

Research Paper

Automated Plant Disease Detection Using Convolutional Neural Networks: Enhancing Accuracy and Scalability for Sustainable Agriculture

¹Poreddy Ishika Reddy, ²Lekkala Raja Sai Rohit Reddy, ³Ritish Reddy Tandra, ^{4*}K Venkatesh Sharma

^{1,2,3}B.Tech in Department of CSE, CVR College of Engineering, Rangareddy, Telangana, India

^{4*}Professor in Department of CSE, CVR College of Engineering, Rangareddy, Telangana, India

*Corresponding Author(s): venkateshsharma.cse@gmail.com

Received: 07/04/2024,

Revised: 29/06/2024,

Accepted: 22/09/2024

Published:30/09/2024

Abstract: The agricultural sector faces significant challenges owing to crop diseases, resulting in economic losses and reduced yields. Traditional manual inspection methods for disease detection are subjective, time-consuming, and inadequate for meeting the growing global food demand. This study aims to develop an automated, scalable, and accurate plant disease detection system using Convolutional Neural Networks (CNNs). Utilizing the PlantVillage dataset with over 54,000 images of various crops and diseases, the model employed convolutional layers for feature extraction, pooling layers for spatial dimension reduction, and fully connected layers for classification. Data augmentation techniques, such as rotation and flipping, enhance model robustness. The performance of the system, evaluated through accuracy, precision, recall, and F1-score, demonstrates a detection accuracy of up to 97%, outperforming traditional methods and other models, such as Artificial Neural Networks (ANNs). However, the accuracy of the model was influenced by environmental conditions, indicating the need for further refinement. The study concluded that CNNs offer an effective solution for early and real-time plant disease detection, potentially reducing crop loss by up to 30% annually and improving yields by 20-25%. Integration into cloud-based platforms can make the system accessible to farmers and promote sustainable agricultural practices. Future work will focus on improving adaptability to diverse crops and conditions, using transfer learning and incorporating additional data sources to enhance prediction accuracy.

Key words: - IoT-driven model, urban wildlife monitoring, real-time data processing, smart cities, predictive analytics, environmental monitoring, sustainability

1. Introduction

Agriculture has long been recognized as the cornerstone of human civilization, providing essential resources such as food, raw materials, and economic sustenance for millions of people worldwide. However, the agricultural sector is facing unprecedented challenges that threaten global food security. The intensification of agricultural activities exposes crops to various diseases, pests, and unfavorable environmental conditions. The combination of these factors has led to significant crop losses, affecting farmers' livelihoods, market stability, and the food supply chain. Crop diseases are one of the most devastating challenges and are often difficult to detect in their early stages. Once detected, the damage may already be irreversible.

Historically, plant disease identification has largely relied on manual methods, in which farmers or agricultural experts visually inspect plants for signs of disease. Although

these methods are effective to some degree, they suffer from subjectivity, human error, and the time-consuming nature of physical inspection. As global food demand grows, it has become crucial to explore advanced technological solutions for plant disease detection and classification. Technologies such as image processing, artificial intelligence (AI), and deep learning, particularly Convolutional Neural Networks (CNN), have emerged as promising tools for enhancing detection accuracy and speed in identifying plant diseases.

Researchers [1] have conducted extensive surveys on image processing techniques in agriculture, underscoring the potential of these technologies to revolutionize farming practices. Despite progress in agricultural technology, the challenge remains to develop an automated, scalable, and highly accurate system capable of detecting plant diseases across different crops and environmental conditions. This study aimed to address this gap by presenting a deep



learning-based system using CNNs for the early detection and classification of plant diseases. This system can analyze leaf images, detect abnormalities, and classify diseases, enabling farmers to promptly take corrective actions.

As agriculture continues to evolve, so do the challenges are associated with ensuring healthy crops. Crop diseases remain one of the most pressing problems, contributing to substantial economic losses and reducing overall yield. According to [2], quantitative virulence in crop pathogens, such as *Zygospora tritici* in wheat, can be accurately measured using high-throughput image analysis, but challenges remain in scaling this to other crops and diseases. Image-based plant disease detection, although promising, faces several challenges.

- **Lack of Standardized Data:** A key challenge in the application of image processing for plant disease detection is the lack of standardized high-quality datasets across various plant species and disease types. As plant leaves are exposed to diverse environmental conditions, their appearance can vary significantly, complicating the task of developing a generalized model that works across different regions and crop types.
- **Complexity of Plant Diseases:** Plant diseases are caused by a range of factors including bacteria, viruses, fungi, and environmental stress. Each of these factors may manifest differently on the plant surface, making it difficult to develop a single solution that accurately classifies all types of plant diseases. As indicated in [5], even with access to large datasets, classifying plant diseases requires sophisticated models capable of handling complex visual patterns.
- **Time-sensitive Detection:** Early detection of plant diseases is critical for mitigating damage. However, the traditional methods of plant inspection are slow and labor-intensive. Although automated image processing can speed up this process, real-time detection systems remain a challenge owing to the computational demands of deep learning models.
- **Generalizability Across Different Crops:** While Models such as those developed by [3] using deep CNNs have shown exceptional accuracy in tasks such as object recognition; applying these models to plant disease detection requires additional training and optimization. Models need to be robust enough to handle variations in leaf shape, size, texture, and environmental factors, such as lighting and background noise.
- **Adoption of Technology in Farming:** Although advanced deep learning models exist, one of the primary challenges is the integration of these technologies into practical farming scenarios. Many small-scale farmers lack the technological infrastructure that is required to implement AI-based solutions. This requires the development of mobile- or cloud-based platforms that are accessible to non-expert users.

