

A High-Accuracy Automated Plant Disease Classification System Using CNN Architectures and Web Deployment

K. Prasath¹, Balika J. Chelliah^{2,*}

^{1,2}Department of Computer Science and Engineering, SRM Institute of Science and Technology, Ramapuram, Chennai, Tamil Nadu, India.

pk9515@srmist.edu.in¹, ballikaj@srmist.edu.in²

Abstract: Agricultural productivity faces significant challenges due to plant diseases caused by microscopic pathogens that are difficult to detect during early developmental stages. Traditional disease detection methods rely heavily on manual inspection by agricultural experts, which is time-consuming, expensive, and prone to human error. This research presents an automated plant disease detection system that leverages deep learning and transfer learning to identify and classify plant diseases with high accuracy. The proposed framework utilises the comprehensive Plant Village dataset, which contains 54,309 images spanning 14 crop species and 38 disease classes. Multiple convolutional neural network architectures, including AlexNet, VGG16, InceptionV3, and MobileNet, were implemented and evaluated using colour, grayscale, and segmented image representations. The models were trained using an 80-20 split between training and test data to ensure robust performance evaluation. Performance metrics, including accuracy, precision, recall, and F1-score, were computed to assess model effectiveness. Among all architectures, AlexNet achieved the highest performance, with 99.56% accuracy, 0.9953 precision, 0.9971 recall, and 0.9961 F1-score. The system is deployed as a web application using HTML, CSS, JavaScript, Keras, and Python, hosted on Heroku. This automated solution enables farmers to detect plant diseases at early stages with minimal cost, facilitating timely intervention and improved crop yield management.

Keywords: Plant Disease Detection; Deep Learning; Transfer Learning; Convolutional Neural Networks; Precision Agriculture; Image Classification; Agricultural Productivity; Crop Yield Management.

Received on: 28/01/2025, **Revised on:** 04/04/2025, **Accepted on:** 11/06/2025, **Published on:** 06/12/2025

Journal Homepage: <https://www.fmdbpublish.com/user/journals/details/FTSHSL>

DOI: <https://doi.org/10.69888/FTSHSL.2025.000512>

Cite as: K. Prasath and B. J. Chelliah, “A High-Accuracy Automated Plant Disease Classification System Using CNN Architectures and Web Deployment,” *FMDB Transactions on Sustainable Health Science Letters*, vol. 3, no. 4, pp. 231–245, 2025.

Copyright © 2025 K. Prasath and B. J. Chelliah, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/), which permits copying, remixing, redistributing, transformation, and building upon the material in any medium so long as the original work is properly cited.

1. Introduction

Agriculture remains the cornerstone of global food security and economic stability, employing approximately 50% of India's workforce and contributing substantially to the nation's gross domestic product. India is the world's largest producer of pulses, rice, wheat, and spices, with agricultural output directly influencing both domestic food availability and international trade. The economic prosperity of farming communities fundamentally depends on the quality and quantity of agricultural produce,

*Corresponding author.

which are intrinsically linked to plant health and disease management throughout the cultivation cycle. Plant diseases represent one of the most critical threats to agricultural sustainability and productivity worldwide. According to the Food and Agriculture Organisation (FAO) of the United Nations, transboundary plant pests and diseases affect food crops, causing significant losses for farmers worldwide and threatening the food security of millions of people. These diseases, predominantly caused by fungal, bacterial, viral, and parasitic pathogens, can devastate entire harvests if not detected and managed promptly. The symptoms of plant diseases manifest in various forms, including leaf discolouration, spot formation, wilting, stunted growth, and tissue necrosis, with leaves being the most affected and visually accessible plant organ for disease diagnosis. The conventional approach to plant disease detection and diagnosis relies primarily on visual inspection by agricultural experts and plant pathologists. This traditional methodology involves periodic field visits during which trained professionals examine plants for visible symptoms of disease. While this approach has been practised for centuries and continues to serve as the primary diagnostic method in many regions, it suffers from several critical limitations that impede effective disease management in modern large-scale agriculture. First, manual inspection is inherently labour-intensive and time-consuming, particularly when monitoring extensive agricultural lands comprising thousands of plants. Second, the accuracy of visual diagnosis depends heavily on the inspector's expertise and experience, introducing subjective variability and the potential for misdiagnosis.

Third, many disease-causing microorganisms are microscopic and invisible to the naked eye during the initial infection stages, becoming apparent only after significant damage has occurred. Fourth, the shortage of trained agricultural experts, especially in rural and developing regions, creates accessibility barriers for farmers seeking timely disease diagnosis. Finally, the cost of hiring experts for regular farm inspections places a financial burden on small- and medium-scale farmers, making preventive disease management economically unfeasible for many agricultural operations. The advent of artificial intelligence (AI) and machine learning (ML) technologies has revolutionised numerous industrial and scientific domains, offering unprecedented opportunities for automation, efficiency enhancement, and intelligent decision-making. Within the agricultural sector, AI-driven solutions have emerged as powerful tools addressing various challenges, including crop yield prediction, soil quality assessment, irrigation optimisation, and pest management. Among these applications, automated plant disease detection and classification using computer vision and deep learning techniques have attracted significant research attention due to their potential to transform traditional agricultural practices and enable precision farming. Deep learning, a specialised subset of machine learning inspired by the structure and function of biological neural networks, has demonstrated remarkable capabilities in image recognition, pattern identification, and feature extraction tasks. Convolutional Neural Networks (CNNs), a class of deep learning architectures specifically designed for processing grid-like data such as images, have achieved human-level or superhuman performance in various visual recognition benchmarks. The hierarchical feature-learning capability of CNNs, which automatically extracts low-level features such as edges and textures in the initial layers and progressively learns higher-level semantic features in deeper layers, makes them particularly suitable for complex image classification tasks, including plant disease identification from leaf images.

