


Implementation of Yolov8 Instance Segmentation for Detection and Segmentation of Leaf Diseases in Horticultural Plants

Da'i Gilang Ramadhan¹, Putra Edi Mujahid²

Information System, Faculty Of Science And Technology, Prima Indonesia University, Jl. Sampul No.3, Sei Putih Baru, Kec. Medan Petisah, Kota Medan, Provinsi Sumatera Utara

Article Info	ABSTRACT
Keywords: YOLOv8, instance segmentation, data augmentation, leaf disease, deep learning.	Leaf diseases are one of the primary factors contributing to reduced agricultural productivity, especially in tropical regions such as Indonesia. Early identification of infected areas is crucial to prevent further disease spread and minimize crop losses. This study proposes the application of YOLOv8 instance segmentation with an in-training augmentation strategy to precisely detect and segment leaf disease areas. The dataset consists of 3,528 diseased leaf images, with 588 original images used for validation and 2,940 augmented images for training. The augmentation process includes mosaic, HSV adjustment, spatial transformations, RandAugment, and other techniques to enhance the model's generalization capability. Model performance was evaluated using mAP, precision, recall, and Intersection over Union (IoU) metrics. Experimental results demonstrate that the model achieved stable performance with high mAP@50, precision and recall approaching optimal values, and an average IoU above 0.5. This approach is effective and has strong potential for implementation in early plant disease detection systems, supporting farmers in making timely and accurate decisions.
This is an open access article under the CC BY-NC license 	Corresponding Author: Da'i Gilang Ramadhan Information System, Faculty Of Science And Technology, Prima Indonesia University, Jl. Sampul No.3, Sei Putih Baru, Kec. Medan Petisah, Kota Medan, Provinsi Sumatera Utara daigilang@gmail.com

INTRODUCTION

Agriculture remains the backbone of the economy in North Sumatra. According to 2021 data, the agricultural sector contributed approximately 13.28% to the national GDP and supplied rice consumed by more than 90% of Indonesia's population, with per capita consumption averaging 160 kg per year [1]. Furthermore, about 32% of Indonesia's population depends on agriculture as their primary source of livelihood [2]. At the commodity level, rice, chili, tomatoes, and other horticultural crops serve not only as staple food sources but also as the main source of rural household income [3]. Red chili consumption in 2021 was recorded at 0.15 kg per capita per month (≈ 40.9 thousand tons per month), with a total consumption of 490.83 thousand tons. Uneven supply distribution has led to price fluctuations, contributing to inflation in the range of 0.01– 0.07% [4].

The greatest threat to crop production is pests and diseases. Globally, plant diseases cause yield losses of up to 30% in staple crops, resulting in economic damage worth hundreds

of billions of dollars [5]. In rice cultivation, disease infections can reduce yields by 20–40% [6]. Such production losses not only undermine food security, but also disrupt farmers' income and the price stability of commodities such as chili and tomatoes [7].

In practice, most farmers and extension workers still rely on visual observation of leaves, stems, or fruits to identify diseases [8]. This manual method requires substantial experience and skill, is time-consuming, inefficient, and highly prone to misdiagnosis. Even at the research level, classical image processing methods still depend on manual feature engineering and struggle to address the variability of natural environments [9]. Laboratory-based observations and visual inspections are also impractical for large-scale agricultural settings [10].

Advancements in artificial intelligence (AI) have created new opportunities. Deep learning models, particularly Convolutional Neural Networks (CNNs), have proven to be far more accurate compared to conventional methods [11] [12] [13]. Numerous studies have employed CNNs, EfficientNet, and Squeeze-and-Excitation (SE) modules for disease classification in rice, maize, and tomato, achieving high levels of accuracy. For example, a modified CNN model achieved an accuracy of 96.08% in rice disease detection [14], while EfficientNet Lite combined with KE-SVM improved potato disease detection accuracy from 79.38% to 87.82% in uncontrolled data and reached 99.54% in controlled data [15]. InsightNet, a MobileNet-based model, successfully classified tomato, bean, and chili leaf diseases with accuracy rates exceeding 97%.

