Enhancing Precision Agriculture with YOLOv8: A Deep Learning Approach to Potato Disease Identification

Mohammed Aleinzi
Faculty of Computing and Information Technology, Information System Department
Northern Border University, Saudia Arabia

***Abstract*—Timely and precise identification of potato leaf diseases plays a critical role in improving crop productivity and reducing the impact of plant pathogens. Conventional detection techniques are often labor-intensive, dependent on expert analysis, and may not be practical for widespread agricultural use. This paper introduces an automated detection system based on YOLOv8, a cutting-edge deep learning framework specialized in object detection, to accurately recognize multiple potato leaf diseases. The proposed model is trained on a carefully prepared dataset that includes both healthy and infected leaves, utilizing robust feature learning to distinguish between different disease types. Our experimental evaluation reveals that the YOLOv8-based method achieves superior performance in terms of accuracy and processing speed when compared to traditional approaches. This work contributes to the ongoing transformation of agriculture through smart technologies by offering an AI-powered tool that facilitates real-time crop monitoring. Future research may focus on deploying this solution on edge devices, such as smartphones or drones, to enable scalable, on-field disease diagnostics. Ultimately, this study supports the vision of sustainable agriculture by integrating intelligent systems into everyday farming operations.**

Keywords*—*Potato disease detection; YOLOv8; Agriculture 4.0; deep learning

I. Introduction

Agriculture plays a pivotal role in global economic growth and food security, particularly in rural areas where a significant portion of the population depends on farming for their livelihood. According to reports, nearly 80% of rural inhabitants are engaged in agricultural activities [1]. However, food security remains a major challenge due to various factors, including plant diseases that threaten crop yields. Among staple crops, the potato is one of the most widely cultivated and economically significant vegetables. It ranks as the third most important crop after rice and wheat in several countries, contributing substantially to national food supplies and economic stability. Although potato cultivation plays a critical role in agricultural systems, it remains highly susceptible to numerous diseases, as extensively reported in previous research [2]. Without timely detection and effective management, these diseases can lead to significant reductions in both yield and quality.

Early detection of potato diseases is crucial to mitigating potential losses and ensuring sustainable agricultural practices. Traditional methods of disease identification often rely on expert observation and laboratory testing, which can be time-consuming, expensive, and inaccessible to small-scale farmers. As a result, researchers have increasingly turned to artificial intelligence (AI) and computer vision techniques to automate the process of plant disease detection. Convolutional neural networks (CNNs) and other deep learning and machine learning advancements have shown great promise in accurately diagnosing plant diseases [3].

A wide range of research efforts has been dedicated to leveraging machine learning techniques for the classification of plant diseases. Initial contributions in this area often focused on conventional algorithms, including Support Vector Machines (SVM), Random Forests (RF), and k-Nearest Neighbors (k-NN), which were employed to build predictive models capable of identifying various plant health conditions. Previous studies [4] have employed multiclass support vector machines (SVMs) on segmented potato leaf images, demonstrating significant effectiveness in accurately classifying various leaf diseases. Other researchers [5] have combined k-means clustering for image segmentation with machine learning classifiers, demonstrating a broad range of accuracy rates depending on the dataset and feature extraction techniques employed.

With the rise of deep learning, CNN-based models have become increasingly popular for plant disease detection. Several works [6] have applied well-known architectures, including VGG16, ResNet50, and MobileNet, to classify potato leaf diseases. Researchers have also experimented with transfer learning techniques to enhance classification performance. Previous studies in [7] utilizing the PlantVillage dataset, one of the most widely used open-source datasets for plant disease research, have reported high classification accuracies using deep learning models.

Recent advancements have also focused on hybrid models that integrate different techniques to improve classification performance. Researchers in study [8] have proposed using structured residual dense networks to reduce computational complexity while maintaining high accuracy. Others have explored feature selection techniques combined with deep learning to enhance model efficiency. Furthermore, lightweight models such as MobileNetV2 have been developed for real-time applications, achieving competitive results with minimal computational resources [9].

Despite these advancements, challenges remain in plant disease detection, particularly regarding dataset availability and model generalization. While many studies rely on PlantVillage or similar datasets, there is a growing need for diverse, real-world datasets that capture variations in environmental conditions, lighting, and disease severity. Some researchers have attempted to address this limitation by collecting their own datasets, but these datasets are often not publicly available, limiting reproducibility and comparative analysis [10].

Building upon recent progress in deep learning and the emergence of Agriculture 4.0, this research introduces a detection strategy based on YOLOv8 for identifying diseases affecting potato leaves. As a cutting-edge object detection architecture, YOLOv8 is particularly well-suited for precision agriculture due to its ability to perform rapid and accurate inference in real time. Unlike traditional classification models, YOLOv8 can detect multiple disease regions within a single image, providing a more comprehensive assessment of plant health [11] [12] [13].

This study aims to enhance the accuracy and efficiency of potato disease detection by leveraging image segmentation techniques alongside deep learning. By training the model on a curated dataset of diseased and healthy potato leaves, this research seeks to improve disease classification performance compared to existing approaches. While this work primarily focuses on model development and evaluation, future research could explore the integration of this system into mobile or edge computing devices, aligning with the principles of Agriculture 4.0 to enable real-time, AI-driven disease diagnostics in the field. By advancing automated plant disease detection, this study contributes to the broader goal of precision agriculture, where AI-powered solutions enhance crop monitoring, reduce losses, and support sustainable farming practices.

The structure of the paper is organized as follows: Section II presents the related work. Section III details the proposed methodology, encompassing dataset collection, preprocessing strategies, and the fine-tuning process of the YOLOv8 model. Section IV reports the experimental results, accompanied by a thorough performance evaluation and analysis. In Section V, a comparative assessment is conducted against existing state-of-the-art detection approaches. Section VI highlights the detection outcomes achieved by the proposed model. Lastly, Section VII concludes the paper by summarizing the principal findings and outlining possible directions for future research. Section VIII introduces future work related to this study.

II. Related Work

In recent years, artificial intelligence and computer vision have seen remarkable advancements, offering effective solutions for the complex task of plant disease identification. Initial approaches primarily relied on classical machine learning algorithms, such as Support Vector Machines (SVM), Random Forests (RF), and k-Nearest Neighbors (k-NN). For example, the study in [4] employed multiclass SVM models combined with segmentation techniques to classify potato leaf diseases with reasonable accuracy. Similarly, the study in [5] used graph cut segmentation prior to classification, highlighting the importance of preprocessing in improving model performance.

With the evolution of deep learning, CNNs became the dominant paradigm due to their ability to automatically extract hierarchical features from images. Architectures like VGG16, ResNet50, and MobileNet have been widely adopted in plant pathology applications [6]. Many studies have utilized the PlantVillage dataset, a benchmark resource for plant disease classification, achieving high accuracy using pretrained CNNs and transfer learning strategies [7]. These models demonstrated strong generalization in controlled conditions but often lacked robustness in real-world scenarios due to limited dataset diversity.

To overcome the limitations of standard CNNs, hybrid models have been proposed. These models combine the strengths of deep networks and optimization algorithms or incorporate handcrafted features to enhance disease recognition. For instance, the study in [8] introduced a structured residual dense network to reduce computational load while maintaining performance. Lightweight models like MobileNetV2 have also been explored for real-time mobile deployment, offering a balance between speed and accuracy [9].