Transfer learning, an advanced deep learning technique that leverages knowledge gained from solving one problem to address related but different problems, has further accelerated progress in plant disease detection research. Pre-trained CNN models such as AlexNet, VGG16, InceptionV3, ResNet, and MobileNet, originally trained on large-scale image datasets like ImageNet containing millions of labelled images across thousands of categories, can be fine-tuned on smaller domain-specific plant disease datasets. This approach significantly reduces training time, computational resource requirements, and the amount of labelled data needed to achieve high classification accuracy, making deep learning solutions more accessible and practical for agricultural applications. Recent research in automated plant disease detection has explored various methodologies, including traditional machine learning approaches using handcrafted features combined with classifiers such as Support Vector Machines (SVMs), Random Forests, and k-Nearest Neighbours, as well as modern deep learning techniques employing CNN architectures of varying complexity. While traditional machine learning methods have shown promising results in controlled experimental settings, they require extensive domain expertise for feature engineering and often struggle to generalise across different plant species, disease types, and imaging conditions. In contrast, deep learning approaches automatically learn discriminative features directly from raw image data, demonstrating superior performance and robustness across diverse datasets and real-world scenarios. Despite significant advances in the field, several challenges persist in developing practical and deployable plant disease detection systems.

These challenges include the need for large, diverse, and well-annotated training datasets covering multiple crop species and disease classes; addressing variability in image quality, lighting conditions, and background complexity in field-captured images; ensuring computational efficiency for real-time disease detection on resource-constrained devices; and developing user-friendly interfaces that enable farmers with limited technical expertise to utilise AI-powered diagnostic tools effectively. The primary objective of this research is to develop a robust, accurate, and accessible automated plant disease detection system using state-of-the-art deep learning and transfer learning techniques. Specifically, this work aims to: (1) implement and evaluate multiple pre-trained CNN architectures including AlexNet, VGG16, InceptionV3, and MobileNet for plant disease classification; (2) conduct comprehensive performance analysis using metrics such as accuracy, precision, recall, and F1-score across different image representations including color, grayscale, and segmented images; (3) identify the optimal model

architecture and training configuration that maximizes classification accuracy while maintaining computational efficiency; (4) deploy the trained model as an intuitive web application accessible to farmers for real-time disease diagnosis; and (5) establish a foundation for future enhancements including recommendation systems for disease-specific fertilizers and treatments.

The key contributions of this research include: (1) comprehensive comparative analysis of four prominent CNN architectures on a large-scale plant disease dataset comprising 54,309 images across 38 disease classes and 14 crop species; (2) achievement of 99.56% classification accuracy using the AlexNet architecture, demonstrating the effectiveness of transfer learning for agricultural disease detection; (3) systematic evaluation of model performance across different data partitioning ratios (80-20, 70-30, 60-40) to assess generalization capability and overfitting tendencies; (4) development and deployment of a functional web-based application that democratizes access to AI-powered plant disease diagnostics; and (5) detailed documentation of methodology, results, and insights to facilitate reproducibility and inspire future research in precision agriculture and computer vision applications. The remainder of this paper is organized as follows: Section II provides a comprehensive review of related literature examining previous research in plant disease detection using machine learning and deep learning approaches; Section III describes the methodology including dataset characteristics, preprocessing techniques, model architectures, and training procedures; Section IV details the data collection process, augmentation strategies, and dataset partitioning; Section V presents experimental results, performance metrics, and comparative analysis of different models; and Section VI concludes the paper with a summary of findings, research contributions, limitations, and directions for future work.

2. Review of Literature

Shrivastava et al. [1] explored transfer learning of a deep CNN for the first time to classify rice plant diseases. Moreover, the experiments were conducted by partitioning the entire dataset into different training-test ratios. The proposed model achieves 91.37% classification accuracy on an 80%-20 % training-test split. Benchmarking the proposed model against the literature is not appropriate because standard labelled rice disease images are unavailable, which is a disadvantage of this paper. Sabrol and Satish [2] follow the greedy approach, in which classification trees are constructed in a top-down, recursive, divide-and-conquer manner. The classification tree uses a top-down approach, starting with the tuples of the training set and their associated class labels. The major drawback of this model is the technique used; more sophisticated classification techniques are available, such as Adaptive Neuro-Fuzzy, Neural Networks, and Genetic Algorithms—support vector machines, etc., for image classification. Elangovan and Nalini [3] discuss various techniques to segment the disease part of the plant. This paper discussed characterisation strategies for extracting elements from contaminated leaves and for classifying plant infections using an SVM classifier. The downside of this paper is that, though SVMs can achieve high precision, they cannot achieve high recall. It will fall flat when applied to a large dataset of various species.

Sladojevic et al. [4] have developed a model that can recognise leaf presence, and Marko Arsenovic has developed a model that recognises sound leaves and 13 distinct illnesses, which can be externally analysed. This model neglected to account for plant diseases when using input images of larger land regions, consolidating aerial photographs of plantations and grape plantations captured by robots and convolutional neural networks for object localisation. Aziz et al. [5] have developed a framework capable of extracting useful image features and classifying disease type. The maximum average accuracy achieved during experimentation is 94%. The drawback of this paper is that it only utilises texture features. It does not include other classification features, such as shape- and colour-based features and texture features, to achieve even better performance. Chollet [6] compares Inception v3 and Xception, showing large gains on the JFT dataset and small gains in performance on ImageNet. Due to its many layers, it can lead to overfitting. They didn't perform hyperparameter tuning. Padol and Yadav [7] extracted both colour and texture features to identify the type of disease affecting the leaf. The proposed system can successfully detect and examine the disease with an accuracy of 88.89%. But here, this model does not consider other classification features, such as shape, and some datasets need to include more colour- and texture-based features to enhance the system's performance.

Chakraborty et al. [8] have used SVM to classify the healthy and unhealthy apple leaves. The proposed system takes at most 3 seconds to classify leaves and achieves 96% accuracy. Although it has high efficiency, the dataset is too small; a larger dataset is required to improve the model, and the model needs to run on a faster computer. Effective algorithms are applied to get better accuracy. Bhagat et al. [9] have classified healthy and unhealthy leaves using the libraries that researchers used to classify the images. The dataset that researchers have tested and validated has a test accuracy of 96.78%. Lack of image quality and features not present in the training data have led to some loss of accuracy in the model. Kirti and Rajpal [10] classify healthy and diseased parts, enabling easier disease detection based on colour. The radial basis function kernel achieves the highest accuracy of 94.1%. The main drawback of this model is that the dataset is large, as the colour and image quality should be high to enhance the model's performance. Tarik et al. [11] have processed images of 2034 unhealthy potato leaves. The system achieves 99.23% accuracy in testing. The model's overall performance is around 98.36%. The drawback of this model is that it only classifies unhealthy potato leaves, and the dataset must be larger [12].

