BIMA
BULLETIN OF INTELLIGENT MACHINES AND ALGORITHMS
Vol. 1 No.4 May 2026
E-ISSN: 3132-5115
DOI:10.65780/bima.v1i4.22

YOLO26n-Based Apple Leaf Disease Detection for Precision Agriculture Using Lightweight Deep Learning and Object Detection

Darsiti[1*], Budiman[2], Dhika Widiyanto[3], Tarsinah Sumarni[4]

Informatic Management[1*], Informatics[2,3,4]
Universitas Teknologi Digital, Bandung, Indonesia[1*]
Universitas Informatikan dan Bisnis Indonesia[2]
Politeknik Sawunggalih Aji, Kutoarjo, Indonesia[3]
Universitas Teknologi Bandung, Bandung, Indonesia[4]
https://digitechuniversity.ac.id/[1*], https://unibi.ac.id/[2], https://polsa.ac.id/[3], https://utb-univ.id/[4]
darsiti@digitechuniversity.ac.id[1*]

Abstract. Early detection of apple leaf diseases is a critical factor in supporting agricultural productivity and minimizing losses caused by plant disease outbreaks. However, manual identification processes still have limitations in terms of accuracy, consistency, and time efficiency. This study aims to develop an apple leaf disease detection model based on object detection using YOLO26n to identify four main classes: Apple__BlackRot, Apple__CedarRust, Apple__Healthy, and Apple__Scab. The dataset was obtained from Kaggle in YOLO format, consisting of 2,754 training images and 687 validation images. The study employs a transfer learning approach with various data augmentation techniques, such as mosaic, mixup, copy-paste, rotation, translation, and HSV transformation, to enhance the model's generalization ability. Evaluation was conducted using the Precision, Recall, mAP50, and mAP50-95 metrics. The results demonstrated that the YOLO26n model achieved an overall Precision of 0.969, Recall of 0.888, mAP50 of 0.959, and mAP50-95 of 0.881. The Apple__BlackRot class exhibited the best detection performance, achieving an mAP50-95 score of 0.987. The inference results also show that the model is capable of accurately localizing diseases through bounding boxes with a high level of confidence. These findings indicate that YOLO26n has great potential as an efficient and accurate lightweight model for the implementation of real-time precision agriculture-based plant disease detection systems.

Key words: Apple Leaf Disease Detection; YOLO26n; Object Detection; Precision Agriculture; Deep Learning

1. INTRODUCTION

Apple leaf diseases are one of the major challenges in the horticultural sector because they can significantly reduce both the quality and productivity of the harvest [1]. Disease outbreaks such as Black Rot, Cedar Rust, and Scab not only affect the physiological condition of the plants but also have a direct impact on the economic value of apples at both the farmer and agribusiness levels [2]. In conventional practice, leaf disease identification is generally still performed manually through visual observation by farmers or agricultural experts [3], [4]. This approach has various limitations, such as reliance on the observer's experience, relatively lengthy inspection times, and a high likelihood of misidentification, especially when symptoms of different diseases share similar visual characteristics [5]. These conditions highlight the need for the development of an automated detection system capable of providing rapid, accurate, and consistent disease identification to support the implementation of precision agriculture [6], [7].

Advances in Artificial Intelligence (AI) technology, particularly Deep Learning, have opened new opportunities for the development of plant disease detection systems based on digital images [8]. Previous studies have shown that Convolutional Neural Network (CNN) models can deliver strong performance in plant disease classification through automatic visual feature extraction [9]. Architectures such as VGGNet, ResNet, DenseNet, and EfficientNet have been widely used for plant leaf classification tasks with high accuracy [10], [11]. However, most previous studies have focused on image classification approaches, where the model only determines the disease category for the entire image without providing specific information on the location of the diseased area. In the context of field applications, this approach has limitations because it cannot pinpoint infected areas in detail on plant leaves.

As computer vision technology advances, object detection approaches are increasingly being applied to address the limitations of conventional image classification [12]. Models based on the YOLO (You Only Look Once) family have become one of the most popular methods because they can perform real-time object detection with a good balance between accuracy and inference speed [13]. Several recent studies report that YOLOv5, YOLOv7, and YOLOv8 have been successfully used to detect various plant diseases with competitive performance [14]. Nevertheless, some studies still face several challenges, such as data imbalance, variations in lighting conditions, image background complexity, and the model's difficulty in distinguishing disease symptoms with similar visual patterns [15]. Additionally, the use of models with high parameter complexity often requires significant computational resources, making them less optimal for implementation on devices with hardware limitations [16]. Given these conditions, this study proposes the use of the YOLO26n model to detect apple leaf diseases. This study employs a transfer learning approach using a pretrained YOLO26n model combined with various data augmentation techniques, such as mosaic augmentation, mixup, copy-paste, rotation, translation, and HSV transformation, to enhance the model's generalization ability regarding variations in apple leaf image conditions.

Unlike previous studies that have primarily focused on image classification, this study focuses on directly detecting the location of diseases using an object detection approach. This approach enables the system not only to recognize the type of disease but also to precisely identify the position of the infected area via a bounding box. Additionally, YOLO26n was selected because it has a relatively small number of parameters, making it more computationally efficient while still maintaining high detection performance.

2. RELATED WORK

Research on the detection of apple leaf diseases using Artificial Intelligence and Deep Learning has advanced rapidly in recent years. Various approaches have been developed to improve the system's ability to automatically recognize plant diseases, both through image classification and object detection methods. Most previous studies indicate that the use of deep learning models can deliver better performance compared to conventional methods based on manual feature extraction, particularly when dealing with the visual complexity of apple leaf disease patterns in natural environments.