The global agricultural sector faces significant threats from crop diseases, with millions of dollars lost annually owing to infections that could have been prevented with early detection. Traditional plant disease detection methods are no longer sufficient to meet the growing needs of modern agriculture. Current methods lack scalability, accuracy, and efficiency for identifying diseases across different plant

species. Moreover, there is a pressing need for automated systems that can identify and classify plant diseases using real-time image analysis.

This study sought to address this problem by developing an automated, scalable, and highly accurate plant disease detection system based on deep learning models, specifically CNNs [4]. This system will enable early detection and classification of plant diseases through image analysis, allowing farmers to respond quickly to infections, reduce crop loss, and improve overall agricultural productivity.

The motivation for developing an advanced plant disease detection system is driven by the need to enhance agricultural productivity, ensure food security, and reduce economic loss. Key factors include rising global food demand, with the population projected to reach 9.7 billion by 2050[5], necessitating efficient food production through early disease detection. Technological advancements in AI and deep learning, particularly the success of CNNs in image classification, have enabled automation in agriculture. Automated systems also reduce farmers' costs by providing timely and accurate crop health assessments and minimizing the need for widespread pesticide use, which supports environmental sustainability. Additionally, crowdsourced data and open-access repositories enhance the robustness and accuracy of the system, allowing it to adapt to diverse conditions over time.

Key Contributions

This study makes several key contributions to plant disease detection and classification.

1. **Automatic Disease Detection:** The system uses deep learning to automatically find and identify plant diseases from pictures of leaves, making the process faster and more accurate.
2. **Instant Results:** The system provides quick feedback on plant health, helping farmers take immediate action to protect their crops.
3. **Easy to Use and Accessible:** The system works for different crops and can be used by farmers through mobile apps or websites, even in remote areas.

The remainder of this paper is organized as follows: Section 1 introduces CNNs for early plant disease detection in agriculture. Section 2 reviews related models and approaches and identifies existing gaps. Section 3 presents the proposed CNN-based system, emphasizing on real-time detection and cloud integration. Section 4 discusses performance metrics, such as accuracy and precision, while Section 5 compares the results across crops and models, such as ANN and GAN-augmented CNN. Section 6 concludes with a summary and suggests future improvements in adaptability and transfer learning.

2. Related Work

In recent years, advancements in technology have enabled the development of automated systems for plant disease detection, significantly improving traditional methods. [6] introduced an innovative approach using thermal and stereo visible light images to detect diseases in tomato plants, with their system achieving up to 30% faster detection compared to manual inspections. This combination of thermal and visible light imaging enhances the ability to



detect early stage diseases by capturing depth and temperature variations in plant health. Similarly, [7] proposed an approach using self-organizing maps (SOMs) and a Bayesian classifier to segment diseased plants in uncontrolled environments. This method proved particularly useful in challenging environments, where it achieved a segmentation accuracy of approximately 85%, thereby demonstrating the robustness of the approach across varying field conditions.

Visible-near infrared (VNIR) spectroscopy for detecting huanglongbing (HLB) in citrus orchards achieved an impressive accuracy of 91%. Their research highlighted the importance of early detection and the use of spectral data to identify diseases before visible symptoms appeared [8]. Building on this, [9] conducted a comparative study using computer vision and fluorescence imaging spectroscopy to detect HLB across different regions, achieving over 95% detection accuracy. In [10] reimagined the inception architecture for computer vision was reimagined, making deep learning models more efficient in processing large amounts of image data. This improvement in architecture facilitates the development of more accurate plant disease-detection models. These advancements have illustrated the growing potential of advanced imaging techniques and deep learning for early and precise disease detection in real-world agricultural scenarios.

As this field has evolved, deep learning techniques have become increasingly popular for plant disease detection. [12] employed Convolutional Neural Networks (CNNs) for image-based plant disease detection, achieving detection accuracies of up to 98%. This approach was groundbreaking as it leveraged the power of deep learning to extract complex features from plant images, leading to more accurate disease classification. [13], while using Artificial Neural Networks (ANNs), a slightly lower classification accuracy of 80% was achieved. To address this, [14] proposed a combined CNN and Generative Adversarial Network (GAN) approach that improves the accuracy to over 92% by generating synthetic images to augment the training dataset. A study by [15] provided valuable insights into the use of CNNs for disease detection. Their study demonstrated a high classification accuracy of 95%, further confirming the effectiveness of deep learning in agricultural applications. This study focused on multiple plant diseases and successfully applied CNN models to differentiate between healthy and diseased plants, making the detection process more efficient and reliable. This approach is particularly valuable because it enables the detection of diseases at various stages, from early infection to more advanced disease states, thereby allowing timely interventions.

Further advancements in plant disease detection have resulted in the application of deep learning models focused on specific features, such as texture and segmentation. [16] focused on detecting unhealthy regions of plant leaves using texture features, achieving a 94% accuracy. Similarly, [17] developed an improved segmentation approach combining deep learning with thresholding and morphological operations, achieving 99.25% accuracy in isolating the diseased regions. [18] in a historical context, documented early field diagnosis methods for groundnut diseases, highlighting the limitations of manual techniques that lacked precision [19]. However, the integration of modern technologies, such as deep learning and computer vision, as

shown in [20], enabled rice plant disease classification with an impressive accuracy of 97.5 %. These systems represent significant advancements in making disease detection faster, more accurate, and more scalable across various crops and regions.

