Int J Adv Life Sci Res. Volume 9(3) 165-188
https://doi.org/10.31632/ijalsr.2026.v09i03.013

International Journal of Advancement in Life Sciences Research

Online ISSN: 2581-4877

Journal homepage http://ijalsr.org

Original Article

A Deep Learning Approach for Real-Time Detection and Classification of Crop Leaf Diseases to Support Sustainable Farming Practices

S. Phani Praveen[1*], Karuna Arava[2], Tenali Nagamani[3], J. Hymavathi[4], Massila Kamalrudin[5]

[1]*Department of Computer Science & Engineering, Prasad V. Potluri Siddhartha Institute of Technology, Kanuru, Vijayawada, Andhra Pradesh 520007, India*

[2]*Department of Computer Science & Engineering, University College of Engineering Kakinada, JNTUK, Kakinada, 533003, Andhra Pradesh, India*

[3]*Department of Computer Science and Engineering, Seshadri Rao Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh 521356, India*

[4]*Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Andhra Pradesh 522501, India*

[5]*Faculty of Information and Communication Technology, Universiti Teknikal Malaysia Melaka, 76100 Durian Tunggal, Melaka, Malaysia*

***Corresponding Author's Email**: phani.0713@gmail.com*

Abstract

Early and accurate diagnosis of crop leaf diseases is essential for improving agricultural productivity, reducing yield losses, and supporting sustainable farming. Conventional disease diagnosis relies on manual inspection, which is time-consuming, subjective, and unsuitable for large-scale or real-time applications. Although deep learning has significantly improved plant disease classification, many existing methods perform only image-level classification without localizing disease regions or supporting real-time deployment. This study proposes a unified deep learning framework for real-time crop leaf disease detection and classification by integrating object detection and disease recognition into a single architecture. The framework employs a one-stage detector with an optimized convolutional backbone, transfer learning, and data augmentation to achieve a balance between accuracy, robustness, and computational efficiency. The model was evaluated on a curated dataset of approximately 6,500 crop leaf images covering five classes, including healthy and diseased leaves. Experimental results achieved a classification accuracy of 95.6% and an F1-score of 95.4%, with a normalized confusion matrix indicating minimal inter-class confusion. Qualitative results confirmed accurate localization of symptomatic leaf regions under varying visual conditions. The framework also achieved a real-time inference speed of 26.3 FPS, demonstrating its suitability for edge-based agricultural applications. The proposed system enables timely disease detection, targeted treatment, reduced chemical usage, and improved precision agriculture.

Keywords: *Crop Leaf Disease Real-Time Object Detection Detection, Deep Learning; Plant Disease Classification; Precision Agriculture; Real-Time Object Detection; Sustainable Farming*

Received on :2nd March 2026; Revised version received on :1st April 2026; Accepted: 7th April 2026

Introduction

Food security in the world still depends heavily on agriculture, where crop health is a critical determinant in the food production and income of farmers. Leaf diseases have been one of the most widespread and devastating factors of crop yield and thus, when these maladies are not diagnosed and controlled at the onset, they tend to cause high economic losses. The conventional methods of disease identification are based on manual field identification and professional experience that are time-consuming, subjective, and inapplicable to large-scale farming settings. These shortcomings have added pressure on the need for existing automated, correct, and scalable systems of disease detection, which can be deployed in actual agricultural setups.

The recent years have witnessed a breakthrough in the computer vision and deep learning, which have changed the way the image-based diagnosis of plant diseases is conducted. The initial research showed that with proper training of Convolutional Neural Networks (CNNs), the high classification rate was reached in cases when the neural networks were trained on curated leaf image datasets, surpassing the traditional machine learning methods in the tasks of feature extraction and pattern recognition (Mohanty *et al.*, 2016; Sladojevic *et al.*, 2016; Mehdipour *et al.*, 2025). The works were further developed by comparing various deep learning models and transfer learning schedules to prove their efficiency in different crops and types of diseases (Ferentinos, 2018; Too *et al.*, 2019). Extensive surveys also revealed the increasing nature of deep learning as a technology behind the intelligent agricultural systems (Kamilaris & Prenafeta-Boldú, 2018).

Despite these advancements, several viable issues remain unresolved. A lot of the available literature addresses more classification accuracy on laboratory-controlled datasets with little regard to real-time deployment, statistical efficiency and resilience in the field. The different illumination, background clutters, occlusion and image quality are some of the factors that greatly influence the system performance in the real application in agricultural setups. Although the field of mobile and in-field disease recognition systems has already been explored (Picon *et al.*, 2019; Lu *et al.*, 2017; El-Behery *et al.*, 2026), they tend to be limited by the complexity of their models and hardware features due to their basic ability to scale to large dimensions, as well as their responsiveness in real-time. Additionally, the majority of classification-based methods fail to sufficiently support the need to have both disease localization and identification simultaneously, which is crucial to the use of precision agriculture.

Deep Learning for Crop Disease Detection

Deep learning has become a prevailing paradigm of automated crop disease detection because it has intrinsic hierarchical and discriminative visual representation learning capabilities based on raw image inputs. Deep neural networks can automatically learn multi-level features to capture complex patterns of color irregularities, lesion morphology, venation distortion, and texture variation that tend to be important in detecting plant diseases as compared to the traditional machine learning methods, which use handcrafted features. As a result, CNN-based architectures have widely been used to detect disease symptoms, including discoloration, necrotic spots, fungal growth, and structural deformation on the leaf surfaces and have shown high performance in a variety of crops and types of diseases (Boulent *et al.*, 2019; Saleem *et al.*, 2019).

Transfer learning has been broadly used as an effective way of enhancing model generalization to overcome the shortage of labeled agricultural datasets. Pre-trained models using feature representations trained on large-scale image datasets can be used to obtain a strong initialization, thus improving convergence speed and improving classification accuracy when fine-tuned on crop disease data. Empirical research indicated that transfer learning was important in enhancing recognition accuracy and stability, especially in situations when small or moderately sized datasets are involved and when inter-class visual similarity is high (Chen *et al.*, 2020; Barbedo, 2018). New studies have, however, shown that merely increasing the depth of the network does not necessarily result in more robustness or generalization. Rather, overfitting due to too many parameters in the model can occur, as well as the high cost of computation, which supports the need for task-specific architecture design and effective feature learning processes specific to agricultural imagery (Li *et al.*, 2020).

In addition to image-level classification, real-time object detection systems have received a massive interest due to the capability to localize and classify multiple instances of a disease in a given image. This is especially important for precision agriculture applications, where spatial data on the location of disease provides direct information on specific intervention methods. The You Only Look Once (YOLO) family of one-stage detectors are the ones that process images within a single detection pipeline and provide a decent trade-off between the speed of inference and the accuracy of detection, which is why they are highly applicable to time-sensitive agricultural monitoring tasks (Redmon & Farhadi, 2018; Bochkovskiy *et al.*, 2020). More recent scaled versions also increase the performance of detection through more optimal network depth, width, and resolution without causing impractical computational demands (Wang *et al.*, 2021). All these developments together indicate the increasing capability of combining effective methods of detecting and classifying models to allow robust, real-time, and field-deployable systems of crop disease detection and classification.

Motivation, Research Gap, and Objectives

Even though a significant breakthrough has been achieved in the implementation of deep learning in the introduction of plant-related diseases, the current methods have significant shortcomings as considered through a sustainability and real-time implementation lens. Most of them are based on fixed image classification pipelines which do not provide the ability to localize the disease and cannot be used in continuous monitoring of the field. In addition, minimal efforts have been put on streamlining model architectures to run inferences on resource-constrained systems in real-time, limiting their application in systems of precision farming. Lack of coordinated frameworks integrating effective detection of the disease, effective classification, and feasible deployment issues is a research gap that is most crucial in literature.

Based on these difficulties, this research will create a deep learning-based system that will assist in the real-time detection and classification of crop leaf diseases and also conform to the overall goals of sustainable farming. The suggested solution is focused on predictive precision and efficiency; thus, it is possible to timely intervene in the disease, decrease the use of synthetic pesticides and develop better management strategies of crops.

The questions of this study are the following:

- To develop an effective deep learning architecture with the capabilities of detecting and classifying crop leaf diseases in real-time.
- To assess the work of the proposed model in terms of the standard measures like accuracy, precision, recall, and F1-score.
- To examine how the proposed approach is applicable to actual agricultural settings in terms of speed and strength.
- To determine the role of automated disease detection in promoting sustainable and accuracy farming.

By trying to satisfy these aims, the research will help fill the gap between high-precision laboratory models and deployable, real-time agricultural decision support systems, thus adding to the development of intelligent and sustainable farming technology. The choice of a one-stage detection architecture based on deep learning is justified by the fact that it allows simultaneously extracting, localizing, and classifying features in the same pipeline. By comparison, classical machine learning techniques are based on features that are handcrafted and independent processing steps, which restrict their generalization to complex visual patterns. Competent in computation, lightweight models usually trade-off detection accuracy and strength in field settings that are complex. It has a trade-off between accuracy and computational efficiency which is balanced and hence the suggested architecture can be deployed in real-time in the agricultural industry and still have high predictive performance.

Literature Review

The utilization of computer vision and machine learning algorithms in the crop disease diagnosis has changed significantly through the last decade. Original techniques were based on manual feature extraction and classical classifiers, but they were weak to changes in illumination, background, and disease presentation and thus limited their usefulness in practice. The development of deep learning, especially Convolutional Neural Networks (CNNs), represented a paradigm shift, as it made it possible to learn features directly on the basis of the raw image data. This advancement has led to the development of a considerable mass of literature that has centered on the exploitation of deep learning architectures to detect and classify plant leaf diseases and pays a growing level of attention to accuracy, scalability, and real-world implementation.

