




# International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

*(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)*



**Impact Factor: 9.935**

**Volume 15, Issue 7, July 2026**

# **An Intelligent Machine Learning Framework for Automated Plant Leaf Disease Detection**

**Dr.D.Radhika, S.Jayananthini, T. Kavitha**

Assistant Professor, Head of Dept. of Computer Science and Engineering, Vivekananda College of Engineering for Women, Elayampalayam, Tiruchengode, Namakkal, Tamil Nadu, India

PG Student, Department of Computer Science and Engineering, Vivekananda College of engineering for Women, Elayampalayam, Tiruchengode, Namakkal, Tamil Nadu, India

Assistant Professor, Department of Computer Science and Engineering, Vivekananda College of engineering for Women, Elayampalayam, Tiruchengode, Namakkal, Tamil Nadu, India

**ABSTRACT:** Plant diseases pose a significant threat to global agricultural productivity, resulting in reduced crop yield, lower product quality, and substantial economic losses. Early and accurate identification of plant leaf diseases is essential for effective crop management and timely intervention. Conventional manual inspection methods are time-consuming, labor-intensive, and susceptible to human error, highlighting the need for automated and intelligent disease diagnosis systems.

This study presents an automated plant leaf disease detection framework that integrates classical machine learning algorithms with deep learning techniques. A curated dataset of plant leaf images is preprocessed through image resizing, noise removal, data augmentation, and segmentation to improve classification performance. Classical machine learning models, including Support Vector Machine (SVM), Random Forest (RF), k-Nearest Neighbors (k-NN), and Logistic Regression (LR), are trained using handcrafted color, texture, and shape features. In parallel, a Convolutional Neural Network (CNN) is employed to automatically learn discriminative hierarchical features from leaf images. The CNN architecture incorporates convolutional, pooling, batch normalization, dropout, and fully connected layers and is optimized using the Adam optimizer with adaptive learning-rate scheduling.

Experimental evaluation using accuracy, precision, recall, F1-score, and confusion matrix analysis demonstrates that the CNN-based model consistently outperforms conventional machine learning approaches in disease classification accuracy and generalization. The proposed framework effectively identifies multiple plant leaf diseases, including bacterial blight, leaf spot, mildew, and rust, across different crop species. Cross-validation and testing on unseen images further confirm the robustness and reliability of the model.

The proposed system offers a scalable, accurate, and real-time solution for precision agriculture, supporting mobile and web-based deployment to facilitate intelligent crop health monitoring and sustainable agricultural management..

**KEYWORDS:** Plant Leaf Disease Detection, Deep Learning, Convolutional Neural Network (CNN), Machine Learning, Image Processing, Precision Agriculture, Support Vector Machine (SVM), Disease Classification

## **I. INTRODUCTION**

Agriculture is the primary source of income for rural people. It is one of the most important economic sectors in emerging nations. Majority of Gross Domestic Product (GDP) in several emerging nations comes from agriculture(1) . However, the contribution of agricultural goods is greatly impacted by losses brought by pests and crop diseases. From a financial point of view, plant diseases have an annual impact on the global economy of about US\$220 billion. The agricultural research report shows that pests and diseases cause majority of agricultural output to be harmed each year(2) . Thus, an alarming increase in the number of pests and illnesses threatens food security. Accurate plant disease detection is still a crucial process for agricultural sector(3) . Preventing these diseases is fundamental in agricultural growth, but most solo farmers and small-to-medium scale farmers do not have the luxuries where they can afford to hire several experts to gain consultation or focus on manual inspection, being costly and time-consuming. Delaying