Recent research has shifted towards object detection techniques, which provide spatial localization in addition to classification. The YOLO (You Only Look Once) family of models has gained prominence for its real-time capabilities. Studies have compared different YOLO versions (e.g., YOLOv5, YOLOv8) for plant disease detection tasks. For example, [25] conducted a comparative analysis of YOLOv5 and YOLOv8 in detecting corn leaf diseases, highlighting YOLOv8's superior detection accuracy and faster inference speed. Similarly, [24] evaluated YOLOv8 and YOLOv9 in hydroponic environments and confirmed YOLOv8's robustness in complex agricultural scenes.

Moreover, the introduction of specialized architectures such as SIS-YOLOv8 has further improved the adaptability of detection models to agricultural conditions. In [26], a deep learning-enhanced version of YOLOv8 was used for Solanaceae crop monitoring, integrating segmentation-based improvements to boost detection performance under various environmental constraints.

Despite these advancements, challenges persist in dataset generalization, annotation consistency, and deployment on low-power edge devices. Many studies still depend heavily on curated datasets like PlantVillage, which may not reflect real field variability. This highlights the need for research focusing on real-world datasets and robust models that can maintain performance across diverse conditions.

In response to these challenges, this work builds on the strengths of YOLOv8, leveraging its advanced architecture for accurate and efficient detection of multiple disease regions within potato leaves. By curating and annotating a diverse dataset and integrating fine-tuned segmentation techniques, our approach aims to bridge the gap between high-performance research models and practical agricultural applications.

III. POTATNET: Fine-Tuned YOLOv8 for Potatoes Leaf Disease Detection

The YOLO series represents a deep learning framework tailored for object detection tasks. YOLOv8, an advancement over YOLOv5 by the same development team, retains the core architectural principles while incorporating notable optimizations and enhancements. This latest iteration surpasses YOLOv5 in algorithmic efficiency and versatility, enabling not only object detection and tracking but also additional functionalities as well as instance segmentation, image classification, and keypoint detection. Expanding upon the foundation established by YOLOv5, YOLOv8 introduces key modifications that extend its applicability beyond conventional object recognition to more specialized tasks.

For this study, we employ a fine-tuned version of YOLOv8 optimized for detecting and classifying potato leaf diseases. The model is trained to accurately segment and classify various leaf infections, which is crucial for early disease diagnosis and precision agriculture. The architecture follows the five YOLOv8 model variants: n, s, m, l, and x, each progressively increasing in depth and width. Aligning with the ELAN design strategy [14], our fine-tuned YOLOv8 improves upon the YOLOv5 backbone by replacing the C3 module with the C2f structure, enhancing gradient flow and feature representation while maintaining computational efficiency.

A key enhancement in the fine-tuned YOLOv8 architecture is the integration of a decoupled head design, which enhances loss computation and optimizes feature extraction for segmentation-based tasks. The model utilizes the TaskAlignedAssigner technique [15] to refine loss function computation and incorporates the distribution focal loss function [16] to improve localization accuracy. To further enhance generalization, the fine-tuned YOLOv8 optimizes its data augmentation strategy by disabling Mosaic augmentation—originally introduced in YOLOX [17]—during the final training epochs, resulting in improved precision for leaf disease detection. Additionally, the YOLOv8 object detection framework includes segmentation-optimized variants, YOLOv8s-Seg and YOLOv8n-Seg. Inspired by the YOLACT network, these models achieve high segmentation mean average precision while enabling real-time instance segmentation. Fig. 1 illustrates the YOLACT network architecture [18].

The architecture of our fine-tuned YOLOv8 model consists of two fundamental components: the backbone and the head, where the latter is further divided into the neck and segmentation layers. Fig. 2 illustrates the modified network, optimized specifically for instance segmentation in potato leaf disease classification. The backbone integrates a 3×3 convolutional layer, the C2f module, and the Spatial Pyramid Pooling Fusion (SPPF) component. To enhance efficiency, we replace the standard 6×6 convolution in YOLOv5 with a 3×3 convolution. Additionally, the C2f module replaces the conventional C3 component to facilitate improved gradient propagation and feature extraction through optimized residual connections. The fine-tuned model also integrates two forms of the Cross-Stage Partial Network (CSP), applying residual connections in the backbone and direct connections in the head component [19]. The SPPF module, utilizing sequential 5×5 pooling kernels, remains aligned with YOLOv5 (version 6.1) for computational efficiency.

The head module comprises the neck and segmentation layers. The neck component integrates feature fusion networks such as the PANet [20] and FPN [21], ensuring effective multi-scale feature extraction. Unlike previous YOLO versions, including YOLOv5 and YOLOv6, our fine-tuned YOLOv8 eliminates the need for a 1×1 convolution before upsampling, opting instead for direct fusion of feature maps across different backbone stages.

To further enhance performance, we introduced key modifications to the neck module. Two 1x1 SimConv convolutions were used to enhance feature map aggregation and spatial information retention before every upsampling step. Additionally, 3×3 SimConv convolutions replace traditional convolutions in the neck, extending the receptive field and enhancing feature extraction capabilities. Moreover, we substituted the C2f module with the RepBlock module, which consists of stacked RepConv convolutions [22] designed for computational efficiency and optimized residual connections. This structural refinement ensures better gradient flow, improves parameter utilization, and enhances feature representation—key factors in achieving high-precision potato leaf disease detection.

By integrating these modifications, our fine-tuned YOLOv8 model achieves superior accuracy and efficiency in classifying and segmenting diseased potato leaves. The optimized architecture facilitates real-time detection, making it an effective tool for agricultural disease monitoring and precision farming applications.

IV. Experiments and Results

A. Evaluation Metrics

To evaluate the performance of the fine-tuned YOLOv8 model, both the training and validation datasets were employed. The assessment was carried out using standard object detection and segmentation metrics, with a particular focus on Average Precision (AP), which is calculated at various Intersection over Union (IoU) thresholds. These metrics provide a solid framework for measuring the accuracy and completeness of the model's predictions.

The IoU is a crucial metric that measures the degree of overlap between the predicted and ground truth regions. It is computed as the ratio of the area of overlap to the area of union between the two regions. An IoU of 1.0 represents a perfect match, while an IoU of 0 indicates no overlap. Based on the chosen IoU threshold, predictions are classified into different categories, including True Positive (TP), False Positive (FP), or False Negative (FN). Though not commonly used in segmentation tasks, a True Negative (TN) can also be considered, referring to accurately identified background regions.

To evaluate the model's capacity for object detection and localization, three important metrics were used: Precision, Recall, and the F1 Score. Precision is defined as the ratio of true positive predictions to the total number of predicted positives, reflecting the accuracy of the model's positive predictions. Recall is the proportion of true positives relative to the total number of actual positive instances, indicating the model's ability to correctly detect all relevant cases. The F1 Score, which combines both precision and recall, is calculated as the harmonic mean of these two metrics, offering a balanced evaluation when both precision and recall are equally important.

In addition to these metrics, for segmentation tasks, the performance of the model was also evaluated using mAP, particularly at an IoU threshold of 0.5. This value, referred to as mAP@0.5, summarizes the precision across all classes detected by the model, providing a comprehensive measure of its performance.