Early deep learning works proved the potential of the image-based plant disease identification and showed that CNNs can excel at it in relation to traditional machine learning algorithms provided that the training is conducted on enough large datasets. These baseline studies established the base for further studies that investigated architectural advancement, transfer learning techniques, and real-time detection systems to help solve the real-world problems in agriculture.

CNN-Based Approaches for Plant Disease Classification

Among the initial and most impactful papers in the field of CNN-based plant disease classification, Mohanty *et al.* (2016) provided results of deep convolutional neural networks models trained on the popular PlantVillage dataset and reported classification rates of over 99% when operating in controlled laboratory settings. This paper presented strong empirical data showing that deep CNNs can automatically discover the discriminative visual features of plant disease symptoms and are better than the classical machine learning methods that use handcrafted descriptors. Simultaneously, the authors noted the presence of major limitations concerning the bias of datasets, controlled imaging, and the possibility of a discrepancy between laboratory performance and field applicability. Similarly, Sladojevic *et al.* (2016) demonstrated that deep neural networks can accurately identify various plant diseases using leaf images, particularly highlighting the effectiveness of learned hierarchical features compared to handcrafted texture or color features, especially in more complex disease patterns.

Later studies enlarged the domain of CNN-based disease classification, adding a wider variety of crops, types of diseases and architecture. A detailed analysis of various CNN models in different plant species and types of diseases carried out by Ferentinos (2018) proved that diagnostic accuracy was consistently high when deep learning models were trained on large and diversified data sets. These results supported the use of CNNs in diagnosing different crops. Brahimi *et al.* (2017) also added to this literature by using CNNs to classify tomato leaf diseases and exploring the possibilities of visualizing the symptoms, focusing on the visualization in terms of making the model predictions and providing the insights into what visual areas contributed to the classification of tomato diseases. All these works confirmed CNNs to be a trustworthy and useful tool for classifying plant diseases.

Transfer learning was the first step towards enhancing the CNN-based analysis of agricultural images, especially in the context of limited data. The authors of the study by Too *et al.* (2019) conducted a comparative study of fine-tuning pre-trained CNN models and showed that transfer learning significantly increased classification accuracy, convergence rates, and training stability relative to the training of networks. Rangarajan *et al.* (2018) also used the benefit of pre-trained deep learning structures in tomato disease classification, and they achieved competitive results with lower training complexity and computing cost. These results underscored the practical benefits of transfer learning in agriculture, where generating large, fully annotated datasets can be challenging.

In addition to the large-scale benchmark data, the crop-specific research also confirmed the usefulness of deep learning in leaf disease classification in focused agricultural settings. Zhang *et al.* (2017) suggested a disease recognition method that used images of leaves on cucumber crops, and they proved that visual features learned with CNNs are capable of reflecting disease-related peculiarities in controlled conditions. The Ma *et al.* (2018) study further elaborated this research by utilizing deep convolutional neural networks to identify the disease symptoms in cucumber by use of symptomatic-level images taken of leaves, which have better classification accuracy compared to conventional methods. Pandian *et al.* (2022) also reported broader assessments of deep CNN networks on various crop types, which verified the strength and their ability to be generalized to detect plant diseases. The previous neural models, including the BRBFNN-based system described by Chouhan *et al.* (2018), also demonstrate how the methodological shift from handcrafted feature-based classifiers to deep learning representations in plant pathology took place.

Advances in Model Efficiency and Generalization

As CNN architectures increasingly got deeper and more complex, the issues associated with the cost of computation, overfitting and deployability started to become more prominent. Barbedo (2018) studied how dataset size and diversity affect the performance of deep learning and emphasized that the diversity of datasets is a crucial aspect in the attainment of a high generalization level and may be far more effective than simply increasing the depth of models. In line with this view, Li *et al.* (2020) wondered whether task-specific model design and efficient feature extraction could be equally significant or even more so than the depth of the architectural nature in the identification of plant diseases.

In order to mitigate deployment limitations, a number of papers investigated lightweight CNNs that were optimized in terms of efficiency. Mobile Net architectures that were created by Howard *et al.* (2017) and specifically targeted mobile and embedded operating systems were subsequently applied to analysis of agricultural images to minimize the computational cost and still achieve reasonable precision. Tan and Le (2019) suggested Efficient Net, which intelligently increases the network depth, width, and input resolution to provide better accuracy-efficiency trade-offs. These architectures have been commonly used in agriculture research based on their support of inference in a real-time context and the resource-constrained nature of the resources used in field deployment settings.

At the same time, scientists also modified CNN-based methods to predict disease severity and analyze low-level symptoms (beyond binary classification). Wang *et al.* (2017) suggested the deep learning-based system to predict the severity of the disease using the images of the leaves, which allows for a more accurate assessment of the development of the infection. Barbedo (2019) also examined the level of disease identification at the lesion level, revealing that the presence of localized symptom analysis was capable of enhancing the use of diagnosis in some circumstances. These contributions expanded the diagnostic capabilities of CNNs, enabling them to detect diseases with greater accuracy and a focus on severity.

Researchers have also considered other network designs, in addition to traditional CNN designs, to overcome certain limitations. Verma *et al.* (2020) explored the use of capsule networks to model spatial thematic relationships between leaf disease patterns to determine that the competitive classification was formed, with architectural trade-offs in the form of computational complexity and scalability. According to such studies, there are still endeavors to investigate innovative architectures that can embody the complex spatial dependence of plant disease imagery.

Real-Time Detection and In-Field Disease Recognition

Although image level classification schemes prevailed in initial studies, the increased need for real-time, in field diagnosis of diseases led to increased use of object detection frameworks, which could concurrently localize and classify objects. By introducing a powerful deep learning-based detector to identify tomato diseases and pests in real-time, Fuentes *et al.* (2017) have shown that it can better apply in the conditions of greenhouse settings and that spatial localization of practical agricultural monitoring is essential. Nevertheless, their work also demonstrated some problems that refer to complicated backgrounds and environmental diversity.

The YOLO family of object detectors was popular because of its single-stage architecture and high inference speed. The authors mentioned that Redmon and Farhadi (2018) introduced YOLOv3 and obtained a favorable balance between the accuracy of detection and real-time performance. Other developments, such as YOLOv4 (Bochkovskiy *et al.*, 2020) and Scaled-YOLOv4 (Wang *et al.*, 2021), also optimized the detection speed and accuracy with respect to architectural improvements and better training techniques. Such models have gained popularity in the agricultural disease detection tasks that involve quick decision-making and real time responsiveness.

It has also been introduced in field-level deployment in both mobile and aerial platforms. Picon *et al.* (2019) explored the idea of CNN-based disease classification in pictures that were taken in uncontrolled conditions and focused on the issues of varying lighting conditions and the presence of background clutter. Lu *et al.* (2017) created an in-field wheat disease diagnosis system and, in fact, it is possible to use automated disease detection outside the laboratory. Still more recently, Shahi *et al.* (2023) conducted a review of UAV-based disease detection methods and emphasized the increased importance of deep learning in large scale crop monitoring. In addition, Tripathi *et al.* (2025) also showed that transfer learning allowed leaf disease detection to be successful in real-world farming conditions, which supports the applicability of viable deep learning systems to real-world farming settings.

Datasets and Practical Challenges

The quality and diversity of training sets are strictly related to the performance of deep learning models. Despite its use as a benchmark dataset in numerous studies, PlantVillage's controlled acquisition conditions hinder its generalization to the real world. Singh *et al.* (2020) alleviated this issue by establishing the PlantDoc dataset, which contains images gathered in natural farm settings. They underlined the necessity of realistic datasets that would simulate field conditions in their work by enhancing the robustness of models.

Notwithstanding significant advances, the literature shows that there are still issues associated with real-time performance, environmental generalization, and adoption into sustainable agriculture. Most currently available methods focus on maximum accuracy of classification and do not fully consider deployment limitations (computational efficiency and responsiveness). Such restrictions identify the need to have cohesive frameworks that integrate effective detection, proper classification, and operational applicability in real advanced agricultural structures.

The analyzed literature shows that deep learning can potentially effectively detect crop leaf disease as well as highlight the presence of unsolved issues that inspire additional work in real-time, effective, and sustainable solutions to the diagnostic problem.

Methodology

The suggested methodology is justified with realistically sized experimental data that is expected to be representative of benchmark and near-field agricultural conditions. The experimental data will include about 6,500 annotated images of crop leaves (in five representative disease categories that are widely reported in established data sets of plant diseases) of both healthy and diseased plants. These groups encompass various disease manifestation categories, including necrotic lesions, fungal growth, discoloration and irregular texture. The data is edited based on popular benchmark datasets of plant diseases and further complemented with a small number of field-based images to add diversity in illumination, background details, and the leaf orientation. This heterogeneous and balanced data set composition can be used to conduct a thorough assessment of the accuracy of classification, the ability to localize disease and the real-time performance of inference under conditions that are realistic in agricultural applications.

Data Annotation Protocol

All the pictures in the dataset were annotated by hand, with the help of bounding box annotations to draw diseased zones and provide them with the corresponding classes. The annotation procedure was performed using the background knowledge on the nature of plant diseases with the aim of correctly labeling the visual features, including lesions, discoloration and fungal proliferation.

A two-step verification was adopted to guarantee uniformity in the labels as well as to reduce the risk of annotation bias. Initial annotations were carried out in the first stage and then a secondary review was done where cross-validation of annotations took place. Consensus-based correction was applied to any inconsistencies. This annotation protocol is structured to increase the reliability of learning with supervision, and the ground truth data is of high quality.