Convolutional Neural Network (CNN)-based approaches have been among the most widely used methods in early research on apple leaf disease classification. Kumar et al. (2024) used a VGG19 model based on transfer learning and achieved a validation accuracy of 98.71% in apple leaf disease classification [17]. This study demonstrated that transfer learning can enhance a model's ability to recognize visual disease patterns even with a limited dataset. However, the image classification approach has limitations because it only produces class labels without being able to specify the exact location of the diseased area on the leaf. A similar limitation was found in the study by Rohith et al. (2025), who used ResNet for apple leaf disease classification and achieved an accuracy of 98.9% [18]. Although it produced high accuracy, this approach was not yet capable of directly localizing diseases, making it less optimal for real-time field monitoring-based implementation.

In addition to conventional CNNs, several studies have begun integrating more complex architectures to enhance feature extraction capabilities. Wang et al. (2023) developed the DEFL model, which combines EfficientNet-B0 and DenseNet121 with Focal Loss and Label Smoothing strategies [19]. This model achieved an accuracy of 99.13% on a natural-environment apple leaf disease dataset. The strength of this research lies in its ability to improve the discrimination of disease features that share similar visual characteristics. However, the high architectural complexity results in greater computational demands, making it less efficient for devices with limited resources.

Subsequent developments have focused on using object detection models based on the YOLO family, capable of both classifying and localizing disease objects. Li et al. (2023) developed BTC-YOLOv5s by integrating BiFPN, Transformer, and CBAM Attention to detect apple leaf diseases in natural environments [20]. This research achieved an mAP of 84.3% and demonstrated fairly good detection capabilities under complex lighting conditions. The advantage of this approach is its ability to perform real-time detection with a relatively lightweight model. However, the mAP performance, which remains below 90%, indicates that the model still faces challenges in distinguishing multi-scale disease objects and complex backgrounds.

Another study by Xu and Wang (2023) proposed ALAD-YOLO, based on YOLOv5s, with the integration of MobileNet-V3, Coordinate Attention, and the Ghost Module to produce a lighter and more efficient model [21]. The model was able to improve accuracy to 90.2% while significantly reducing FLOPs. Nevertheless, the study still faces a trade-off between model accuracy and efficiency. This indicates that developing lightweight models with high detection performance remains a major challenge in the field of plant disease detection.

In a more recent study, Liu and Li (2024) developed an A-Net based on YOLOv5 with the integration of Wise-IoU, an attention mechanism, and a RepVGG module [22]. The model achieved an mAP50 of 92.7%, demonstrating improved localization capabilities and detection stability compared to standard YOLOv5. Meanwhile, Li et al. (2024) introduced YOLO-Leaf, which combines Dynamic Snake Convolution (DSConv), BiFormer Attention, and IF-CIoU to enhance the model's generalization to variations in lighting and disease object sizes [23]. This model achieved an mAP50 of 95.69% on the FGVC8 dataset. Although they deliver high performance, most of these studies employ fairly complex architectures that require relatively significant computational resources.

In addition to CNN and YOLO approaches, several recent studies have begun to explore Vision Transformers (ViT). Ullah et al. (2024) developed AppViT, which combines convolutional blocks with multi-head self-attention to enhance global and local feature extraction capabilities [24]. The model achieved a precision of 96.38% using only approximately 1.3 million parameters. This research demonstrates the significant potential of transformers in the field of smart agriculture; however, implementing transformers generally requires a more complex training process compared to conventional CNNs.

Based on this literature review, several research gaps can be identified. First, most studies still focus on image classification tasks and thus cannot provide detailed information on the location of the disease. Second, object detection-based research generally still faces challenges in balancing detection accuracy, model complexity, and computational efficiency. Third, there is still limited research using the latest generation of YOLO architectures with a lightweight approach that maintains high performance on real-world apple leaf disease datasets. Therefore, this study proposes the use of a YOLO26n object detection model to automatically detect apple leaf diseases using the APPLE Leaf Diseases dataset.

3. METHODS

This research method is designed to develop an apple leaf disease detection system based on object detection using the YOLO26n model, as shown in Figure 1. The research process begins with dataset preparation, data configuration, and model training, and proceeds through the evaluation and inference stages to assess the model's ability to accurately detect and classify apple leaf diseases. The approach used integrates transfer learning, data augmentation, and evaluation based on object detection metrics to produce a robust model with good generalization capabilities across various apple leaf image conditions.