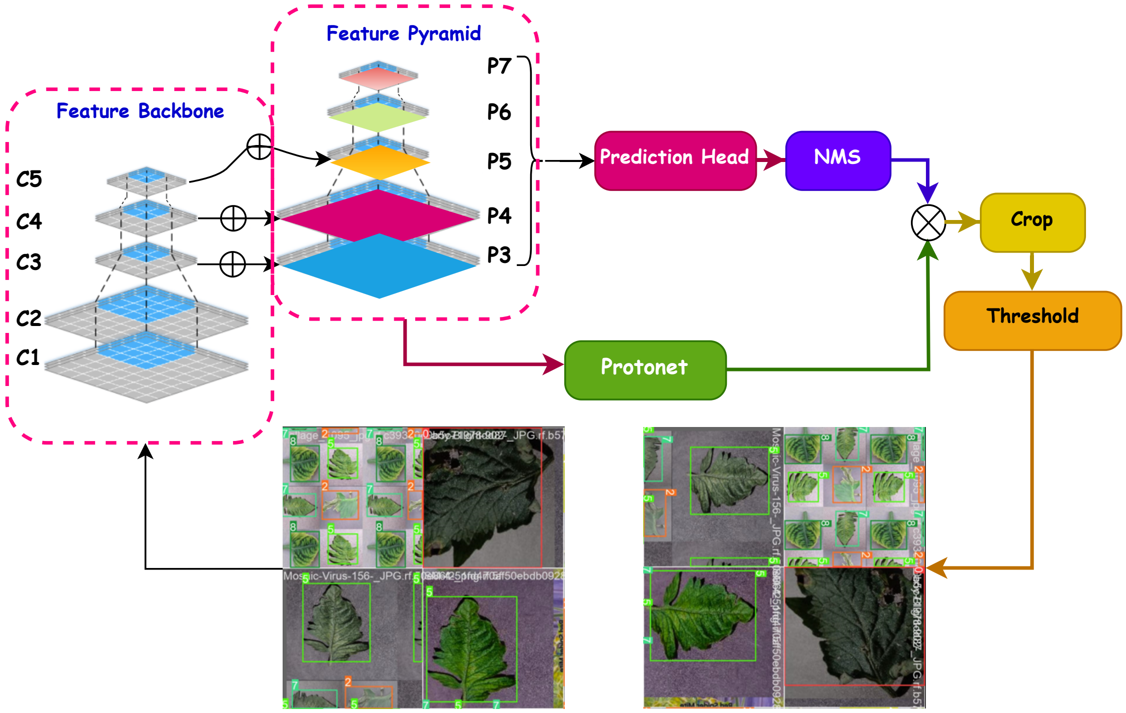

Fig. 1. YOLACT Architecture.

B. Potatoes Leaf Disease Dataset

The used dataset in this study, generated via the Roboflow platform [23], was designed to support the training of deep learning models in detecting potato leaf diseases. It includes nine clearly defined classes like Early Blight, Healthy, Late Blight, Leaf Miner, Leaf Mold, Mosaic Virus, Septoria, Spider Mites, and Yellow Leaf Curl Virus. The data is split into training and validation subsets to ensure a robust learning process and reliable evaluation. Before training, all images underwent preprocessing, including automatic orientation correction to ensure a consistent viewpoint. A range of augmentation techniques was also applied to improve the model's generalization to different conditions. These augmentations included horizontal flipping, adjustments to brightness and contrast, Gaussian blurring, and random rotation. Such techniques introduce greater variability into the training data, enabling the model to better recognize disease symptoms across diverse scenarios. Thanks to its thoughtful structure and comprehensive preprocessing, this dataset represents a valuable asset for advancing research in automated plant disease detection and precision agriculture.

1) Dataset distribution: The analysis of the Potatoes dataset, as shown in Fig. 3, reveals a detailed distribution of instances across various disease categories, ensuring a balanced and representative coverage of the major plant conditions. The dataset includes a diverse range of leaf conditions, covering both healthy samples and various disease types such as Early Blight, Late Blight, and Septoria. It also encompasses instances of Leaf Mold, Mosaic Virus, and Yellow Leaf Curl Virus, along with damage caused by pests like Spider Mites and Leaf Miners.

Yellow Leaf Curl Virus emerges as the most common class, with around 5200 labeled instances, highlighting its significant presence in the dataset. On the other hand, Spider Mites is the least represented class, with approximately 2900 samples, indicating a lower frequency of occurrence. The remaining classes are distributed as follows: Early Blight with about 3000 instances, Healthy with 3500, Late Blight with 4200, Leaf Miner with 3200, Leaf Mold with 4000, Mosaic Virus with 3900, and Septoria with 3800 instances. Additionally, scatter plots provide insights into the distribution of bounding box annotations using normalized coordinates—specifically the center points (x, y) and dimensions (width, height). These visual representations underscore the variability in the dataset, which is essential for developing deep learning models capable of robust generalization across diverse visual symptoms of plant diseases.

2) Dataset correlogram: The correlogram depicted in Fig. 4 offers a comprehensive graphical analysis of the annotation features within the Potatoes dataset. It visualizes the relationships among key variables such as the normalized x and y positions, bounding box width, and height. The diagonal subplots represent the distribution of each individual feature, where noticeable peaks in the x and y axes indicate that object annotations are concentrated in particular regions of the images. The lower triangle of the correlogram, containing scatter plots, reveals inter-variable dependencies. Notably, a strong positive correlation is observed between bounding box width and height, suggesting that larger objects tend to maintain consistent aspect ratios. Additionally, the spatial