For tomato crops, the TomatoDet method, which integrates Swin DETR, Meta-ACON, and IBiFPN modules, achieved a mean average precision (mAP) of 92.3%, while also accelerating detection speed. Other researchers developed E-TomatoDet, which combines transformer architectures and multi-kernel modules for tomato leaf disease classification, reaching mAP50 of 97.2%.

METHOD

This research is quantitative with an experimental and descriptive approach. The primary focus is the development and evaluation of a YOLOv8-based deep learning model for segmenting diseased areas of plant leaves in digital images. Analysis was conducted using a leaf image dataset that had been masked with diseased areas and haze using data augmentation techniques to test model performance improvements.

The population in this study consists of all horticultural plant leaf images that are vulnerable to disease attacks in the region of North Sumatra and Indonesia. The research sample comprises a collection of diseased leaf images along with infection area masks (labels) obtained from open-source datasets (Kaggle) as well as digitally augmented data (rotation, flipping, zooming, and others) to increase the variety of training and validation data.

In addition to classification methods, object detection algorithms such as YOLO and SSD have been widely employed for plant disease identification. For instance, YOLOv7 applied to tea leaf diseases achieved an accuracy of 97.3%, with a precision of 96.7% and a recall of 96.4%. Meanwhile, YOLOv4 combined with data augmentation reported an accuracy of 99.99%. An enhanced variant, SerpensGate YOLOv8, integrates Dynamic Snake Convolution and Super Token Attention modules, resulting in a 3.3% improvement in mAP compared to the standard

YOLOv8. For segmentation tasks, a modified YOLOv8-Seg that incorporates Ghost and BiFPN modules, along with the strategy of disabling mosaic augmentation during the final 10 epochs, demonstrated improved leaf segmentation accuracy. In other domains, YOLO LeafNet, a customized model derived from YOLOv5/v8, achieved a precision of 0.985, recall of 0.980, and mAP50 of 0.990. A comparative study between YOLOv4 and YOLOv3 revealed that YOLOv4 attained higher accuracy and mAP (98%) while reducing computation time (29 s vs. 105 s). Furthermore, lightweight models such as SNMPF (based on ShuffleNetV2 and SimAM) successfully detected maize leaf diseases with 98.40% accuracy, while maintaining a compact model size of only 1.56 MB.

Nevertheless, most studies have primarily focused on image-level classification or bounding box detection rather than detailed segmentation of infection areas. This research gap is particularly evident in tropical crops specific to Indonesia. Recent studies have indicated that data augmentation and YOLOv8-based instance segmentation can enhance the identification of diseased regions. However, farmers in practice still rely heavily on personal experience, which often leads to misdiagnosis and improper treatment.

In response to these conditions, this study proposes the application of a YOLOv8 instance segmentation model with data augmentation strategies to detect and segment disease areas on rice, chili, and tomato leaves. This approach is expected to achieve precise spatial detection, provide clear visualizations of infection areas for farmers and extension workers, and accelerate field diagnosis. Consequently, disease management can be carried out earlier and more effectively, thereby supporting crop productivity and strengthening local food security.

1. Research Instrument

Dataset: Plant leaf images (JPG/PNG format) along with binary masks of disease areas (PNG). Software:

- a. Python (Jupyter/Kaggle Notebook)
- b. Main libraries: Ultralytics YOLOv8, OpenCV, Albumentations, NumPy, Matplotlib, scikit-learn.
- c. Hardware: GPU (NVIDIA Tesla P100 / NVIDIA Tesla T4) through cloud colaboratory services.
- d. Evaluation Metrics: Intersection over Union (IoU), Dice coefficient, and mean Average Precision (mAP).

2. Research Procedure

Dataset Preparation:

- a. Collecting diseased leaf image datasets and masks from open sources (Kaggle, GitHub, etc.).
- b. Performing augmentation on images and masks to expand the diversity of training data.

Data Preprocessing:

- a. Converting masks to YOLOv8 segmentation polygon format.
- b. Structuring folders according to YOLOv8 standards (images/train, images/val, labels/train, labels/val).
- c. Normalizing image and mask sizes to a resolution of 640x640 pixels.