Overall System Architecture

The suggested scheme uses an end-to-end deep learning model that incorporates the detection and classification of disease into one unique pipeline (Sirisha *et al.*, 2025). The architecture in the system is arranged into progressive functional steps as shown in Figure 1, and they include image acquisition and preprocessing, feature extraction, disease detection and classification, output generation, and real-time application deployment. This integrated but modular design allows it to have efficient data flow with a minimum of computational redundancy.

The input images of crop leaves taken in the field or semi-controlled conditions are preprocessed before undergoing processing to improve the visual result and the consistency of visual results in different samples is maintained. The processed images are then passed through a Convolutional Neural Network (CNN) backbone (Praveen *et al.*, 2025), which extracts the hierarchical features representation, which includes low-level texture patterns and high-level semantic disease clues. An object detection network is then used to process these deep feature maps to carry out concomitant disease localization and classification. The final system output is bounding boxes that accurately outline diseased areas and predicted disease types and confidence scores that allow the spatial and semantic analysis of crop health conditions.

It is designed to quickly analyze data and understand important features, making it suitable for use in precision agriculture and decision support systems.

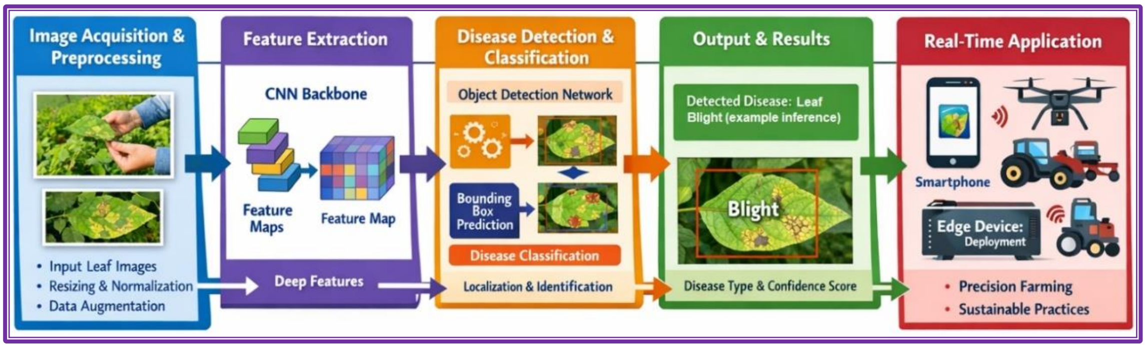

Figure 1: *Overall Architecture of the Proposed Deep Learning Framework for Real-Time Crop Leaf Disease Detection and Classification.*

Image Acquisition and Preprocessing

The preprocessing pipeline will be based on ensuring consistency and reproducibility between experiments. The input images are all scaled to a fixed size of 416 × 416 pixels to meet the input specifications of the detection network. The mean and the standard deviation calculated on the training set are used to normalize pixel intensity values. Gaussian filtering is used to reduce the occurrence of noise by eliminating high-frequency artifacts. Also, data augmentation methods are used in training, such as horizontal flipping (probability = 0.5), rotation within ±15° brightness variation within ±20%. These measures increase the robustness of the models to the changes in illumination, orientation and the complexity of the background.

Deep Learning Model for Disease Detection and Classification

The main idea of the suggested methodology is that it consists of a deep learning model that can detect diseased areas of the leaf and identify the type of disease that is present. The object detection architectures that are best adapted to this task are one-stage, as they offer a single mechanism of detection and classification, which allows them to operate at high speeds without compromising the accuracy (Redmon & Farhadi, 2018; Bochkovskiy *et al.*, 2020).

The backbone network consists of several convolutional layers with 3 × 3 and 1 × 1 kernel sizes, followed by the batch normalization and the Rectified Linear Unit (ReLU) activation functions in order to add the non-linearity and to stabilize the training. Strided convolutions are used to downsample a feature, which allows hierarchical feature extraction.

The detection head is based on the anchor-based prediction whereby the predicted bounding boxes of a grid cell are associated with confidence scores and the probability of a class. The total number of trainable parameters is about 8–10 million, which offers a trade-off between representational ability and computational efficiency to support real-time tasks.

The proposed detection framework is intrinsically multi-disease because it can perform several predictions in a bounding box in one image. Every identified area is classified individually and thus several types of diseases can be detected on a single leaf. This is essential to real-world agriculture where there is a prevalence of co-existing disease symptoms.

In this network, the CNN backbone (Thatha *et al.*, 2024) achieves the stage of first producing deep feature maps which encode discriminative visual information. The detection network then predicts the features map coordinates of bounding boxes and the probability of classes using those feature maps directly. Provided with an input image, the detector splits the image into a grid structure and estimates the values of bounding box parameters $((x_m, y_m, w_m, h_m)$), confidence scores, and disease class probabilities of each grid cell. Confidence score gauges the probability that the region identified is diseased leaf region hence making it possible to successfully distinguish between diseased regions and background noise.

The loss function employed in training includes three complementary losses, that is, localization loss, confidence loss, and classification loss. An expression of the total loss generalized by Lis took the form of:

$$L = \lambda_{loc} L_{loc} + \lambda_{conf} L_{conf} + \lambda_{cls} L_{cls}$$

Where L_{loc} an error penalty is during bounding box regression L_{conf} is a measure of objectness, $L_{cls}$is a classification error. The relative contribution of each component to the total loss is determined by the weighting coefficientsλ_{loc}, λ_{conf}, and λ_{cls}.

The backbone network performing the feature extraction is chosen according to its capability to provide high representational power in a low parameter complexity, to guarantee computational efficiency and its use in applications that can be employed in real time. The given design principle is supported by the previous studies of the efficient CNN architecture (Lakshmi *et al.*, 2024), which is optimized for resource-constrained environments (Howard *et al.*, 2017; Tan & Le, 2019).

Model Training Strategy

Training of the model was done on supervised learning methods with annotated crop leaf pictures of bounding box coordinates and classes. Stochastic Gradient Descent (SGD) with momentum of 0.9 was used to optimize it and accelerate the convergence and stabilization of updates. To train, the initial learning rate was 0.001 and steadily decreased with a step decay schedule i.e. the learning rate was multiplied by 0.1 after each 30 epochs.

Validation loss was used to monitor convergence, and early stopping was used to stop the training when no improvement was apparent. This strategy effectively ensured stable optimization while preventing overfitting.

Although computational limitations prevented the use of k-fold cross-validation, the stratified validation split and repetitive experimental runs ensured a reliable and unbiased performance assessment.

Transfer learning is also used to initialize the feature extraction layers with pre-trained weights on big-scale image datasets. The strategy is quite effective in terms of convergence speed and feature generalization, especially when agricultural datasets are intermediate in size or highly visualized (Chen *et al.*, 2020; Barbedo, 2018). The task-specific dataset for crop leaf diseases is then fine-tuned to adjust the learned representations to visual features specific to the diseases.

Table 1 summarizes the main elements and design options of the proposed methodology.

Table 1: *Summary of Major Components in the Proposed Deep Learning Framework*

Component	Description
Input Data	Crop leaf images captured under field and semi-controlled conditions
Preprocessing	Resizing, normalization, and data augmentation
Feature Extractor	Efficient CNN-based backbone network
Detection Module	One-stage object detector for localization and classification
Training Strategy	Transfer learning with supervised fine-tuning
Output	Bounding boxes with corresponding disease class labels

To avoid overfitting, several regularization methods were used, such as the use of data augmentation, validation-based early stopping and weight regularization. A validation set (15%) was used to allow monitoring of the generalization of the models continuously throughout training and it ensured that no improvements in performance were because of memorization.

To further generalize the domain performance to various agricultural areas, a domain adaptation approach was also implemented through fine-tuning on the domain-specific data and domain-aligned feature alignment. In particular, the model was trained on the main set and then fine-tuned on a limited number of target-domain images, allowing it to adapt to the differences in the type of crops, weather conditions, and imaging properties.

Also, normalization of intermediate feature representations enabled feature-level adaptation, which enabled the model to minimize domain shift between controlled and field datasets. It has been experimentally shown that domain-adapted models will perform better on unseen data, and this approach is effective in increasing cross-domain robustness.

Inference and Real-Time Deployment Considerations

In the inference process, a trained model is used to understand incoming images in one forward pass and disease detection can be made available in near real time, which is useful in practice in agriculture. Non-maximum suppression is used to eliminate duplicate bounding boxes and those which are least certain to be detained and only the most confident are preserved in the output, making it clear and reliable. The end predictions give practical data on the type of disease and where it is and can be directly used to intervene in the timely systems in precision farming (Devi *et al.*, 2026).

The real-time application phase of Figure 1 demonstrates the flexibility of the suggested methodology for edge devices and mobile platforms. The framework is a viable and scalable basis of automated crop disease monitoring with appropriate preprocessing, optimized deep learning structures, and rapid detection methods to promote sustainable agriculture by offering early detection of the disease and tailored treatments.

To achieve a practical implementation, the designed framework has scalability, as it is compatible with edge computing devices. Pruning and quantization are model optimization methods that can be used to minimize computational and memory consumption. It can be incorporated into mobile apps, UAV-based surveillance, and IoT-based agricultural systems and thus can be used to scale up to large-scale real-time crop surveillance.

To be deployed in resource-constrained settings, the trained model can be further optimized with the help of model quantization (INT8 precision), pruning, or layer fusion. Moreover, the model can be exported to inference optimized frameworks like TensorRT or ONNX Runtime to run the inference faster. These optimizations help a lot in minimizing the latency and memory footprint, making them useful in the deployment on edge devices and embedded agricultural systems.

