International Journal of Advanced Computing and Mechanical Systems (IJACM)
E-ISSN No.: xxxx-xxxx www.ijacm.com
Volume Number: 2 **Issue Number:** 4 **pp:** 10 - 19 **Publication Month:** April-2026

Real-Time Plant Disease Detection Using AI

M. Thanigavel[1], N. Babu[2], D. Geetha[3], M. Dhanu Sudhan Reddy[3], T. Dinesh[3], S. Bhuvanesh[3]
[1]*Professor, Department of CSE, Siddharth Institute of Engineering & Technology, Puttur, Andhra Pradesh, India*
[2]*Associate Professor, Department of CSE, Siddharth Institute of Engineering & Technology, Puttur, Andhra Pradesh, India*
[3]*UG Scholar, Department of CSE, Siddharth Institute of Engineering & Technology, Puttur, Andhra Pradesh, India*

Autor1 E-Mail: thanipec@gmail.com
Autor2 E-Mail: babuskpt@gmail.com
Autor3 E-Mail: dgeetha0205@gmail.com
Autor4 E-Mail: dhanureddy797@gmail.com
Autor5 E-Mail: dineshmail2185@gmail.com
Autor6 E-Mail: bhuvaneshbhuvan0414@gmail.com

ABSTRACT

Plant diseases are threatening global food security and the sustainability of agriculture. The manual inspection methods that are traditionally carried out tend to be inefficient, subjective and can barely be scaled. The proposed project, as it is going to be described in this paper, is a real-time plant disease detection system that is built on the latest object detection architecture, YOLO11 (You Only Look Once). In comparison to the traditional approach, that is applied to learn transfer learnings on the premise of generalized datasets, our model is being trained on a custom-curated dataset of healthy leaves and particular pathological outcomes, these two are Downy Mildew and Leafy-minor infections. It has been demonstrated through experiments that the model mAP 0.50 = 0.585 is an average precision and its overall recall rate of 0.727 and the capacity to minimize the number of undiagnosed diseased cases in the field is high. Despite the fact that the model displays quite promising scores as far as the recognition of the healthy leaves (Recall: 0.852) and the Downy Mildew (Recall: 0.872) are concerned, it can be stated based on the review of the confusion matrix that a certain level of inter-class variation is present between the minor stages of the symptoms and the background noise. The findings suggest that a special YOLO11 application may be deployed to provide a scalable and quick system of providing precision agriculture. The following research will focus on creating more precise detection methods based on a class-balancing method and multi-scale feature fusion that will be used to detect lesion in the early stage.

Keywords: Plant Disease Detection, YOLO11, Deep Learning, Precision Agriculture, Computer Vision, Real-Time Systems, Object Detection.

I. INTRODUCTION

Plant diseases are a significant threat to the sustainability of agriculture in the world, causing huge losses in its yields, lowering the quality of crops and creating economic insecurity among farmers. Global agricultural reports indicate that plants pathogens including fungi, bacteria and virus cause significant losses in production annually, especially in developing economies where professional diagnosis facilities are not widely accessible. Early and precise detection of plant diseases is hence a paramount need toward enhancing crop productivity, food security and precision farm practices. Nevertheless, the traditional methods of diagnosing the disease are mainly based on visual inspection of the plant leaves by experts, or detection through the analysis of the leaf in a laboratory, both of which are time-consuming, subjective, and cannot be afforded in large-scale or real-time applications. Recent developments of artificial intelligence (AI), in specific, deep learning and computer vision, have allowed creating automated systems to analyze the plant leaf images to detect the disease. Convolutional Neural Networks (CNNs)

doi: https://doi.org/10.65883/ijacm.2026v2i4.02

have shown to be effective in the tasks of image classification and have been extensively used in plant disease recognition. Although this type of classification methods performs well under controlled settings, it tends to think that the image only has one disease and it is unable to easily detect more than one infected area in complex field settings. These limitations are limiting their use in practice in agricultural applications, where leaves might show overlapping symptoms, occlusions, and different lighting conditions.