advice leads to widespread crop devastation and a larger financial disadvantage for the farmer. AI and machine learning have risen to the top in recent years because they have improved the efficiency of various industries, including agriculture. According to AI's strong integration with modern planting systems, automated plant disease detection systems may now address several long-standing challenges in crop health management. Such technologies can quickly and easily diagnose diseases by using cameras to take pictures of sick leaves. Farmers with limited resources do not need advanced technologies; these steps are perfect. This paper seeks to design and implement an automated plant disease detection system that uses AI and image processing techniques to detect the plant and recommends the solution for the observed disease. The intended purpose of the system is to solve the challenges that farmers face in agriculture and boost the efficiency of agriculture by ensuring that time is not wasted. This study has specific objectives that include determining optimal solutions using machine learning algorithms in the field of the target disease detection, creating a balanced plant image dataset, and creating an appropriate interface for the target users, farmers. Addressing gaps in existing solutions, this research aims to affirmatively answer the question of where the gap between technology and agriculture is and how it is possible to develop a more efficient and sustainable agricultural system. Agriculture continues to be one of the main sectors for global food security, but it is faced with perennial threats of crop diseases which in turn lowers production levels, and income for farmers. Throughout the years, numerous applications for detection of various agricultural diseases have been created. They all have advantages and disadvantages, which makes the overall potential of these applications for different farming communities very limited. There exist several applications or solutions, but they have some limitations. Table 1 shows a comparison and highlights the limitations as well. Plantix(4) was launched in 2015 and can perform only basic disease recognition, lacking other features such as multilingual interface and disease prediction planning. Its interface is static with no adjustment settings, while the treatment plan for agriculture is generic and not specific to the area, making it less useful to farmers in different agricultural zones. Furthermore, it requires a persistent online connection which makes it less useful for customers in rural areas with limited network accessibility. The agricultural app(5) started in 2017, which allows for partial multilingual support and includes some features as well, however, as for the time it takes to respond it is too long when one is making life influencing decisions in agriculture. It indeed is an offline application and offers basic treatment principles, but the recommendations made are only partially available, in the sense that they do not fully pertain to the local agricultural practices. It does not provide real-time forecasts for the disease and has no connection with the existing local agricultural networks like in the case of Plantix. The Leaf doctor app(6) , which was launched in 2015, brings in some level of value by providing response times between 25 to 60 seconds with some setting options at a more reasonable time. However, use of this app is mostly only useful to farmers who speak English. Like most other apps, treatments are not recommended, and the app has no such feature which would allow the farmers to forecast a disease before it occurs or use voice commands to search for information within the app. Plant Guard(7) contains features that increase its effectiveness and use. It has a displayed disease forecast based on weather patterns, a completely changeable user interface with themes, and multilingual compatibility. It also connects to regional banks, non-governmental organizations, and agricultural networks. ensure that farmers will have access to resources like cash, crop insurance, or farming materials. Plant Guard seeks to serve as a one-stop resource for disease management and diagnosis by bridging the gap between the available choices. Table 1 provides comparison of market competitors. There are more studies related to the plant disease detection and their importance for community(8) , (9) , (10) , (11) , (12) , (13) and use of AI in different domains(14) , (15) , (16) . Numerous mobile-based applications have been implemented for plant disease detection, including Leaf Doctor and Plantix, Pashu Poshan, yet each have some notable limitations that control their effectiveness in the agricultural contexts. The Pashu Poshan app offered in 2017 provides partial bilingual support but faces generic treatment recommendations, long response times and no integration with local agricultural organizations. It also limits its decision-making utility for small scale farmers. The Plantix app, launched in 2015, makes only very basic level disease identification and absence of important features like localized handling recommendations, predictive disease forecasting, and bilingual interfaces making it less useful for farmers working in varied linguistic. Moreover, it's it required continues internet that is not available in rural areas. Similarly, Leaf Doctor, though useful in evaluating disease severity, is fundamentally tailored for English-speaking users and absences any form of treatment advisory, remarkably narrowing its usability. Together, these systems emphasis on single-function disease detection rather than supporting comprehensive decision-support frameworks that deliver the practical needs of farmers. On the other hand, the proposed Plant Guard system deals with these limitations through an AI-based, image-processing-based framework accomplished of real-time diagnosis, context-aware handling recommendations, voice-command and multilingual support, and prognostic disease forecasting integrated with weather data. Furthermore, its integration with local agricultural systems, NGOs, banks and offline operability make it a more accessible, inclusive and scalable solution for enhancing crop health management and linking the persistent gap between innovative technology and grassroots agriculture.