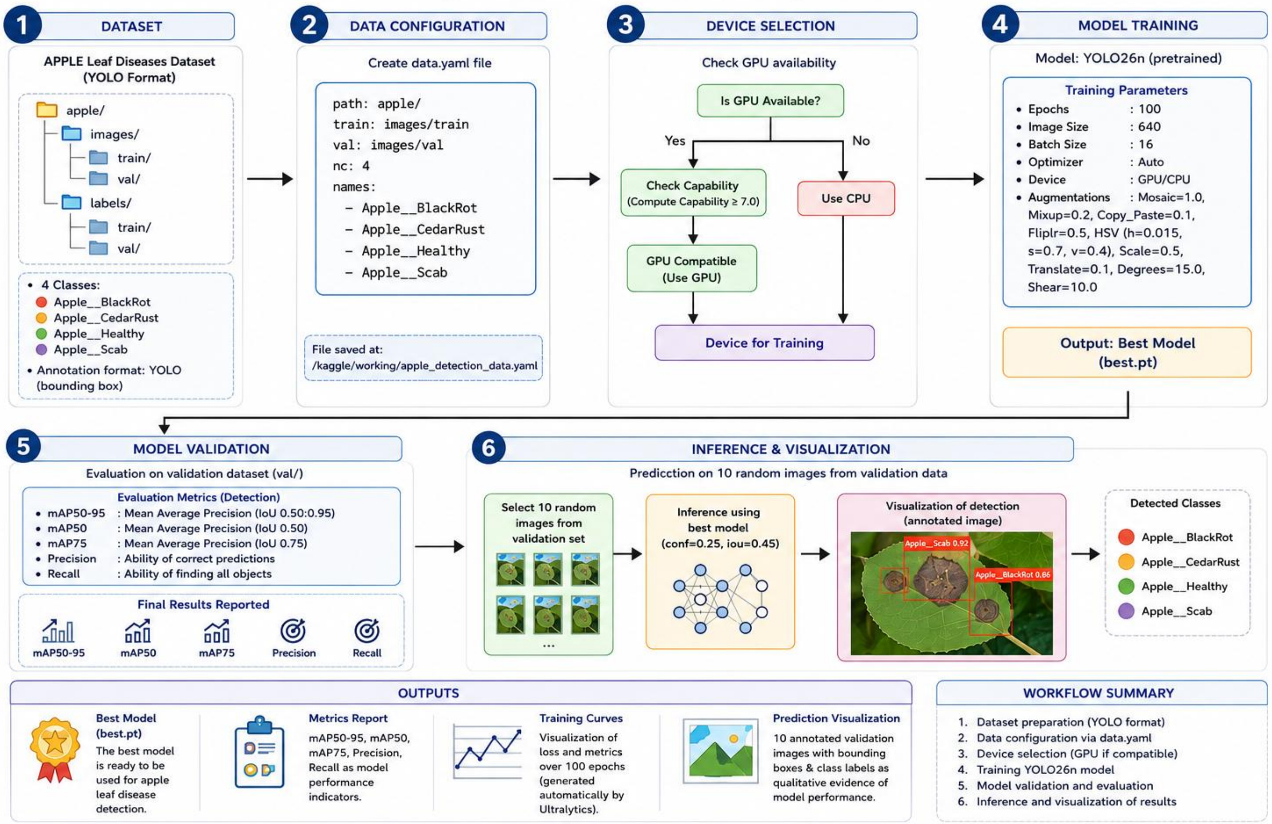

```
path: apple/
train: images/train
val: images/val
nc: 4
names:
  - Apple__BlackRot
  - Apple__CedarRust
  - Apple__Healthy
  - Apple__Scab
```

Fig.1. Proposed research workflow for apple leaf disease detection using YOLO26n, including dataset preparation, data configuration, device selection, model training, validation, inference, and visualization of detection results within the precision agriculture framework.

This study proposes a framework for detecting apple leaf diseases based on the YOLO26n object detection model to identify four classes: Apple__BlackRot, Apple__CedarRust, Apple__Healthy, and Apple__Scab. The dataset was obtained from Kaggle under the name APPLE Leaf Diseases Dataset and is provided in a YOLO-compatible format, allowing for direct integration with the Ultralytics framework without the need for additional annotation conversion. This dataset consists of the directories images/train, images/val, labels/train, and labels/val, where each image is associated with corresponding bounding box annotations stored in .txt files. In this annotation scheme, each bounding box represents a single apple leaf. For the disease classes (BlackRot, CedarRust, and Scab), the bounding boxes encompass leaves showing visible disease symptoms, whereas for the Healthy class, the bounding boxes correspond to the entire leaf without visible signs of infection. Consequently, the detection task is performed at the leaf level, allowing the model to simultaneously localize and classify healthy and diseased leaves within an image. Based on the dataset scanning process, 2,754 images were used for training, with three corrupted images automatically excluded, while 687 validation images containing 805 annotated leaf instances were used for performance evaluation. Since the original dataset only provided training and validation partitions and did not include an independent test set, model evaluation was conducted on the validation subset in accordance with the dataset protocol. Although this evaluation strategy is widely applied in object detection studies, future research should incorporate an independent test dataset or cross-validation procedures to further validate the generalization ability and robustness of the proposed model under diverse field conditions.

The next step involved dataset configuration through the creation of a data.yaml file. This file serves as the interface between the dataset and the Ultralytics YOLO framework. The configuration specifies the locations of the training dataset (train), validation dataset (val), the number of classes (nc = 4), and the corresponding apple leaf disease class names. This configuration enables the YOLO26n model to systematically access and process the dataset during training, validation, and inference stages.

Next, a device selection process is performed to determine whether to use a GPU or CPU. The system first checks for GPU availability using the PyTorch library. If a GPU is available and has a minimum compute capability of 7.0, the training process is run using the GPU to accelerate computations. In this study, a Tesla T4 GPU was used, which supports CUDA-based training. The use of a GPU is crucial for accelerating the forward propagation and backpropagation processes in the YOLO26 network, especially when handling thousands of images with a resolution of 640 × 640 pixels.