Model Training:

- Initializing the YOLOv8 instance segmentation model with pretrained weights.
- Training the model on augmented training data and validating on original data.
- Using predetermined epochs, batch size, and training parameters.

Model Evaluation:

- Evaluating performance using IoU, Dice, and mAP metrics on validation data.
- Saving prediction results, prediction visualizations, and per-image scores into a CSV file.
- Comparing model performance with and without data augmentation.

Analysis and Interpretation

Interpreting the evaluation results, including IoU heatmap visualization per image and the distribution of Dice/IoU values.

RESULTS AND DISCUSSION

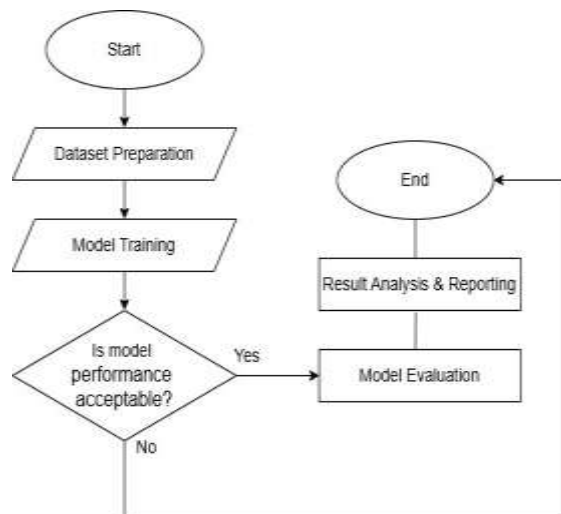


Figure 1. Research Stages

Dataset Description and Initial Analysis

The dataset used in this study is the Leaf Disease Segmentation Dataset (Kaggle, fakhrealam9537/leaf-disease-segmentation-dataset), consisting of a total of 3,528 diseased plant leaf images. The dataset is divided into two main groups: original data (588 images, used as validation data) and augmented data (2,940 images, used as training data). Each image is paired with a digitally annotated disease mask, providing the exact location and shape of the infected areas on the leaves. To ensure consistency, all data were normalized to a resolution of 640×640 pixels using bilinear interpolation. The dataset preparation process also adhered to the standard YOLOv8 segmentation folder format, where images and masks were systematically organized into separate *imgcs* and *mxsks* folders for each subset of *rxin* and *vxl*. In addition, data augmentation was systematically applied using strategies such as rotation, flipping, zooming, and color transformations to enable the model to learn from diverse visual

variations that may occur in real-world conditions. Through this procedure, the resulting dataset became highly representative and robust for training deep learning- based segmentation models.

Implementasi Model YOLOv8 Instance Segmentation



Figure 2. Mask and Segmentation of Plant Disease Areas

In this study, the YOLOv8n-seg model was employed as the primary backbone due to its efficient architecture, consisting of 151 layers and more than 3 million parameters, making it well-suited for experiments under limited computational resources (e.g., GPU Tesla P100). YOLOv8n-seg is a variant of YOLOv8 specifically designed for instance segmentation, combining fast object detection with the ability to distinguish and separate each object instance within an image.

Its architecture is divided into three main components: the backbone, responsible for extracting visual features from images; the neck, which fuses information across multiple resolution levels using FPN/PAN mechanisms; and the segmentation head, which predicts bounding boxes, classes, and mask coefficients for each instance. This model was selected not only for its efficiency but also for its flexibility in handling various complexities and variations in diseased leaf images, including differences in size, color, and distribution of infection areas within the dataset. Furthermore, YOLOv8n-seg leverages a prototypical mask (ProtoNet) approach, enabling the spatial prediction of disease area masks with high accuracy. Unlike mere image classification, this approach provides richer information for analysis and decision-making in plant disease management.