The system uses images for both training and testing. The colour image is preprocessed and undergoes k-means clustering segmentation. The texture feature is extracted from the GLCM method. The accuracy of this method is 90%. The disadvantages of this model are that researchers get output for only a few samples, and the samples that are not processed are from different categories, which reduces the model's performance [21]. This paper focuses on the algorithm used, which correctly distinguishes between the healthy part of a Chilli plant, the area affected by cucumber mosaic, and the background [22]. A high accuracy is obtained for both the classification of background images and the healthy parts of the plant. However, only 57.1% accuracy was achieved in classifying cucumber mosaic. Fewer training datasets affected by cucumber mosaic virus have reduced the model's accuracy [14]. Artificial Intelligence is used for automation and to improve efficiency. The disease is already detectable at an early stage, and it can be diagnosed using deep learning techniques. Here, researchers have used the VGG-16 architecture to train a model [13]. The accuracy is 96%. The disease identified at an earlier stage has a good outcome, but due to poor-quality images or a lack of datasets, the loss will be around 86%, which is a bad sign.

Saranya et al. [15] helped farmers analyse plant growth and automatically detect diseases at minimal cost. Therefore, a system which uses image processing techniques to detect diseases at an earlier stage by an ANN algorithm. The imported dataset contains images with missing plant information. Ramesh and Vydeki [16] proposed a system to classify the disease-affected parts of the rice leaf and the healthy parts of the rice leaf. The system achieved 99% accuracy on blast-infected images and 100% on normal images during training. The testing-phase accuracy is 90% and 86% for infected and healthy images, respectively. To ensure high image quality, images should be captured with high-saturation pixels. Faster computers are required to enhance the performance [17]. The evolution of plant disease detection research reveals a clear trajectory from traditional image processing and manual feature engineering toward automated deep learning approaches offering superior accuracy and generalisation [18]. However, persistent challenges remain, including the need for large, diverse datasets, performance validation under real-world field conditions, computational efficiency for mobile deployment, integration with agricultural decision-support systems, and the development of user interfaces accessible to farmers with varying levels of technical literacy [19]. The present research builds upon this rich foundation, implementing and comparing multiple state-of-the-art CNN architectures on a comprehensive multi-crop disease dataset to advance toward practical, deployable plant disease detection systems [20].

3. Data Collection

The foundation of any robust deep learning model lies in the quality, diversity, and comprehensiveness of the dataset used for training and validation. For this research, the researchers utilised the Plant Village dataset, one of the most widely recognised and extensively used open-access repositories for plant disease identification research. The PlantVillage dataset was specifically designed to enable the development of mobile disease diagnostics and automated plant health monitoring systems, making it an ideal choice for our study.

3.1. Data Source and Dataset Overview

The PlantVillage dataset is an open-access repository that contains a comprehensive collection of plant health images captured under controlled laboratory conditions. This dataset was created through a collaborative effort involving agricultural experts, plant pathologists, and computer vision researchers to provide a standardised benchmark for plant disease detection algorithms. The dataset comprises 54,309 high-resolution images spanning 14 different crop species, including Apple, Blueberry, Cherry, Grape, Orange, Peach, Bell Pepper, Potato, Raspberry, Soybean, Squash, Strawberry, and Tomato. These crop species were selected based on their economic importance, global cultivation prevalence, and susceptibility to various diseases. The dataset encompasses a wide spectrum of plant health conditions, including both healthy and diseased leaf images. Specifically, it contains images of 17 fungal diseases, four bacterial diseases, two mould (oomycete) diseases, two viral diseases, and 1 disease caused by pests. Additionally, 12 crop species have images of healthy leaves unaffected by disease, serving as the baseline for comparison. This comprehensive coverage of disease types and crop species makes the PlantVillage dataset particularly valuable for developing generalised plant disease detection models that can be applied across multiple agricultural contexts. The images in the PlantVillage dataset were captured using standardised photography protocols to ensure consistency in image quality, lighting conditions, and background uniformity. All images were taken against plain backgrounds (typically white or light-coloured) to minimise noise and facilitate easier feature extraction during training. This controlled environment approach, while potentially limiting real-world applicability, provides an excellent foundation for initial model development and algorithm validation.

3.2. Image Characteristics and Quality

Each image in the PlantVillage dataset is a high-resolution colour photograph with dimensions that vary but typically range from 256 x 256 pixels to higher resolutions. The images are stored in JPEG format, which provides a good balance between image quality and file size. The colour depth of the images is 24-bit RGB (Red, Green, Blue), allowing for accurate

representation of disease symptoms that often manifest as colour changes in plant tissues. The image characteristics are particularly important for deep learning models, as they directly influence the model's ability to extract meaningful features. The PlantVillage images exhibit several key characteristics that make them suitable for deep learning applications. First, the images clearly show distinct disease symptoms, including leaf spots, discolouration, necrosis, wilting, and other pathological manifestations. Second, controlled lighting conditions ensure that colour information is consistent across images, reducing variability that might arise from different illumination conditions. Third, the plain backgrounds eliminate confounding factors that might distract the model from focusing on the relevant leaf features. For our experiments, researchers utilised colour images as the primary input format, as preliminary analysis indicated that colour information is crucial for distinguishing between different disease types. Many plant diseases manifest as specific colour patterns, for example, yellowing (chlorosis), browning (necrosis), or coloured spots and lesions. Grayscale images, while computationally less expensive to process, would lose this critical colour information. Researchers also experimented with segmented images, isolating the leaf region from the background. Still, researchers found that colour images without segmentation provided the best balance between preprocessing complexity and model performance.

3.3. Class Distribution and Dataset Composition

The PlantVillage dataset contains 38 distinct classes, representing different combinations of crop species and health conditions. The class distribution is relatively balanced, though some classes contain more images than others due to differences in disease prevalence and the availability of sample specimens during dataset creation. Each class represents a unique combination of plant species and disease type, such as "Tomato Early Blight," "Apple Cedar Apple Rust," "Potato Late Blight," or "Grape Healthy." The distribution of images across classes is an important consideration for training deep learning models, as class imbalance can lead to biased predictions that favour the majority classes. In our dataset, the number of images per class ranges from approximately 1,000 to 2,000, providing sufficient samples for effective learning. However, to address residual class imbalance and further enhance model robustness, researchers employed data augmentation techniques, as described in the following subsection. The diversity of disease types and crop species in the dataset presents both opportunities and challenges. On the one hand, it allows the model to learn generalised features applicable across different plant species and disease conditions. On the other hand, it increases the complexity of the classification task, as the model must distinguish between 38 classes, some of which may share visually similar symptoms.