The core phase of the research is the training process of the YOLO26n model (YOLO26 nano pretrained model) using a transfer learning approach. The initial model (pretrained weights) is obtained from previous training results on a general dataset, thereby accelerating the convergence process and enhancing feature extraction capabilities. The experiments were conducted using the YOLO26n architecture within the Ultralytics framework. Model training was performed for 100 epochs using input images resized to 640 × 640 pixels and a batch size of 16. Transfer learning was employed by initializing the network with pretrained YOLO26n weights to accelerate convergence and improve feature extraction capability. To enhance model robustness and generalization performance, several data augmentation techniques were applied during training, including Mosaic augmentation, MixUp, Copy-Paste augmentation, horizontal flipping, HSV color transformations, scaling, translation, rotation, and shearing. The optimization process incorporated a predefined learning rate strategy, weight regularization, and confidence-based object filtering to ensure stable training and reliable detection performance. During inference, object predictions were evaluated using confidence and IoU thresholds to determine valid detections and suppress overlapping bounding boxes. All experiments were executed on a workstation equipped with an AMD Ryzen 9 8940HX processor, an NVIDIA GeForce RTX 5050 GPU, and 32 GB DDR5-5600 RAM. This hardware configuration provided sufficient computational resources for deep learning training and validation while enabling efficient processing of high-resolution agricultural image data. The detailed experimental configuration was designed to facilitate reproducibility and support future comparative studies involving lightweight object detection models for precision agriculture applications.

During the optimization process, YOLO26 uses a loss function consisting of box loss, classification loss, and Distribution Focal Loss (DFL). Generally, the total loss can be represented by Equation (1).

$$L_{\{total\}} = L_{\{box\}} + L_{\{cls\}} + L_{\{DFL\}} \tag{1}$$

where $L_{\{box\}}$ represents the bounding box regression loss, $L_{\{cls\}}$ denotes the classification loss, and $L_{\{DFL\}}$ corresponds to the Distribution Focal Loss used to improve bounding box localization accuracy. The combination of these loss components enables the model to simultaneously optimize object localization and disease classification performance. Model performance was evaluated using several standard object detection metrics, namely Precision, Recall, mAP50, and mAP50-95. Precision measures the proportion of correctly detected objects among all predicted objects and is defined as:

Model evaluation is performed using several key metrics in object detection, namely Precision, Recall, mAP50, and mAP50-95. The Precision value is used to measure the proportion of true positive predictions relative to all positive predictions, as shown in Equation (2).

$$Pecision = \frac{TP}{TP + FP} \tag{2}$$

where (TP) denotes the number of true positive detections and (FP) represents the number of false positive detections. Meanwhile, Recall is used to measure the model's ability to identify all objects present in the dataset, as shown in Equation (3).

$$Recall = \frac{TP}{TP + FN} \tag{3}$$

where (FN) indicates the number of false negative detections. In object detection evaluation, the quality of the bounding box is measured using Intersection over Union (IoU), which compares the predicted area with the ground truth area, as shown in Equation (4).

$$IoU = \frac{Area(Overlap)}{Area(Union)} \tag{4}$$

where $Area(Overlap)$ denotes the predicted bounding box and $Area(Union)$ represents the ground-truth bounding box. Furthermore, the model's overall performance is measured using Mean Average Precision (mAP). The mAP50 value represents the average precision at an IoU threshold of 0.50, while mAP50-95 calculates the average performance across IoU thresholds from 0.50 to 0.95 in steps, resulting in a more rigorous and comprehensive evaluation.

4. RESULTS AND DISCUSSIONS

4.1. Results

The training process of the YOLO26 model on the APPLE Leaf Diseases dataset showed a stable and consistent convergence pattern over 100 epochs, as shown in Figure 2. The box loss value decreased gradually from 1.151 in the first epoch to 0.2513 in the 100th epoch. This decrease indicates that the model's ability to localize leaf objects and disease areas becomes increasingly accurate as the training process progresses. Additionally, the classification loss dropped significantly from 4.395 to 0.1102, indicating an improvement in the model's ability to distinguish the four main classes: Healthy, Black Rot, Cedar Rust, and Scab. The most consistent decrease was also observed in the Distribution Focal Loss (DFL), which dropped from 0.03371 to 0.008424, indicating that the quality of bounding box coordinate predictions became increasingly precise. Overall, the loss curve shows that the model did not experience extreme fluctuations or divergence, so the optimization process can be categorized as stable and effective for the task of apple leaf disease detection based on object detection.

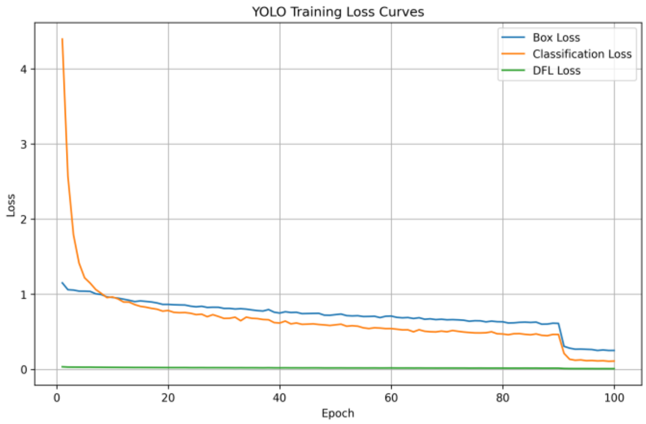

Fig.2. YOLO26n training loss curves over 100 epochs, illustrating the progressive reduction of Box Loss, Classification Loss, and Distribution Focal Loss (DFL), which indicates stable learning behavior and effective model convergence during the training process.