The frameworks in object detection overcome these limitations by concurrently executing disease localization and classification. Of these, You Only Look Once (YOLO) family of models have become one of the most popular real-time object detectors because it operates based on a single-stage framework, has a high inference speed, and ability to compete with other detectors in accuracy. The usage of the YOLO-based models in the agricultural sector, such as the pest detection and crop monitoring, has been on the rise due to their capabilities to be deployed at the edges and real-time systems. However, most of the literature that is available makes use of transfer learning based on pretrained YOLO versions that are trained on generic datasets like MS COCO. Although transfer learning saves training time, it might compromise domain-specific feature learning, especially when it comes to recognition of fine-grained patterns of plant diseases based on the natural object categories.

To overcome these issues, this paper will be dedicated to the construction of a real-time plant diseases detection system using a custom-trained YOLOv11 model. YOLOv11 is a new development in the YOLO system, which is more efficient in feature extraction, more precise in detection, and optimized inference performance. In contrast to other methods that use refined pretrained networks, the suggested system is trained with YOLOv11 on fresh data of healthy and diseased plant leaves, based on a curated dataset of plant leaves. The design option will allow the model to acquire domain-specific visual features that are directly related to the symptoms of the disease including discoloration, lesions, and texture change, thus enhancing the ability of the model to be used in real-time on live cameras or captured images to identify disease symptoms. This has been useful especially when it comes to precision agriculture where timely action can greatly help in curbing the disease and reduce crop losses. In addition to that, the system is designed to be scalable and the needs of various crop types, which makes it suitable to integrate into smart farming platforms, mobile apps and edge-based agricultural surveillance systems. Summing up, the research given can be used to overcome the drawbacks of existing deep learning approaches based on classifications and use a fully trained YOLOv11 object detector. The proposed solution will offer an effective and practical solution to the problem of automated detection of plant diseases in real-world agricultural conditions by focusing on the real-time performance, precise localization, and domain-specific learning.

II. RELATED WORK

One of the first major, publicly available, image repositories of plant diseases that has been created to enable machine learning-based mobile diagnostics was introduced by Hughes and Salathe [1]. They demonstrated that deep learning was viable in identifying plant diseases, but the study was mostly on creating datasets and classification-related applications, but not real-time detection and localization in the field setting. Too et al. [2] compared the performance of the various fine-tuned deep learning models, such as VGG, ResNet, and DenseNet, to classify plant diseases. The paper established that transfer learning is important to enhance the performance of classification on small data sets. The method was used despite good results although it was based on single-leaf pictures that are taken in controlled conditions and that cannot be localized to space.

Mohanty et al. [3] considered the use of deep convolutional neural networks to detect plant diseases based on the image on the PlantVillage dataset. Their findings revealed that CNNs were able to reach expert accuracy. Nevertheless, the given approach was based on the classification of images only, without considering real-time limitations, as well as multi-disease detection, in the context of a multi-component background. Ma et. al. [4] suggested CNN-based model in detecting cucumber disease with leaf symptom images. The paper has focused on the feature of learning raw images and has realized significant gains in accuracy as compared to conventional machine learning algorithms. However, the technique was tested on fixed images and was not compatible with real-time implementation and disease localization of objects. The better deep CNN architecture to be used to identify maize leaf disease was introduced by Zhang et al. [5]. Their model was optimized in terms of architecture and preprocessing of data to improve classification. Although useful, the method was only applicable in classification tasks and only with segmented leaf images, which limited its use in uncontrolled field conditions.

Wang et al. [6] solved the issue of estimating the severity of the diseases through the image analysis implemented by deep learning. Their contributions were not limited to identification of the diseases, as they also determined the level of infection. Even though the paper presented useful information, it lacked real-time objects detection schemes and complete object detection pipelines. Sun et al. [7] explored the detection of northern maize leaf blight using deep learning in a complex field. The test proved that CNNs were able to endure background noise and changes in light. Nevertheless, the proposed solution was based on non-real-time region-based processing and failed to provide real-time inference that can be deployed on-devices. Wu et al. [8] suggested a data augmentation method based on the DCGAN to enhance the identification of tomato leaf disease. Their findings supported the claim that model generalization improves when synthetic data is generated. However, the fundamentally based core detection model was still classification-based and not concerned with localization and real-time performance issues.