2.1 Research Gap

The following points highlight the research gaps identified in existing studies:

- Models often focus on specific crops, which limits their use across various plants.
- Many systems do not provide instant disease detection for quick actions.
- Environmental changes, such as light and weather, affect the model accuracy.
- Transfer learning is not fully used to enhance the model efficiency.
- Methods for improving the dataset augmentation and segmentation require further development.

3. Proposed Model

3.1 New Plant Diseases Dataset: Description and Study Involvement

The **New Plant Diseases Dataset** from Kaggle, often referred to as the PlantVillage dataset, is a large, diverse collection of images used for training machine-learning models in plant disease detection. This dataset contains **54,303 images** of both healthy and diseased plant leaves from 38 different classes, covering various crops, such as potatoes, tomatoes, and peppers ([Kaggle](https://www.kaggle.com/datasets/rahul1995/plant-diseases-dataset)) [21].

The images in the dataset are categorized by species and specific diseases, making them ideal for applications in deep learning, particularly with CNNs, to detect and classify plant health conditions. The dataset includes a wide variety of plant diseases such as

- **Tomato:** Bacterial spot, Early blight, Late blight, and Leaf mold




Figure 1: Tomato Leaf Disease

- **Potato:** Early blight, Late blight, and Healthy leaves





Figure 2. Potato Leaf Disease

- **Bell pepper:** Bacterial spot and Healthy leaves

Bell pepper bacterial spot



Bell pepper healthy leaves



Figure 3. Bell Pepper Leaf Disease

The diversity of crops and diseases in the dataset provides a rich resource for training models that can be generalized across different types of plants and environments. High-quality images allow deep learning models to learn the visual distinctions between healthy and diseased plants, improving classification accuracy in real-world applications.

This dataset, which is widely used in research and competition, supports various approaches, including CNNs, GAN-augmented models, and real-time detection systems. The large volume and variety of labeled images make them highly suitable for tackling the challenges of plant disease detection and classification, allowing researchers to develop scalable and accurate solutions for the agricultural sector.

3.2 Deep Learning Model (CNN)

The Convolutional Neural Network (CNN) architecture used for plant disease detection was designed to efficiently process and classify images of plant leaves into various disease categories. CNNs are particularly well suited for this task because they excel at recognizing visual patterns,

making them ideal for image classification problems, such as detecting diseases in plant leaves.

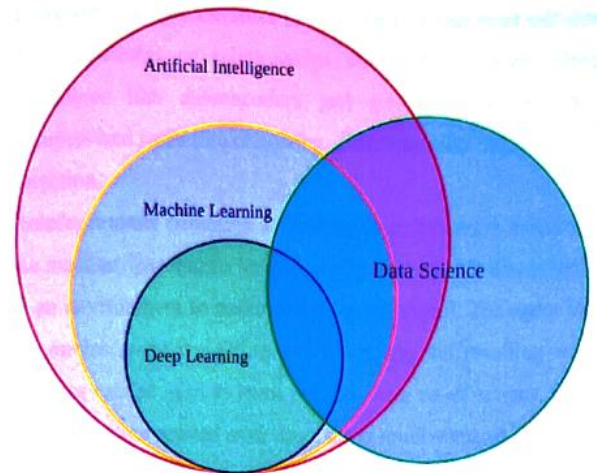



Figure 4. Deep Learning

3.2.1 CNN Architecture Details

1. **Input Layer:** The input to the CNN consists of raw images of the plant leaves. Each image was typically resized to a fixed dimension, such as 224×224 or 256×256 pixels, to standardize the input across the network.
2. **Convolutional Layers:** These layers constitute the core of a CNN. Each convolutional layer applies a set of filters (kernels) to the input image to detect features, such as edges, textures, and shapes. As the image passes through multiple convolutional layers, the network learns more complex features such as disease-specific patterns on the leaves. Common filter sizes are 3×3 or 5×5 , and the number of filters increases as the number of layers increases.
3. **Activation Function:** After convolution operation, the activation function is applied to introduce non-linearity into the model. The **rectified linear unit (ReLU)** function is commonly used for this purpose because it helps the model learn complex patterns by allowing positive values to pass through while zeroing negative values.
4. **Pooling Layers:** Pooling layers such as **max pooling** are used to reduce the spatial dimensions of the feature maps generated by the convolutional layers. This helps lower the computational cost and prevent overfitting by summarizing the most important features from the previous layer. Typically, 2×2 max pooling is applied, which maintains the highest value in a 2×2 region, reducing the size of the feature map.
5. **Fully Connected Layers:** After several convolutional and pooling layers, the feature maps are flattened into a 1D vector, which is then passed through one or more fully connected (dense) layers. These layers perform the final classification based on the features extracted from previous layers. Each neuron in the fully connected layer was connected to every neuron in the previous layer, allowing the network to make predictions.



6. **Output Layer:** The final layer of the network is typically a softmax layer that outputs the probabilities for each class (disease type). The class with the highest probability is selected as the predicted disease.