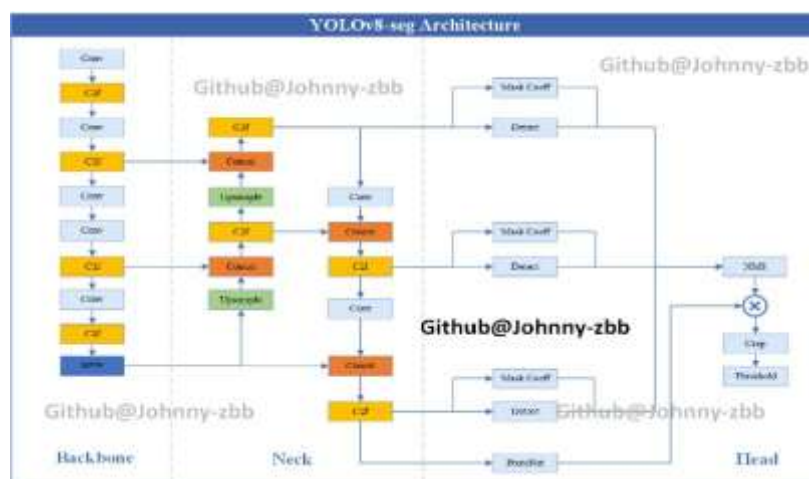


Figure 3. YOLOv8-seg Architecture

Training Configuration and Performance

In-training augmentation

Based on the sequence of processes in the notebook, the initial stage of training configuration begins with the application of a series of augmentation transformations designed to enhance the generalization capability of the YOLOv8n-seg model across various conditions and image characteristics. These augmentation techniques not only expand the diversity of the training data but also assist the model in addressing challenges such as variations in lighting, differences in scale, and the spatial distribution of diseases on the leaves.

- Mosaic augmentation combines four images into one composite frame, then is disabled during the last 10 epochs to maintain training stability before convergence.
- HSV adjustment modifies hue, saturation, and value to make the model more robust against variations in lighting and color.
- Spatial transformations include translation, scaling, rotation, and zoom to simulate field image capture from different angles and distances.
- RandAugment adds visual diversity through controlled random transformations, preventing the model from overfitting to fixed augmentation patterns.
- Erasing augmentation randomly removes parts of the image, forcing the model to focus on remaining contextual information.
- Light augmentations such as Blur, MedianBlur, ToGray, and CLAHE enhance local details or contrast, making diseased areas more distinguishable.
- Disabled techniques (MixUp, CutMix, Copy-Paste) are excluded as they may alter object shapes and mask boundaries, potentially reducing segmentation accuracy.

Model Training

The YOLOv8n-seg model was trained using a Tesla P100 GPU with batch size, number of epochs, and learning rate configurations that were carefully adjusted to maximize the utilization of limited resources while maintaining training stability. This process considered computational efficiency while ensuring the consistency of learning.

The dataset used was neatly separated into train and validation subsets according to the standard YOLOv8 format, complete with folder structure and metadata arrangements, and employed an image resolution of 640×640 pixels to guarantee consistent input. During the training process, the model's performance was continuously evaluated using a series of metrics that describe its object detection and segmentation capabilities. The main monitored metrics include box loss, class loss, and mask loss.

- Box loss represents the prediction error of object locations compared to their actual positions in the image. The smaller the box loss value, the more precise the model is in determining the bounding box coordinates.
- Class loss measures the degree of error in classifying objects into the correct class. This value is closely related to the model's ability to distinguish between object categories.
- Mask loss is used in segmentation tasks to assess the accuracy of predicted masks, so that the model not only identifies the location of objects but also accurately differentiates the boundaries of the object area from the background.

- b. In addition to loss, the model's performance was also evaluated using the mAP (mean Average Precision) metric, which represents the average precision across all classes at different recall values. For the case of mAP@50, the calculation follows the equation:

$$mAP@50 = \frac{1}{N} \sum_{i=1}^N AP_i, \text{ dengan } AP_i = \int_0^1 p_i(r) dr$$

where $p_i(r)$ is the precision function with respect to recall for i -th class, and NN is the total number of classes. The value of AP_i is obtained from the area under the precision– recall curve (AUC-PR) for each class. For mAP50–95, the calculation is performed by averaging the mAP values at Intersection over Union (IoU) thresholds ranging from 0.50 to 0.95, with an increment of 0.05. This metric provides a stricter evaluation since it accounts for varying levels of difficulty in detecting and localizing objects.