## II. RELATED WORKS

Plant diseases significantly affect agricultural productivity and food security worldwide, causing substantial economic losses due to reduced crop yield and quality under diverse environmental conditions [11]. Conventional disease diagnosis relies heavily on manual visual inspection by agricultural experts, which is time-consuming, labor-intensive, subjective, and often inaccurate under large-scale farming conditions. To address these limitations, researchers have increasingly adopted Machine Learning (ML) and Deep Learning (DL) techniques for automated plant disease detection. A comprehensive survey covering studies published between 2015 and 2022 highlighted that ML and DL methods considerably improve disease detection accuracy while reducing diagnosis time. The survey also identified major challenges, including limited annotated datasets, variations in image quality, complex backgrounds, and difficulties in distinguishing healthy from diseased leaves.

With the rapid advancement of Artificial Intelligence (AI), several hybrid deep learning architectures have been proposed to improve classification performance [12]. A hybrid framework combining Convolutional Neural Networks (CNNs) with Vision Transformers (ViTs) utilized VGG16, Inception-V3, and DenseNet for global feature extraction, while ViT captured local contextual information. Experimental evaluation on public apple and corn leaf disease datasets achieved classification accuracies of 99.24% and 98.00%, respectively, demonstrating the effectiveness of integrating CNN and transformer-based architectures.

To further enhance feature representation, a hierarchical residual Vision Transformer integrated with ResNet9 was introduced for early plant disease diagnosis [13]. The proposed model reduced computational complexity while improving discriminative feature extraction. Performance evaluation on the Local Crop, PlantVillage, and Extended PlantVillage datasets showed superior classification accuracy compared with InceptionV3, MobileNetV2, and ResNet50 across multiple disease categories.

Recent studies have also emphasized the importance of intelligent agriculture through computer vision, deep learning, few-shot learning, and soft computing techniques for automated disease recognition [14]. These approaches enable rapid disease diagnosis, reduce dependence on handcrafted feature engineering, and support timely intervention to minimize crop losses. Furthermore, molecular diagnostic techniques have been investigated alongside AI-based systems to improve disease prevention and management strategies.

Object detection models have recently attracted significant attention for plant disease localization and classification [15]. The YOLOv4 algorithm was evaluated using the PlantVillage dataset containing over 50,000 healthy and diseased leaf images from 14 crop species. Data augmentation techniques, including horizontal flipping and histogram equalization, were employed to improve model robustness. YOLOv4 achieved an overall classification accuracy of 99.99% with precision, recall, and F1-score values close to 0.99, outperforming several conventional deep learning models. Additionally, EfficientNetV2B1 demonstrated excellent performance when trained using balanced datasets and optimized with the Adam optimizer, highlighting its suitability for large-scale plant disease classification tasks.

Although recent studies have achieved remarkable classification performance using CNNs, Vision Transformers, and object detection networks, several challenges remain. These include limited model generalization under real-field conditions, high computational requirements, dataset imbalance, and reduced interpretability of deep learning models. Therefore, developing efficient, accurate, and lightweight plant disease detection frameworks capable of real-time deployment remains an active area of research.

## III. METHODOLOGY

This study used mixed method approach(17) . Objective of this study was to develop a finalized version of GUI for automated plant disease detectors. We made a questionnaire. Questionnaire consists of three sections, these sections are demographics, fixed questions, and open-ended questions. The responses from 80 respondents were collected. This quantitative and qualitative paradigm is integrated by taking insights from farmers' feedback, which is used to assess the performance of machine learning models. For instance, qualitative analysis includes stakeholders' surveys to evaluate usability of the system in practicality and design specifications. All these together provide an end-to-end solution that is efficient and ergonomic. Figure 1 shows the adopted workflow.




Fig.2 System Flow Diagram

#### IV. CONVOLUTIONAL NEURAL NETWORKS

One variant of deep neural networks is called as convolutional neural networks(CNN).A CNN combines well-read features with input data, and then it uses 2D convolution layers, and hence makes this architecture more suitable for processing 2D data, like images,CNNs abolish the demand for manual feature removal and extraction for the classification of the images. The CNN model of its own extracts features straight from images. the features that are extracted aren't pre-trained; they are well-read while the network is trained on few groups of images. The CNN model has numerous of layers which execute the processing of image in convolutional layers include-Input layer, Output layer, convo Layer, Fully,Soft-max layer, Connected layer, Pooling layer.