As shown in Figure 3, the Performance Metrics graph, the model's detection performance improved significantly from the start of training until it reached a state of convergence in the final epoch. The Precision value increased from 0.625 in the first epoch to 0.960 in the 100th epoch. This indicates that the majority of predictions generated by the model are correct detections with a relatively low false positive rate. Meanwhile, the Recall value increased from 0.610 to 0.895, indicating that the model's ability to comprehensively identify disease objects has improved, although a small portion of objects remain undetected. The model's key performance is evident in the mAP50 value, which increased from 0.653 to 0.951, while mAP50-95 increased from 0.480 to 0.877. These values indicate that the model has excellent detection performance both at the standard IoU threshold and under stricter evaluation criteria. Peak performance was achieved at the 73rd epoch with an mAP50-95 value of 0.880, indicating an optimal balance between the model's generalization ability and stability.

Fig.3. Training performance curves of the YOLO26n model showing the progression of Precision, Recall, mAP50, and mAP50-95 over 100 training epochs, indicating stable convergence and improved detection performance during the learning process.

Based on the Training Overview graph, it is evident that the model's learning process occurs in several main phases. During the initial training phase (epochs 1–20), there is a very sharp decrease in loss accompanied by a rapid improvement in evaluation metrics. This indicates that the model successfully learned the basic features of apple leaf disease patterns during the early stages of training. Subsequently, in the middle phase (epochs 20–70), the decrease in loss proceeds more slowly but steadily, while the Precision, Recall, and mAP metrics continue to improve gradually. This phase reflects the process of refining more complex feature representations, particularly in distinguishing visual characteristics between leaf diseases that share similar textures or colors. In the final training phase (epochs 70–100), the model's performance begins to reach a state of saturation or convergence, marked by relatively small changes in metrics. Although there were slight fluctuations in the final epochs, evaluation scores remained in the high range with mAP50 around 0.95 and mAP50–95 around 0.87–0.88. This indicates that the YOLO26 model possesses good generalization capabilities on the validation dataset without significant signs of overfitting.

The overall experimental results demonstrate that the proposed YOLO26n model exhibits strong capabilities in detecting apple leaf diseases within the object detection framework. The use of a YOLO-formatted dataset with well-structured bounding box annotations contributes to stable model convergence and effective feature learning during the training process. As presented in Table 1, this model achieved a Overall Precision of 0.969, a Recall of 0.888, an mAP50 of 0.959, and an mAP50-95 of 0.881 on the validation dataset. These findings indicate that this model is capable of accurately identifying and localizing disease symptoms across various classes of apple leaves. Furthermore, the high detection performance obtained in this study highlights the potential of YOLO26n as a reliable and computationally efficient solution to support automated disease monitoring in precision agriculture environments.

TABLE 1. Performance Evaluation Results of YOLO26n for Apple Leaf Disease Detection on the Validation Dataset Using Precision, Recall, mAP50, and mAP50-95 Metrics

Kelas	Images	Instances	Precision (P)	Recall (R)	mAP50	mAP50-95
All Classes	687	805	0.969	0.888	0.959	0.881
Apple__BlackRot	124	124	0.997	1.000	0.995	0.987
Apple__CedarRust	72	97	0.931	0.830	0.924	0.792
Apple__Healthy	347	417	0.986	0.885	0.980	0.893
Apple__Scab	144	167	0.962	0.838	0.935	0.851

The evaluation results summarized in Table 1 demonstrate the effectiveness of the proposed YOLO26n model in detecting and classifying apple leaf diseases. Using 687 validation images containing 805 annotated disease instances, the model achieved an overall Precision of 0.969, Recall of 0.888, mAP50 of 0.959, and mAP50-95 of 0.881. The high Precision score indicates that the majority of detected objects correspond to actual disease instances, reflecting a low rate of false positives. Meanwhile, the Recall score indicates that the model successfully identified most of the disease instances present in the validation set, although a small number of instances remained undetected. The mAP50 and mAP50-95 scores further confirm the model's strong capability in disease classification and bounding box localization under both standard and stricter evaluation criteria. Overall, these results demonstrate that the proposed model achieves a robust balance between detection accuracy and localization precision, making it suitable for practical application in smart crop monitoring systems.

A more in-depth analysis of performance by class reveals variations in detection accuracy across disease categories. These differences are primarily related to the visual characteristics of each disease class, including lesion size, color distribution, texture patterns, and the number of training samples available for model training. Nevertheless, all classes achieved high detection performance, indicating that the proposed model successfully learned distinguishing features to differentiate healthy leaves from various disease conditions. Class-specific evaluation results are discussed in the following section to provide a deeper understanding of the strengths and limitations of the proposed approach.

In the Apple__BlackRot class, the model demonstrated the best performance compared to other classes, with a Precision of 0.997, a Recall of 1.000, an mAP50 of 0.995, and an mAP50-95 of 0.987. These results show that all Black Rot disease objects in the validation data were detected with high accuracy without any missed detections. The high performance in this class indicates that the visual characteristics of Black Rot exhibit a sufficiently consistent pattern that is easily learned by the model, in terms of color, texture, and the shape of the disease area on apple leaves.

In the Apple__CedarRust class, the model's performance is relatively lower compared to other classes, particularly with a Recall of 0.830 and an mAP50-95 of 0.792. Although the Precision remains high at 0.931, these results show that while the model tends to produce accurate predictions when detecting Cedar Rust, there are still some instances that go undetected. This condition is likely influenced by the more complex visual variations of Cedar Rust or the relatively smaller amount of data compared to the Healthy class. With only 97 instances in the validation dataset, the feature representation of this disease may not yet be fully optimized to improve the model's generalization ability.

