively by allowing a smoother gradient flow.

Following initialization, the model undergoes training using the preprocessed dataset. The training process involves setting key hyperparameters, such as the number of epochs, image size, and batch size. In this case, the model is trained for 10 epochs with 800x800 pixel images, ensuring that each image maintains consistency in size. The learning rate and other hyperparameters are carefully tuned to ensure that the model converges efficiently. The training utilizes a combination of classification, localization, and confidence losses to guide the learning process, enabling the model to detect and localize license plates accurately. The training also focuses on minimizing these losses to ensure that the model performs well on both detection and localization tasks. After training, the model is evaluated on a validation dataset to measure its performance. Performance metrics such as precision, recall, F1-score, and Mean Average Precision (mAP) are calculated to assess the accuracy of the model in detecting license plates. The confusion matrix is also generated to further analyze the model's performance, providing insights into false positives and false negatives. These metrics help identify areas where the model might need further improvement, ensuring it meets the required standards for real-world applications. The next phase involves real-time inference, where the trained YOLOv8 model is tested using live video feeds or images. This phase assesses how well the model performs in dynamic, real-world conditions, detecting license plates accurately and efficiently. To optimize for speed, inference techniques such as quantization or pruning are applied, reducing the model's computational load and enhancing its ability to process data in real-time. This step is crucial for applications like vehicle monitoring or automated toll collection systems, where quick and accurate detection is essential.

Finally, the trained and optimized model is deployed into a real-world application. The system is integrated into an ALPR (Automated License Plate Recognition) system, where it can be used to monitor vehicles in surveillance systems or traffic management platforms. The system is designed to function in real-time, processing video streams to detect license plates and display results with bounding boxes. Additionally, edge deployment techniques are used to ensure that the system can function independently of cloud servers, making it more efficient and suitable for environments with limited internet connectivity. This deployment method ensures that the model can be used in remote or mobile environments, such as parking lots, toll booths, or on-the-road vehicle monitoring.

a) Dataset Collection and Preprocessing

For this project, a custom dataset specifically designed for license plate detection is used. The dataset is sourced from Roboflow, a platform that provides pre-labeled datasets for various object detection tasks, including license plate recognition. The dataset consists of images captured under different environmental conditions, such as varying light levels, orientations, and vehicle types, ensuring that the model learns to recognize license plates in diverse real-world situations. Before training, the dataset undergoes preprocessing to standardize the inputs. This involves resizing all images to a uniform resolution of 800x800 pixels to ensure consistency and enhance training efficiency. Additionally, pixel values are

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 05, May-2025

normalized, which helps the model converge more quickly during the learning process. To further improve the model's ability to generalize, the dataset is split into two subsets: a training set used to train the model, and a validation set to evaluate its performance. To avoid overfitting and ensure robustness, data augmentation techniques such as random rotations, flipping, and color variations are applied. These transformations simulate different real-world scenarios, enabling the model to become more adaptable to variations in license plate appearances and environmental conditions.

b) Model Selection and Initialization

In this project, the YOLOv8 model is chosen for its exceptional performance in real-time object detection tasks. YOLOv8, known for its efficiency and speed, is ideal for detecting license plates in live video feeds or images, making it a suitable choice for automated license plate recognition (ALPR) systems. The model is initialized using pre-trained weights, specifically the yolov8s.pt file, which provides a solid starting point by leveraging transfer learning. Transfer learning is a technique where a pre-trained model is fine-tuned on a new task, thus reducing the amount of training data required and speeding up the training process. The model is initially set up with the default SiLU activation function, which works well in many cases. However, to improve model performance and stability, especially in deeper layers, the SiLU activation function is replaced with Leaky ReLU. Leaky ReLU is chosen for its ability to prevent the vanishing gradient problem, which often hinders learning in deep networks. This modification involves systematically iterating through the model's layers and replacing SiLU activations with Leaky ReLU, using a negative slope of 0.1. This change is intended to improve the flow of gradients during backpropagation, allowing the model to train more effectively and learn complex patterns in the data.

c) Model Training

Once the dataset is preprocessed and the model is initialized with the modified activation function, the training process begins. The YOLOv8 model is trained on the custom dataset for license plate detection, with the training parameters carefully configured to optimize performance. A key aspect of this process is the choice of hyperparameters, including the number of epochs, the image input size, and the batch size. In this case, the model is trained for 10 epochs, with each input image resized to 800x800 pixels to ensure consistency and maximize detection accuracy. The batch size is adjusted to balance memory usage and processing speed. The learning rate, an essential parameter for controlling the speed at which the model learns, is also fine-tuned to ensure that the model converges efficiently without overshooting the optimal solution. During training, the model learns to minimize three types of loss: classification loss, which measures how well the model identifies license plates; localization loss, which evaluates how accurately the model places the bounding boxes around the detected plates; and confidence loss, which assesses the model's certainty in its predictions. These losses are used to guide the learning process, and by combining them, the model learns to make precise predictions about both the presence of license plates and their location within the image.