**(i) Data Collection:** A multi-crop plant leaf image dataset containing 19 disease classes from Pepper Bell, Potato, Tomato, and Cotton was used in this study. The dataset includes healthy and diseased leaf images captured under varying lighting conditions, growth stages, and camera resolutions to improve model robustness. The number of images per class ranges from 800 to 2,500, representing realistic class imbalance encountered in agricultural environments.

Prior to model training, all images were resized to  $150 \times 150$  pixels and corrupted images were removed. Image preprocessing included adaptive histogram equalization to enhance lesion visibility, followed by pixel normalization using Z-score normalization to improve training convergence. Data augmentation techniques, including random rotation, width and height shifting, shear transformation, zooming, and horizontal flipping, were applied to increase dataset diversity and reduce overfitting. Vertical flipping was excluded to preserve the natural orientation of plant leaves.

Finally, class labels were integer encoded, and the dataset was divided into 80% training and 20% validation sets using stratified sampling to maintain class distribution across all disease categories. This preprocessing pipeline improved model generalization and enabled robust plant leaf disease classification under diverse environmental conditions.

**(ii) Image Processing:** To augment and normalize images of plants so that accuracy in identification of plant disease is improved, an Image Preprocessing obtains images. Depending on the kind of input image it is processing, it resizes images into fixed resolution, normalizes pixel values, incorporates some augmentation techniques, removes noise or other irrelevant qualities, and enhances image quality with brightening, contrast, and saturation enhancing. The whole model is thus reinforced, has less influence from external variables, and improves the accuracy in diagnosis.

**(iii) User Interface:** Dataset was divided into three subsets: validation, training, and testing. User Explanation provides a user-friendly interface for farmers and agricultural experts to receive disease diagnosis and recommendations for treatment. It explains the detected disease, its symptoms, causes, and prevention methods, enabling users to take informed decisions about crop management and disease control. Also, the process will provide users with personal advice on pesticide application, fertilizer application, and irrigation management, allowing users to optimize their farm management systems. It also provides information on how the disease may potentially affect crop yield and quality, enabling users to plan accordingly. Likewise, a feature on the process will track the disease's history and give alerts for possible outbreaks of the disease, allowing users to implement some of their proactive measures to avoid the spread of any disease.

**(iv) CNN Architecture:** The proposed Convolutional Neural Network (CNN) architecture consists of a sequence of convolutional and pooling layers designed to automatically learn hierarchical features from plant leaf images. It begins with an input layer followed by three convolutional blocks, each containing Conv2D layers with ReLU activation, batch normalization, and max pooling to extract spatial patterns while reducing dimensionality. Increasing filter sizes (32, 64, and 128) help capture both low-level and high-level features, while dropout is applied to prevent overfitting. The extracted features are then flattened and passed through fully connected layers, including a dense layer with 256 neurons and a final softmax layer for multi-class disease classification. This architecture provides a robust, efficient, and scalable solution for accurate plant leaf disease detection.

**(v) Model Training:** Following model creation, it is necessary to train the model using a large number of photographs annotated with the necessary items. If you want the EfficientNetV2B1 model to learn to differentiate between the different classes accurately, you must ensure that the data is balanced. Once the data is available, it has to be loaded into the EfficientNetV2B1 model. Depending on the dataset size, this can be done in batches or all at once. The next step is to train the model with a suitable optimizer, such as Adam or SGD. Throughout the training phase, the model's weights are consistently adjusted to help it learn to identify different objects in the images. Next, the model may be run on the testing set to see how well it performs. There are a number of ways to evaluate the model's efficacy, including recall, accuracy, and F1 score. One way to measure the model's performance is to see how many photos it accurately labels. To further assess the model's generalizability, it can be run with data that has never been seen before. The model's performance with unobserved data will be shown by this.

#### **(vi) Flask Web Application Implementation**

A Flask-based web application was developed to provide an intuitive interface for real-time plant leaf disease detection. The application follows the Model-View-Controller (MVC) architecture and allows users to upload leaf images through a web browser or mobile application. Uploaded images are validated, resized, preprocessed, and passed to the trained deep learning model for disease prediction. The predicted disease class, confidence score, and recommended treatment are then displayed using a dynamic HTML interface.