In the Apple__Healthy class, the model achieved excellent performance with a Precision of 0.986, a Recall of 0.885, an mAP50 of 0.980, and an mAP50-95 of 0.893. These results indicate that the model is capable of distinguishing healthy leaves from disease-infected leaves with a very high level of accuracy. The high performance in the Healthy class is also influenced by the largest number of instances, namely 417 objects, so that the model has more data variation to study the characteristics of healthy leaves in greater depth.

Meanwhile, in the Apple__Scab class, the model achieved a Precision of 0.962, a Recall of 0.838, an mAP50 of 0.935, and an mAP50-95 of 0.851. These values indicate that the model is capable of detecting Scab disease effectively, although the lower recall rate compared to the BlackRot and Healthy classes suggests that some Scab objects remain undetected. This may be due to the visual similarity of Scab patterns to other diseases or the considerable variation in lesion shapes on apple leaves.

The inference results show that the YOLO26n model is capable of detecting and classifying apple leaf diseases with a high level of confidence under various leaf image conditions, as shown in Figure 4. Most objects were correctly identified as the Apple__Healthy, Apple__Scab, and Apple__BlackRot classes with confidence values ranging from 0.76 to 0.96. The bounding boxes generated by the model also appear quite precise in marking the leaf areas that are the detection targets, both on healthy leaves and those infected with disease. These results indicate that the model is able to effectively learn the visual characteristics of each class, including differences in texture, color, and disease spot patterns on the leaf surface. Furthermore, the high and consistent confidence scores across most samples indicate that the model possesses good generalization capabilities regarding variations in leaf shape, image capture angles, and lighting conditions. These findings demonstrate that the YOLO26n-based approach has the potential to be implemented as an automated, real-time apple leaf disease detection system to support the implementation of precision agriculture.

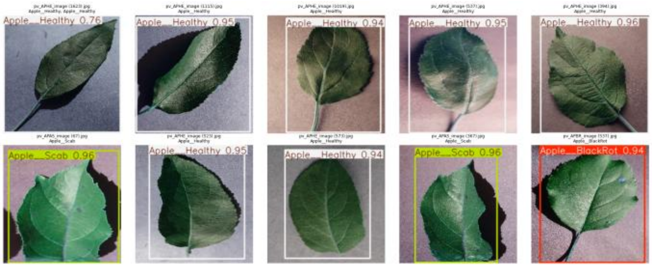

Fig.4. Visualization of YOLO26n inference results on apple leaf images showing successful detection and classification of Apple__Healthy, Apple__Scab, and Apple__BlackRot classes with high confidence scores and accurate bounding box localization.

Overall, the evaluation results show that the YOLO26 model has very strong object detection capabilities for classifying apple leaf diseases in the context of precision agriculture. Its high performance across nearly all classes indicates that the model is able to effectively learn the visual features of the diseases while maintaining good bounding box localization quality under various conditions, including both healthy and infected apple leaves.

4.2. Discussions

The research results show that the proposed YOLO26n model is capable of delivering highly competitive performance in detecting apple leaf diseases, with a Precision of 0.968, Recall of 0.887, mAP50 of 0.958, and mAP50-95 of 0.880. These findings demonstrate that an object detection-based approach using a modern lightweight architecture can achieve a good balance between detection accuracy, computational efficiency, and model generalization capabilities in an agricultural environment. These results further reinforce the argument that the development of plant disease detection systems requires not only high classification accuracy but also the ability to precisely localize diseased areas to support the implementation of precision agriculture.

Previous studies have demonstrated the effectiveness of convolutional neural network (CNN) architectures for apple leaf disease classification, often reporting classification accuracies exceeding 98% [17], [18], [19]. However, these studies focused on image-level classification tasks, where the objective is to assign a disease label to an entire image without identifying the spatial location of disease symptoms. Consequently, the reported classification accuracies reflect a different problem setting and should not be directly compared with object detection metrics such as Precision, Recall, mAP50, and mAP50-95.

To address this limitation, the present study adopts an object detection approach based on YOLO26n, which simultaneously performs disease classification and localization through bounding box prediction. By providing information on both the disease category and its location within the leaf image, the proposed approach is more suitable for practical precision agriculture applications that require accurate disease monitoring and targeted intervention. Therefore, the performance of the proposed model is discussed in relation to previous object detection studies rather than classification-based approaches.

Among object detection methods, Li et al. (2023) reported an mAP of 84.3% using BTC-YOLOv5s[20], while Xu and Wang (2023) achieved a detection accuracy of 90.2% with ALAD-YOLO [21]. Liu and Li (2024) further improved detection performance, achieving an mAP50 of 92.7% using A-Net based on YOLOv5. In comparison, the proposed YOLO26n model achieved an mAP50 of 95.8%, demonstrating strong detection capability while maintaining a lightweight architecture suitable for efficient deployment.

Another advantage of this study lies in the model's efficiency. YOLO26n has only approximately 2.37 million parameters with a computational complexity of 5.2 GFLOPs. Compared to previous models that heavily integrated complex attention mechanisms, additional transformers, or multi-branch feature extraction, the proposed model remains capable of delivering high performance without excessively increasing computational complexity. These findings demonstrate that optimizing modern lightweight architectures can be an effective solution for implementing plant disease detection systems on resource-constrained devices, such as edge computing devices or mobile agriculture systems.

Furthermore, class-specific evaluation results show that the model performs exceptionally well in detecting BlackRot disease, achieving an mAP50-95 of 0.987 and a Recall of 1.000. However, performance on the CedarRust class remains lower compared to other classes, with an mAP50-95 of 0.792. This indicates that variations in visual patterns, lesion sizes, and a smaller dataset remain the primary challenges in detecting apple leaf diseases. These findings are consistent with the research by Zhang et al. (2023) and Gong et al. (2023), which stated that small-sized disease objects and complex background conditions can reduce a model's ability to perform multi-scale detection [25], [26].