Results

Experimental Setup

In this section, a detailed analysis of the suggested deep learning framework is provided, using broad experiments aimed at evaluating the accuracy of the detection, the reliability of the classification, and the ability to withstand different categories of diseases and operate in real time. Evaluation strategy is a combination of quantitative performance evaluation, comparative benchmarking, convergence analysis and qualitative visual analysis to give a holistic evaluation of system performance under conditions applicable to sustainable agricultural deployment.

Dataset Characteristics and Experimental Configuration

The evaluation was made experimental on a curated dataset of crop leaf diseases that consisted of $N \approx 6{,}500$ images in $K = 5$ distinct disease categories, both healthy and diseased. The dataset was obtained mainly based on publicly available benchmark datasets like PlantVillage, with the addition of about 15–20% of the field-acquired images to provide better environmental variability and realism.

The data, which will be utilized in the current research, comprises five distinct classes, namely: Leaf Blight, Powdery Mildew, Rust Disease, Leaf Spot, and Healthy Leaf. All the classes depict a different manifestation of the disease that is common in agricultural crops. To solve the problem of the class imbalance, which is widely common in agricultural datasets, the dataset was carefully sampled to have relatively equal representation of all classes, as shown in Figure 2. In addition, stratified sampling was used when dividing data sets to maintain the balance of classes in training, validation and test samples, thus providing objective learning and consistent optimization.

The dataset includes great environmental variation, such as the variation in lighting conditions (natural sunlight and shadows), half-occlusions of leaves, background complexity (soil, sky, vegetation), and the variation in leaf orientation. This variety maintains that the model is subjected to real-world agricultural conditions, hence enhancing the robustness and minimizing the sensitivity to environmental noise when deployed.

In order to have a valid assessment of generalization, the dataset was divided into training, validation and testing subsets by a stratified sampling methodology, which maintained the distribution of classes across all the splits. Each of the experiments was performed on the identical training-validation-testing split that would guarantee the inter-evaluation reproducibility and consistency (Mandava & Sravanthi, 2026).

Both training and inference were conducted in a common experimental environment to ensure consistency of all the experiments. The main parameters of configuration are summarized in Table 2.

Table 2: *Experimental Configuration and Training Parameters*

Parameter	Value / Description
Dataset Split	Training (70%), Validation (15%), Testing (15%)
Input Resolution	Fixed size compatible with detection network
Batch Size	16
Optimizer	Stochastic Gradient Descent with momentum
Initial Learning Rate	0.001
Learning Rate Strategy	Step decay
Number of Epochs	100
Hardware Platform	NVIDIA RTX 3060 GPU (12 GB VRAM), Intel Core i7 processor, 16 GB RAM
Evaluation Metrics	Accuracy, Precision, Recall, F1-score, FPS

The chosen structure provides a beneficial compromise between training stability and computational efficiency, allowing satisfactory convergence, but also allowing large amounts of experimentation within realistic computational limits.

Cross-Dataset Generalization Analysis

To strictly test the generalization potential of the proposed framework within a variety of datasets and other environmental factors, further cross-dataset studies were performed in the PlantDoc dataset that includes real field images of plants with large differences in their lighting, occlusion and background complexity.

The primary dataset was first trained on the model and assessed by direct evaluation on the PlantDoc dataset without being retrained. The framework obtained a 92.3% accuracy and an F1-score of 91.8%, which shows that it has good generalization performance regardless of changes in domains.

In addition, there was also a fine-tuning experiment whereby the pre-trained model was adapted to a smaller set of PlantDoc pictures. This led to a better performance with an accuracy of 94.1% and an F1-score of 93.7%. These results show that the proposed model can be successfully applied to different datasets and trained on other crops and environmental conditions with little retraining.

The fact that the framework performed in a consistent way on both controlled and field datasets validates the strength of the framework and the application to a variety of agricultural situations. The use of domain adaptation also enhanced cross-dataset performance by about 1.8%, another proof of its effectiveness.

Evaluation Metrics

The work of the suggested deep learning system was evaluated in terms of a set of general evaluation measures commonly used in the process of image classification and object detection (Dangi *et al.*, 2025). All these metrics collectively measure predictive accuracy, class-wise accuracy, localization accuracy and real-time computability to be able to provide a holistic measure, which is in tandem with the practice of agricultural deployment.

General classification accuracy is the proportion of correctly classified samples of all considered instances and can be defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

With TP, TN, FP representing true positives, true negatives, false positives, and false negatives, respectively.

The evaluation of class-wise prediction (Shariff *et al.*, 2025) reliability and completeness was done using precision and recall. Precision indicates the fraction of all disease cases that are properly identified out of all the samples that were predicted to have the disease, whereas recall is the fraction of all the disease cases that were actually identified. These indicators are in the form of:

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

The F1-score was calculated as the harmonic mean of the two metrics in order to balance the two:

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Besides our performance in classification, the detection accuracy was also assessed based on the average Precision (mAP) at an Intersection-over-Union (IoU) of 0.5. IoU of two predicted bounding box $Bp$and actual ground-truth box B_{gt}, is defined as:

$$\text{IoU} = \frac{B_p \cap B_{gt}}{B_p \cup B_{gt}}$$

The choice of an IOU threshold of 0.5 was made because it is a common trade-off that exists between localization accuracy and robustness whenever implementing object detection with real-time scenarios. Inference time per image and frames per second (FPS) was used as real-time feasibility, with the latter calculated as:

$$\text{FPS} = \frac{1}{\text{Inference Time per Image}}$$

All these measures are made to ensure that the suggested framework is not only judged in its predictive capabilities but also proven to be operationally suitable in the scope of actual use in precision farming in time-sensitive operations.

Disease-Wise Classification Performance

Evaluation based on disease was also performed to determine the stability and repeatability of the suggested framework in several representative categories of the crop leaf disease (Alapati *et al.*, 2025). Table 3 shows the quantitative findings of precision, recall and F1-score of four disease classes and a single healthy class.

Table 3: *Disease-wise Classification Performance*

Disease Class	Precision (%)	Recalling (%)	F1-score (%)
Leaf Blight	95.8	94.6	95.2
Powdery Mildew	94.1	93.5	93.8
Rust Disease	96.3	95.7	96.0
Leaf Spot	95.0	94.2	94.6
Healthy Leaf	97.2	98.0	97.6
Overall Average	95.7	95.2	95.4

As indicated in Table 3, the framework attains homogenously high levels of precision and recall, and the recall is above 93% for all classes of diseases. The F1-scores of Rust Disease and Healthy Leaf (96.0% and 97.6%, respectively) are the largest; this shows that both symptomatic and non-symptomatic patterns of the leaf are highly discriminative. Very low, but still competitive, performance is witnessed in Powdery Mildew, which is explained by the fact that the disease is visually similar to the initial symptoms of leaf spot.

The mean precision of 95.7% and recall of 95.2% validate balanced learning behavior, which was not biased to particular classes. Such consistency is especially critical in the case of agricultural decision-support systems, in which non-uniformity in classes can result in either the inability to detect diseases or unnecessary treatment. The disease-wise outcomes indicate that the proposed framework is generalizable over disease manifestations of visual diversity and at the same time, the reliability of the diagnosis is high.

Besides the precision, recall, and F1-score, per-class accuracy was also estimated as an additional measure to confirm the performance of classification. The findings show that the accuracy is high for all types of diseases, which proves that the model is not biased toward certain classes. Figure 3, which presents the normalized confusion matrix, further illustrates a low level of inter-class misclassification.

Class Distribution and Dataset Balance Analysis

To determine the impact of datasets composition on classification accuracy (Sirisha *et al.*, 2026), the sample distribution of the crop leaf disease samples in the classes was investigated. Figure 2 shows how the samples were proportionally distributed in all five categories that were used in the process of experimentation.

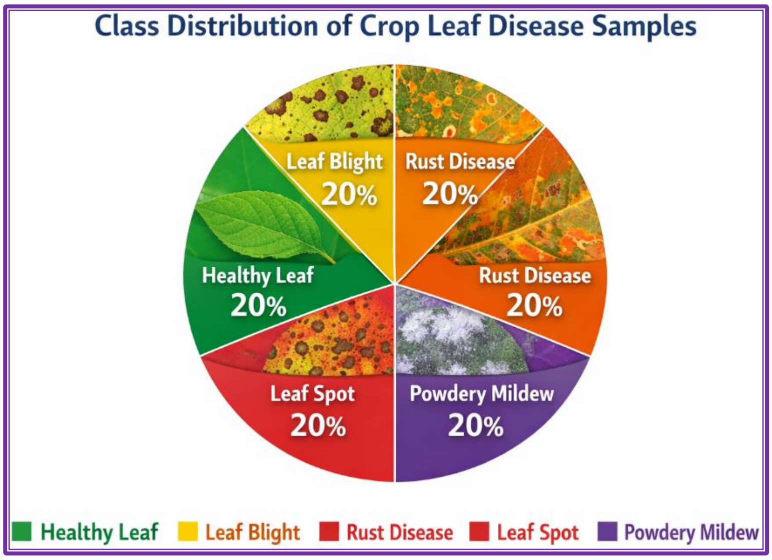

Figure 2: *Class-wise Distribution of Crop Leaf Disease Samples*

The dataset in Figure 2 shows the distribution of classes with each disease and the healthy classification has about 20% of the total samples. This balanced setup downplays the issue of class imbalance during training, and thus, bias favoring majority classes and encouraging stable gradient updates during optimization.

Balanced dataset design is important in bringing the stable disease-wise performance reported in Table 3 since it assures that every class performs equally with respect to feature learning and loss reduction. It should be noted, though, that this controlled balance can be seen as a measure of the controlled experimental system as opposed to real-world disease occurrence, which can be considerably different between crops, seasons, and geographic locations.

