International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Hyperspectral Imaging for Early Detection of Tomato Bacterial Wilt

Rohit Krishnan Kandathil[1], Anuush M B[2], Anshu Prasad Sah[3], Ayush Binil Nair[4], Shalu Murali[5]

[1,2,3,4]Student, [5]Assistant Professor

[1,2,3,4,5]Department of Computer Science and Engineering, Amrita Vishwa Vidyapeetham, Amritapuri

Abstract

Early and accurate detection of plant diseases is essential for improving crop health and minimizing yield losses. This study explores the application of hyperspectral imaging (HSI) combined with deep learning techniques for the early-stage detection of tomato bacterial wilt, a severe and fast-spreading crop disease. Spectral reflectance data collected from hyperspectral scans of tomato leaves were preprocessed using Savitzky-Golay smoothing, dimensionality reduction via Principal Component Analysis (PCA), and normalized before classification using a lightweight one-dimensional Artificial Neural Network (1D-ANN). The proposed model achieved a validation accuracy of 96% and a test accuracy of 89.77%, demonstrating high precision and robustness in distinguishing healthy and infected plants, even at asymptomatic stages. These findings highlight the potential of integrating HSI and deep learning for non-invasive, real-time plant disease diagnosis, contributing to the advancement of precision agriculture systems.

Keywords: Hyperspectral Imaging(HSI), Artificial Neural Networks (ANNs), UAV-based Disease Surveillance, Spectral Normalization

1.Introduction

Plant diseases plaguing the agricultural sector continue to undermine food security, cut crop yields, and bring tremendous economic losses to the world. Staple crops such as rice and tomato are especially susceptible to pathogens, with timely diagnosis being central to controlling infections. Conventional methods of detection based on manual examination and biochemical assays are laborious, retrospective, and reliant upon highly trained analysts. Such weaknesses underscore the importance of developing automated, non-destructive, and scalable methods for timely detection of disease to act upon.

New technologies in machine learning and computer vision have changed plant pathology in that RGB imaging and CNNs have provided cost-effective means to classify diseases [4], [5]. The models, though, identify the diseases only after symptoms are apparent, which hinders their applications to early intervention. In addition, use of RGB channels hinders detection of internal physiological alterations that are antecedent to detectable symptoms [4].

To overcome these shortcomings, hyperspectral imaging (HSI) is an incredibly useful alternative, which is able to capture over hundreds of spectral bands in the visible and near infrared range. The spectral resolution makes identification of biochemical markers including chlorophyll breakdown, water imbalance, and breakdown in cellular structure possible even before symptoms are visibly apparent [2], [3], [9]. Combined with deep learning frameworks such as Artificial Neural Networks (ANNs), HSI systems hold great promise to achieve accurate, early, non-destructive plant disease diagnosis [1], [2].

This research outlines a two-phase study in plant disease identification through deep learning, as well as image analysis methods: Phase 1 is intended for rice plant disease classification from RGB images, utilizing a multi-output CNN with an EfficientNetB0 architecture. The network distinguishes plant type from disease class in parallel with high accuracy, making it feasible to use in applications with mobile or UAV platforms [5], [12].

Phase 2 involves the application of hyperspectral imaging to identify tomato bacterial wilt, caused by Ralstonia solanacearum, in early, symptomless stages. Based on preprocessed spectra, utilizing a trained 1D ANN optimised with Principal Component Analysis (PCA), phase 2 illustrates how HSI can be included in precision agriculture pipelines [1], [2], [4].

By contrasting hyperspectral with traditional RGB-based methods, this research emphasizes the technological evolution from reactive to proactive plant disease surveillance. The decision to use RGB imaging for rice and hyperspectral imaging for tomato is based on practical and biological considerations. Rice leaf diseases typically manifest with visible symptoms—such as lesions, spots, and discoloration—that can be effectively detected through RGB images, making it suitable for low-cost, scalable deployment via mobile or UAV platforms. In contrast, tomato bacterial wilt, caused by Ralstonia solanacearum, often progresses internally before visual symptoms emerge. Therefore, hyperspectral imaging, which can detect subtle biochemical changes at the early stages of infection, is more appropriate for timely detection and intervention in tomato crops. The work highlights how deep learning and spectral imaging increasingly contribute to developing more sustainable agriculture with far-reaching consequences in real-time in-the-field deployment, food security, and environmental responsibility [3], [10].