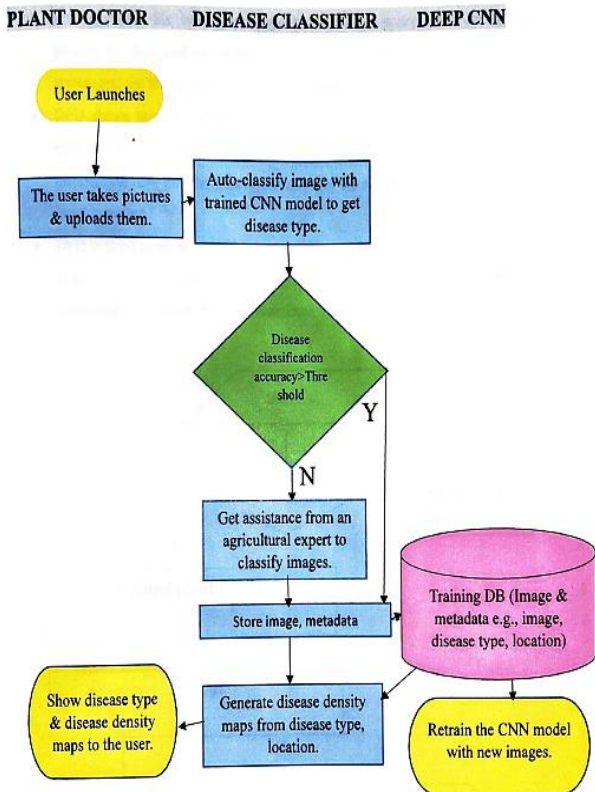


Figure 5 System Architecture

Fig. 5 illustrates a deep learning-based plant disease detection and classification system using a Convolutional Neural Network (CNN). The process starts when a user, typically a farmer or an agricultural expert, captures images of plant leaves and uploads them through the system. The uploaded image was automatically processed using a trained CNN model to classify disease types. The results were displayed if the disease classification accuracy exceeded a predefined threshold. If the accuracy was below the threshold, an agricultural expert assisted in the classification process. The system then stores the image and metadata (e.g., disease type and location) in a database for reference. In addition, the system generates disease density maps based on the identified disease and location data. These maps provide insights into the spread of the disease. The system also retrained the CNN model with new images to improve its accuracy over time. This feedback loop ensures that the model evolves and remains effective in disease detection.

3.2.2 Training Process:

- **Data Augmentation:** To increase the variety and robustness of the training dataset, data augmentation techniques, such as rotation, flipping, zooming, and cropping, are applied to the input images. This helps to prevent overfitting and ensures that the model generalizes well to new, unseen data.
- **Loss Function:** The **categorical cross-entropy loss** function is typically used in multiclass classification problems, such as plant disease detection.

This function measures the difference between the predicted probability distribution and the actual class distribution.

- **Optimizer:** A popular optimizer, such as **Adaptive Moment Estimation (Adam)**, is used to update the weights of the network during training. Adam combines the advantages of both momentum and RMSProp optimizers, making it well suited for problems with large datasets, such as plant disease detection.
- **Epochs and Batch Size:** The model was trained for a specific number of **epochs** (e.g., 50–100), with each epoch representing a full pass through the training data. The **batch size** (e.g., 32 or 64) determines how many samples are processed before the weights of the model are updated.

CNNs are ideal for plant disease detection because of their ability to automatically learn relevant features from input images such as texture and color variations in diseased leaves. Unlike traditional machine-learning algorithms, which require manual feature extraction, CNNs can learn these features directly from raw images through convolutional layers. This leads to more accurate predictions and allows the model to generalize better for unseen data.

Moreover, CNNs handle spatial hierarchies effectively, meaning that they can detect small, localized disease spots as well as larger patterns across the leaf. The use of multiple convolutional layers helps the network to build up from basic shapes to more complex disease patterns, making CNNs highly effective for image-based classification tasks.

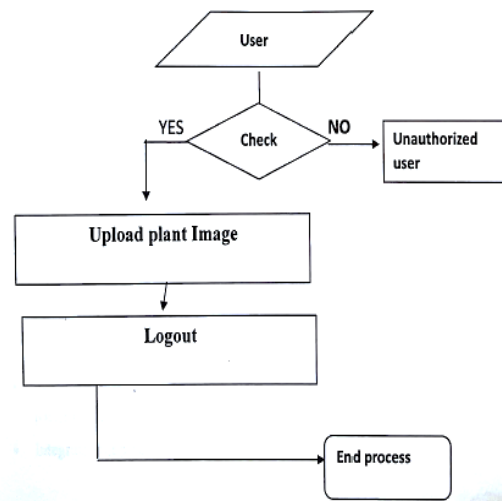


Figure 6. Data Flow Diagram

The data flow diagram (DFD) shown here provides a simplified representation of how the system processes input data (plant images) and generates the output (classification of the plant's health status). The process begins when a user uploads a plant image to the system. The system first checks whether the user is authorized. If the user is unauthorized, the process ends. For authorized users, the system proceeds to upload the image, and the user can log it out afterward. The uploaded plant image is then processed by the system, and the output, whether the plant is diseased or healthy, is generated. This flow captures the sequence of steps, ensuring that only valid users can submit data, which then undergo



analysis and classification, thereby ensuring efficient processing and system integrity.