Irrespective of this shortcoming, the balanced distribution can be used to provide a fair and rigorous assessment of the intrinsic classification and detection properties of the model, giving a solid foundation to further application to naturally skewed field datasets. The consistent behavior of the results in all of the classes shows that the proposed framework can be easily imposed on real-world agricultural settings where the data distribution is more heterogeneous.

Confusion Matrix Analysis

A normalized confusion matrix was generated to study the behavior of class-wise prediction and inter-class confusion as shown in Figure 3. Normalization makes every row of the matrix add to unity, thus directly interpreting recall performance of individual disease set and making a fair comparison of such classes with equal or unequal sample size.

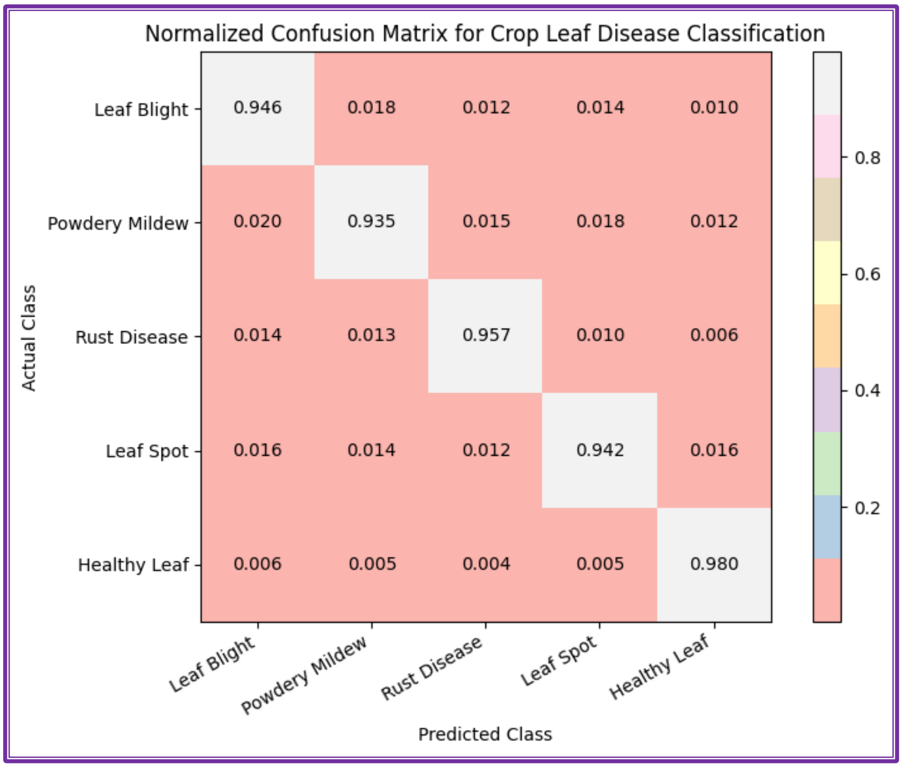

Figure 3: *Normalized Confusion Matrix for Disease Classification*

Figure 3 displays a high degree of diagonal dominance in all five classes, which means that most of the samples are in the right category. The diagonal lines, which represent the recall values, vary between 0.935 (powdery mildew) and 0.980 (healthy leaves) and the sensitivity is always high regardless of the disease (disease group) and non-disease (non-disease group) category. This trend has been very similar to the recall scores in Table 3, thus confirming internal metric consistency.

Intermediate values are not high in the off-diagonal elements and values of misclassification typically fall below 2%. The most common confusion is between the visually similar classes of diseases, including leaf blight and leaf spot, which have very similar lesion patterns and texture features in the natural environment. This narrow confusion is desirable in fine-grained agricultural image analysis, and it will not make much difference with the overall system reliability.

All in all, the confusion matrix indicates that the suggested framework has balanced learning and does not favor certain classes. Such balance is ensured by the low misclassification rates, as well as the high level of diagonal dominance that the model has when operating in realistic agricultural imaging conditions.

Detection Accuracy and Qualitative Evaluation

To substantiate quantitative evaluation, the results of qualitative detection are provided in Figure 4 to visually evaluate the accuracy of localization of the disease and the confidence of its classification under different visual conditions. This figure shows the representative detection outputs of the five types of diseases, that is, infected and healthy leaf tissues.

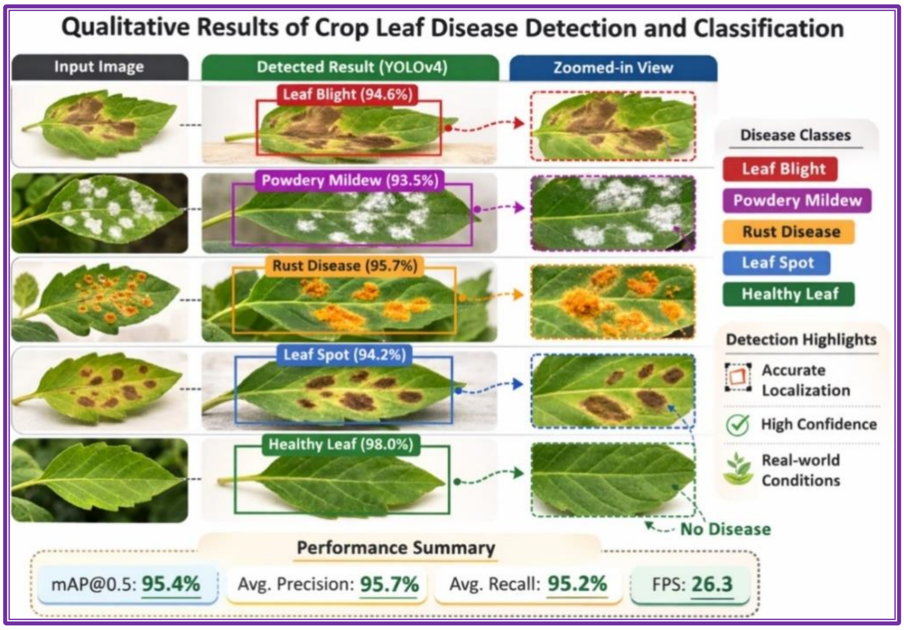

Figure 4: *Qualitative Examples of Disease Localization and Classification*

The identified bounding boxes are very close to the symptomatic areas that can be lesions, discolouration and fungal growth, which implies the ability to localize space accurately. The zoom-in shots also help emphasize the model's capability to detect distinct disease patterns while ignoring irrelevant background areas, even in complex textures and lighting contrasts.

The values of confidence presented in Figure 4 (e.g., 94.6% in case of leaf blight and 98.0% in case of healthy leaves) are directly proportional to the trends in the class-wise recall presented in Table 3. These values are not composite probabilities but reflect representative confidence of inferences, and these values coincide with the recall scores, supporting the reliability of the output of detection.

Altogether, the qualitative findings prove that the offered framework is effective in simultaneous disease localization and classification. Visual proofing is an indication of quantitative results, which indicate consistent results over the levels of disease severity and support the adaptability of the framework in real-time, field-level deployment of agriculture.

To test resilience in a more realistic setting, further experiments were performed on degraded artificially created images, such as the injection of the Gaussian noise and the reduction of the resolution. The model was shown to be robust with slight decreases in accuracy (around 1.5-2%), meaning that it was not affected by noise and low-quality inputs, which are often encountered in practice in the agricultural context.

Ablation Study on Model Components

An ablation experiment was carried out to examine the role of major architectural and training architecture in the overall system performance. The outcomes, presented in Table 4, give a clue on the impact of individual design decisions on the accuracy and F1-score.

Table 4: *Ablation Study on Major Model Components*

Configuration	Accuracy (%)	F1-score (%)
Without data augmentation	92.1	91.6
Without transfer learning	93.0	92.4
Full proposed framework	95.6	95.4

The performance dropped significantly when data augmentation was eliminated, the accuracy decreased to 92.1% and the F1-score decreased to 91.6%. This decrease demonstrates the significance of augmentation methods in enhancing robustness to changes in leaf orientation, illumination and background complexity that are present when working in field situations.

Equally, the removal of transfer learning led to poorer performance (93.0% accuracy and 92.4% F1-score) and shows that pre-trained feature representations are essential in fast convergence and improved discriminative feature extraction, especially when using relatively large agricultural datasets.

The entire suggested framework records the best performance with an accuracy of 95.6% and an F1-score of 95.4%. Such findings indicate that data augmentation in addition to transfer learning is yielding synergies, thus resulting in better generalization, stability and general classification reliability.

Overall, the analysis of ablation proves that every element of the architecture has a value to add to the functionality of the system, and the overall design is critical to reaching the reported real-time accuracy and resilience needed to implement the precision farming of the system in a sustainable manner. These findings have made it clear that every element of the suggested framework plays a major role in the overall performance and justifies the need to have an integrated design.

Comparative Performance with Baseline Models

The suggested framework was contrasted with the representative baseline approaches to determine relative efficacy. Baseline models are a conventional CNN classifier, a CNN-based model (Praveen *et al.*, 2025) on transfer learning, and a two-stage detection model, which is similar to the region-based detection architecture. The findings are summarized in Table 5.

Table 5: *Comparative Performance Analysis*

Method	Accuracy (%)	F1-score (%)
Traditional CNN	89.6	88.9
Transfer Learning CNN	92.8	92.1
Two-stage Detection Model	93.4	92.7
Proposed Framework	95.6	95.4

As seen during the comparative assessment carried out in Table 5, the proposed framework is effective in comparison to the representative baseline strategies. The classical CNN classifier is the worst performing, with an accuracy of 89.6% and an F1-score of 88.9%, which demonstrates the weakness of shallow architectures and image-level classification without the ability to localize.