3.4. Data Augmentation Techniques

Data augmentation is a critical technique in deep learning that artificially expands the training dataset by creating modified versions of existing images. This approach helps to improve model generalisation, reduce overfitting, and enhance the model's ability to handle variations in real-world scenarios. For our plant disease detection system, researchers implemented several data augmentation techniques to increase the effective size and diversity of our training dataset. The augmentation techniques applied include: (1) Rotation: Images were randomly rotated by angles ranging from 15° to $+15^\circ$ to simulate different orientations of leaves during image capture. (2) Horizontal and Vertical Flipping: Images were randomly flipped horizontally and vertically to account for different leaf orientations and camera positions. (3) Zoom: Random zoom operations (ranging from 0.9 to 1.1) were applied to simulate different distances between the camera and the leaf. (4) Width and Height Shifts: Images were randomly shifted horizontally and vertically by up to 10% of the image dimensions to simulate off-centre leaf positioning. (5) Brightness Adjustment: Random brightness variations were introduced to simulate different lighting conditions that might be encountered in field settings. (6) Shear Transformation: Slight shear transformations were applied to simulate different viewing angles. These augmentation techniques were applied randomly during training, ensuring the model encountered a diverse set of image variations at each epoch. This approach significantly enhances the model's robustness and generalisation to new, unseen images captured under different conditions.

3.5. Dataset Partitioning and Validation Strategy

Proper dataset partitioning is essential for training robust machine learning models and for obtaining reliable performance estimates. For our experiments, researchers adopted a systematic approach to dataset partitioning that balances the need for sufficient training data with the requirement for comprehensive model evaluation. Researchers divided the Plant Village dataset into training and test sets using an 80-20 split, allocating 80% (approximately 43,447 images) to the training set and 20% (approximately 10,862 images) to the test set. This split was stratified, ensuring that each class maintained the same proportional representation in both the training and test sets. Stratified splitting is crucial for maintaining class balance and ensuring that the model is evaluated on a representative sample of all disease types. To further validate our model and to assess its generalisation capability, researchers also experimented with different training-testing ratios, including 70-30 and 90-10 splits. These experiments allowed us to evaluate how the amount of training data affects model performance and to identify potential overfitting issues. The results across these different partitioning strategies were consistent, indicating that our models did not suffer from significant overfitting. In addition to the train-test split, researchers used 5-fold cross-validation during model

development. Cross-validation involves dividing the training data into five equal subsets (folds), training the model on four folds, and validating it on the remaining fold. This process is repeated five times, with each fold serving as the validation set once. The performance metrics are then averaged across all five folds to obtain a more robust estimate of model performance. Cross-validation helps to ensure that the model's performance is not dependent on a particular train-test split and provides greater confidence in the model's generalisation ability. The careful attention to data collection, preprocessing, validation, and partitioning strategies laid a solid foundation for the development of our plant disease detection system. By leveraging the comprehensive PlantVillage dataset and implementing rigorous data-handling procedures, researchers ensured that our deep learning models were trained on high-quality, diverse, and representative data, resulting in superior performance in plant disease classification.

4. Methodology

This section presents a comprehensive description of the methodological framework used to develop the automated plant disease detection system. The methodology encompasses dataset selection and preprocessing, the implementation of multiple CNN architectures via transfer learning, training procedures, evaluation metrics, and the system deployment architecture.

4.1. Dataset Selection and Characteristics

The research utilises the Plant Village dataset, a publicly accessible open-source repository specifically curated for plant disease recognition research. This dataset serves as the foundation for training and evaluating deep learning models due to its comprehensive coverage of multiple crop species and disease types. The Plant Village dataset contains 54,309 high-resolution digital images of healthy and diseased plant leaves, captured under controlled laboratory conditions with uniform backgrounds and consistent lighting, to minimise confounding variables during initial model development. The dataset comprises 14 distinct crop species: Apple, Blueberry, Cherry, Grape, Orange, Peach, Bell Pepper, Potato, Raspberry, Soybean, Squash, Strawberry, and Tomato. These species represent economically significant crops with substantial global cultivation and trade value. Disease coverage includes 26 distinct disease classes caused by various pathogenic organisms: 17 fungal diseases, representing the most common category of plant pathogens; four bacterial diseases, including bacterial spot and blight variants; two oomycete diseases (water moulds); two viral diseases affecting plant tissues; and one pest-induced disease. Additionally, the dataset includes images of healthy leaves from 12 crop species, providing essential negative examples for training balanced classifiers that can distinguish diseased from healthy plant tissue.

4.2. Image Preprocessing Pipeline

Effective preprocessing is crucial for standardising input data and enhancing model training efficiency. The preprocessing pipeline implemented in this research consists of several sequential stages:

- **Image Resizing and Normalisation:** All images are resized to uniform dimensions compatible with the input requirements of the selected CNN architectures. For AlexNet, images are resized to 227×227 pixels, while VGG16 and InceptionV3 require 224×224 and 299×299 pixels, respectively. Pixel intensity values originally in the range $[0, 255]$ are normalised to $[0, 1]$ by dividing by 255, facilitating faster convergence during gradient-based optimisation.
- **Colour Space Transformations:** The dataset is processed in three distinct representations to evaluate their impact on classification performance:
 - Colour images in the RGB colour space, preserving all spectral information,
 - Grayscale images reduce computational complexity while retaining spatial and intensity patterns,
 - Segmented images where diseased regions are isolated from healthy tissue and background using threshold-based or clustering segmentation algorithms.
- **Data Augmentation:** To enhance model generalisation and prevent overfitting, geometric and photometric augmentation techniques are applied during training, including random rotation (15°), horizontal and vertical flipping, random cropping with subsequent resizing, brightness adjustment (20%), and contrast variation. These augmentations artificially expand the training dataset while exposing models to realistic variations encountered in field conditions.

4.3. Deep Learning Architectures

Four prominent CNN architectures are implemented using transfer learning, leveraging pre-trained weights from ImageNet and fine-tuning on the plant disease dataset:

- **AlexNet Architecture:** AlexNet consists of five convolutional layers followed by three fully connected layers, totalling approximately 60 million parameters. The architecture employs ReLU (Rectified Linear Unit) activation functions defined as:

$$f(x) = \max(0, x) \quad (1)$$

This introduces nonlinearity while mitigating vanishing gradients. Max pooling layers with 3x3 kernels and a stride of 2 reduce the spatial dimensions while retaining dominant features. Dropout regularisation with probability 0.5 is applied to fully connected layers to prevent overfitting.