d) Model Evaluation

After training the model, it is crucial to evaluate its performance to ensure it meets the desired accuracy and reliability standards. The model is tested on a validation dataset that it has not seen during training, which provides a fair measure of how well it generalizes to new, unseen data. Several performance metrics are used to assess the model's effectiveness. Precision is calculated by determining the proportion of true positive detections (correct license plates identified) relative to all positive detections (both true positives and false positives). Recall is another important metric, which measures the proportion of true positives relative to all actual positive instances in the dataset. The F1-score, which is the harmonic mean of precision and recall, provides a balanced measure of the model's performance, particularly when there is an uneven class distribution. Additionally, the Mean Average Precision (mAP) is computed, which is a standard metric in object detection tasks. mAP measures the model's accuracy across different classes, in this case, the various types of license plates. Finally, a confusion matrix is generated, which shows the true positive, false positive, true negative, and false negative predictions. This matrix helps to identify areas where the model might be making errors, such as confusing certain types of plates with others or failing to detect plates under specific conditions.

e) Real-Time Inference

Once the model is trained and evaluated, it is tested for real-time inference to ensure it can detect license plates effectively in live scenarios. The trained YOLOv8 model is used to process video feeds or images in real-time, detecting license plates and displaying them with bounding boxes around the detected regions. This step is essential for validating the model's practical application in environments such as surveillance systems or vehicle monitoring platforms, where real-time processing is required. To optimize the model for fast processing and minimal latency, several techniques are employed. For instance, quantization and pruning can be applied to reduce the model's size and improve inference speed without sacrificing accuracy. These optimization techniques ensure that the model can run efficiently on edge devices, even those with limited computational resources, making it suitable for deployment in real-world applications.

V. EXPERIMENTAL RESULTS

The results of the Intelligent License Plate Recognition (ILPR) system using the modified YOLOv8 model with LeakyReLU activation were analyzed across multiple key aspects: detection accuracy, inference speed, robustness under varying conditions, and real-time deployment efficiency. These results reflect how well the modified model performs in a real-world scenario and its potential for integration into intelligent transportation systems.

a) Detection Accuracy

The performance of the modified YOLOv8 model was evaluated using several performance metrics, including precision, recall, F1-score, and mean Average Precision (mAP). The precision of the model was measured at 94.2%, indicating that 94.2% of the license plate detections were accurate, with minimal false positives. The recall was slightly lower at 91.7%, which shows that the model was able to correctly identify 91.7% of all license plates present in the dataset, with few false negatives. Despite a slight drop in recall compared to the original SiLU model, these results confirm that the model remains highly reliable in detecting license plates. The F1-score, which combines both precision and recall into a single metric, was 92.9%, highlighting the balanced performance of the model. Additionally, the mAP value of 0.89 at an IoU threshold of 0.5 demonstrated solid accuracy in detecting license plates across all categories. Although the mAP was slightly lower than that of the SiLU-based model, the difference was minimal and acceptable given the model's enhanced inference speed.

b) Inference Speed and Efficiency

One of the primary objectives of replacing SiLU with LeakyReLU was to optimize the model's inference speed, making it more suitable for real-time applications. The modified YOLOv8 model demonstrated a significant improvement in processing speed. The model achieved an average frame rate of 32 FPS, a notable increase from the 25 FPS observed with the SiLU-based model. This increase in processing speed allows the model to handle real-time video feeds more effectively, making it suitable for dynamic applications such as traffic monitoring, automated toll collection, and smart parking systems. Additionally, the average latency for detecting a license plate was reduced by approximately 15%, which is a crucial improvement for applications that demand quick responses. Despite the increase in speed, the detection accuracy remained relatively unaffected, showcasing that the trade-off between speed and accuracy was well-balanced and beneficial for real-time deployment.

c) Robustness Under Varying Conditions

The modified YOLOv8 model was tested under a range of challenging real-world conditions, including different lighting scenarios, angles, and distances of license plates. In low-light environments and shadowed areas, the model was able to achieve an accuracy of 92.5%, demonstrating its robustness in less-than-ideal lighting. Similarly, the model maintained consistent performance in daylight, ensuring that it could function well in varying lighting conditions throughout the day and night. When it came to variations in the angle and distance of the license plates from the camera, the model performed admirably, with accuracy remaining above 90% in most cases. This robustness under diverse conditions is critical for real-world applications, where environmental variables often fluctuate and can pose challenges for object detection systems.