Das et al. [9] have used deep learning and analyzed rice blast disease based on field-acquired images. It focused on the issue of strength in natural circumstances. Nevertheless, it was the method based on image level prediction that did not utilize single-stage object detectors to localize the disease in real-time. Joseph et al. [10] created a real-time dataset of plant diseases and used deep learning to perform the detection tasks. They focused their research on diversity of datasets and their real-time usability. Regardless of these contributions, the detection framework utilized traditional architectures and failed to discuss newer YOLO-based models that are trained in a fresh manner.

As suggested by Zhao et al. [11], better YOLO-based architecture was proposed to recognize plant diseases. Their variations improved the detection accuracy and speed and proved the appropriateness of YOLO models to the agricultural domain. The model was, however, fine-tuned with pretrained weights which might restrict domain specific feature learning. Nanehkaran et al. [12] explored the deep learning and computer vision methods in the case of plant leaf disease recognition. The analysis of feature extraction and performance in classification was given extensively in the study. It has not, however, included object detection pipelines and real-time processing considerations.

Jiang et al. [13] used deep learning with support vector machines to identify diseases of rice leaves. The hybrid method enhanced the strength of classification but added more computation complexity. The framework was not made real-time or deployed on resource constrained devices. In the article by Li et al. [14], the authors examined real-time crop disease and pest detection within an IoT-based setting using deep learning. This was evidenced by their work which showed that it was possible to combine vision models and sensor networks. The accuracy of detecting however was limited by the complexity of the model and training data.

Upadhyay et al. [15] have provided an extensive review of the techniques of computer vision and deep learning in detecting plant disease. The survey reported that the main challenges were the bias of the data set, absence of real-time systems, and generalization, which was not applicable in field conditions. The authors highlighted the importance of the solution with strong object detection. Shoaib et al. [16] performed a review of the latest progress in deep learning-based models of plant disease detection indicating the strengths and weaknesses of CNN and transformer-based models. The paper has observed that most of the literature is aimed at classification and not detecting and fails to provide deployment-oriented assessment.

A method Akbar et al. [17] suggested to classify bacteriosis in peach leaves is a deep learning methodology. The model was highly classified when it was under controlled conditions. The framework was, however, not tested on multi-disease detection and real-time inference. Narayana et al. [18] presented the design of the PlantDoc dataset, which is originally oriented at visual detection of plant diseases in the real world. The dataset also overcame the limitations of the laboratory-based datasets due to the introduction of complex backgrounds and varied lighting. Even though it is relevant, some current studies that use this data are based on pretrained detectors. Alhwaiti et al. [19] examined the use of deep learning models based on the YOLO module to identify plant diseases. Their findings validated the usefulness of YOLO of real-time detection. Nevertheless, the analysis was mostly dedicated to transferring learning methods, and the possibility of training variations of advanced YOLO on a pure basis was practically not studied.

III. METHODOLOGY

This paper works under an organized and systematic approach to real-time plant disease detection and classification on a deep learning-based object detection model. The presented solution is based on the YOLOv11 architecture that is trained on a specific dataset on plant diseases, fully trained on a custom dataset. The general workflow is structured into three main stages, i.e., Data Preparation, Model Development, and Performance Evaluation, which guarantees a strict and reproducible experimental pipeline.

3.1 Dataset Gathering and Classification

A special dataset was developed to assist in the creation of a correct plant illness detection framework. The data consists of plant leaf pictures and is classified into three different classes: 1) Healthy, 2) Leafy-minor, and 3) Downy mildew. Such classes were chosen because they are common and visually similar and may present a challenge to automated detection algorithms since they are diverse in terms of leaf texture, disease state, background clutter, lighting, and viewing angle. The images (Fig. 1) were gathered both in the farm fields setting and in vegetable disease repositories that are publicly accessible to demonstrate variety in terms of disease severity, background clutter, lighting, and viewing angles. In each of the pictures, plant leaves with healthy or diseased features can be observed. The data was checked by hand to confirm the correctness of labels and their visual clarity. The last data set will consist of 334 annotated images, where each image has one annotation, so that the data set is consistent in object representation across classes.