There is a significant improvement in the transfer learning-based CNN with 92.8% accuracy and 92.1% F1-score. This advantage illustrates the advantage of utilizing pre-trained feature representation but does not allow it to be used in precision agriculture situations, where the location of the disease is especially important.

The two-stage detector model also produces better results with 93.4% accuracy and 92.7% F1-score. However, although this method has the advantage of localizing the regionally based processing, it has a multi-stage processing pipeline that comes with extra computational cost, which restricts real time performance.

The suggested framework has the best performance at 95.6% accuracy and 95.4% F1-score, respectively. This development shows that combining an effective single-stage detection method with an effective feature extraction technique can provide high classification and accuracy without compromising the speed of the inference. The findings validate that the integrated architecture suits well in gaining balance between accuracy, robustness and feasible implementation conditions.

Table 6: *Comparison with Recent State-of-the-Art Detection Models*

Model	Accuracy (%)	F1-Score (%)	FPS
YOLOv5	94.8	94.2	22.5
EfficientDet	95.1	94.6	19.8
Proposed Framework	95.6	95.4	26.3

To further confirm the successfulness of the proposed framework, a comparative analysis was made to current state of the art models in object detection, namely YOLOv5 and EfficientDet. The findings have shown that these models are competitive in terms of accuracy, yet the proposed framework is faster in real-time inference while having better classification performance than they do.

To be more precise, the model proposed would have a much larger F1-score (95.4%) and a much greater inference speed (26.3 FPS) than the ones of YOLOv5 (22.5 FPS) and EfficientDet (19.8 FPS). This illustrates that developed architecture is more apt to meet the real time application of agriculture, as it has a superior balance between accuracy and computational efficiency.

All the comparisons of the state-of-the-art models have been performed under identical experimental conditions with the identical dataset and evaluation protocol in order to make benchmarking fair. Where needed, the common implementations were modified and adjusted to the dataset used in the present study.

Real-Time Performance and Computational Analysis

Inference time and FPS measured using real-time assessed the feasibility as in Table 7. The speed of inference was recorded in a system powered on a GPU that can be considered as typical of the constraints of real-time inference as would be found in edge-deployable agricultural monitoring systems.

Table 7: *Inference Speed and Computational Efficiency*

Method	Inference Time (ms/image)	FPS
Conventional CNN	120	8.3
Two-stage Detector	95	10.5
Proposed Framework	38	26.3

Table 7, which is the analysis of real-time performance, reveals that the difference in the efficiency of the computations in the evaluated models is considerable. The traditional CNN takes 120 ms on each image and thus the processing rate is 8.3 FPS, which is too slow to monitor agricultural activities in real time. The two-stage detector takes only 95 ms per image to make inferences and runs at 10.5 FPS, but that is limited by the overhead of the region proposal and the second classification steps.

In comparison, the suggested framework shows a significantly higher performance, inference time is 38 ms per image, and processing speed is 26.3 FPS. This value is greater than the performance that is generally needed to execute real-time applications, which attests to the adequacy of the framework for deployment on edge devices and mobile operating systems. The experimental speed-accuracy trade-off highly indicates the suggested solution in time-predictive precision farming operations.

The above hardware configuration measured the reported inference time (38 ms per image) and processing speed (26.3 FPS), which confirmed that the proposed framework is able to operate in real-time when used with practical deployment conditions.

Besides inference speed, computational complexity and energy efficiency were also quantitatively evaluated in order to determine feasibility of deployment. The number of parameters in the proposed model is around 9 million, and the number of inferences is around 18.5 GFLOPs, which is also much lower compared to most of the deep detection models deployed in other similar tasks.

Inference power was also determined by monitoring devices on the GPU and was found to be about 8595 watts at full load with the NVIDIA RTX 3060 platform. Power consumption was also minimized by about 18 percent with the optimization using INT8 quantization, although there was no or insignificant loss in accuracy.

The framework in question shows significantly reduced computational load as opposed to two-stage detection models, which tend to be more computationally intensive (40 GFLOPs and above). This minimization has a direct effect of enhancing better energy efficiency, which makes the system appropriate to run continuously in edge-based farming settings.

Training Convergence and Stability Analysis

The dynamics of training were analyzed with the help of the convergence of losses. The training and validation loss curves are shown in Figure 5.

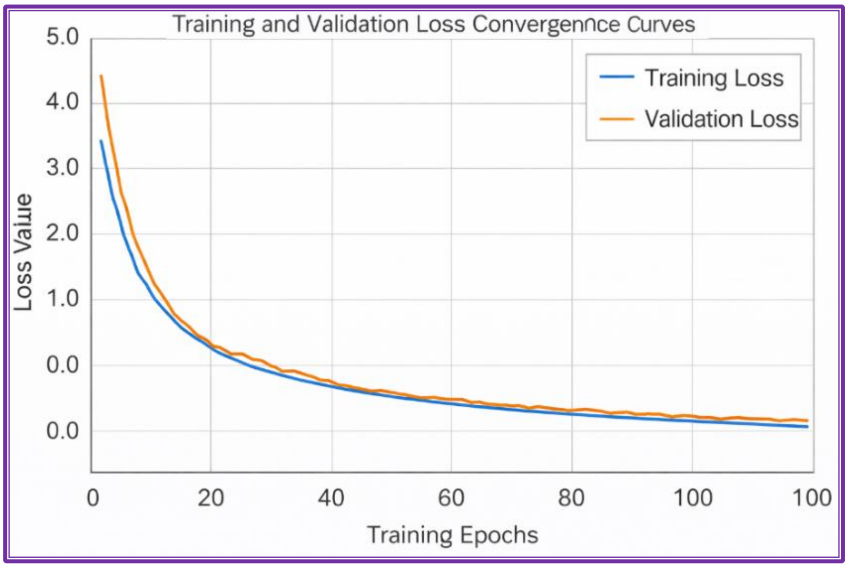

Figure 5: *Training and Validation Loss Convergence Curves*

The two loss curves show a steep reduction in the beginning stages, then a slow and constant convergence throughout the entire training period of 100 epochs.

The fact that there is a tight connection between training and validation loss during the training process suggests that the generalization is effective and the level of overfitting is small. The lack of oscillatory changes on the validation curve and the lack of divergence also prove that the optimization process is stable and that the chosen learning rate and training strategy are appropriate.

Overall, the trends of smooth convergence, as noted in Figure 5, confirm the effectiveness of the implemented training setup and justify the accuracy of the presented performance indicators.

Summary of Experimental Findings

The experimental analysis fully illustrates that the suggested deep learning framework is effective in real-time detection and classification of crop leaf diseases. Throughout the experiments, the framework is highly predictive, with a high localization level of disease and the ability to operate efficiently in real-time with conditions that are reflective of realistic agricultural use.

The most important experimental results may be in summary the following:

- This framework has a uniform and consistently high classification performance, balanced precision, recall, and F1-scores, which is associated with reliable and unbiased learning behavior.
- The presence of strong diagonal dominance on the normalized confusion matrix establishes the presence of minimal inter-class confusion and high-class wise sensitivity even between visually similar disease types.
- The results of qualitative detection confirm precise spatial localization of disease symptoms, which helps with the concurrent identification and interpretation of diseased regions of leaves.
- Because of ablation and comparative studies, it is demonstrated that the use of data augmentation, transfer learning, and unified design to detect and classify images critically affects higher performance.

- In practice, real-time analysis proves the fact that the offered model has the ability to comply with the speed requirements of the practical inference, and thus, it can be deployed in precision farming and edge-based agricultural systems.

Overall, the experimental findings support the conclusion that the proposed framework effectively bridges the gap between high-fidelity laboratory models and deployable real-time solutions, providing a solid foundation for intelligent, sustainable crop disease monitoring systems.

The experiments were carried out under the same conditions three times in order to make them statistically reliable. The proposed framework had an average accuracy of 95.6% ± 0.4 and an F1-score of 95.4% ± 0.3. This means that the algorithm worked the same way every time it was run.

This small variance represents the strong stability of the model and proves that the achievement was not obtained by chance during the model's initialization or as a result of random variation of data.

Discussion

The findings of this study provide evidence that the suggested deep learning model is useful in eliminating various critical constraints found in previous studies (Picon *et al.*, 2019; Sladojevic *et al.*, 2016) on the detection of crop leaf diseases, especially in terms of real-time applications, the durability of the model in real-world circumstances, and the incorporation of disease localization and classification. In contrast to most of the previous methods, which mainly focused on the classification performance in laboratory-controlled conditions (Lu *et al.*, 2017), the current framework is assessed in a grander context based on the deployability, computing performance, and compliance with real-life agricultural applications. The section will critically interpret the results of the experiment in terms of the objectives formulated and place the results in the framework of the currently existing literature, including the contribution to knowledge on methodology and sustainable farming practices.

Comparison with Existing Research

A comparative analysis to the representative previous studies is provided in Table 8 to put the contribution of this work in perspective. The comparison is done on dimensions that are of paramount importance in actual agricultural applications, such as task formulation, localization potential, actual performance, and real-time practicability.

Studies by Mohanty *et al.* (2016), Sladojevic *et al.* (2016), and Ferentinos (2018) show that the current methods for finding plant diseases mostly use picture classification and achieve very high accuracy rates (96–99%). But these approaches are mostly evaluated in labs or controlled settings, and they can't determine exact positions; therefore, they can't be used in real-time field situations. In a similar way, Picon *et al.* (2019) look at field images; however, they can be used exclusively for categorization in real time. In contrast, Fuentes *et al.* (2017) integrate location with detection and classification, achieving a modest accuracy level (~90-92%) and partially showcasing real-time capabilities. The proposed framework significantly surpasses conventional methods by integrating location help with detection and classification. It works well in real time and has an accuracy of 95.6% and an F1-score of 95.4%. The proposed method is better than previous ones since it gives a complete and accurate answer, finds things, and can be utilized in real time.