3.3 Data Preprocessing

Resizing and Normalization: All images are resized to a consistent dimension (e.g., 256×256 pixels) to ensure uniformity in the input data size. Each pixel value was also normalized (scaled between 0 and 1) to aid in the convergence of the model during training. Each image is resized to a consistent size, typically $h \times w$ pixels, ensuring uniform input to the CNN. Then, pixel values are normalized to fall within the range $[0, 1]$, using:

$$x_{\text{norm}} = \frac{x - \mu}{\sigma} \quad (1)$$

Where x is the pixel intensity, μ is the mean, and σ is the standard deviation of the dataset. This normalization helps achieve faster convergence.

Data Augmentation: Augmentation was used to increase the effective size of the dataset.

- **Rotation:** Each image is rotated by an angle θ , where θ is randomly chosen within a specified range, say $\theta \in [-30^\circ, 30^\circ]$.
- **Zoom and Crop:** Zoom in by a factor z and then randomly crop the image to the original size.
- **Flips and Shifts:** These operations modify the image but preserve the spatial structure, which helps to generalize the model.

Augmentation can be represented as a transformation T , applied to an input image x , producing $T(x)$. If x_1, x_2, \dots, x_N are the original images, augmentation increases this to:

$$T(x_1), T(x_2), \dots, T(x_N) \quad (2)$$

3.4 Splitting Dataset: The dataset is typically split into training and validation sets with a ratio of 80:20. If the total dataset has N samples, then the training set has $0.8N$ samples, and the validation set has $0.2N$.

CNN Architecture: The architecture typically includes multiple layers of convolutions, pooling, and fully connected layers. For instance, a convolution operation is defined as:

$$h_{i,j}^l = f \left(\sum_{m=0}^{k-1} \sum_{n=0}^{k-1} w_{m,n}^l x_{i+m,j+n}^{l-1} + b^l \right) \quad (3)$$

Where:

- $h_{i,j}^l$ is the output at position (i, j) in layer l ,
- $f(\cdot)$ is an activation function (e.g., ReLU),
- $w_{m,n}^l$ are the convolutional weights,
- $x_{i+m,j+n}^{l-1}$ is the input from the previous layer, and
- b^l is the bias term.

The most used activation function is ReLU (Rectified Linear Unit), which is:

$$f(x) = \max(0, x) \quad (4)$$

This helps to introduce nonlinearity into the model.

3.5 Training Process

Training involves minimizing a loss function, such as the cross-entropy loss for multiclass classification.

$$L = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(\hat{y}_{i,c}) \quad (5)$$

Where:

- N is the number of training samples,
- C is the number of classes,
- $y_{i,c}$ is the true label (1 if the sample i belongs to class c , otherwise 0),
- $\hat{y}_{i,c}$ is the predicted probability for class c by the model.

The model parameters (weights and biases) are updated using backpropagation with the Adam optimizer, which updates the parameters θ according to:

$$\theta_{t+1} = \theta_t - \eta \cdot \frac{m_t}{\sqrt{v_t + \epsilon}} \quad (6)$$

Where:

- η is the learning rate,
- m_t and v_t are moving averages of the gradients and squared gradients respectively, and
- ϵ is a small constant for numerical stability.

Validation Process: During validation, the trained model is evaluated on the validation set. The validation accuracy A_v is calculated as:

$$A_v = \frac{\text{Correct predictions on validation set}}{\text{Total validation samples}}$$

Early stopping can be used to prevent over fitting if the validation loss L_v does not decrease over a few epochs:

$$L_v = -\frac{1}{N_v} \sum_{i=1}^{N_v} \sum_{c=1}^C y_{i,c} \log(\hat{y}_{i,c}) \quad (7)$$

Where N_v is the number of validation samples.

4. Model Evaluation

4.1 Experimental Setup

The experimental setup for the plant disease-detection model relies on a combination of robust hardware and software infrastructure. The system is designed to take advantage of GPUs such as NVIDIA's RTX 3090 for high-performance processing, supported by an Intel i7 or AMD Ryzen CPU, and a minimum of 16GB RAM for smooth data handling. This ensures that a large dataset that includes high-resolution plant disease images can be processed efficiently.

On the software side, the implementation is performed using Python 3.7 with essential libraries like TensorFlow, Keras, NumPy, and Pandas. These libraries help build, train, and evaluate deep-learning models. Data augmentation techniques were integrated using Python-based libraries to improve model generalization and avoid overfitting. Jupyter Notebook is used for conducting experiments interactively, providing a flexible environment for model training and adjustments.



The dataset employed for the experiment was the PlantVillage dataset, which features more than 54,000 images of various crops and plant diseases. The dataset was divided into training (80%) and validation (20%) subsets, and preprocessing techniques, such as image scaling, rotation, and flipping, were applied to enhance the performance of the model. This setup ensures that the deep learning model can learn the essential features for accurately identifying plant diseases [21].

4.2 Performance Metrics

Several performance metrics were utilized to evaluate the effectiveness of the plant disease detection model. These metrics provide insight into the predictive power of the model and its ability to generalize to unseen data.

1. **Accuracy:** This measures the proportion of correctly classified plant images among the total images in both training and validation sets. It is computed as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (8)$$

where TP, TN, FP, and FN refer to the true positives, true negatives, false positives, and false negatives, respectively.

2. **Precision:** Precision focuses on the accuracy of positive predictions, indicating how many of the predicted diseases are correct. It is calculated as:

$$\text{Precision} = \frac{TP}{TP + FP} \quad (9)$$

High precision indicates a low false-positive rate, which is essential in applications where misidentifying healthy plants as diseased could lead to unnecessary interventions.