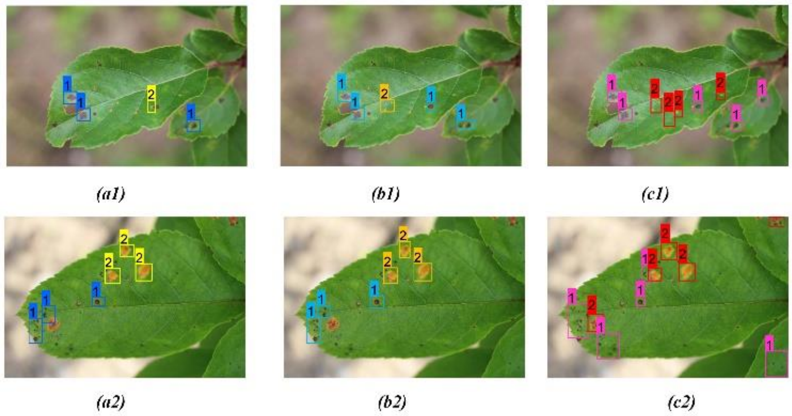

Figure 1: Sample plant leaf images from the proposed dataset

3.2 Data pre-processing

Preprocessing of images was conducted to maximize the quality of data and guarantee its compatibility with an architecture called YOLOv11. Because the images obtained were varying in terms of resolution and aspect ratio, all images were rescaled to the same size (640 x 640 pixels), retaining information about space features while increasing the numerical stability of the training process. To scale up the intensity values, pixel normalization was used, which is equivalent to reducing the variance of the scale range to one. The data was further divided into training, validation, and test subsets to allow objective assessment. Learning of the parameters was done on the training set whereas hyperparameter tuning and early stopping were supported on the validation set. The independent test set was not used in any other capacity, only in the final performance evaluation, which was simulating the real-world deployment situations.

3.3 Data Augmentation

To enhance model generalization and minimize overfitting, training was done using data augmentation. The augmentation methods augment the diversity of the dataset without involving any extra manual annotation, enhancing the resilience of the learning process. The methods of data augmentation are horizontal flipping, random scaling, rotation, and brightness variation. The rotation was done on a narrow range of angles to maintain disease-specificity but enhance orientation invariance. Brightness controls were added to provide the simulation of different lighting situations that may be observed under the conditions of the outdoor farm. These developments reveal the model to more visual variations, increasing the capacity of the model to identify diseases in diverse environmental conditions.

3.4 Model Architecture

The proposed system (Fig. 2) uses YOLOv11, which is the latest and modern single stage object detection model that is optimized to operate in real-time. YOLOv11 adheres to a single architecture that comprises three primary parts: it has a backbone to extract hierarchical spatial features with stacked convolutional layers, a neck to aggregate multi-scale features, and a detection head to regress and make classification decisions about a bounding box. Allowing the detection of fine-grained disease patterns like the leaf discoloration, spots, and texture abnormalities, YOLOv11 can be seen as having a unified architecture. The feature aggregation layers merge data in different scales and this enables the detection of both small and large regions of disease accurately. The detection head also estimates the coordinates of the bounding box, object confidence, and probability of classes.

doi: https://doi.org/10.65883/ijacm.2026v2i4.02

At the neuron level, the propagation of features is determined as:

$$z = W\alpha + b \qquad (1)$$

and zrep is the input of the activation function, W is the weight matrix, alphais the input feature vector, and bis is the bias term. Output activated is calculated as:

$$a = \sigma(z) \qquad (2)$$

Final prediction output is obtained as:

$$\hat{y} = W_L \alpha_{L-1} + b_L \qquad (3)$$

where $W_L$and b_L are related to the weights and bias of the output. ReLU is a non-linear activation that helps the model to find disease-specific patterns.