The training results are visualized in the form of loss curves and mAP curves, which facilitate real-time monitoring of the learning trends. The decreasing loss and increasing mAP trends indicate an improvement in model performance. As a preventive measure against overfitting, an early stopping strategy was applied by halting the training process when performance on the validation data showed a decline or stagnation. This approach not only saves computational time but also preserves the model's generalization quality on unseen data.

Results and Evaluation

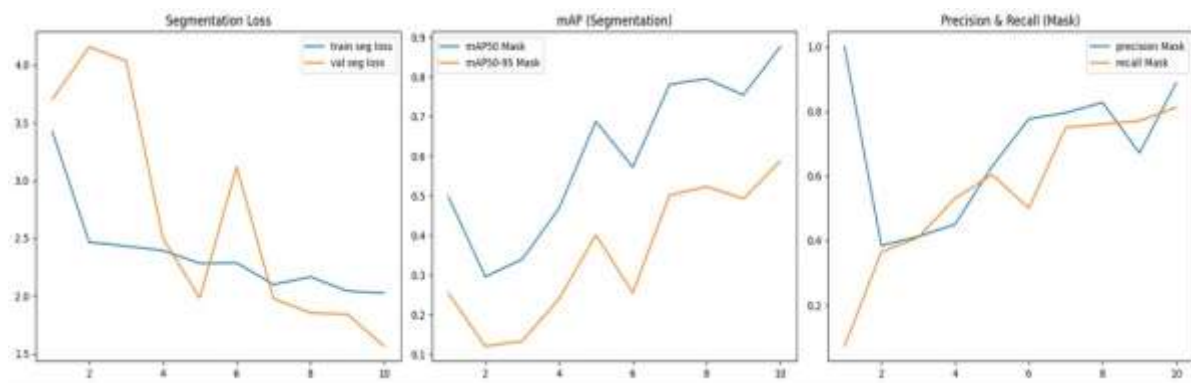


Figure 4. Segmentation Loss Curve, mAP Segmentation, and Precision & Recall Mask

The following figure presents the training curves of the model comprehensively, providing deep visual insights into the development of the model's performance over time, ranging from the efficiency of the learning process to the achievement of optimal evaluation metrics.

- Left Graph (Segmentation Loss): Shows a significant decrease in train loss and val loss, with an intensive learning phase at the beginning that gradually converges. The small gap between the two indicates minimal overfitting and stable learning.
- Middle Graph (mAP Segmentation): Both mAP@50 and mAP@50–95 consistently increase with each epoch. mAP@50 reflects detection accuracy with looser tolerance, while mAP@50–95 demonstrates robustness under stricter precision criteria. Their steady rise indicates overall improvement in segmentation performance.
- Right Graph (Precision & Recall Mask): Precision and recall steadily approach optimal values with parallel trends, showing a balance between minimizing false

positives and capturing relevant objects. This pattern reflects reliable segmentation performance even under diverse data conditions.