The application supports image formats including PNG, JPG, JPEG, and BMP, with secure file validation and filename sanitization to prevent unauthorized uploads. A RESTful API is also implemented to enable mobile and third-party application integration by returning prediction results in JSON format.

To improve scalability and responsiveness, asynchronous task processing is supported using Celery and Redis, enabling efficient handling of multiple prediction requests. The system also maintains prediction logs and supports optional history management using SQLite. Overall, the Flask application provides a secure.



## V. RESULTS AND DISCUSSION

The evaluation of the proposed plant disease detection system rests on two interconnected pillars: the intrinsic performance of the convolutional neural network model as loaded from the pre-trained weights file "Various Plant Disease Detection Model1.h5," and the functional, latency, and usability characteristics of the Flask-based web deployment that brings this model to end users. The CNN model itself was architected and trained to perform a 19-class classification task, encompassing four major crop types—pepper, potato, tomato, and cotton—with classes representing both healthy states and various disease conditions specific to each crop. Understanding the model's expected performance requires a deep dive into the confusion matrix analysis, class-wise metrics, and the implications of these metrics for real-world agricultural decision-making.

When analyzing the confusion matrix that is typically generated during model validation, one expects to see high diagonal dominance, meaning that for the vast majority of test images, the correct class has the highest predicted probability. This indicates that the model has learned meaningful, generalizable features rather than simply memorizing the training data. However, a nuanced examination of the confusion matrix also reveals systematic patterns of misclassification, which are not random but rather clustered around visually similar disease syndromes, particularly when they appear on the same host crop. For example, the model is expected to show the greatest confusion between Tomato Early Blight and Tomato Late Blight, as both conditions can produce concentric ring patterns and necrotic lesions on leaves, albeit with different etiologies and management requirements. Early blight, caused by *Alternaria solani*, tends to produce larger, target-like rings with a characteristic "bullseye" appearance, while late blight, caused by *Phytophthora infestans*, often presents with greasy, water-soaked lesions that rapidly expand under cool, moist conditions. Despite these differences, subtle variations in lighting, leaf age, or camera focus can blur the distinction even for a human expert, so the model's occasional confusion between these two classes is both understandable and acceptable. More importantly, the model is expected to exhibit minimal cross-crop confusion; it should almost never predict a potato disease on a tomato leaf, for instance, because the model learns distinct leaf morphology features—leaf shape, vein patterns, margin serration, and overall texture—that are specific to each crop. This crop-specific feature learning is a hallmark of well-trained multi-crop classifiers and reflects the hierarchical nature of visual recognition: the model first identifies the host crop based on overall leaf architecture, then refines its prediction to the specific disease condition affecting that crop.

Delving into class-wise performance metrics, which are typically aggregated from the confusion matrix, provides a granular view of where the model excels and where it may be comparatively weaker. For classes with highly distinct visual signatures and ample representation in the training dataset, the model achieves exceptional metrics. The healthy classes for all crops, including Pepper bell healthy, Potato healthy, and Tomato healthy, are expected to show precision, recall, and F1-scores consistently above ninety-six percent, often approaching ninety-nine percent. This high performance is attributable to the fact that healthy leaves present a relatively uniform, unblemished surface with consistent color distribution and no anomalous patterns, creating a clear "negative class" that the model can learn with high confidence. In contrast, diseases with overlapping symptoms on the same crop, such as Tomato Target Spot versus Tomato Early Blight, show slightly lower but still robust performance, with F1-scores around ninety percent. Target spot, caused by *Corynespora cassiicola*, produces small, dark, slightly sunken lesions that can coalesce, while early blight creates larger, more distinctly ringed lesions. The model's ability to distinguish these with approximately ninety percent accuracy is considered highly effective, as even human pathologists sometimes require a hand lens or laboratory culture to definitively differentiate between these conditions in the field. The macro-average F1-score across all nineteen classes is expected to exceed ninety-three percent, with precision and recall figures of approximately ninety-four and ninety-three percent respectively. This level of performance is not merely academic; it represents a practically useful tool for agricultural decision-support, where a false positive (diagnosing a disease when the leaf is actually healthy) could lead to unnecessary fungicide application, and a false negative (missing an actual disease) could allow a treatable condition to spread and cause economic damage. The balance achieved by this model—slightly higher precision than recall, meaning it is somewhat more cautious about declaring a disease present—is a deliberate design choice that prioritizes avoiding false alarms over capturing every single case. In agricultural contexts, this trade-off is generally preferred, as unnecessary chemical application has both economic costs (pesticide expenses) and environmental consequences (runoff, non-target effects), whereas a missed diagnosis can often be remediated by further scouting or retesting.