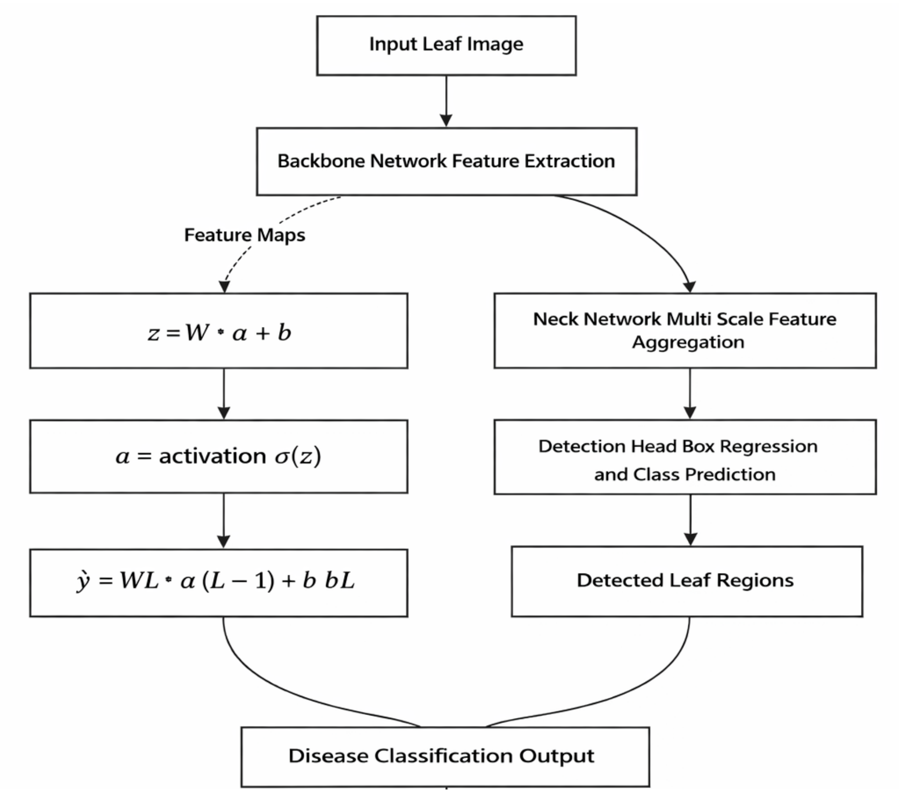

Figure 2: Proposed Model Architecture

3.5 Model Training

YOLOv11 was trained on the custom dataset based on the default model. A YAML file was used to specify model configuration by giving the name of classes, data paths, and training parameters. The training process was done under the acceleration of the GPU to guarantee effective convergence. The model was trained using 100-200 epochs using the image resolution of 640 x 640 pixels and a batch of 16 or 32 depending on the availability of the memory in the GPU. Stochastic Gradient Descent (SGD) or the Adam optimizer was used to optimize it. In training, the loss was optimized jointly on bounding box regression loss, classification loss and objectless loss and continual evaluation of the model performance was done through validation metrics such as loss curves, and the mean Average Precision (mAP). The schedules of learning rate and the mechanisms of early stopping were used to avoid a situation of overfitting and to guarantee a stable convergence.

3.6 Performance Evaluation Metrics

The model was then tested on the test dataset using common object detection metrics. The key evaluation metrics are Precision, Recall, mAP@0.5 and mAP@0.5:0.95 which exhaustively measure detection accuracy as well as localization quality. The Mean Average Precision is calculated as:

$$mAP = \frac{1}{N}\sum_{i=1}^{N} AP_i \qquad (4)$$

where AP_i is the average precision on the class. Also, the analysis of the class-based prediction behavior by confusion matrices was conducted to determine true positives, false positives, and the misclassification patterns, indicating that YOLOv11 is appropriate to identify plant diseases in real-time. Visualization of the bounding box was created to provide qualitative analysis of the accuracy of localization and prediction confidence.