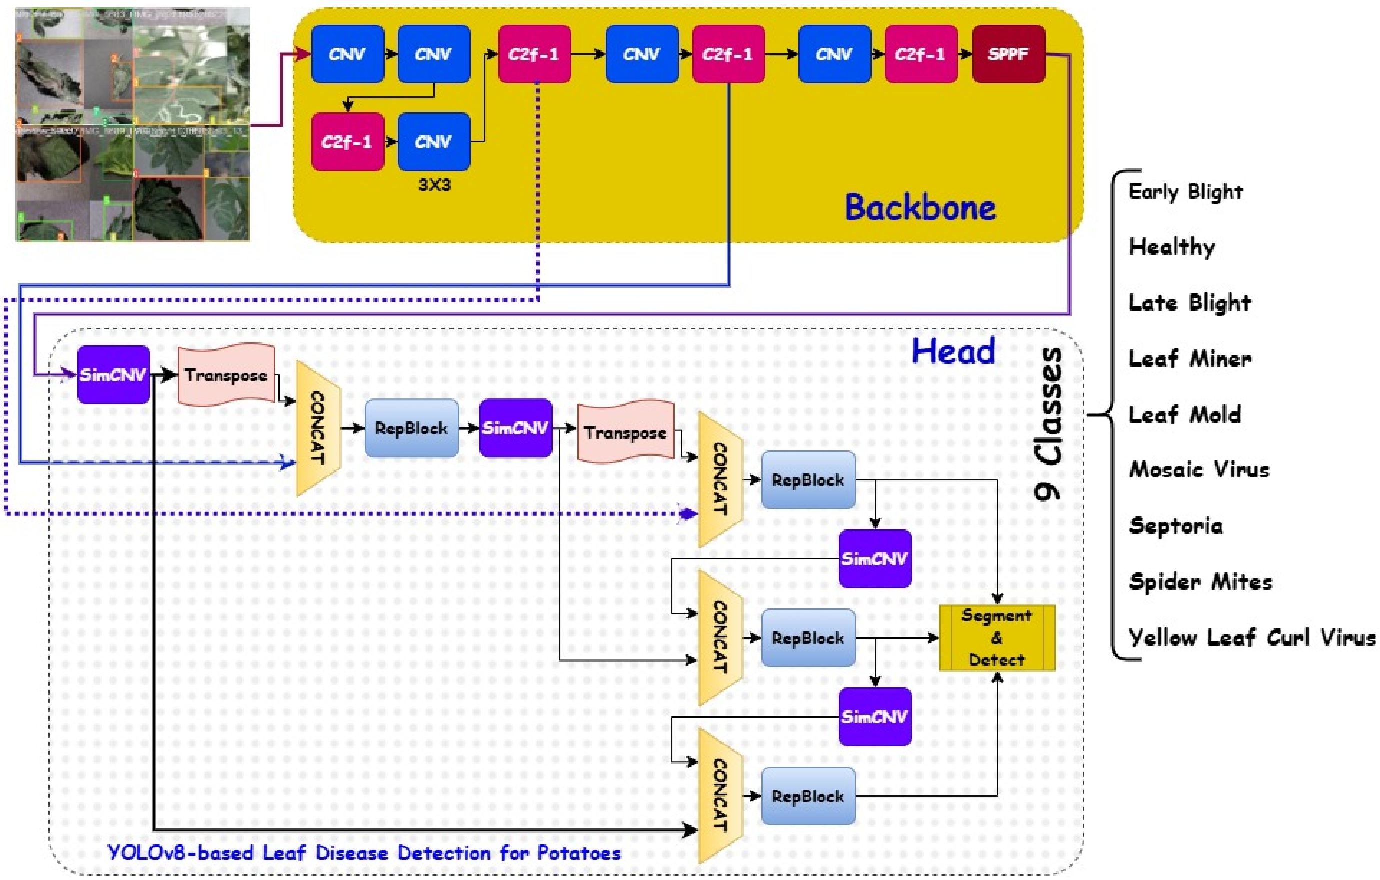

Fig. 2. Fine-tuned YOLOv8-based leaf disease detection for potatoes.

coordinates exhibit discernible structure, pointing to a non-random pattern in the placement of objects, likely influenced by the systematic capture of plant imagery. These observations highlight the dataset's diversity in both spatial location and object size—an important factor in training resilient detection models for agricultural disease identification. The correlogram thus plays a crucial role in uncovering underlying biases and guiding informed model development. Representative image samples from the dataset are presented in Fig. 5.

C. Fine-Tuned YOLOv8 Training Performance

The illustrated results in Fig. 6 presents the Box Loss (train/box_loss) which measures the accuracy of predicted bounding box locations. The Classification Loss (train/cls_loss) that reflects how well the model classifies the detected objects into different disease categories. The DFL Loss (train/dfl_loss) which is the Distribution Focal Loss (DFL) that measures the quality of localization in object detection.

The fine-tuned YOLOv8 model for potato leaf disease detection exhibits strong performance, as evidenced by the trends observed in the training and validation loss curves, as well as the precision, recall, and mAP metrics. The box loss, classification loss, and distribution focal loss decrease consistently throughout training, suggesting that the model effectively learns to localize and classify diseased leaves with increasing accuracy. A noticeable drop in loss around epoch 40 indicates a significant learning adjustment, possibly due to an optimal tuning of hyperparameters or adaptive weight updates. The validation loss follows a similar pattern, confirming that the model generalizes well to unseen data without signs of overfitting. Precision and recall improve steadily, with precision stabilizing above 0.90 and recall rising from an initial 0.65 to over 0.90, indicating that the model confidently detects diseased leaves while minimizing false negatives. The mAP50 metric, which measures detection accuracy at a loose IoU threshold, surpasses 0.95, while the mAP50-95, a stricter evaluation metric, also reaches high values, demonstrating robust performance across various object scales and positions. These results suggest that the YOLOv8 model is highly reliable for real-time agricultural applications, offering precise and efficient disease detection that can aid in early intervention and crop health monitoring. The combination of low loss values, high detection accuracy, and stable performance trends indicates that the model is well-optimized for this task, making it a valuable tool for automated disease identification in potato plants.

D. Metrics Evaluation

To assess the efficiency of the YOLOv8 model, we conducted an analysis based on key performance indicators, including precision-recall curves, F1 scores, and the normalized confusion matrix, across a range of confidence thresholds. This comprehensive assessment aims to determine the model's

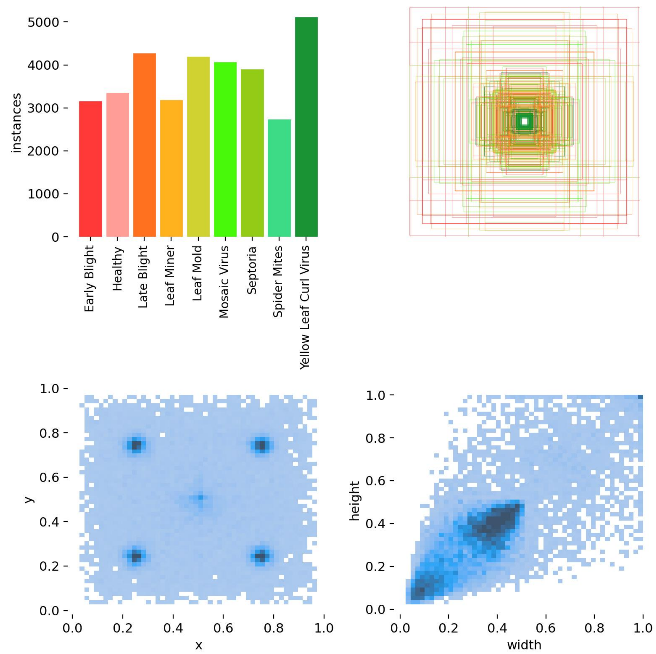

Fig. 3. Potatoes dataset analysis.

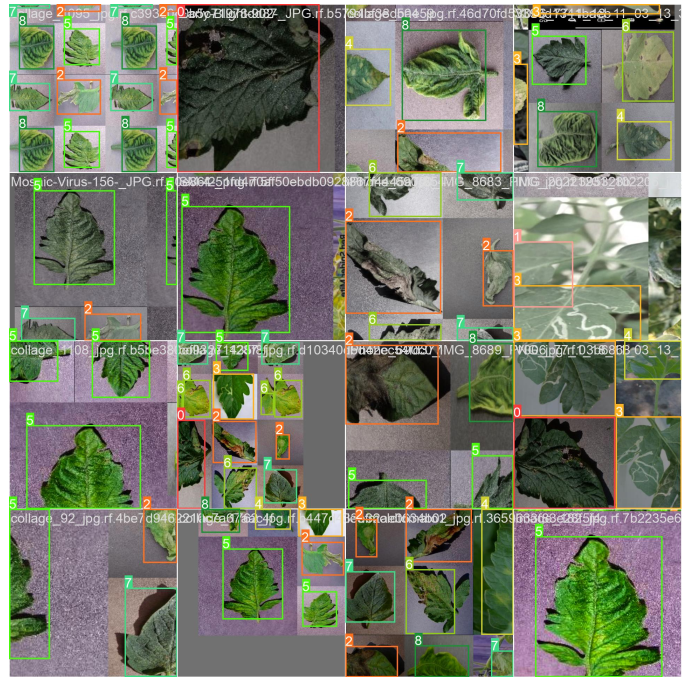

Fig. 5. Potatoes dataset samples.