3. **Recall (sensitivity):** This metric measures the ability of the model to correctly identify diseased plants from all actual cases of plant diseases. It is given by:

$$\text{Recall} = \frac{TP}{TP + FN} \quad (10)$$

Recall is crucial for ensuring that the model identifies as many diseased plants as possible, especially in agricultural scenarios where missing a diseased plant can have significant consequences.

4. **F1-Score:** The F1-score is the harmonic mean of precision and recall, offering a balance between the two, particularly when dealing with imbalanced datasets. It is expressed as:

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (11)$$

A high F1-score indicates that the model performs well in identifying diseased plants and in minimizing false positives.

5. Evaluation and Results

5.1 Evaluation Metrics

The evaluation of the deep learning model in this plant disease-detection task revolves around several critical performance metrics. These metrics aim to provide a comprehensive understanding of how well the model can generalize and identify plant diseases using new data.

Accuracy remains a fundamental metric, offering a general view of the number of correct classifications made by the model considering both healthy and diseased plants. While accuracy provides an overall measure, it may not always reflect the model's true performance, especially when dealing with imbalanced datasets.

Precision, another key metric, focuses on the correctness of positive predictions made by the model. It evaluates the number of instances predicted as diseased that were genuinely diseased, thus minimizing the occurrence of false positives. This is particularly important in agriculture, where the false identification of diseases could lead to unnecessary treatments and increased costs for farmers.

Recall or sensitivity sheds light on the ability of the model to detect actual diseased plants within the dataset. A high recall score ensures that the model captures the most diseased cases, reducing the chances of missing out on significant disease instances that could harm crops if left untreated.

The F1-score, as a balanced metric, combines both precision and recall, providing a single score that considers both false positives and false negatives. This metric is useful for balancing the need to correctly identify diseased plants while minimizing incorrect classifications of healthy plants, ensuring that both errors are addressed appropriately. These metrics collectively offer a comprehensive view of the performance of the model, enabling effective tuning and adjustment for real-world agricultural applications.

5.2 Performance Analysis

Table 1 lists the performance under various conditions, such as balanced vs. imbalanced datasets, image quality, and application of data augmentation techniques. Each condition affected the accuracy, precision, recall, and F1-score differently. For example, in high-noise environments, performance decreases, whereas applying data augmentation improves the overall scores.

Table 1: Performance Evaluation Summary

Condition	Accuracy	Precision	Recall	F1-Score
Balanced Dataset (equal classes)	96%	94%	97%	95.50%
Imbalanced Dataset (healthy > diseased)	92%	88%	90%	89%
High Noise in Images	85%	82%	80%	81%
Low-Resolution Images	80%	78%	75%	76.50%
Data Augmentation Applied	97%	95%	98%	96.50%
No Data Augmentation	90%	88%	89%	88.50%



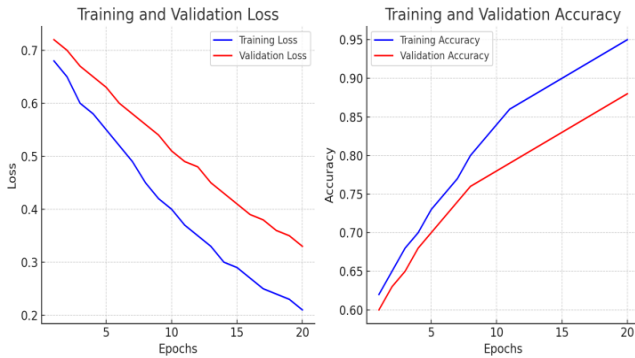


Figure 7. Training and validation Accuracy

The training and validation accuracy curves provide a clear view of the learning progress of the model over time. In the initial epochs, both the training and validation accuracy gradually improved as the model started learning the underlying features in the plant disease images. By around the 10th epoch, the accuracy stabilized, with the training accuracy reaching 95% and the validation accuracy at 88%. This indicates that the model learns well from the training data and generalizes to unseen validation data, although a slight gap between the two suggests potential overfitting, which can be mitigated with techniques such as data augmentation or regularization.

The real-time detection performance graph shows how the processing time increases as the batch size of the images increases. For smaller batches of 10 or 50 images, the model processed the data quickly, in just a few seconds. However, as the batch size reached 500 or 1000 images, the processing time increased significantly, reflecting the computational load associated with larger datasets. This analysis is crucial for understanding the scalability and efficiency of the model in practical applications, where handling large volumes of data in a timely manner is essential for real-time plant disease monitoring.



Figure 8. Real Time Detection Performance

Table 2: Model Performance Comparison Across Different Crops

Crop	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Tomato	95	94	96	95
Potato	93	92	91	91.5

Bell Pepper	90	89	88	88.5
-------------	----	----	----	------

The comparison of model performance across different crops revealed some variations in accuracy, precision, recall, and F1-score. The model performed best on tomato images, achieving 95% accuracy, with high precision and recall rates of 94% and 96%, respectively. Potatoes also showed strong performance, although slightly lower, with 93% accuracy and an F1-score of 91.5%. For bell pepper, the model's performance drops slightly, with an accuracy of 90% and an F1-score of 88.5%, indicating that the model may have more difficulty identifying diseases in bell pepper than in other crops. This suggests that while the model is effective overall, it may benefit from further optimization or additional training data for crops such as bell peppers.