- **VGG16 Architecture:** VGG16 implements a deeper architecture with 16 weight layers, consisting of 13 convolutional layers using exclusively 3x3 filters and three fully connected layers. The uniform small filter size enables learning of complex feature hierarchies while maintaining manageable parameter counts per layer. The architecture totals approximately 138 million parameters, requiring substantial computational resources but offering enhanced representation capacity.
- **InceptionV3 Architecture:** InceptionV3 employs inception modules that perform multiple parallel convolutional operations with different filter sizes (1x1, 3x3, 5x5) and concatenate the results, enabling multi-scale feature extraction within a single layer. This architecture achieves approximately 24 million parameters by efficiently factorising convolutions and aggressively using dimensionality reduction. Auxiliary classifiers integrated at intermediate layers help mitigate vanishing gradients during backpropagation in deep networks.
- **MobileNet Architecture:** MobileNet utilises depth-wise separable convolutions, decomposing standard convolutions into depth-wise and point-wise operations, significantly reducing computational cost and parameter count. This efficiency-focused architecture contains approximately 4.2 million parameters, making it suitable for deployment on mobile and embedded devices with limited computational capabilities (Figure 1).




Figure 1: System architecture of our proposed model

4.4. Transfer Learning Implementation

Transfer learning leverages knowledge encoded in pre-trained models, accelerating convergence and improving performance on domain-specific tasks with limited training data. The implementation follows a systematic approach:

- **Weight Initialisation:** Convolutional layer weights are initialised with values pre-trained on ImageNet, containing 1.2 million images across 1,000 object categories. These weights capture general visual features, including edges, textures, shapes, and object parts, and are transferable to plant disease recognition.

- **Architecture Modification:** The original classification layers designed for 1,000 ImageNet classes are replaced with custom fully connected layers terminating in a softmax output layer with 38 neurons corresponding to the 38 plant disease classes in the dataset.
- **Fine-Tuning Strategy:** During the initial training epochs, freeze the convolutional layer weights and train only the newly added classification layers. Subsequently, all layers are unfrozen, enabling end-to-end fine-tuning with a reduced learning rate to prevent catastrophic forgetting of pre-trained features.

4.5. Training Configuration

The models are trained using the following configuration parameters and procedures. Loss Function: categorical cross-entropy loss is employed, defined as:

$$L = -\sum_{i=1}^C y_i \log(\hat{y}_i) \quad (2)$$

Where C represents the number of classes (38), y_i denotes the true label (one-hot encoded), and $\hat{y}_i$ represents the predicted probability for class i . Optimisation Algorithm, Adam (Adaptive Moment Estimation) optimiser is utilised, computing adaptive learning rates for each parameter based on first and second moment estimates of gradients:

$$\theta_{t+1} = \theta_t - \frac{\alpha \hat{m}_t}{\sqrt{\hat{v}_t + \epsilon}} \quad (3)$$

Where θ represents model parameters, α is the learning rate (initially set to 0.001), and $\hat{m}_t$ and $\hat{v}_t$ are bias-corrected first. Second-moment estimates, with ϵ a small constant to prevent division by zero. Training Protocol, models are trained for 50 epochs with a batch size of 32, employing early stopping monitoring validation loss to prevent overfitting. Learning rate scheduling reduces the learning rate by a factor of 0.1 whenever the validation loss plateaus for five consecutive epochs.

4.6. Evaluation Metrics

Comprehensive performance assessment employs multiple metrics capturing different aspects of classification performance:

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \quad (4)$$

$$\text{Precision} = \frac{TP}{TP+FP} \quad (5)$$

$$\text{Recall} = \frac{TP}{TP+FN} \quad (6)$$

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (7)$$

Where TP, TN, FP, and FN represent true positives, true negatives, false positives, and false negatives, respectively.

4.7. System Architecture and Deployment

The complete system architecture consists of three integrated components:

- **Frontend Interface:** Developed using HTML5, CSS3, and JavaScript, providing an intuitive web interface enabling users to upload plant leaf images for disease diagnosis. Responsive design principles ensure compatibility across desktop, tablet, and mobile devices.
- **Backend Processing:** Implemented in Python using the Keras deep learning framework with TensorFlow backend. The server receives uploaded images, preprocesses them, performs model inference, and returns classification results with associated confidence scores.
- **Cloud Deployment:** The application is deployed on the Heroku cloud platform, providing scalable infrastructure, automated deployment pipelines, and global accessibility without requiring users to install specialised software or maintain local computational resources.

The modular architecture facilitates future enhancements, including the integration of recommendation engines that suggest disease-specific treatments and fertilisers, the development of mobile applications for Android and iOS platforms, the

incorporation of multilingual support for regional languages, and the implementation of real-time disease monitoring with temporal trend analysis.

5. Result and Discussion

The performance evaluation of our plant disease detection system involved a comprehensive analysis using multiple deep learning architectures and various performance metrics. This section presents detailed results from our experiments, comparative analysis among different models, and a discussion of the implications of our findings.

5.1. Experimental Setup and Training Configuration

All experiments were conducted on a high-performance computing system equipped with NVIDIA GPU acceleration to facilitate efficient training of deep neural networks. The training process utilised the Adam optimiser with an initial learning rate of 0.001, which was dynamically adjusted using a learning rate scheduler that reduced the learning rate by a factor of 0.1 when validation loss plateaued for more than five consecutive epochs. The batch size was set to 32 images, and training was run for a maximum of 50 epochs, with early stopping to prevent overfitting. Early stopping monitored the validation loss and terminated training if no improvement was observed for 10 consecutive epochs. The loss function used was categorical cross-entropy, which is appropriate for multiclass classification. Data augmentation was applied in real-time during training to increase the effective size of the training dataset and improve model generalisation. All models were initialised with pre-trained weights from ImageNet, and transfer learning was employed by freezing the initial convolutional layers and fine-tuning the deeper layers along with the fully connected classification layers.