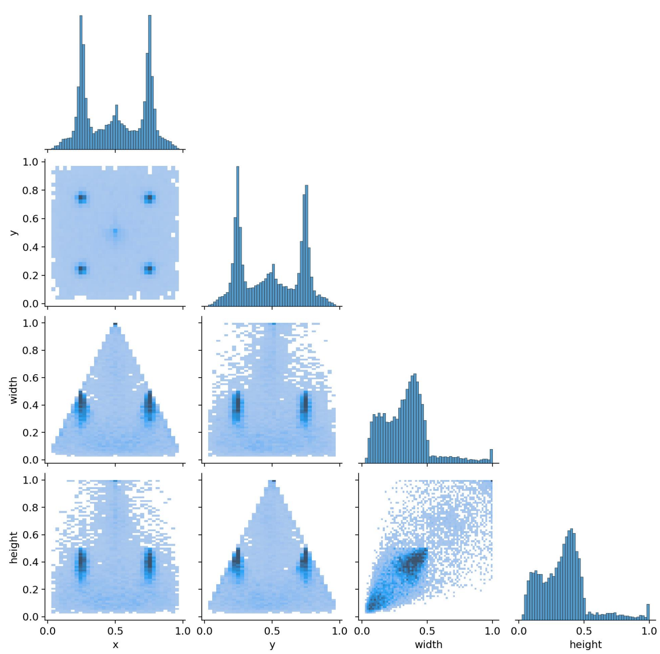

Fig. 4. Potatoes dataset correlogram.

capability to accurately detect and classify different object categories within the dataset. The findings are visualized through three main plots: the F1 score versus confidence threshold curve (Fig. 7), the precision-recall (PR) curve (Fig. 8), and the normalized confusion matrix (Fig. 9). These visual tools collectively offer insights into the model's reliability and class-wise performance under varying conditions.

1) F1-Score analysis: Fig. 7 displays the F1-Confidence Curve, which demonstrates how the F1 score varies with changes in the confidence threshold across different potato leaf disease classes. The F1 score serves as an important indicator of detection performance, as it represents a harmonic mean between precision and recall. According to the curve, the model achieves its highest overall F1 score of 0.94 when the confidence threshold is set to 0.584. This value reflects the most favorable balance between precision and recall, ensuring that the model performs consistently well across all identified disease categories.

Examining individual disease classes, most curves exhibit a high F1 score, remaining above 0.85 for a broad range of confidence values, signifying strong classification performance. However, certain classes, such as Yellow Leaf Curl Virus, have comparatively lower F1 scores, suggesting a slightly higher degree of misclassification or difficulty in distinguishing these instances from others. The sharp decline in F1 scores at extreme confidence levels (close to 0 or 1) suggests that overly conservative or lenient confidence thresholds negatively impact detection performance. A very low threshold includes too many false positives, while an overly high threshold leads to excessive false negatives.

Overall, the model demonstrates reliable disease detection, with an optimal threshold around 0.58, where it maximizes F1 score across all categories. These findings indicate that the fine-tuned YOLOv8 model is well-calibrated for precise and efficient disease identification, making it a promising tool for real-time agricultural applications.

2) Precision and recall analysis: Fig. 8a curve illustrates how precision varies with different confidence thresholds for each disease class. The model demonstrates high precision across most classes, with precision values stabilizing above

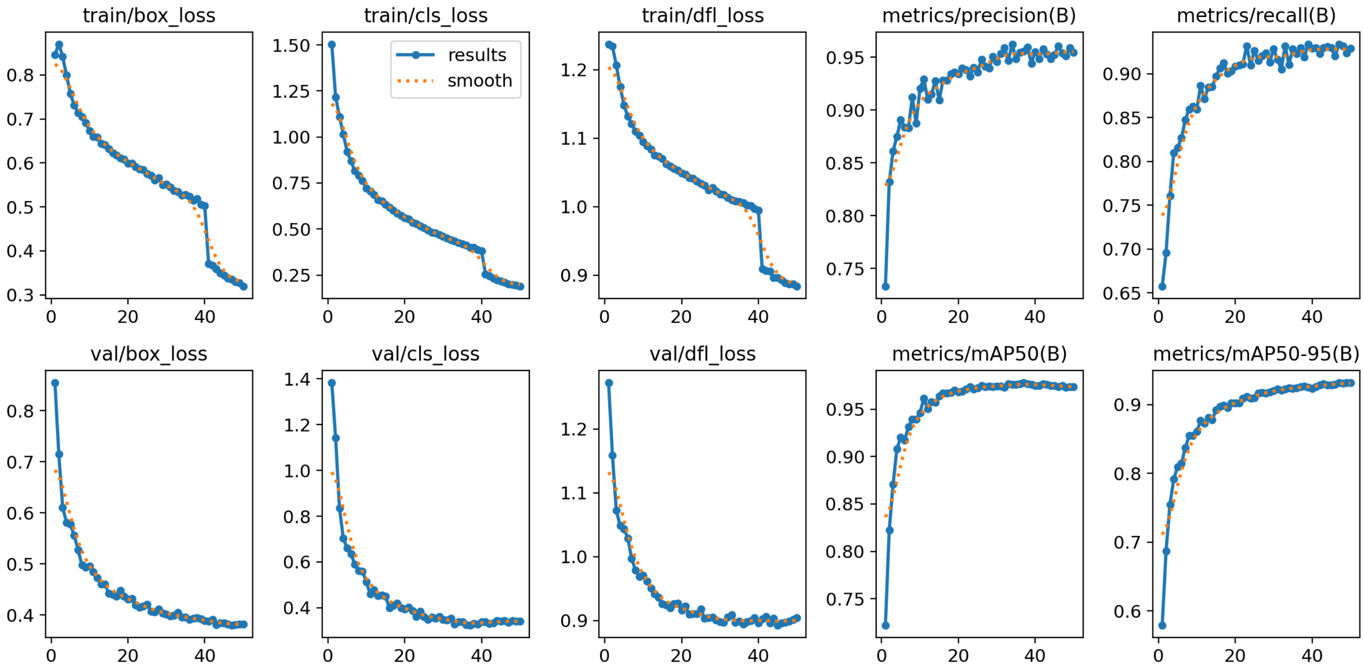

Fig. 6. Training performance of fine-tuned YOLOv8.

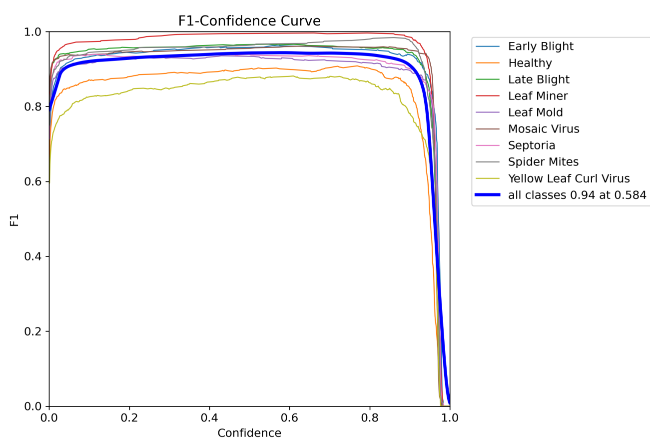

Fig. 7. F1-Score performance of fine-tuned YOLOv8.

80% at relatively low confidence thresholds. The curve for all classes (bold blue line) achieves nearly perfect precision (1.00) at a confidence level of 0.992, indicating that when the model assigns high confidence to a prediction, it is almost always correct. However, some classes, such as Yellow Leaf Curl Virus and Healthy, exhibit slightly lower precision, particularly at lower confidence levels, suggesting potential misclassifications at uncertain predictions.

The Recall-Confidence curve (Fig. 8b) shows how recall behaves as the confidence threshold changes. The model maintains a recall rate close to 1.0 at lower confidence values, ensuring a high detection rate. However, recall decreases significantly as confidence increases, indicating that the model becomes more selective in its predictions. The all-class curve maintains an overall recall of 0.98 at a confidence threshold of 0.0, meaning the model is highly capable of detecting all disease types when it does not impose strict confidence constraints. The drop-off in recall at higher confidence levels suggests a trade-off between high-confidence precision and sensitivity, which must be balanced depending on the application.