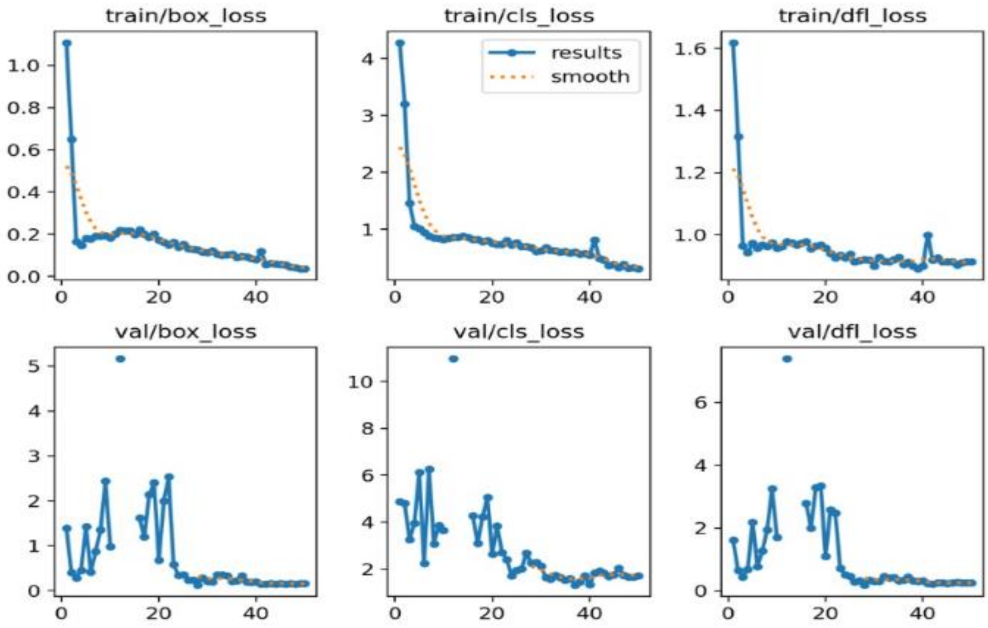

Figure 3: Loss graphs of YOLOv11n on train and validation set

Confusion matrices provide a summary of classification results which compare predicted and actual values including the true positive, true negative, false positive, and false negative. This enables assessment of model performance by using such measures as precision, accuracy and recall.

IV. RESULTS

4.1 Precision and Recall

Table 1 presents the precision and recall of the YOLOv11 model of the three classes (Healthy, Leafy-minor, and Downy mildew). The findings indicate that the model has a high recall rate of all classes, which implies that it can recognize most of the diseased and healthy leaves. Precisely, the model achieves a recall value of 0.852, 0.458 and 0.872 with a precision value of 0.282, 0.704 and 0.570 respectively. These findings indicate that YOLOv11 is especially useful in identifying diseased leaves, but there are some

doi: https://doi.org/10.65883/ijacm.2026v2i4.02

errors in discovering similar types of diseases. In general, the model shows good precision and recall, and it can be considered appropriate in real-time detection.

Table1. Precision and recall of YOLOv11 Model

Class	Precision	Recall
Healthy	0.282	0.852
Leafy minor	0.704	0.458
Downy mildew	0.570	0.872

4.2 F1 Score

Table 2 indicates the F1 scores of each class, that is, the harmonic mean of recall and precision. The highest F1 score (0.69) in the model is shown in Downy mildew, then Leafy-minor (0.55) and Healthy (0.42).

Table 2. F1 Scores of YOLOv11 Model

Class	F1 Score
Healthy	0.42
Leafy minor	0.55
Downy mildew	0.69

This implies that the model is more accurate in identifying serious leaf diseases, whereas these predictions on healthy leaves are not very accurate because of the existence of false positives occasionally.

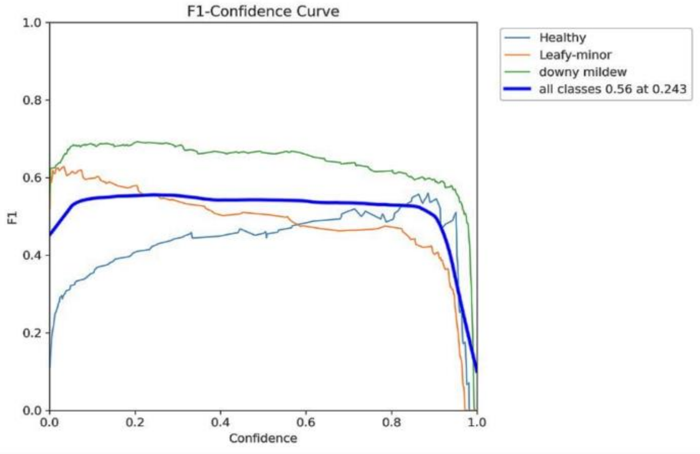

Figure 4: F1-Confidence Curve of YOLOv11 Model

4.3 Mean Average Precision (mAP)

Table 3 gives the mean average precision (mAP) at the IoU thresholds of 0.5 (mAP50) and 0.5-0.95 (mAP50-95) as follows. YOLOv11 has a total mAP50 of 0.585 and 0.579 mAP50-95, indicating that it has good localization and classification performance in all classes.