5.2. Performance Metrics and Evaluation Criteria

To comprehensively evaluate the performance of our implemented models, researchers utilised multiple performance metrics that capture different aspects of classification performance. The primary metrics include:

- **Accuracy:** The proportion of correctly classified images out of the total number of images. Accuracy provides an overall measure of model performance, but it can be misleading when there is class imbalance.
- **Precision:** The proportion of true positive predictions out of all positive predictions made by the model. Precision indicates the proportion of predicted disease cases that are correct.
- **Recall (Sensitivity):** The proportion of true positive decisions out of all actual positive cases. Recall measures the model's ability to identify all instances of a particular disease.
- **F1 Score:** The harmonic mean of precision and recall, providing a balanced measure that accounts for both false positives and false negatives.
- **Training and Validation Loss:** Measures of how well the model fits the training data and generalises to unseen validation data.
- **Training Time per Epoch:** The computational time required to complete one full pass through the training dataset, which is important for practical deployment considerations.

5.3. Comparative Performance Analysis of Deep Learning Models

The researchers implemented and evaluated four state-of-the-art convolutional neural network architectures: AlexNet, VGG16, MobileNet, and Inception V3. Each architecture has distinct characteristics in terms of depth, parameter count, and computational requirements:

Performance with 80-20 Train-Test Split: Using an 80-20 split of the PlantVillage dataset (80% for training and 20% for testing), researchers obtained the following accuracy results on colour images:

- **AlexNet:** 99.56%
- **Inception V3:** 98.42%
- **VGG16:** 97.89%
- **MobileNet:** 97.02%

AlexNet achieved the highest classification accuracy of 99.56%, demonstrating superior performance compared to the other architectures. This result is particularly noteworthy given that AlexNet is a relatively older and simpler architecture compared to VGG16 and Inception V3. The excellent performance of AlexNet can be attributed to several factors: (1) the controlled nature of the PlantVillage dataset with plain backgrounds makes it well-suited for AlexNet's feature extraction capabilities, (2)

the transfer learning approach with Pre-trained ImageNet weights provided a strong initialisation, and (3) the architecture's balance between model complexity and generalisation ability proved optimal for this specific task (Table 1).

Table 1: Performance metrics for different deep learning models

Model	Accuracy	Precision	Recall	F1 Score	Time (s/epoch)
AlexNet	99.56%	0.9953	0.9971	0.9961	545
Inception V3	98.42%	0.9821	0.9856	0.9838	782
VGG16	97.89%	0.9765	0.9802	0.9783	892
MobileNet	97.02%	0.9678	0.9724	0.9701	565

Inception V3 achieved the second-highest accuracy of 98.42%, demonstrating strong performance with its multi-scale feature extraction through inception modules (Table 2).

Table 2: Model performance across different train-test splits

Model	70-30 Split	80-20 Split	90-10 Split
AlexNet	99.48%	99.56%	99.61%
Inception V3	98.31%	98.42%	98.54%
VGG16	97.76%	97.89%	98.02%
MobileNet	96.89%	97.02%	97.18%

The architecture's ability to capture features at different scales is particularly beneficial for plant disease detection, where disease symptoms can manifest at various sizes and patterns. VGG16, despite its deeper 16-layer architecture, achieved an accuracy of 97.89%. While VGG16 is known for its excellent performance on various image classification tasks, its relatively lower performance in our experiments may be attributed to potential overfitting, given its large number of parameters (approximately 138 million) relative to the dataset size. MobileNet achieved an accuracy of 97.02%, which, while lower than the other architectures, is still highly competitive. MobileNet is specifically designed for mobile and embedded applications with limited computational resources. Its slightly lower accuracy is a trade-off for significantly reduced computational requirements and faster inference times, making it an excellent choice for deployment on mobile devices:

- **Detailed Performance Metrics:** Table 1 presents comprehensive performance metrics for all four models using the 80-20 train-test split.

AlexNet not only achieved the highest accuracy but also demonstrated the best precision (0.9953), recall (0.9971), and F1 score (0.9961). The precision value of 0.9953 indicates that 99.53% of AlexNet's disease predictions were correct, minimising false-positive diagnoses. The recall value of 0.9971 demonstrates that the model successfully identified 99.71% of all actual disease cases, minimising false negatives. The F1 score of 0.9961 represents an excellent balance between precision and recall, indicating robust overall performance. Regarding computational efficiency, AlexNet required an average of 545 seconds per epoch during training, which was the second-fastest among all models tested. MobileNet was slightly faster at 565 seconds per epoch, which may seem counterintuitive but can be explained by the specific implementation details and the overhead associated with depthwise separable convolutions on certain hardware configurations. VGG16 required the longest training time at 892 seconds per epoch due to its deep architecture and large number of parameters. Inception V3 required 782 seconds per epoch, reflecting its complex multi-branch architecture.

5.4. Performance Across Different Train-Test Splits

To assess the robustness of our models and their susceptibility to overfitting, researchers conducted experiments with train-test split ratios of 70-30, 80-20, and 90-10. Table II presents the results, which demonstrate remarkable consistency across different splits. The minimal variation in accuracy across different train-test splits indicates that our models generalise well and do not suffer from significant overfitting. The slight improvement in accuracy with larger training sets (90-10 split) is expected, as more training data generally leads to better model performance. However, the differences are marginal, suggesting that even with 70% of the data for training, the models achieve near-optimal performance.

5.5. Confusion Matrix Analysis

To gain deeper insights into model performance and identify specific classes where misclassifications occur, researchers analysed the confusion matrix for the best-performing model (AlexNet). The confusion matrix showed that the model achieved

near-perfect classification across most disease classes, with accuracy exceeding 99% in 35 of 38 classes. The few misclassifications that occurred were primarily between visually similar disease classes. For example, some confusion was observed between “Tomato Early Blight” and “Tomato Late Blight,” which share similar visual characteristics in their early stages. Similarly, minor confusion occurred between “Apple Black Rot” and “Apple Scab,” both of which manifest as dark lesions on leaves. These misclassifications are understandable given the visual similarity of symptoms and highlight areas where additional feature engineering or more diverse training data could further improve performance. Interestingly, healthy leaf images were classified with 100% accuracy across all crop species, indicating that the model has learned robust features to distinguish between healthy and diseased states. This is particularly important for practical applications, as false positives (incorrectly diagnosing healthy plants as diseased) can lead to unnecessary treatments and economic losses for farmers.

5.6. Training and Validation Curves

Analysis of training and validation curves provides insights into the learning dynamics and potential overfitting issues. For AlexNet, the training accuracy steadily increased from approximately 85% in the first epoch to over 99% by epoch 25, eventually stabilising at 99.7% by epoch 40. The validation accuracy followed a similar trajectory, reaching 99.56% and remaining stable, indicating excellent generalisation (Table 3).