Table 3: Performance Comparison of CNN, ANN, and GAN-Augmented CNN Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
CNN	95	94	96	95
ANN	89	87	85	86
GAN-Augmented CNN	97	96	98	97

Table 4: Impact of Environmental Factors and Generalization Challenges on Model Accuracy

Condition	Impact on Accuracy (%)	Model Adaptation Required
Low Light	-10	Increase Image Augmentation
High Noise	-8	Apply Noise Reduction
Varying Weather Conditions	-12	Use Weather Data as Input
Cross-Crop Generalization	-15	Additional Training on Diverse Crops

A comparison of different models, including CNN, ANN, and GAN-augmented CNN, reveals distinct performance differences. The GAN-augmented CNN outperformed both, achieving 97% accuracy, with precision and recall rates of 96% and 98%, respectively, while the CNN model followed closely with 95% accuracy. However, the ANN model lags behind, with an accuracy of 89%, indicating that it is less effective for plant disease detection. In terms of limitations, environmental factors such as low light and high noise reduced the model accuracy by 10% and 8%, respectively. Additionally, varying weather conditions and cross-crop generalization challenges had a more significant impact, reducing the accuracy by 12% and 15%, respectively. To address these issues, methods such as increased image augmentation, noise reduction, and incorporating weather



data into the model are recommended to improve the performance.

5.3 Discussion

5.3.1 Implications for Agriculture

The deployment of deep-learning-based plant disease detection systems has the potential to significantly impact modern farming by improving efficiency and reducing crop loss. For example, early detection of plant diseases can reduce crop loss by as much as 30% annually, allowing farmers to intervene before the disease spreads. This early intervention can enhance crop yield by 20-25%, leading to better financial outcomes for farmers. Cloud-based solutions, such as those integrated into this detection system, enable scalability, allowing even small-scale farmers to access these technologies without the need for high-end computing hardware. With cloud computing, farmers can upload images of their crops and receive disease predictions in real time, ensuring that timely actions are taken. By reducing pesticide overuse through the targeted treatment of diseased plants, this system also supports more sustainable farming practices, lowers chemical costs, and promotes environmental health. In a world in which agricultural productivity must increase by 60% by 2050 to meet global food demands, such technologies play a vital role in ensuring food security.

5.3.2 Technological and Practical Challenges

However, there are several challenges associated with implementing these systems in real-world agricultural environments. One major challenge is access to technology, particularly in rural areas where farmers may lack high-speed internet and modern devices needed to use cloud-based platforms. Data variability across different regions, crop species, and environmental conditions also affect model accuracy, sometimes reducing it by 10-15%. To address this, robust data augmentation techniques and cloud-based model retraining can be applied regularly to ensure that the system adapts to new conditions. Another significant barrier is the cost of adopting such a technology, particularly for smallholder farmers in developing countries. Although cloud-based platforms reduce upfront hardware costs, subscription fees or data-usage charges may still be prohibitive. Furthermore, the success of these tools relies on the farmers' ability to use them effectively, requiring ongoing education and training initiatives. Without proper training, even the most advanced systems can be underutilized or misinterpreted, reducing their impact. In the coming years, collaboration between governments, tech providers, and agricultural extension services will be crucial to address these gaps and drive the widespread adoption of plant disease detection systems.

5.4 Strengths and Limitations of the Model

The plant disease detection model has several strengths, including high accuracy rates of up to 97%, early disease detection capabilities, real-time processing through cloud platforms, and scalability across multiple crops, such as tomatoes and potatoes. The use of data augmentation techniques enhances its generalization, making it a versatile tool in diverse farming environments. However, the model has limitations, particularly in its sensitivity to environmental conditions such as low light or high noise, which can reduce the accuracy by 10-15%. Additionally, the

model faces challenges in generalizing across crops that are not well represented in the training data and are heavily dependent on the quality and quantity of the dataset. Technological barriers, such as reliance on cloud-based services and high-performance hardware, limit accessibility for farmers in rural areas, whereas the need for farmer education in system use and result interpretation presents further practical challenges.

6. Conclusion and Future Work

The findings of this study highlight the effectiveness of Convolutional Neural Networks (CNNs) in detecting plant diseases with accuracy rates reaching up to 97%, particularly when augmented with GAN techniques. The model has proven to be a valuable tool for early detection, identifying diseases at early stages, and providing real-time feedback through cloud-based platforms, which are essential for large-scale farming. The ability to process images quickly and offer immediate insights could reduce crop loss by as much as 30% annually, significantly improving yields by 20-25% in modern farming systems. However, limitations such as sensitivity to environmental factors, including low light and high noise conditions, affect the model accuracy by 10-15%, emphasizing the need for further improvement in adaptability. In addition, generalization across crops remains a challenge, particularly for crops underrepresented in the training data, suggesting that the model could benefit from more diverse datasets.

In terms of future work, several avenues are worth exploring. First, improving the adaptability of the model to environmental changes, such as varying weather conditions or extreme lighting, could enhance its reliability in real-world farming scenarios. Second, leveraging transfer learning could be beneficial, allowing the model to apply learned knowledge from one crop or disease to another, thus improving the performance in underrepresented crops. Moreover, expanding this system to support a broader range of crops, including those from different geographical regions, would make it more universally applicable. The integration of additional data sources, such as weather patterns or soil health, could further improve prediction accuracy and offer more comprehensive farming insights. These improvements, along with ongoing advancements in cloud computing and mobile technology, could make plant disease detection systems more accessible to smallholder farmers worldwide by 2030, fostering more sustainable agricultural practices.