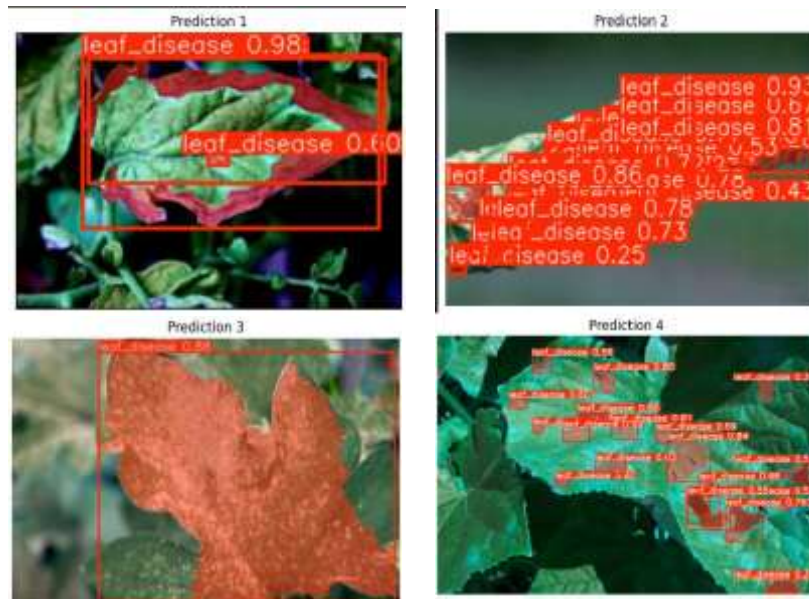


Figure 5. Example of Model Segmentation Prediction Results

- Prediction 1: The model successfully identified the diseased area with very high confidence (0.98) and produced precise segmentation, covering the entire infected region even though the background had a color similar to the target object. This highlights the model's ability to distinguish fine textures and subtle color patterns that separate healthy leaves from infected ones.
- Prediction 2: The model detected multiple disease instances with varying confidence levels, ranging from 0.25 to above 0.90. Several overlapping bounding boxes indicate potential duplicate predictions within the same area, possibly caused by irregular disease patterns or visually similar features in neighboring regions. This case emphasizes the challenge of handling duplicate predictions in instance detection models.
- Prediction 3: Displays segmentation of a large diseased region with confidence of 0.88, reflecting the model's capability to recognize disease patterns spread evenly across the leaf surface. This result also demonstrates the model's robustness to variations in leaf shape and size, as well as its consistency in segmenting larger infected areas.
- Prediction 4: Shows multiple detections in a single image across different disease region sizes, ranging from small spots to larger infected areas. Confidence values span from 0.34 to 0.85, indicating the model's sensitivity to varying levels of infection severity. This distribution suggests that the model can capture fine details while still detecting more significant regions, although further optimization is needed to improve confidence score consistency in small-scale areas.

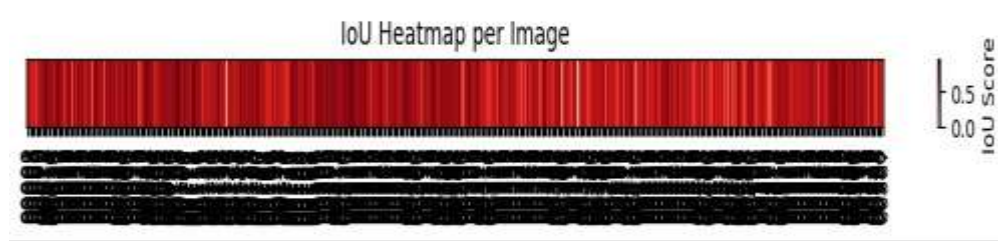


Figure 6. IoU Heatmap per Image

The IoU heatmap visualization shows the distribution of Intersection over Union scores across the test dataset. Each column represents an image, where darker red indicates higher IoU (strong mask–ground truth alignment), and lighter shades mark lower scores, often in cases with complex backgrounds or subtle disease patterns.

Most images achieve IoU above 0.5, confirming reliable segmentation, though variations reveal differing difficulty levels—such as small lesions, diverse patterns, or visual disturbances (e.g., shadows and reflections). This analysis highlights challenging cases and provides direction for model refinement, where low-scoring data can benefit from additional augmentation or retraining. The efficient YOLOv8n-seg architecture, enhanced with ProtoNet, ensures accurate spatial segmentation, robustness, and readiness for real-world implementation.

CONCLUSION

The implementation of the YOLOv8n-seg model with appropriate augmentation configurations and an efficient architecture has been proven capable of performing leaf disease segmentation with high and consistent accuracy, even on complex datasets. Evaluation using mAP, precision, recall, and IoU heatmap demonstrates stable performance under various lighting conditions and background complexities. The integration of ProtoNet enhances the precision of disease boundary delineation, while the application of in-training augmentation improves model adaptability and reduces overfitting. These findings indicate that the proposed approach holds strong potential for implementation in early plant disease detection systems, providing valuable support for decision-making processes in the agricultural domain.