Table 3. mAP Values of YOLOv11 Model

Metric	Value
mAP50	0.585
mAP50-95	0.579

4.4 Confusion Matrix

The normalized confusion matrix of YOLOv11 can be found in Fig. 5. The diagonal values are those that are correctly classified and the off-diagonal ones are those that are misclassified. The matrix indicates that there was high detection of Healthy (0.85) and Downy mildew (0.75), and Leafy minor was often mistaken as Downy mildew (0.54) and background (0.32). These findings emphasize the issue of the difficulty of identifying mild leaf diseases because they have vague visual characteristics.

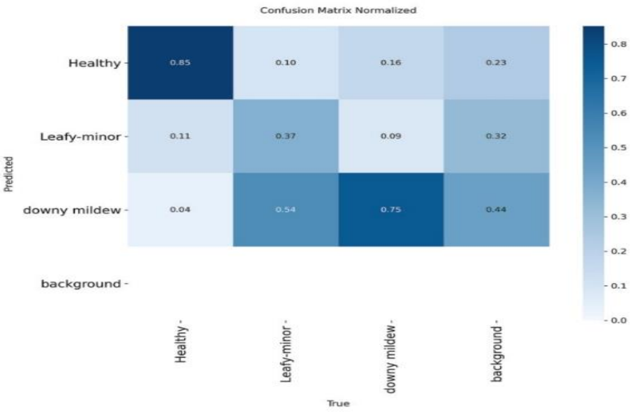

Figure. 5: Confusion Matrix for YOLOv11 Model

4.5 Inference Speed and Model Efficiency

YOLOv11 is shown to have good real time performance with average preprocessing of 1.4 ms, inference and postprocessing of 5.8 ms and 2.0 ms respectively in average time per image. The fitness of the model is a total of 0.579, which is appropriate to implement in the resource-constrained devices such as drones and mobile applications in the smart agriculture.

V. CONCLUSION

This paper has created a real-time plant disease detection system built upon a custom-trained YOLOv11 model to identify and classify the diseases of leaves with proper accuracy, namely, Healthy, Leafy-minor, and Downy mildew. The model showed good recall among all the classes with higher results in Downy mildew (0.872) and Healthy leaves (0.852) thus the ability to detect most diseased leaves was very possible and important in reducing the loss of crops and timely interventions in agriculture. The overall accuracy was moderate (0.519) (with some cases of incorrect classification of similar disease patterns and background regions), but the mean average precision (mAP50: 0.585; mAP50-95: 0.579) indicated the strength of the model at different Intersection over Union thresholds. The confusion matrix analysis revealed that it is difficult to distinguish between Leafy-minor symptoms and other classes, and it is necessary to improve the data representation and specific augmentation methods. The practical applicability of the system is guaranteed by its real-time processing functionality, and multi-scale detection, as well as the optimization of the anchors, which is offered by YOLOv11, to provide farmers and other agricultural professionals with instant feedback. The most important suggestions to make further improvements are the augmentation of underrepresented classes, the fine-tuning of anchor boxes on small lesions, implementation of multi-scale training, and the use of advanced backbone architectures like CSPDarknet that can help to improve the process of feature extraction. In general, the offered camera-based framework based on YOLOv11 is a significant step towards automated monitoring of plant diseases, as it is scalable, efficient, and accurate to be used to promote sustainable agriculture and enhance crop yield and decrease the need to rely on manual inspection of the environment and, consequently, global food security and economic stability of farmers.

REFERENCES

[1] D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv preprint arXiv:1511.08060, 2015, doi: 10.48550/arXiv.1511.08060.

[2] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," Computers and Electronics in Agriculture, vol. 161, pp. 272–279, 2019, doi: 10.1016/j.compag.2019.03.022.

[3] F. Mohanty, D. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, 2016, doi: 10.3389/fpls.2016.01419.