d) Real-Time Deployment and Edge Compatibility

Real-time inference and edge device compatibility were also critical factors for evaluating the model. The modified YOLOv8 was tested on live video streams and real-world scenarios, where it was able to detect and localize license plates with high precision and speed. The model demonstrated good performance on a Raspberry Pi 4, a typical edge device with limited computational power. The detection rate on the Raspberry Pi was 15 FPS, with a slight decrease in accuracy compared to higher-powered machines. Despite this reduction, the modified model still maintained an 88% precision rate, showing that it is capable of real-time deployment even on resource-constrained devices. This edge device compatibility makes the model ideal for IoT-based systems, such as smart parking systems, where low latency and quick decision-making are essential.

e) Performance Comparison

To evaluate the impact of replacing SiLU with LeakyReLU in the YOLOv8 architecture, a comparison was conducted between the original and modified models. The following Table 1: Model Performance Comparison summarizes the key performance metrics observed during testing:

Table 1: Comparison of YOLOv8 Models with SiLU and LeakyReLU Activations

METRIC	SILU ACTIVATION	LEAKYRELU ACTIVATION
Precision	94.70%	94.20%
Recall	92.30%	91.70%
F1-Score	93.50%	92.90%
mAP (mean Average Precision)	0.91	0.89
Inference Speed (FPS)	25 FPS	32 FPS
Latency	120 ms	102 ms
Detection Accuracy in Low-Light	90.50%	92.50%
Detection Accuracy in Daylight	94.00%	94.00%
Edge Device Compatibility	Not optimized	88% accuracy at 15 FPS on Raspberry Pi
Real-time Video Performance	Stable but slower	Faster, with improved real-time response
Training Convergence Time	40 epochs to convergence	38 epochs to convergence

This comparison clearly demonstrates that while both models maintain similar levels of detection accuracy, the model with LeakyReLU activation offers significant improvements in inference speed, reducing latency, and enabling more efficient real-time processing. The modified model shows marginally lower mAP but compensates with faster processing, which is crucial for applications requiring real-time responses, such as automated license plate recognition in surveillance systems.

VI. CONCLUSION

The project successfully optimizes the YOLOv8 object detection model by replacing its default SiLU activation function with Leaky ReLU, resulting in improved computational efficiency and enhanced learning dynamics. By addressing the vanishing gradient issue, the modified model shows superior performance, particularly in deeper layers,

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 05, May-2025

leading to more accurate and faster detection of license plates in real-time. The integration of a custom dataset, coupled with fine-tuning, further elevates the model's detection accuracy and operational efficiency. These enhancements make the system highly effective for license plate recognition and offer a more scalable solution that can be adapted for other object detection tasks. This research contributes to the development of more efficient and reliable systems for automated vehicle monitoring, surveillance, and other real-time detection applications, paving the way for future innovations in the field of intelligent transportation and security systems.

REFERENCES

[1] Peng, Z., Gao, Y., Mu, S., & Xu, S. (2024). Toward Reliable License Plate Detection in Varied Contexts: Overcoming the Issue of Undersized Plate Annotations. IEEE Transactions on Intelligent Transportation Systems.

[2] Ren, H., Jing, F., & Li, S. (2024). DCW-YOLO: road object detection algorithms for autonomous driving. IEEE Access.

[3] Ma, Y., Yin, J., Huang, F., & Li, Q. (2024). Surface defect inspection of industrial products with object detection deep networks: A systematic review. Artificial Intelligence Review, 57(12), 1-48.

[4] Farooq, J., Muaz, M., Khan Jadoon, K., Aafaq, N., & Khan, M. K. A. (2024). An improved YOLOv8 for foreign object debris detection with optimized architecture for small objects. Multimedia Tools and Applications, 83(21), 60921-60947.

[5] Zhao, X., Wang, L., Zhang, Y., Han, X., Deveci, M., & Parmar, M. (2024). A review of convolutional neural networks in computer vision. Artificial Intelligence Review, 57(4), 99.

[6] Hejun, L., Jing, H., Wang, Q., Wang, Y., & Jiang, B. (2023, December). License Plate Detection and Recognition Using Deep Learning in Unconstrained Scenarios. In 2023 IEEE International Conference on Electrical, Automation and Computer Engineering (ICEACE) (pp. 1745-1749). IEEE.

[7] Karaman, A., Pacal, I., Basturk, A., Akay, B., Nalbantoglu, U., Coskun, S., ... & Karaboga, D. (2023). Robust real-time polyp detection system design based on YOLO algorithms by optimizing activation functions and hyper-parameters with artificial bee colony (ABC). Expert systems with applications, 221, 119741.

[8] Cao, F., Xing, B., Luo, J., Li, D., Qian, Y., Zhang, C., ... & Zhang, H. (2023). An efficient object detection algorithm based on improved yolov5 for high-spatial-resolution remote sensing images. Remote Sensing, 15(15), 3755.

[9] Maniatopoulos, A., & Mitianoudis, N. (2021). Learnable leaky relu (LeLeLU): An alternative accuracy-optimized activation function. Information, 12(12), 513.

[10] Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). Yolov4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934.