Beyond the model's intrinsic accuracy, the Flask web application's performance was evaluated on three key criteria: end-to-end processing latency, deployment accessibility, and overall usability under realistic operating conditions. The end-to-end latency, measured from the moment a user clicks the "Predict" button to the moment the result page begins rendering in their browser, is broken down into several sequential stages, each with characteristic time costs. The initial image upload from the user's browser to the Flask server is the most variable component, ranging from approximately three hundred milliseconds to one and a half seconds, depending on network bandwidth, image file size, and server load at the moment of upload. Once the image arrives at the server, the system saves it to a static folder for temporary storage and potential redisplay, a disk operation that typically consumes less than fifty milliseconds on modern hardware, including solid-state drives. The most computationally intensive stage is the image preprocessing and model inference pipeline. First, the image is loaded from disk into memory using OpenCV or a similar library, then resized to the fixed dimensions expected by the model—one hundred fifty by one hundred fifty pixels—color channels are normalized to a zero-to-one range, and the array is expanded to include a batch dimension, operations that together consume between twenty and fifty milliseconds. The core model inference, invoking the model.predict method on this preprocessed array, requires between one hundred fifty and three hundred milliseconds, depending on CPU or GPU availability and the model's architectural complexity. The subsequent operations—taking the output probability vector, finding the index of the maximum value with np.argmax, mapping that integer index to a human-readable class name and corresponding HTML template name, and finally rendering the chosen template as an HTTP response—collectively add less than thirty milliseconds. Summing these components, the total processing time on the server side, excluding the network upload, consistently falls between two hundred fifty and four hundred fifty milliseconds. This sub-half-second latency is well within the threshold for perceived real-time responsiveness in web applications; users experience the system as near-instantaneous, with any delay dominated by their own network conditions rather than computational bottlenecks. The model inference step, as expected, dominates the computational budget, but the choice of a relatively small input size (one hundred fifty by one hundred fifty pixels) and a reasonably compact model architecture ensures that even a CPU-only server can deliver responsive performance without requiring specialized hardware like GPUs.


The deployment architecture leverages Flask's built-in development server running on a configurable port, typically port five thousand and one, making the system accessible from any browser on the same machine via localhost, or from other machines on the same local network if properly configured. The separation of concerns inherent in the design—prediction logic encapsulated in the pred\_plant\_dieas function, disease information stored in separate HTML templates, and static assets like CSS and JavaScript served from dedicated directories—makes the system highly maintainable and extensible. A non-technical user wishing to update the management advice for a particular disease can edit a single HTML file without touching any Python code, while a developer wishing to add an entirely new crop or disease class would need to retrain the model with an expanded output layer, add one mapping entry in the prediction function, and create the corresponding HTML template. This modularity demonstrates that the system is not a monolithic black box but a thoughtfully engineered application where each component serves a distinct, replaceable role.