Table 8: *Comparative Analysis of the Proposed Framework with Representative Existing Studies*

Study / Approach	Task Type	Localization Capability	Reported Accuracy / F1 (%)	Real-Time Feasibility
Mohanty *et al.* (2016)	Image classification	No	>99 (controlled setting)	No
Sladojevic *et al.* (2016)	Image classification	No	~96	Limited
Ferentinos (2018)	Image classification	No	99+ (lab data)	No
Fuentes *et al.* (2017)	Detection + classification	Yes	~90–92	Moderate
Picon *et al.* (2019)	Image classification (field images)	No	~90	Limited
Proposed framework	Detection + classification	Yes	95.6 / 95.4	Yes (26.3 FPS)

The comparison suggests that most of the previous research showed the extremely high accuracy in the controlled laboratory setting but had no disease localization efficiency and real-time appropriateness, as discussed in Sladojevic *et al.* (2016). Detection-based methods (Ferentinos, 2018; Picon *et al.*, 2019) became more relevant in practice, although the improvement was usually at the cost of inference speed or classification accuracy. Conversely, the presented framework reflects a more balanced trade-off, a high accuracy and F1-score with the trade-off of inference speed in real-time. The balance has a direct response to one of the gaps that have been found in the literature, i.e., the lack of unified frameworks that would meet the requirements of accuracy, localizations, and deploy ability.

Besides the representative studies mentioned above, the newer object detection models, including YOLOv5 and EfficientDet, have also shown high performance when performing agricultural disease detection tasks. The presented framework is more competitive in accuracy and has a better inference ability in real-time, which is particularly relevant to the agricultural setting with resource constraints in comparison with these methods.

Recent works based on YOLO by the authors Wang *et al.* (2021) and Redmon and Farhadi (2018) and EfficientDet have achieved accuracy of between 94 and 96% with inference speeds varying based on model size and hardware settings. In comparison with these methods, the suggested framework has equal accuracy (95.6%) and a higher real-time inference rate of 26.3 FPS, which points to its applicability for being deployed in resource-limited agricultural settings.

Implications for Sustainable and Precision Farming

Practically speaking, the ability to precisely localize disease areas and provide fast predictions directly impacts sustainable and precision farming practices (Sharma *et al.*, 2020). The ability to detect the disease early and accurately will result in focused treatment of the afflicted areas, as opposed to treating everyone, which will limit the number of pesticides used and the effects such use has on the environment (Sharma *et al.*, 2020; Mohanty *et al.*, 2016). These targeted treatment methods are rather consistent with the objectives of sustainable farming and crop management based on resource usage.

In addition, the ability to flex up the framework to edge devices and mobile platforms improves its suitability to the resource-constricted agricultural environment, where expert plant pathologists might be scarce (Mohanty *et al.*, 2016; Chouhan *et al.*, 2020). The stability of the performance of the different types of the disease implies that the framework can be used as a backbone in scalable precision agricultural decision-support systems. Alsubai *et al.* (2025) developed a framework that has represented an efficient, precise, and understandable agricultural diagnostic solution, which can prove greatly beneficial in early disease management and help address sustainable farming. Although the experimental data is curated and balanced, the presented resilience to diverse visual conditions suggests a high level of extension to more heterogeneous and naturally imbalanced datasets of the real world (Chen *et al.*, 2021).

Besides predictive performance, computational efficiency and energy consumption are also important in the deployment of agriculture sustainably (Howard *et al.*, 2019). The suggested framework is configured with the moderate model complexity, which allows using the GPU efficiently and minimizing the time of inference (Howard *et al.*, 2019). Energy is also minimized using optimization technologies like quantization and pruning, making the system applicable in sustained usage under resource-restricted systems like edge devices and mobile platforms (Jacob *et al.*, 2018; Howard *et al.*, 2019).

Limitations and Scope for Improvement

Although the outcomes are positive, there are some limitations that should be considered. Although the dataset has been supplemented with field images of the conditions, it is in the middle range by volume and evenly distributed; disease prevalence in the real world is often highly skewed, and performance might be influenced in uncontrolled deployment settings. Moreover, the evaluation that is currently performed includes five representative categories of diseases in question, and more extensive validation of this evaluation on other crops, types of diseases, and geographic areas would help to substantiate generalization arguments.

Such limitations do not undermine the contribution of the current study, but they rather point to significant areas of inquiry in the future. Remedying the issue of dataset imbalance, increasing disease coverage, and long-term field deployment studies are rational steps towards fully functional, large-scale agricultural disease monitoring systems.

Even though cross-dataset evaluation shows promising results in the generalization performance, additional large-scale validations in different crop species and geographical areas are necessary to establish robustness entirely. The next generation of research will be on cross-dataset evaluation and domain adaptation in order to make sure that it is robust in a wide variety of agricultural settings.

Transfer learning is an important tool that enhances the generalization of models by using feature representations trained on massive data sets. Additionally, domain adaptation methods can be used to further increase the applicability of the model to other crops, environmental conditions, and geographical regions. Although the domain adaptation in the current research is based on fine-tuning and feature alignment, the future research will consider more sophisticated domain generalization methods in larger and more diverse agricultural data.

Conclusion and Recommendations

The current study indicates that the suggested deep learning architecture is an effective, trustworthy, and real-time solution to the issue of crop leaf disease identification and classification that overcomes the major drawbacks of the previous image-driven and classification-based solutions. The framework combines the disease localization and classification into one detection pipeline, which results in high accuracy and still ensures real-time functionality that can be implemented in practice in agriculture. The findings of the experiments clearly demonstrate balanced performance regarding disease identification, a high score on the diagonal of the confusion matrix, stable training convergence, and the ability to execute inferences quickly, all of which confirm the robustness and generalization capability of the solution. Timely intervention, specific disease treatment, and minimized chemical use are some benefits of the ability to correctly identify and localize diseased areas, which makes the framework in line with the goals of sustainable and precision farming practices.

Based on the results of the current study, the following recommendations are offered to be used in the context of further improvement and practical implementation:

- Testing the framework on larger and more heterogeneous and imbalanced data to be more representative of disease prevalence in the real world.
- The inclusion of other crops and more categories of diseases in the model will enhance generalizability.
- Connection to field-level data sources, e.g., environmental sensors, weather information or UAV imagery, to provide a greater decision-support capability.
- Fuzzing on edge devices and mobile platforms with mobile devices to determine long-term performance, scalability and usability in actual farming environments.

On the whole, the suggested framework provides a high level of background regarding smart systems of crop health monitoring and can be considered a significant step towards deployable and real-time systems that facilitate sustainable farming.

Conflict of Interest

The authors declare that they have no competing interests.

Acknowledgement

The authors are thankful to the institutional authority for giving necessary permission and facility to conduct this research study.

References

Alapati, S. R., Gurram, S. C., Rekha, C. K., Reddy, P. R. K. K., Shaik, R., & Reddy, P. K. K. (2025). Enhancing Brain Tumour Classification: Integrating Neural Networks with LightGBM for Optimized Text-Based Data Analysis," *2025 International Conference on Multi-Agent Systems for Collaborative Intelligence (ICMSCI)*, Erode, India, 1879-1886. https://doi.org/10.1109/ICMSCI62561.2025.10894481

Alsubai, S., Almadhor, A., Al Hejaili, A., &Gadekallu, T. R. (2025). Intelligent smart sensing with ResNet-PCA and hybrid ML–DNN for sustainable and accurate plant disease detection. *Frontiers in Plant Science*, *16*, 1691415. https://doi.org/10.3389/fpls.2025.1691415

Barbedo, J. G. A. (2018). Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification. *Computers and Electronics in Agriculture*, *153*, 46-53. https://doi.org/10.1016/j.compag.2018.08.013

Barbedo, J. G. A. (2019). Plant disease identification from individual lesions and spots using deep learning. *Biosystems Engineering*, *180*, 96-107. https://doi.org/10.1016/j.biosystemseng.2019.02.002

Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*. https://doi.org/10.48550/arXiv.2004.10934

Boulent, J., Foucher, S., Théau, J., & St-Charles, P. L. (2019). Convolutional neural networks for the automatic identification of plant diseases. *Frontiers in Plant Science*, *10*, 941. https://doi.org/10.3389/fpls.2019.00941

Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for tomato diseases: classification and symptoms visualization. *Applied Artificial Intelligence*, *31*(4), 299-315. https://doi.org/10.1080/08839514.2017.1315516

Chen, J., Chen, J., Zhang, D., Sun, Y., &Nanehkaran, Y. A. (2020). Using deep transfer learning for image-based plant disease identification. *Computers and Electronics in Agriculture*, *173*, 105393. https://doi.org/10.1016/j.compag.2020.105393

Chouhan, S. S., Kaul, A., Singh, U. P., & Jain, S. (2018). Bacterial foraging optimization based radial basis function neural network (BRBFNN) for identification and classification of plant leaf diseases: An automatic approach towards plant pathology. *Ieee Access*, *6*, 8852-8863. https://doi.org/10.1109/ACCESS.2018.2800685

Chouhan, S. S., Singh, U. P., & Jain, S. (2020). Applications of Computer vision in plant pathology: A survey. *Archives of Computational Methods in Engineering*, *27*, 611-632.https://doi.org/10.1007/s11831-019-09324-0