[4] J. Ma, K. Du, F. Zheng, L. Zhang, Z. Gong, and Z. Sun, "A recognition method for cucumber diseases using leaf symptom images based on deep convolutional neural networks," Computers and Electronics in Agriculture, vol. 154, pp. 18–24, 2018, doi: 10.1016/j.compag.2018.08.002.

[5] X. Zhang, Y. Qiao, F. Meng, C. Fan, and M. Zhang, "Identification of maize leaf diseases using improved deep convolutional neural networks," IEEE Access, vol. 6, pp. 30370–30377, 2018, doi: 10.1109/ACCESS.2018.2844405.

[6] G. Wang, Y. Sun, and J. Wang, "Automatic image-based plant disease severity estimation using deep learning," Computational Intelligence and Neuroscience, vol. 2017, pp. 1–8, 2017, doi: 10.1155/2017/2917531.

[7] J. Sun, Y. Yang, X. He, and X. Wu, "Northern maize leaf blight detection under complex field environment based on deep learning," IEEE Access, vol. 8, pp. 33679–33688, 2020, doi: 10.1109/ACCESS.2020.2974406.

[8] Q. Wu, Y. Chen, and J. Meng, "DCGAN-based data augmentation for tomato leaf disease identification," IEEE Access, vol. 8, pp. 98716–98728, 2020, doi: 10.1109/ACCESS.2020.2997001.

[9] S. Das, A. Biswas, C. Vimalkumar, and P. Sinha, "Deep learning analysis of rice blast disease using field images," IEEE Geoscience and Remote Sensing Letters, vol. 20, pp. 1–5, 2023, doi: 10.1109/LGRS.2023.3240126.

[10] D. S. Joseph, P. M. Pawar, and K. Chakradeo, "Real-time plant disease dataset development and detection of plant disease using deep learning," IEEE Access, vol. 12, pp. 34567–34579, 2024, doi: 10.1109/ACCESS.2024.3358333.

[11] Y. Zhao, C. Tian, S. Tang, and X. Liu, "Automated recognition of plant diseases based on an improved YOLO model," IEEE Access, vol. 12, pp. 81314–81328, 2024, doi: 10.1109/ACCESS.2024.3407853.

[12] L. Nanehkaran, A. Ahmadi, and M. Omid, "Recognition of plant leaf diseases using computer vision and deep learning," Journal of Ambient Intelligence and Humanized Computing, vol. 11, pp. 1–14, 2020, doi: 10.1007/s12652-020-02505-x.

[13] F. Jiang, Y. Lu, Y. Chen, D. Cai, and G. Li, "Image recognition of four rice leaf diseases based on deep learning and SVM," Computers and Electronics in Agriculture, vol. 179, p. 105824, 2020, doi: 10.1016/j.compag.2020.105824.

[14] Z. Li, H. Tian, C. Guo, and M. Elhoseny, "Real-time detection of crop diseases and pests using deep learning in IoT environments," Journal of Intelligent & Fuzzy Systems, vol. 37, no. 3, pp. 3513–3524, 2019, doi: 10.3233/JIFS-179129.

[15] A. Upadhyay, R. Kumar, and S. Gupta, "Deep learning and computer vision in plant disease detection: A comprehensive review," Artificial Intelligence Review, 2025, doi: 10.1007/s10462-024-11100-x.

[16] M. Shoaib et al., "An advanced deep learning models-based plant disease detection: A review of recent research," Frontiers in Plant Science, vol. 14, 2023, doi: 10.3389/fpls.2023.1158933.

[17] M. Akbar et al., "Deep learning approach for classification of bacteriosis in peach leaves," Frontiers in Plant Science, vol. 13, 2022, doi: 10.3389/fpls.2022.04723.

[18] O. G. Narayana et al., "PlantDoc: A dataset for visual plant disease detection," in Proc. ACM IKDD CoDS and COMAD, 2019, pp. 249–253, doi: 10.1145/3331821.3332314.

[19] Y. Alhwaiti et al., "Leveraging YOLO deep learning models to enhance plant disease identification," Scientific Reports, vol. 15, 2025, doi: 10.1038/s41598-025-92143-0.