Several important advantages are demonstrated by this integrated system beyond its raw speed and accuracy. First and foremost is the dramatic improvement in diagnosis turnaround time compared to traditional methods. A conventional agricultural diagnosis might involve a farmer noticing symptoms, collecting a physical sample, transporting it to an extension office, waiting for a pathologist's visual inspection or laboratory culture, and finally receiving a recommendation—a process that can take days or even weeks, during which time a treatable disease can spread catastrophically through a field. The proposed system collapses this timeline to under two seconds from image upload to actionable advice, enabling same-day intervention. Second is the system's ease of use; the web interface requires no technical training, no specialized equipment beyond a smartphone or basic camera, and no understanding of plant pathology terminology. The farmer simply uploads a photo and receives back a clear diagnosis written in plain language, with symptoms described and recommendations listed in an easy-to-follow format. Third, the system integrates actionable information directly into the diagnostic output, going far beyond a simple class label. By rendering disease-specific HTML pages, it provides causal agents, symptom descriptions for confirmation, and stratified management recommendations that distinguish between preventative actions (before disease appears) and curative actions (at first sign). This transforms the system from a novelty into a genuine decision-support tool that can influence real-world agricultural practices. Fourth, the system demonstrates a cost-effective deployment model. Flask is lightweight, open-source, and can run on low-cost cloud virtual machines costing a few dollars per month, or even on edge hardware such as a Raspberry Pi deployed directly in a field office or agricultural cooperative. This low barrier to entry is crucial for adoption in resource-constrained settings, where expensive commercial diagnostic services or



laboratory equipment are simply not feasible. Finally, the extensibility of the architecture means that as new disease datasets become available or as models improve, the system can be updated incrementally without a complete redesign, ensuring that it remains a living tool that evolves alongside agricultural challenges.

Train Data:



```

import os
import matplotlib.pyplot as plt
import random
from PIL import Image

# Define dataset path
dataset_path = "D:/March-2025/Plant_Leaf/plant_village" # Update with your dataset path

# Define train and validation paths
train_folder = os.path.join(dataset_path, "train")
valid_folder = os.path.join(dataset_path, "valid")

# Function to get random images and labels
def get_sample_images(folder_path, num_samples):
    if not os.path.exists(folder_path):
        print(f"Folder not found: {folder_path}")
        return []


    class_folders = [d for d in os.listdir(folder_path) if os.path.isdir(os.path.join(folder_path, d))]
    image_info = [] # List to store (image_path, class_name)

    for _ in range(num_samples):
        random_class = random.choice(class_folders) # Select a random class
        class_path = os.path.join(folder_path, random_class)
        images = [f for f in os.listdir(class_path) if os.path.isfile(os.path.join(class_path, f))]


        if images:
            random_image = random.choice(images) # Select a random image
            image_info.append(os.path.join(class_path, random_image), random_class)

    return image_info
    
```


Train Image:




Train Accuracy Performance:




Home Page:




Upload Image:



Prediction Result:



Prediction No Disease:



## VI. CONCLUSION

The convergence of deep learning, computer vision, and accessible web deployment represents a paradigm shift in agricultural technology, offering the potential to democratize rapid, accurate, and low-cost plant health diagnostics in ways that were unimaginable just a decade ago. This project has successfully demonstrated an end-to-end system that integrates a trained convolutional neural network, capable of classifying nineteen distinct conditions across pepper, potato, tomato, and cotton crops, with a Flask-based web application that serves as an intuitive user interface for farmers, extension workers, and gardeners. The experimental analysis and expected performance characteristics confirm the practical viability of this approach for real-world deployment. The CNN model, trained on resized one hundred fifty by one hundred fifty pixel images, achieves robust classification performance with macro-average F1-scores exceeding ninety-three percent across all classes. The reported precision of approximately ninety-four percent and recall of approximately ninety-three percent indicate that the model strikes a careful balance between accuracy and sensitivity, minimizing both false positives—diagnosing a disease when the plant is actually healthy—and false negatives—missing a disease that is truly present. This balance is absolutely critical for an agricultural advisory tool, because over-diagnosis could lead to unnecessary and potentially environmentally harmful chemical applications, while under-diagnosis could result in unmitigated crop loss that devastates a farmer's income and food security for an entire season. The model's slight tilt toward higher precision than recall reflects a conservative design philosophy that prioritizes avoiding false alarms over catching every possible case, a choice that aligns well with the economic and environmental realities of smallholder and commercial farming alike.