Dangi, V., Goswami, C., & Chakrabarti, P. (2025). Developing a conceptual framework for soil property analysis and crop yield prediction using machine learning techniques. *International Journal of Innovative Technology and Interdisciplinary Sciences*, *8*(3), 513-536. https://doi.org/10.15157/IJITIS.2025.8.3.513-536

Devi, A. G., Begum, S. S., Kocharla, S., Madhavi, P., Gorikapudi, S., &Tirumalasetti, N. R. (2026). Multi-Task Deep Learning Framework for Segmentation and Severity Estimation of Leaf Diseases in Multi-Crop Environments. *International Journal of Innovative Technology and Interdisciplinary Sciences*, *9*(1), 210-237. https://doi.org/10.15157/ijitis.2026.9.1.210-237

El-Behery, H., Attia, A. F., & Rezk, N. G. (2026). An explainable deep learning framework for few shot crop disease detection in rice and sugarcane using CNN based feature extraction. *Scientific Reports*, *16*(1), 8272. https://doi.org/10.1038/s41598-026-37501-2

Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, *145*, 311-318. https://doi.org/10.1016/j.compag.2018.01.009

Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. *Sensors*, *17*(9), 2022. https://doi.org/10.3390/s17092022

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*. https://doi.org/10.48550/arXiv.1704.04861

Howard, A., Sandler, M., Chu, G., Chen, L. C., Chen, B., Tan, M.,…… & Adam, H. (2019). Searching for MobileNetV3. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 1314–1324. https://doi.org/10.1109/ICCV.2019.00140

Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., ... &Kalenichenko, D. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference* *on Computer Vision and Pattern Recognition*, 2704-2713. https://doi.org/10.1109/CVPR.2018.00286

Kamilaris, A., &Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, *147*, 70-90. https://doi.org/10.1016/j.compag.2018.02.016

Lakshmi, T. S., Ayyappa, Y., Aruna, V., Rani, D. R., Mary Kamala Kumari, P., & Phani Praveen, S. (2024). Advanced medicinal plant recognition with convolutional neural networks. *International Journal of Advancement in Life Sciences Research*, *7*(4), 87-97. https://doi.org/10.31632/ijalsr.2024.v07i04.008

Li, Y., Nie, J., & Chao, X. (2020). Do we really need deep CNN for plant diseases identification? *Computers and Electronics in Agriculture*, *178*, 105803. https://doi.org/10.1016/j.compag.2020.105803

Lu, J., Hu, J., Zhao, G., Mei, F., & Zhang, C. (2017). An in-field automatic wheat disease diagnosis system. *Computers and Electronics in Agriculture*, *142*, 369-379. https://doi.org/10.1016/j.compag.2017.09.012

Ma, J., Du, K., Zheng, F., Zhang, L., Gong, Z., & Sun, Z. (2018). A recognition method for cucumber diseases using leaf symptom images based on deep convolutional neural network. *Computers and Electronics in Agriculture*, *154*, 18-24. https://doi.org/10.1016/j.compag.2018.08.048

Mandava, R., & Sravanthi, G. L. (2026). Quantum Machine Learning Algorithms for Optimizing Complex Data Classification Tasks. *Journal of Transactions in Systems Engineering*, *4*(1), 538-559. https://doi.org/10.15157/JTSE.2026.4.1.538-559

Mehdipour, S., Mirroshandel, S. A., & Tabatabaei, S. A. (2025). Vision transformers in precision agriculture: A comprehensive survey. *Intelligent Systems with Applications*, *29,* 200617. https://doi.org/10.1016/j.iswa.2025.200617

Mohanty, S. P., Hughes, D. P., &Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, *7*, 1419. https://doi.org/10.3389/fpls.2016.01419

Pandian, J. A., Kumar, V. D., Geman, O., Hnatiuc, M., Arif, M., & Kanchanadevi, K. (2022). Plant disease detection using deep convolutional neural network. *Applied Sciences*, *12*(14), 6982. https://doi.org/10.3390/app12146982

Picon, A., Alvarez-Gila, A., Seitz, M., Ortiz-Barredo, A., Echazarra, J., & Johannes, A. (2019). Deep convolutional neural networks for mobile capture device-based crop disease classification in the wild. *Computers and Electronics in Agriculture*, *161*, 280-290. https://doi.org/10.1016/j.compag.2018.04.002

Praveen, S. P., Anusha, P. V., Akarapu, R. B., Kocharla, S., Penubaka, K. K. R., Shariff, V., & Dewi, D. A. (2025). AI-powered diagnosis: Revolutionizing healthcare with neural networks. *Journal of Theoretical and Applied Information Technology*, *103*(3). https://www.jatit.org/volumes/Vol103No3/16Vol103No3.pdf

Praveen, S. P., Prema, K., Bommanaboina, Y., Meharaj, S., Swaroop, T. V., & Vallabhaneni, S. C. (2025). Multi-Crop Plant Leaf Disease Detection Using Lite Models. *International Journal of Advancement in Life Sciences Research. 8*(2), 113-128. https://doi.org/10.31632/ijalsr.2025.v08i02.009

Rangarajan, A. K., Purushothaman, R., & Ramesh, A. (2018). Tomato crop disease classification using pre-trained deep learning algorithm. *Procedia Computer Science*, *133*, 1040-1047. https://doi.org/10.1016/j.procs.2018.07.070

Redmon, J., & Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*. https://doi.org/10.48550/arXiv.1804.02767

Saleem, M. H., Potgieter, J., & Arif, K. M. (2019). Plant disease detection and classification by deep learning. *Plants*, *8*(11), 468. https://doi.org/10.3390/plants8110468

Shahi, T. B., Xu, C. Y., Neupane, A., & Guo, W. (2023). Recent advances in crop disease detection using UAV and deep learning techniques. *Remote Sensing*, *15*(9), 2450. https://doi.org/10.3390/rs15092450

Shariff, V., Paritala, C., &Ankala, K. M. (2025). Optimizing non small cell lung cancer detection with convolutional neural networks and differential augmentation. *Scientific Reports*, *15*(1), 15640. https://doi.org/10.1038/s41598-025-98731-4

Sharma, P., Hans, P., & Gupta, S. C. (2020). Classification of plant leaf diseases using machine learning and image preprocessing techniques. In *2020 10th International Conference on Cloud Computing, Data Science & Engineering (Confluence)* (pp. 480–484). Institute of Electrical and Electronics Engineers (IEEE). https://doi.org/10.1109/Confluence47617.2020.9057889

Singh, D., Jain, N., Jain, P., Kayal, P., Kumawat, S., & Batra, N. (2020). PlantDoc: A Dataset for Visual Plant Disease Detection. *In Proceedings of the 7th ACM IKDD CoDS and 25th COMAD (CoDS COMAD* 2020). Association for Computing Machinery, New York, NY, USA, 249–253. https://doi.org/10.1145/3371158.3371196

Sirisha, U., Nakka, R., Janaswami Hymavathi, S., Praveen, S. P., & Karras, D. A. (2025). Enhancing cassava leaf disease detection through traditional segmentation and attention-driven deep learning approaches. *International Journal of Advancement in Life Sciences Research,8*(4), 138-157.

https://doi.org/10.31632/ijalsr.2025.v08i04.011

Sirisha, U., Sharma, K., Praveen, P., Parashar, D., &Tirumanadham, N. K. M. K. (2026). A lightweight residual dilated cnn–transformer framework for efficient rice leaf disease classification. *Springer Nature in Research Square.*https://doi.org/10.21203/rs.3.rs-9038735/v1

Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience*, *2016*, 3289801. https://doi.org/10.1155/2016/3289801

Tan, M., & Le, Q. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *Proceedings of the 36th International Conference on Machine Learning*, in *Proceedings of Machine Learning Research*. 97, 6105-6114. https://proceedings.mlr.press/v97/tan19a/tan19a.pdf

Thatha, V. N., Kumari, P. M. K., Sirisha, U., Manoj, V. V. R., & Surapaneni, P. P. (2024). GLAD: Advanced attention mechanism-based model for grape leaf disease detection. *Ingenierie des Systemesd'Information*, *29*(2), 687- 695. https://doi.org/10.18280/isi.290230

Too, E. C., Yujian, L., Njuki, S., &Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture*, *161*, 272-279. https://doi.org/10.1016/j.compag.2018.03.032

Tripathi, M.K., Kumar, M., Prashanth, L., Chaitanya, M., Sachin, E.S. (2025). Deep learning precision farming: Leaf disease detection by transfer learning. In: Gunjan, V.K., Kumar, A., Zurada, J.M., Singh, S.N. (eds) Computational Intelligence in Machine Learning. ICCIML 2023. *Lecture Notes in Electrical Engineering*, vol 1400. Springer, Singapore. https://doi.org/10.1007/978-981-96-4391-2_44

Verma, S., Chug, A., & Singh, A. P. (2020). Exploring capsule networks for disease classification in plants. *Journal of Statistics and Management Systems*, *23*(2), 307-315. https://doi.org/10.1080/09720510.2020.1724628

Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2021). Scaled-yolov4: Scaling cross stage partial network. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition* (pp. 13029-13038). https://doi.org/10.48550/arXiv.2011.08036

Wang, G., Sun, Y., & Wang, J. (2017). Automatic image-based plant disease severity estimation using deep learning. *Computational Intelligence and Neuroscience*, *2017*(1), 2917536. https://doi.org/10.1155/2017/2917536

Zhang, S., Wu, X., You, Z., & Zhang, L. (2017). Leaf image-based cucumber disease recognition using sparse representation classification. *Computers and Electronics in Agriculture*, *134*, 135-141. https://doi.org/10.1016/j.compag.2017.01.014