Overall, the results of this study provide an important contribution to the development of AI-based apple leaf disease detection systems. First, this study demonstrates that lightweight models such as YOLO26n can achieve high detection performance without requiring high model complexity. Second, this study shows that the integration of transfer learning and appropriate data augmentation can improve the model's generalization ability regarding variations in apple leaf images. Third, this study reinforces the direction of developing real-time object detection-based smart agricultural systems that are more practical for direct implementation in the field.

TABLE 1. Comparative analysis of previous studies and the proposed YOLO26n model for apple leaf disease detection, highlighting differences in methods, task types, and performance metrics across classification and object detection approaches.

Research	Method	Task Type	Key Results	Strengths	Weaknesses
[17]	VGG19 Transfer Learning	Classification	Accuracy 98.71%	High accuracy	Does not detect the location of the disease
[19]	DEFL (EfficientNet + DenseNet)	Classification	Accuracy 99.13%	Very strong feature extraction	High model complexity
[18]	ResNet	Classification	Accuracy 98.9%	Robust in classification	Does not support object localization
[20]	BTC-YOLOv5s	Object Detection	mAP 84.3%	Real-time detection	Accuracy is still limited
[21]	ALAD-YOLO	Object Detection	Accuracy 90.2%	Lightweight and efficient	Trade-off akurasi dan efisiensi
[22]	A-Net YOLOv5	Object Detection	mAP50 92.7%	Better bounding box regression	More complex architecture
[23]	YOLO-Leaf	Object Detection	mAP50 95.69%	High generalization	Requires more computation
Current Research	**YOLO26n**	**Object Detection**	**Precision 0.968, Recall 0.887, mAP50 0.958, mAP50-95 0.880**	**Lightweight, high accuracy, efficient for real-time detection**	**CedarRust performance still needs improvement**

5. CONCLUSIONS

This study contributes to the growing body of research on intelligent plant disease detection by providing an empirical evaluation of the YOLO26n architecture for apple leaf disease detection under an object detection framework. Rather than introducing a new detection architecture, this study demonstrates the applicability of a recent lightweight YOLO model for simultaneous disease classification and localization in apple leaf images. The findings show that YOLO26n achieved a Precision of 0.968, Recall of 0.887, mAP50 of 0.958, and mAP50-95 of 0.880, indicating that lightweight object detection models can provide competitive performance while maintaining computational efficiency. From a practical perspective, the results suggest that YOLO26n represents a viable solution for real-time disease monitoring in precision agriculture environments. The study therefore contributes empirical evidence regarding the effectiveness of modern lightweight object detection architectures for agricultural disease detection tasks and provides a benchmark for future comparisons involving more advanced or customized detection models.

REFERENCES

[1] E. Szpyrka, S. Migdal-Pecharroman, and P. Książek-Trela, 'Biological Strategies and Innovations in Pest Control and Fruit Storage in Apple Orchards: A Step Towards Sustainable Agriculture', *Agronomy*, vol. 15, no. 10, p. 2373, Oct. 2025, doi: 10.3390/agronomy15102373.

[2] K. I. Roumeliotis, R. Sapkota, M. Karkee, N. D. Tselikas, and D. K. Nasiopoulos, 'Plant Disease Detection through Multimodal Large Language Models and Convolutional Neural Networks', Apr. 29, 2025, *arXiv*: arXiv:2504.20419. doi: 10.48550/arXiv.2504.20419.

[3] A. Upadhyay *et al.*, 'Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture', *Artif Intell Rev*, vol. 58, no. 3, p. 92, Jan. 2025, doi: 10.1007/s10462-024-11100-x.

[4] E. D. Ruby, G. Amirthayogam, G. Sasi, T. Chitra, A. Choubey, and S. Gopalakrishnan, 'Advanced Image Processing Techniques for Automated Detection of Healthy and Infected Leaves in Agricultural Systems', *Mesopotamian Journal of Computer Science*, vol. 4, no. 1, pp. 44–52, Jan. 2024, doi: https://doi.org/10.58496/MJCSC/2024/006.

[5] M. John, I. Bankole, O. Ajayi-Moses, T. Ijila, T. Jeje, and L. Patil, 'Relevance of Advanced Plant Disease Detection Techniques in Disease and Pest Management for Ensuring Food Security and Their Implication: A Review', *American Journal of Plant Sciences*, vol. 14, Nov. 2023, doi: 10.4236/ajps.2023.1411086.

[6] M. Majdalawieh, C. Martins, M. Radi, M. Alaraj, and S. Khan, 'Precision agriculture in the age of AI: A systematic review of machine learning methods for crop disease detection', *Smart Agricultural Technology*, vol. 12, p. 101491, Dec. 2025, doi: 10.1016/j.atech.2025.101491.

[7] G. Mohyuddin, M. A. Khan, A. Haseeb, S. Mahpara, M. Waseem, and A. M. Saleh, 'Evaluation of Machine Learning Approaches for Precision Farming in Smart Agriculture System: A Comprehensive Review', *IEEE Access*, vol. 12, pp. 60155–60184, 2024, doi: 10.1109/ACCESS.2024.3390581.

[8] Y. Wu, L. Chen, N. Yang, and Z. Sun, 'Research Progress of Deep Learning-Based Artificial Intelligence Technology in Pest and Disease Detection and Control', *Agriculture*, vol. 15, no. 19, p. 2077, Oct. 2025, doi: 10.3390/agriculture15192077.