REFERENCE

- A. Ansari *c" xl*, "Evaluating the effect of climate change on rice production in Indonesia using multimodelling approach," *Hcliyon*, vol. 9, no. 9, p. e19639, Sep. 2023, doi: 10.1016/j.heliyon.2023.e19639.
- F. R. Moeis, T. Dartanto, J. P. Moeis, and M. Ikhsan, "A longitudinal study of agriculture households in Indonesia: The effect of land and labor mobility on welfare and poverty dynamics," *World Dcv. Pcrspcc*., vol. 20, p. 100261, Dec. 2020, doi: 10.1016/j.wdp.2020.100261.
- G. Olaitan Onasanya *c" xl*, "Techniques of Using Peripheral Blood Mononuclear Cells as the Cellular System to Investigate How of the Bovine Species (Indian Zebu-Jersey Crossbreds) Responds to in vitro Thermal Stress Stimulation (Thermal Assault/Heat

- Shock),” 2024. doi: 10.5772/intechopen.109431.
- D. M. Rizzo, M. Lichtveld, J. A. K. Mazet, E. Togami, and S. A. Miller, “Plant health and its effects on food safety and security in a One Health framework: four case studies,” *Onc Hcxl. Ou''look*, vol. 3, no. 1, p. 6, Dec. 2021, doi: 10.1186/s42522-021-00038-7.
- G. Latif, S. E. Abdelhamid, R. E. Mallouhy, J. Alghazo, and Z. A. Kazimi, “Deep Learning Utilization in Agriculture: Detection of Rice Plant Diseases Using an Improved CNN Model,” *Plxn''s*, vol. 11, no. 17, p. 2230, Aug. 2022, doi: 10.3390/plants11172230.
- J. Liu and X. Wang, “Plant diseases and pests detection based on deep learning: a review,” *Plxn'' Mc''hods*, vol. 17, no. 1, p. 22, Dec. 2021, doi: 10.1186/s13007-021-00722-9.
- A. Y. Ashurov c''xl, “Enhancing plant disease detection through deep learning: a Depthwise CNN with squeeze and excitation integration and residual skip connections,” *Fron''. Plxn'' Sci.*, vol. 15, Jan. 2025, doi: 10.3389/fpls.2024.1505857.
- M. Shoaib c''xl, “An advanced deep learning models-based plant disease detection: A review of recent research,” *Fron''. Plxn'' Sci.*, vol. 14, Mar. 2023, doi: 10.3389/fpls.2023.1158933.
- X. Wang and J. Liu, “An efficient deep learning model for tomato disease detection,” *Plxn'' Mc''hods*, vol. 20, no. 1, p. 61, May 2024, doi: 10.1186/s13007-024-01188-1.
- M. J. A. Soeb c''xl, “Tea leaf disease detection and identification based on YOLOv7 (YOLO-T),” *Sci. Rcp.*, vol. 13, no. 1, p. 6078, Apr. 2023, doi: 10.1038/s41598-023-33270-4.
- E. A. Aldakheel, M. Zakariah, and A. H. Alabdall, “Detection and identification of plant leaf diseases using YOLOv4,” *Fron''. Plxn'' Sci.*, vol. 15, Apr. 2024, doi: 10.3389/fpls.2024.1355941.
- L. Li, S. Zhang, and B. Wang, “Plant Disease Detection and Classification by Deep Learning—A Review,” *IEEE Accss*, vol. 9, pp. 56683–56698, 2021, doi: 10.1109/ACCESS.2021.3069646.
- A. Kamilaris and F. X. Prenafeta-Boldú, “Deep learning in agriculture: A survey,” *Compu''. lcc''ron. Agric.*, vol. 147, pp. 70–90, Apr. 2018, doi: 10.1016/j.compag.2018.02.016.
- Y. Miao, W. Meng, and X. Zhou, “SerpensGate-YOLOv8: an enhanced YOLOv8 model for accurate plant disease detection,” *Fron''. Plxn'' Sci.*, vol. 15, Jan. 2025, doi: 10.3389/fpls.2024.1514832.
- P. Wang c''xl, “Leaf Segmentation Using Modified YOLOv8-Seg Models,” *Lific*, vol. 14, no. 6, p. 780, Jun. 2024, doi: 10.3390/life14060780.