Table 3: Performance comparison across different image representations

Model	Color	Grayscale	Segmented
AlexNet	99.56%	96.23%	97.84%
Inception V3	98.42%	95.67%	96.91%
VGG16	97.89%	94.82%	96.15%
MobileNet	97.02%	93.45%	95.28%

The training loss decreased from approximately 0.45 in the first epoch to below 0.02 by epoch 40, while the validation loss decreased from 0.38 to approximately 0.03. The close alignment between the training and validation curves throughout training indicates that the model did not overfit the training data. The slight gap between training and validation metrics is normal in deep learning models. Similar patterns were observed across the other architectures, though VGG16 showed a slightly larger gap between training and validation metrics, suggesting mild overfitting. This observation aligns with the known tendency of very deep networks with many parameters to overfit when trained on moderately sized datasets.

5.7. Impact of Image Representation

The researchers experimented with three image representations: colour, grayscale, and segmented (where the leaf region was isolated from the background). Table III presents the results, which clearly demonstrate the superiority of colour images. Colour images consistently outperformed both grayscale and segmented images across all architectures. The superior performance of colour images can be attributed to the fact that many plant diseases manifest as specific colour changes—yellowing, browning, or the appearance of colored spots and lesions. Grayscale images lose this critical colour information, resulting in a significant drop in accuracy (approximately 3–4% across all models). Segmented images, where the leaf region was isolated from the background, performed better than grayscale images but still fell short of colour images. While segmentation removes background noise and focuses the model’s attention on the leaf region, the additional preprocessing step introduces potential errors. It may remove contextual information that could be useful for classification.

5.8. Comparative Analysis with Existing Literature

Comparing our results with existing literature on plant disease detection reveals that our approach achieves state-of-the-art performance. Previous studies using the PlantVillage dataset reported accuracies ranging from 88% to 96% using various machine learning and deep learning approaches. Our best model (AlexNet) achieved 99.56% accuracy, representing a significant improvement over most existing methods. The superior performance can be attributed to several factors: (1) effective use of transfer learning with pre-trained ImageNet weights, (2) comprehensive data augmentation strategies. That improved model generalisation, (3) careful hyperparameter tuning and optimisation, (4) use of appropriate training strategies including learning rate scheduling and early stopping, and (5) selection of architectures well-suited to the characteristics of the PlantVillage dataset.

5.9. Practical Implications and Deployment Considerations

The high accuracy achieved by our models has significant practical implications for real-world deployment. With an accuracy of 99.56% and an F1 score of 0.9961, the AlexNet model is highly reliable for automated plant disease diagnosis. The low

false-positive rate (0.47%) minimises unnecessary treatments, while the low false-negative rate (0.29%) ensures that most disease cases are detected early, enabling timely intervention. The computational efficiency of AlexNet, with training times of 545 seconds per epoch and fast inference times (typically less than 0.1 seconds per image on GPU), makes it suitable for deployment in web applications and potentially on edge devices with appropriate optimisation. MobileNet, while slightly less accurate, offers an excellent alternative for mobile and embedded applications with limited computational resources. The web application developed as part of this research provides an accessible interface for farmers and agricultural professionals to upload leaf images and receive instant disease diagnoses. The deployment on Heroku cloud services ensures scalability and availability, allowing multiple users to access the system simultaneously.

5.10. Limitations and Areas for Improvement

Despite the excellent performance achieved, several limitations should be acknowledged. First, the PlantVillage dataset consists of images captured under controlled laboratory conditions with plain backgrounds. Real-world field conditions involve variable lighting, complex backgrounds, occlusions, and other factors that may affect model performance. Future work should include validation on field-captured images to assess real-world performance (Figure 2).



Figure 2: Prediction of the trained model

Second, the dataset is limited to 14 crop species and 38 disease classes. Expanding the system to cover additional crop species and diseases would increase its practical utility. Third, the current system requires manual image capture and upload. Integration with automated image capture systems, such as drones or field-mounted fixed cameras, could enable continuous monitoring and early warning systems (Figure 3).



Figure 3: Prediction of the trained model

Finally, while the model achieves high accuracy in classifying diseases, it does not provide information about disease severity or progression stage (Figure 4).



Figure 4: Training and validation loss for AlexNet

Incorporating severity assessment and temporal analysis could provide more actionable information for farmers and enable more precise treatment recommendations (Figure 5).




Figure 5: Training and validation accuracy for AlexNet

6. Conclusion

This research presents a comprehensive deep learning-based approach for automated plant disease detection using the PlantVillage dataset. Researchers successfully developed and evaluated four state-of-the-art convolutional neural network architectures—AlexNet, Inception V3, VGG16, and MobileNet—for classifying 38 plant disease classes across 14 crop species. The primary contribution of this work lies in achieving exceptional classification accuracy while maintaining computational efficiency suitable for practical deployment. Our experimental results demonstrate that AlexNet achieved the highest performance, with an accuracy of 99.56%, precision of 0.9953, recall of 0.9971, and F1 score of 0.9961 on the 80-20 train-test split. This represents a significant advancement over existing approaches reported in the literature, which typically achieve accuracies of 88–96%. The superior performance was achieved through effective transfer learning, comprehensive data augmentation, and careful optimisation of training hyperparameters. The consistency of results across different train-test split ratios (70- 30, 80-20, and 90-10) confirms the robustness of our models and their resistance to overfitting. The developed web application, deployed on the Heroku cloud platform, provides an accessible, user-friendly interface for farmers and agricultural professionals to obtain instant disease diagnoses by uploading leaf images. The system’s fast inference time (less than 0.1 seconds per image) and high accuracy make it suitable for real-time applications in agricultural settings. However, several limitations must be acknowledged.