[9] M. Yogeshwari and G. Thailambal, 'Automatic feature extraction and detection of plant leaf disease using GLCM features and convolutional neural networks', *Materials Today: Proceedings*, vol. 81, pp. 530–536, Jan. 2023, doi: 10.1016/j.matpr.2021.03.700.

[10] M. K. A. Mazumder, M. M. Kabir, A. Rahman, M. Abdullah-Al-Jubair, and M. F. Mridha, 'DenseNet201Plus: Cost-effective transfer-learning architecture for rapid leaf disease identification with attention mechanisms', *Heliyon*, vol. 10, no. 15, Aug. 2024, doi: 10.1016/j.heliyon.2024.e35625.

[11] J. G. Kotwal, R. Kashyap, and P. Mohd. Shafi, 'Artificial Driving based EfficientNet for Automatic Plant Leaf Disease

Classification', *Multimed Tools Appl*, vol. 83, no. 13, pp. 38209–38240, Apr. 2024, doi: 10.1007/s11042-023-16882-w.

[12] N. Manakitsa, G. S. Maraslidis, L. Moysis, and G. F. Fragulis, 'A Review of Machine Learning and Deep Learning for Object Detection, Semantic Segmentation, and Human Action Recognition in Machine and Robotic Vision', *Technologies*, vol. 12, no. 2, p. 15, Feb. 2024, doi: 10.3390/technologies12020015.

[13] A. Vijayakumar and S. Vairavasundaram, 'YOLO-based Object Detection Models: A Review and its Applications', *Multimed Tools Appl*, vol. 83, no. 35, pp. 83535–83574, Oct. 2024, doi: 10.1007/s11042-024-18872-y.

[14] M. Abedin, M. A. A. Khan, and N. Sultana, 'Comparative Analysis of You Only Look Once (Yolo) Models V5, V8, V10, and V11 for Automated Detection of Lemon Leaf Canker Disease', Aug. 13, 2025, *Social Science Research Network, Rochester, NY*: 5390866. doi: 10.2139/ssrn.5390866.

[15] M. Saini and S. Susan, 'Tackling class imbalance in computer vision: a contemporary review', *Artif Intell Rev*, vol. 56, no. 1, pp. 1279–1335, Oct. 2023, doi: 10.1007/s10462-023-10557-6.

[16] G. Menghani, 'Efficient Deep Learning: A Survey on Making Deep Learning Models Smaller, Faster, and Better', *ACM Comput. Surv.*, vol. 55, no. 12, p. 259:1-259:37, Mar. 2023, doi: 10.1145/3578938.

[17] A. Kumar, L. Nelson, and S. Gomathi, 'Transfer Learning of VGG19 for the Classification of Apple Leaf Diseases', in *2024 2nd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, Jan. 2024, pp. 1643–1648. doi: 10.1109/IDCIoT59759.2024.10467967.

[18] D. Rohith, P. Saurabh, and D. Bisen, 'An integrated approach to apple leaf disease detection: leveraging convolutional neural networks for accurate diagnosis', *Multimed Tools Appl*, vol. 84, no. 32, pp. 40307–40342, Sep. 2025, doi: 10.1007/s11042-025-20735-z.

[19] R. Wang, F. Chen, X. Zhu, and X. Zhang, 'Identifying apple leaf diseases using improved EfficientNet', *Transactions of the Chinese Society of Agricultural Engineering*, vol. 39, no. 18, pp. 201–210, Sep. 2023, doi: 10.11975/j.issn.1002-6819.202306140.

[20] H. Li, L. Shi, S. Fang, and F. Yin, 'Real-Time Detection of Apple Leaf Diseases in Natural Scenes Based on YOLOv5', *Agriculture*, vol. 13, no. 4, p. 878, Apr. 2023, doi: 10.3390/agriculture13040878.

[21] W. Xu and R. Wang, 'ALAD-YOLO:an lightweight and accurate detector for apple leaf diseases', *Front. Plant Sci.*, vol. 14, Aug. 2023, doi: 10.3389/fpls.2023.1204569.

[22] Z. Liu and X. Li, 'An improved YOLOv5-based apple leaf disease detection method', *Sci Rep*, vol. 14, no. 1, p. 17508, Jul. 2024, doi: 10.1038/s41598-024-67924-8.

[23] T. Li, L. Zhang, and J. Lin, 'Precision agriculture with YOLO-Leaf: advanced methods for detecting apple leaf diseases', *Front. Plant Sci.*, vol. 15, Oct. 2024, doi: 10.3389/fpls.2024.1452502.

[24] W. Ullah *et al.*, 'Efficient identification and classification of apple leaf diseases using lightweight vision transformer (ViT)', *Discov Sustain*, vol. 5, no. 1, p. 116, Jun. 2024, doi: 10.1007/s43621-024-00307-1.

[25] X. Gong and S. Zhang, 'A High-Precision Detection Method of Apple Leaf Diseases Using Improved Faster R-CNN', *Agriculture*, vol. 13, no. 2, p. 240, Feb. 2023, doi: 10.3390/agriculture13020240.

[26] Y. Zhang, G. Zhou, A. Chen, M. He, J. Li, and Y. Hu, 'A precise apple leaf diseases detection using BCTNet under unconstrained environments', *Computers and Electronics in Agriculture*, vol. 212, p. 108132, Sep. 2023, doi: 10.1016/j.compag.2023.108132.