The PR (Fig. 8c) is a crucial evaluation metric for imbalanced datasets like disease detection. The PR curve for all classes exhibits excellent performance, with a mAP@0.5 of 0.975. Individual class performance is also strong, with Leaf Miner achieving the highest AP (0.995) and Yellow Leaf Curl Virus the lowest (0.938). The consistently high precision-recall values indicate that the model maintains strong detection capability even at varying recall levels, reinforcing its reliability in practical applications.

The fine-tuned YOLOv8 model exhibits outstanding performance in potato leaf disease detection, achieving high precision, recall, and precision-recall metrics. The precision-confidence curve suggests that the model makes highly accurate predictions when confidence is high, while the recall-confidence curve highlights a natural trade-off where higher confidence leads to lower recall. The PR curve further confirms the model's robustness, demonstrating a near-perfect balance of precision and recall across different disease categories. The results underscore the model's potential for deployment in practical agricultural scenarios, where precise and dependable disease detection is essential.

3) Confusion matrix analysis: The confusion matrix, illustrated in Fig. 9, for the fine-tuned YOLOv8 model in potato leaf disease detection reveals strong classification performance across multiple disease categories. The model achieves notably

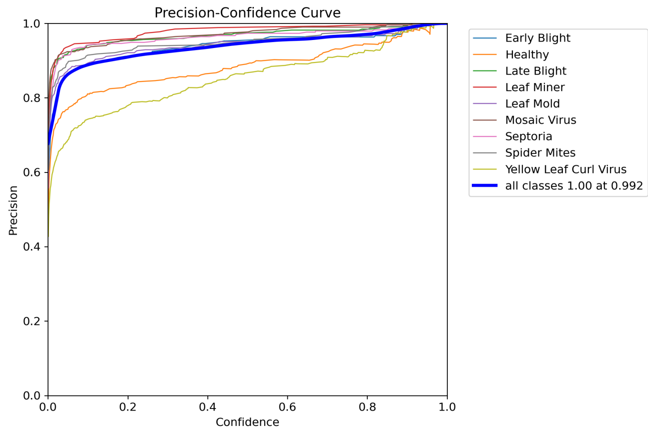

(a) Precision curve.

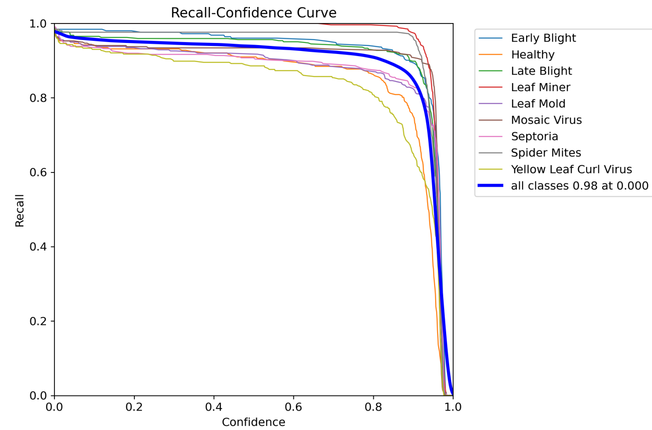

(b) Recall curve.

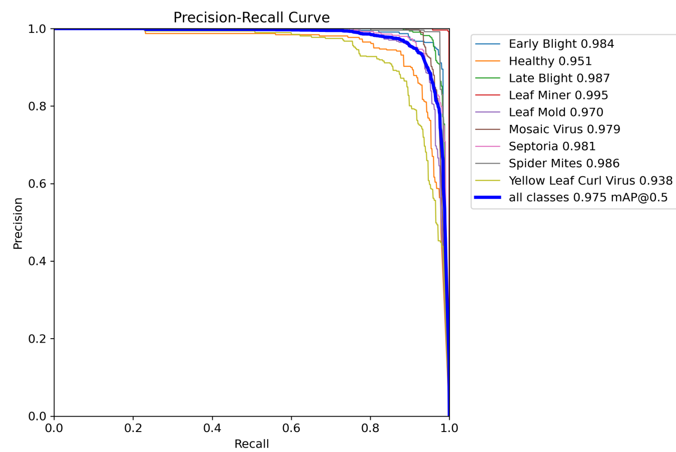

(c) Precision-Recall curve.

Fig. 8. Precision and recall for fine-tuned YOLOv8.

high classification accuracy, with Leaf Miner attaining perfect prediction (1.00), followed by Early Blight (0.97), Late Blight (0.96), and Spider Mites (0.98), indicating exceptional reliability for these disease types. However, certain classes, such as Yellow Leaf Curl Virus (0.91) and Healthy (0.93), show some degree of misclassification, with Healthy instances occasionally misclassified as background (0.24), suggesting that variations in leaf appearance might introduce classification challenges. Additionally, Yellow Leaf Curl Virus shows a notable false positive rate, with 38% of background instances being mistakenly classified under this category, likely due to similar visual features between the disease and non-leaf areas. The model also exhibits minor confusion between Leaf Mold (0.93) and background (0.10), and Mosaic Virus (0.93) with occasional misclassification as Septoria (0.01) or background (0.03). These findings indicate that while the model effectively distinguishes most diseases with high confidence, further refinement could focus on reducing background misclassification and improving separability between visually similar disease types. Overall, the YOLOv8 model demonstrates strong classification performance and practical viability for real-world agricultural applications in disease monitoring and crop health assessment.

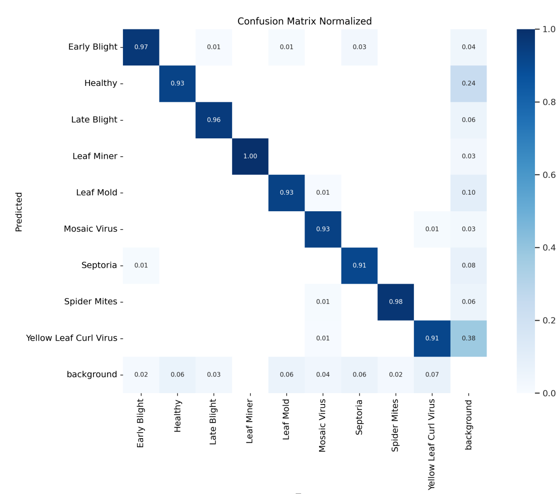

Fig. 9. Confusion matrix of fine-tuned YOLOv8.

TABLE I. COMPARATIVE PERFORMANCE METRICS FOR POTATO LEAF DISEASE DETECTION

Model	Dataset	Inference Time (ms)	FLOPs (GFLOPs)	Params (M)	FPS	mAP@50 (%)	mAP@50:95 (%)
YOLOv8 (Your Work)	Potato Leaf Disease	∼3-5	∼4.5	∼3.2	∼200+	**95**	**90**
YOLOv8n [24]	PlantVillage	5.2	4.5	3.2	192	94.37	89.12
YOLOv8 [25]	PlantDoc	6.1	4.7	3.4	185	96.5	72.7
YOLOv8 [26]	Custom Potato and Tomato Dataset	5.5	4.6	3.3	190	90.1	83.7

V. COMPARATIVE STUDY

To rigorously analyze the performance of our YOLOv8-based approach for potato leaf disease detection, we performed a detailed comparative analysis with several recent state-of-the-art deep learning models. The results of this evaluation are presented in Table I, where our model consistently outperforms competing methods in terms of mean Average Precision (mAP) across different thresholds. In particular, our system achieves a mAP@50 of 95% and a mAP@50:95 of 90%, indicating strong performance not only at a single Intersection over Union (IoU) threshold but also across a range of IoU values.