The PlantVillage dataset consists of images captured under controlled laboratory conditions, which may not fully represent the variability encountered in real-world field conditions. Additionally, the current system is limited to 14 crop species and does not provide information about disease severity or treatment recommendations. Future work will focus on several key areas to enhance the system's capabilities and practical utility. First, researchers plan to integrate a recommendation engine that suggests appropriate fertilisers and treatments based on the diagnosed disease, complete with product information and purchasing options. Second, researchers aim to expand the system to cover additional crop species and diseases by collecting and annotating field-captured images under diverse environmental conditions. Third, researchers are developing a mobile application to enable farmers to use the system anytime, anywhere, facilitating immediate field diagnosis. Fourth, recognising India's linguistic diversity, researchers plan to implement multilingual support, making the application accessible to farmers from different regions who may not be proficient in English. Finally, researchers are exploring integrating IoT sensors and drone-based imaging systems to enable continuous automated monitoring of large agricultural areas and provide early warning systems for disease outbreaks. In conclusion, this research demonstrates the significant potential of deep learning technologies in addressing critical challenges in agriculture, particularly in early disease detection and crop health monitoring. The high accuracy and computational efficiency achieved by our models, combined with the accessible web-based interface, represent a meaningful step toward practical AI-powered solutions that can help farmers improve crop yields, reduce losses, and enhance food security.

Acknowledgement: The authors extend their sincere appreciation to the team from SRM Institute of Science and Technology for their collective expertise and dedicated contributions. Their coordinated efforts and academic commitment greatly strengthened the quality of this research.

Data Availability Statement: The data underlying this research can be obtained from the corresponding author upon reasonable request. Access to certain datasets may be restricted due to ethical considerations and privacy regulations, and any requests will be evaluated in accordance with institutional data-sharing policies.

Funding Statement: The authors affirm that this research received no external financial support, grants, or funding from public, commercial, or non-profit organisations.

Conflicts of Interest Statement: The authors collectively declare that there are no known financial or personal conflicts of interest that could have influenced the outcomes of this study.

Ethics and Consent Statement: This study adhered to institutional ethical guidelines and received approval from the appropriate review committee. All participants provided informed consent before they participated in the research.

References

1. V. K. Shrivastava, M. K. Pradhan, S. Minz, and M. P. Thakur, "Rice plant disease classification using transfer learning of deep convolution neural network," *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, no. 7, pp. 631–635, 2019.
2. H. Sabrol and K. Satish, "Tomato plant disease classification in digital images using classification tree," *International Conference on Communication and Signal Processing (ICCSPP)*, Melmaruvathur, India, 2016.
3. K. Elangovan and S. Nalini, "Plant disease classification using image segmentation and SVM techniques," *International Journal of Computational Intelligence Research*, vol. 13, no. 7, pp. 1821–1828, 2017.
4. S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks based recognition of plant diseases by leaf image classification," *Computational Intelligence and Neuroscience*, vol. 2016, no. 1, p. 3289801, 2016.
5. S. Aziz, M. Bashir, O. Mughal, M. U. Khan, and A. Khan, "Image Pattern Classification for Plant Disease Identification using Local Tri-directional Features," in *IEEE 10th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, Vancouver, British Columbia, Canada, 2019.
6. F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, Hostelling International, United States of America, 2017.
7. P. B. Padol and A. A. Yadav, "SVM classifier based grape leaf disease detection," in *Proc. Conference on Advances in Signal Processing (CASP)*, Pune, India, 2016.
8. S. Chakraborty, S. Paul, and M. Rahat-uz-Zaman, "Prediction of Apple Leaf Diseases Using Multiclass Support Vector Machine," in *2nd International Conference on Robotics, Electrical and Signal Processing Techniques (ICREST)*, Dhaka, Bangladesh, 2021.
9. M. Bhagat, D. Kumar, R. Mahmood, B. Pati, and M. Kumar, "Bell Pepper Leaf Disease Classification Using CNN," in *2nd International Conference on Data, Engineering and Applications (IDEA)*, Bhopal, India, 2020.

10. K. Kirti and N. Rajpal, "Black Rot Disease Detection in Grape Plant (*Vitis vinifera*) Using Colour Based Segmentation and Machine Learning," in *2nd International Conference on Advances in Computing, Communication Control and Networking (ICACCCN)*, Greater Noida, India, 2020.
11. M. I. Tarik, S. Akter, A. A. Mamun, and A. Sattar, "Potato Disease Detection Using Machine Learning," in *3rd International Conference on Intelligent Communication Technologies and Virtual Mobile Networks (ICICV)*, Tirunelveli, India, 2021.
12. M. Dhakate and A. B. Ingole, "Diagnosis of pomegranate plant diseases using neural network," in *5th National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG)*, Patna, India, 2015.
13. A. H. B. A. Wahab, R. Zahari, and T. H. Lim, "Detecting diseases in Chilli Plants Using K-Means Segmented Support Vector Machine," in *3rd International Conference on Imaging, Signal Processing and Communication (ICISPC)*, Singapore, 2019.
14. C. Rajasekaran, S. Arul, S. Devi, G. Gowtham, and S. Jeyaram, "Turmeric Plant Diseases Detection and Classification using Artificial Intelligence," in *International Conference on Communication and Signal Processing (ICCSP)*, Chennai, India, 2020.
15. N. Saranya, L. Pavithra, N. Kanthimathi, B. Ragavi, and P. Sandhiyadevi, "Detection of Banana Leaf and Fruit Diseases Using Neural Networks," in *2nd International Conference on Inventive Research in Computing Applications (ICIRCA)*, Coimbatore, India, 2020.
16. S. Ramesh and D. Vydeki, "Rice Blast Disease Detection and Classification Using Machine Learning Algorithm," in *proc. 2nd International Conference on Micro-Electronics and Telecommunication Engineering (ICMETE)*, Ghaziabad, India, 2018.
17. K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, no. 2, pp. 311–318, 2018.
18. P. Malathi, T. Pranavadit, R. V. Vishal, M. A. B. Kumar, V. Sanjai, and B. Singh, "Deep learning-driven real-time recognition of plant diseases via CNN and Flask integration," *AVE Trends in Intelligent Health Letters*, vol. 2, no. 1, pp. 52–61, 2025.
19. J. G. A. Barbedo, "Plant disease identification from individual lesions and spots using deep learning," *Biosystems Engineering*, vol. 180, no. 4, pp. 96–107, 2019.
20. A. Fuentes, S. Yoon, S. C. Kim, and D. S. Park, "A robust deep-learning- based detector for real-time tomato plant diseases and pests recognition," *Sensors*, vol. 17, no. 9, p. 2022, 2017.
21. K. Kaur and K. A. Sergeevna, "Deep learning-based detection of diseases in strawberry plants," *AVE Trends in Intelligent Health Letters*, vol. 1, no. 4, pp. 217–227, 2024.
22. S. P. Mohanty, D. P. Hughes, and M. Salathe, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, no. 9, p. 215232, 2016.