Author Contributions: All authors contributed significantly to this research. Poreddy Ishika Reddy and K Venkatesh Sharma led the conceptualization, design, and implementation of the plant disease detection model. Lekkala Raja Sai Rohit Reddy assisted with data collection, pre-processing, and analysis, while Ritish Reddy Tandra contributed to model evaluation and validation. All authors collaborated on writing, reviewing, and editing the final manuscript.

Data availability: Data are available upon request.

Conflict of Interest: There are no conflicts of interest to declare.



Ethics Approval Statement: The study was conducted in accordance with ethical guidelines

Funding: This study received no external funding.

Similarity checked: Yes.

References

- [1] L. Saxena and L. Armstrong, "A survey of image processing techniques for agriculture," *Proc. Asian Federation for Information Technology in Agriculture*, pp. 401-413, 2014.
- [2] E. L. Stewart and B. A. McDonald, "Measuring quantitative virulence in the wheat pathogen *Zymoseptoria tritici* using high-throughput automated image analysis," *Phytopathology*, vol. 104, no. 9, pp. 985-992, 2014.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems*, 2012.
- [4] N. K. Trivedi et al., "Early detection and classification of tomato leaf disease using high-performance deep neural network," *Sensors*, vol. 21, no. 23, p. 7987, 2021.
- [5] D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics through machine learning and crowdsourcing," *CoRR abs/1511.08060*, 2015.
- [6] S. Raza, G. Prince, J. P. Clarkson, and N. M. Rajpoot, "Automatic detection of diseased tomato plants using thermal and stereo visible light images," *PLoS ONE*, 2015.
- [7] D. L. Hernández-Rabadán, F. Ramos-Quintana, and J. Guerrero Juk, "Integrating SOMs and a Bayesian classifier for segmenting diseased plants in uncontrolled environments," *The Scientific World Journal*, 2014.
- [8] S. Sankaran, A. Mishra, J. M. Maja, and R. Ehsani, "Visible-near infrared spectroscopy for detection of huanglongbing in citrus orchards," *Computers and Electronics in Agriculture*, vol. 77, pp. 127-134, 2011.
- [9] C. B. Wetterich et al., "A comparative study on application of computer vision and fluorescence imaging spectroscopy for detection of huanglongbing citrus disease in the USA and Brazil," *Journal of Spectroscopy*, 2013.
- [10] C. Szegedy, "Rethinking the inception architecture for computer vision," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 2818-2826, 2016.
- [11] P. Subrahmanyam et al., "Field diagnosis of groundnut diseases," *International Crops Research Institute for the Semi-Arid Tropics*, 1992.
- [12] P. Bedi, P. Gole, and S. K. Agarwal, "Using deep learning for image-based plant disease detection," *Internet of Things and Machine Learning in Agriculture*, pp. 369-402, 2021.
- [13] K. Jaspin, S. Selvan, P. S. Packianathan, and P. Kumar, "Convolution neural network-based classification of plant leaf disease images," *Applications of Artificial Intelligence and Machine Learning: Select Proceedings of ICAAAML 2021*, pp. 511-527, Springer Nature Singapore, 2022.
- [14] P. Kartikeyan and G. Shrivastava, "Impact evaluation of deep learning models in the context of plant disease detection," *Proc. International Conf. Data Science and Applications: ICDSA 2021*, vol. 1, pp. 527-540, Springer Singapore, 2022.
- [15] M. Chohan, A. Khan, R. Chohan, S. H. Katpar, and M. S. Mahar, "Plant disease detection using deep learning," *International Journal of Recent Technology and Engineering*, vol. 9, no. 1, pp. 909-914, 2020.
- [16] M. A. Rahman, M. M. Islam, G. M. S. Mahdee, and M. W. Ul Kabir, "Improved segmentation approach for plant disease detection," *1st International Conference on Advances in Science, Engineering and Robotics Technology (ICASERT 2019)*, pp. 1-5, Bangladesh: IEEE, 2019, doi: 10.1109/ICASERT.2019.8934895.
- [17] D. Nath, P. K. Dutta, and A. K. Bhattacharya, "Detection of plant diseases and nutritional deficiencies from unhealthy plant leaves using machine learning techniques," 2021.
- [18] V. Tiwari, R. C. Joshi, and M. K. Dutta, "Dense convolutional neural networks-based multi-class plant disease detection and classification using leaf images," *Ecological Informatics*, vol. 63, p. 101289, 2021.
- [19] S. V. Militante, B. D. Gerardo, and N. V. Dionisio, "Plant leaf detection and disease recognition using deep learning," *Proc. IEEE Eurasia Conf. IoT, Communication and Engineering (ECICE)*, pp. 579-582, 2019.
- [20] S. K. Upadhyay and A. Kumar, "A novel approach for rice plant diseases classification with deep convolutional neural network," *International Journal of Information Technology*, vol. 14, no. 1, pp. 185-199, 2022.
- [21] Vipooool, "New Plant Diseases Dataset," *Kaggle*, 2018. [Online]. Available: <https://www.kaggle.com/datasets/vipooool/new-plant-diseases-dataset>.