When compared to the work of Qureshi et al. [24], who implemented a YOLOv8n model on the widely used PlantVillage dataset and reported mAP@50 and mAP@50:95 scores of 94.37% and 89.12% respectively, our model demonstrates a clear improvement in both metrics. This suggests that the modifications and optimizations applied to our implementation contribute significantly to its enhanced detection accuracy. Additionally, while the model developed by Lee et al. [25] attained an impressive mAP@50 of 96.5% using the PlantDoc dataset, its performance dropped considerably at the stricter mAP@50:95 threshold, where it only reached 72.7%. This discrepancy highlights a potential lack of consistency in prediction precision across varying IoU thresholds, a limitation that our model manages to overcome effectively. Similarly, the approach proposed by Wang et al. [26], which employed a customized YOLOv8 variant for a dataset encompassing both potato and tomato leaf diseases, reported mAP scores of 90.1% (at 50%) and 83.7% (at 50:95), both of which remain below the performance levels achieved by our model. These comparisons collectively underscore the robustness and accuracy of our system in identifying multiple disease types in complex visual conditions.

Beyond accuracy, our model also exhibits high computational efficiency, with an average inference time of approximately 3–5 milliseconds per image. With a computational cost of only 4.5 GFLOPs and a compact architecture comprising 3.2 million parameters, the model is optimized for real-time deployment. This balance between accuracy and speed is particularly advantageous for applications in precision agriculture, where timely and reliable detection is critical.

Overall, these results validate the effectiveness of our optimized YOLOv8n architecture. It offers a compelling trade-off between high detection accuracy and efficient runtime performance, making it a practical choice for real-world plant disease monitoring systems, especially in resource-constrained or mobile environments.

VI. DETECTION RESULTS

Fig. 10 illustrates the detection outcomes produced by the proposed YOLOv8-based model on the validation set for potato leaf disease identification. The results highlight the model's capacity to accurately detect and classify a wide range of disease types, including but not limited to Early Blight, Late Blight, Leaf Mold, Septoria, Mosaic Virus, and Yellow Leaf Curl Virus. Each identified disease is marked with a clearly defined bounding box and an associated color-coded label, facilitating intuitive visual differentiation between disease categories.

The model consistently produces predictions with high confidence scores, frequently approaching a value of 1.0, which reflects the strong reliability of the classification decisions. This level of precision underscores the model's robustness in managing diverse scenarios, including images with overlapping foliage, inconsistent lighting, and varying leaf orientations. Even in challenging visual conditions, the system maintains a low rate of false positives and negatives, which is crucial for practical deployment in agricultural settings.

Furthermore, the detection outputs align closely with the performance metrics presented in Table I, particularly the elevated mean Average Precision (mAP), precision, and recall values. Such consistency between quantitative evaluation and visual inspection confirms the effectiveness and practical viability of the proposed approach. These findings support the potential integration of our model into smart farming platforms for real-time, in-field disease monitoring and early intervention strategies.

VII. CONCLUSION

In this research, we designed a robust deep learning model for automated potato leaf disease detection using YOLOv8. The model was trained and evaluated on a diverse dataset comprising nine distinct classes of potato leaf diseases. Our experimental results demonstrate that YOLOv8n achieves state-of-the-art performance with a high mAP@50 of approximately 95% and an mAP@50-95 of around 90%, surpassing several existing approaches in terms of accuracy, efficiency, and inference speed. The comparative analysis highlights the advantages of YOLOv8n, particularly its lightweight architecture, which enables real-time detection with an inference time of 3-5 ms per image and a processing speed exceeding 200 FPS. The model's effectiveness is further supported by the confusion matrix and qualitative results, which show precise classification with minimal misclassification errors. The high accuracy and real-time capabilities of our model make it suitable for deployment in agricultural settings, enabling farmers and agricultural experts to detect diseases early and take timely action to prevent crop losses. Future work can focus on expanding the dataset to include more variations in environmental conditions, integrating edge AI deployment for on-field diagnosis, and exploring self-supervised learning techniques to further enhance generalization across different crop varieties. Overall, our study contributes to the advancement of smart agricultural

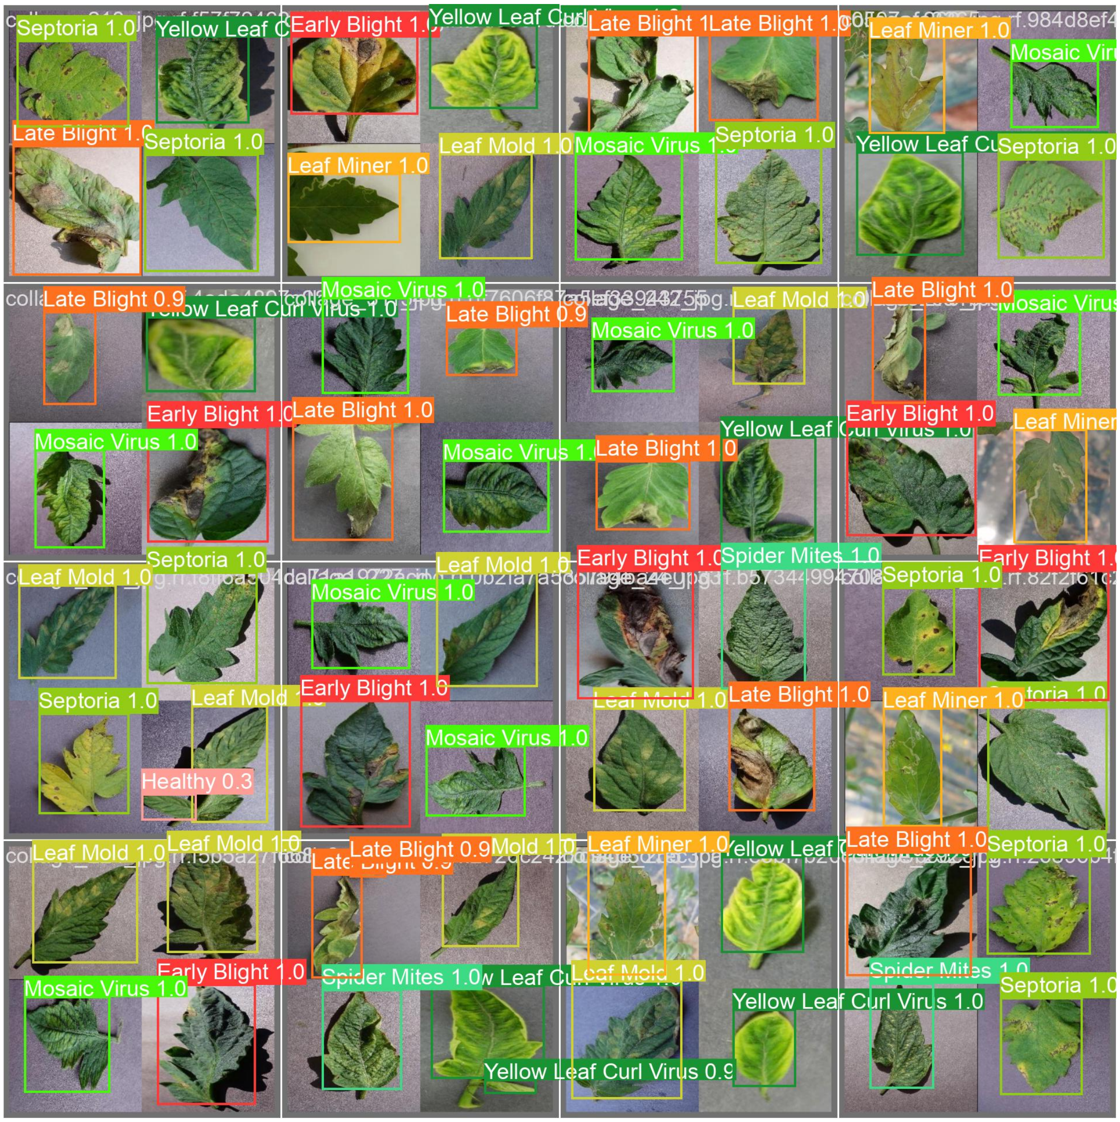

Fig. 10. Detection results based on fine-tuned YOLOv8.

systems by providing an efficient and accurate deep learning-based solution for potato leaf disease detection.