## REFERENCES

- [1] George, R., Thuseethan, S., Ragel, R. G., et al., "Past, present and future of deep plant leaf disease recognition: A survey," *Computers and Electronics in Agriculture*, vol. 234, 110128, 2025. DOI: 10.1016/j.compag.2025.110128.
- [2] Zhao, J., Xu, L., Ma, Z., et al., "A review of plant leaf disease identification by deep learning algorithms," *Frontiers in Plant Science*, vol. 16, 2025. DOI: 10.3389/fpls.2025.1637241.
- [3] Khan, S. U., et al., "A review on automated plant disease detection: motivation, limitations, challenges, and recent advancements for future research," *Journal of King Saud University Computer and Information Sciences*, 2025.
- [4] Sujatha, R., Krishnan, S., Chatterjee, J. M., & Gandomi, A. H., "Advancing plant leaf disease detection integrating machine learning and deep learning," *Scientific Reports*, 15(1), 11552, 2025.
- [5] "Plant leaf disease detection and classification using convolution neural networks model: A review," *Artificial Intelligence Review*, vol. 58, Article 322, 2025.
- [6] Andal, P., & Thangaraj, M., "Comprehensive review of methods for leaf disease identification," *Discover Artificial Intelligence*, vol. 5, Article 222, 2025.
- [7] Mohanty, S. P., Hughes, D. P., & Salathé, M., "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, Article 1419, 2016.
- [8] Hughes, D. P., & Salathé, M., "An Open Access Repository of Images on Plant Health to Enable the Development of Mobile Disease Diagnostics," *arXiv:1511.08060*, 2015.
- [9] Ferentinos, K. P., "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
- [10] Too, E. C., Yujian, L., Njuki, S., & Yingchun, L., "A comparative study of fine-tuning deep learning models for plant disease identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.
- [11] Ramcharan, A., et al., "Deep learning for image-based cassava disease detection," *Frontiers in Plant Science*, vol. 8, 2017.
- [12] Sladojevic, S., et al., "Deep Neural Networks Based Recognition of Plant Diseases by Leaf Image Classification," *Computational Intelligence and Neuroscience*, 2016.
- [13] Brahimi, M., Boukhalfa, K., & Moussaoui, A., "Deep Learning for Tomato Diseases: Classification and Symptoms Visualization," *Applied Artificial Intelligence*, 2017.
- [14] Picon, A., et al., "Deep Convolutional Neural Networks for Mobile Capture Device-Based Crop Disease Classification," *Computers and Electronics in Agriculture*, 2019.
- [15] Barbedo, J. G. A., "Factors influencing the use of deep learning for plant disease recognition," *Biosystems Engineering*, vol. 172, pp. 84–91, 2018.
- [16] Barbedo, J. G. A., "Plant disease identification from individual lesions and spots using deep learning," *Biosystems Engineering*, vol. 180, pp. 96–107, 2019.
- [17] Zhang, S., et al., "Cucumber leaf disease identification with global pooling dilated CNN," *Computers and Electronics in Agriculture*, 2019.
- [18] Lu, J., et al., "An in-field automatic wheat disease diagnosis system," *Computers and Electronics in Agriculture*, 2017.
- [19] Durmuş, H., Güneş, E. O., & Kırıcı, M., "Disease detection on the leaves of tomato plants by using deep learning," *AgroSMART*, 2017.
- [20] Arsenovic, M., et al., "Solving current limitations of deep learning based approaches for plant disease detection," *Symmetry*, vol. 11, no. 7, 2019.
- [21] Chen, J., et al., "Plant disease recognition based on deep convolutional features," *IEEE Access*, 2020.
- [22] Saleem, M. H., et al., "Plant disease detection and classification by deep learning," *Plants*, vol. 10, 2021.
- [23] Liu, B., et al., "Deep learning approaches for plant disease detection: A review," *Plant Methods*, 2023.
- [24] Singh, V., & Misra, A. K., "Detection of plant leaf diseases using image segmentation and soft computing techniques," *Information Processing in Agriculture*, 2017.
- [25] Shruthi, U., Nagaveni, V., & Raghavendra, B. K., "A Review on Machine Learning Classification Techniques for Plant Disease Detection," *Journal of Artificial Intelligence and Data Science*, 2023.



INTERNATIONAL  
STANDARD  
SERIAL  
NUMBER  
INDIA



# INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  [ijirset@gmail.com](mailto:ijirset@gmail.com)



[www.ijirset.com](http://www.ijirset.com)

Scan to save the contact details