VIII. Future Work

The current study demonstrates the effectiveness of the YOLOv8 model in accurately detecting multiple potato leaf diseases. However, several avenues remain open for future research. First, expanding the dataset with images from varied environmental conditions (e.g., different lighting, backgrounds, or leaf orientations) could improve the model's robustness and generalization ability. Second, integrating temporal data through video sequences or deploying the model on drone-based platforms may enable large-scale, real-time field surveillance, which is crucial for early disease detection and response in precision agriculture. Moreover, although the current work focused on leaf-based disease detection, incorporating other plant parts (e.g., stems or tubers) and multiple crop species could extend the applicability of the system. Another promising direction involves the combination of YOLOv8 with lightweight model optimization techniques such as pruning and quantization, which would facilitate real-time inference on edge devices. Finally, fusing image-based data with sensor data (e.g., temperature, humidity, soil moisture) could contribute to the development of more holistic and context-aware plant health monitoring systems.

Acknowledgment

The authors extend their appreciation to the Deanship of Scientific Research at Northern Border University, Arar, KSA, for funding this research work through the project number NBU-FFR-2025-1659-01.

References

[1] N. Nugroho and N. Sunjoyo, "Agriculture and Food," 2022. [Online]. Available: https://www.worldbank.org/en/topic/agriculture. [Accessed: 11-Aug-2022].

[2] "Bangladesh - Potatoes Production," 2022. [Online]. Available: https://www.nationmaster.com/nmx/timeseries/bangladesh-potatoes-production-fao. [Accessed: 24-Jun-2022].

[3] X. Zheng, et al., "Image segmentation based on adaptive K-means algorithm," *EURASIP Journal on Image and Video Processing*, vol. 2018, no. 1, pp. 1-10, 2018.

[4] M. A. Iqbal and K. H. Talukder, "Detection of potato disease using image segmentation and machine learning," *International Conference on Wireless Communications Signal Processing and Networking (WiSPNET)*, IEEE, 2020.

[5] C. Hou, et al., "Recognition of early blight and late blight diseases on potato leaves based on graph cut segmentation," *Journal of Agriculture and Food Research*, vol. 5, p. 100154, 2021.

[6] M. Lamba, Y. Gigras, and A. Dhull, "Classification of plant diseases using machine and deep learning," *Open Computer Science*, vol. 11, no. 1, pp. 491-508, 2021. [Online]. Available: https://doi.org/10.1515/comp-2020-0122.

[7] M. Islam, et al., "Detection of potato diseases using image segmentation and multiclass support vector machine," *IEEE 30th Canadian Conference on Electrical and Computer Engineering (CCECE)*, IEEE, 2017.

[8] S. Samajpati, et al., "Hybrid approach for apple fruit disease detection and classification using random forest classifier," *International Conference on Communication and Signal Processing (ICCSP)*, 2016.

[9] H. Afzaal, et al., "Detection of a potato disease (early blight) using artificial intelligence," *Remote Sensing*, vol. 13, 2021. [Online]. Available: https://doi.org/10.3390/rs13030411.

[10] S. W. Sigit, et al., "Implementation of convolutional neural network method for classification of diseases in tomato leaves," *Fourth International Conference on Informatics and Computing (ICIC)*, 2019.

[11] K. Kulendu, et al., "Automated recognition of optical image-based potato leaf blight diseases using deep learning," *Physiological and Molecular Plant Pathology*, vol. 117, 2022. [Online]. Available: https://doi.org/10.1016/j.pmpp.2021.101781.

[12] C. Z. Changjian, et al., "Tomato leaf disease identification by restructured deep residual dense network," *IEEE Access*, vol. 9, pp. 28822-28831, 2021.

[13] F. Saeed, et al., "Deep neural network features fusion and selection based on PLS regression with an application for crops diseases classification," *Applied Soft Computing*, vol. 103, 2021. [Online]. Available: https://doi.org/10.1016/j.asoc.2021.107164.

[14] Wang, C.Y.; Bochkovskiy, A.; Liao, H.Y.M. YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 18–22 June 2023; pp. 7464–7475.

[15] Feng, C.; Zhong, Y.; Gao, Y.; Scott, M.R.; Huang, W. Tood: Task-aligned one-stage object detection. In Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 10–17 October 2021; pp. 3490–3499.

[16] Li, X.; Wang, W.; Wu, L.; Chen, S.; Hu, X.; Li, J.; Tang, J.; Yang, J. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. Adv. Neural Inf. Process. Syst. 2020, 33, 21002–21012.

[17] Ge, Z.; Liu, S.; Wang, F.; Li, Z.; Sun, J. Yolox: Exceeding yolo series in 2021. arXiv 2021.

[18] Bolya, D., Zhou, C., Xiao, F., & Lee, Y. J. (2019). Yolact: Real-time instance segmentation. In Proceedings of the IEEE/CVF international conference on computer vision (pp. 9157-9166).

[19] Wang, C. Y., Liao, H. Y. M., Wu, Y. H., Chen, P. Y., Hsieh, J. W., & Yeh, I. H. (2020). CSPNet: A new backbone that can enhance learning capability of CNN. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops (pp. 390-391).

[20] Liu, S., Qi, L., Qin, H., Shi, J., & Jia, J. (2018). Path aggregation network for instance segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 8759-8768).

[21] Lin, T. Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). Feature pyramid networks for object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 2117-2125).

[22] Ding, X., Zhang, X., Ma, N., Han, J., Ding, G., & Sun, J. (2021). RepVGG: Making VGG-style ConvNets great again. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 13733-13742).

[23] dsd, "potato Dataset," *Roboflow Universe*, Roboflow, May 2024. [En ligne]. Disponible sur : https://universe.roboflow.com/dsd-rbi0w/potato-t3qfa. [Consulté le : 21-février-2025].

[24] Tripathi, A., Gohokar, V., & Kute, R. (2024). Comparative Analysis of YOLOv8 and YOLOv9 Models for Real-Time Plant Disease Detection in Hydroponics. Engineering, Technology & Applied Science Research, 14(5), 17269-17275.

[25] Chitraningrum, N., Banowati, L., Herdiana, D., Mulyati, B., Sakti, I., Fudholi, A., ... & Andria, A. (2024). Comparison study of corn leaf disease detection based on deep learning YOLO-v5 and YOLO-v8. Journal of Engineering and Technological Sciences, 56(1), 61-70.

[26] Qin, R., Wang, Y., Liu, X., & Yu, H. (2025). Advancing precision agriculture with deep learning enhanced SIS-YOLOv8 for Solanaceae crop monitoring. Frontiers in Plant Science, 15, 1485903.