2.Related Works

Progress over the last ten years in automating plant disease detection using machine learning (ML) and deep learning (DL) techniques has been tremendous. Researchers have experimented with various imaging modalities such as RGB, thermal, multispectral, and hyperspectral to come up with reliable and scalable diagnostic frameworks for precision agriculture.

2.1. RGB-Based Plant Disease Detection

One of the most widely used and accessible modalities is RGB imaging as a consequence of the widespread availability of standard digital camera hardware and smartphone devices. Convolutional Neural Networks (CNN) have demonstrated very good performance in classifying plant disease symptoms from visible plant images.

Li et al. [6] (2021) performed an extensive review of deep learning for plant disease diagnosis, with special focus placed upon how handcrafted feature-based approaches have changed to end-to-end CNNs. It draws attention to the role of transferring knowledge along with visualization techniques in increasing classification performance, although recognizing the deficiency of early detection as well as deployment limitations in fields.

Islam et al. [7] in their work (2023) presented "DeepCrop," which is based on CNN, with ResNet50 used for classifying diseases coupled with web-based diagnosis for farmers. In addition to receiving 98.98%

accuracy, use of curated datasets such as PlantVillage restricts generalization under variable field scenarios.
Jung et al. [8] (2023) also developed a multi-stage classification pipeline with an "unknown" class to learn to generalize to unseen disease types. Even though this makes classifications more robust, scalability to crop types as well as to environmental variation is still limited.
These studies make CNN-based RGB-based systems reliable devices for detection of apparent diseases, but reactive in nature, sensing only after damage to physiology.

2.2. **Hyperspectral Imaging and Spectral Analysis**

Hyperspectral imaging delivers revolutionary insight by allowing detection of early plant stress through reflectance pattern analysis across many narrow bands. The technique detects biochemical and structural plant tissue changes, such as chlorophyll breakdown and water loss, that are in many cases undetectable by RGB-based sensors.
Wasswa et al. [3] in their work in 2023 performed a systematic review of 176 studies conducted in plant disease detection based on image-based AI techniques. The increasing role of lightweight DL models, and standardized data sets, as well as benchmark sets for hyperspectral analysis, were highlighted by them.
Jiang et al. [4] presented in 2025 that model fusion methods integrating SVM, Random Forest, and XGBoost classifiers over hyperspectral data enhanced black vegetable pesticide residue detection accuracy. The model's intricacy, however, created challenges for actual field implementation.
Zhang et al. [1] in 2025 integrated hyperspectral imaging with 1D-ANNs and GAN-based data augmentation to detect rice bacterial blight symptoms. The model had greater than 97% accuracy in pure samples but underperformed with mixed-pathogen datasets.
In another work, Bhargava et al. [2] (2024) wrote about different spectral preprocessing and feature selection methods, such as Savitzky-Golay smoothing and PCA, that are essential to hyperspectral classification. They highlighted the use of explainable AI (XAI) and data augmentation for generalizability.

2.3. **Integrations and Limitations**

While RGB-based models offer cost-effective and accessible solutions, they fall short in early-stage detection and generalization across environments. Conversely, hyperspectral systems provide detailed spectral insights necessary for proactive diagnosis but are hindered by high computational costs, specialized hardware requirements, and limited scalability.
The reviewed literature underscores the need for integrated, multimodal systems that combine the scalability of RGB-based models with the precision of hyperspectral imaging. Furthermore, future research should focus on real-time deployment through edge computing, UAV-based monitoring, and Internet of Things (IoT) integration for automated agricultural disease surveillance [10].

3.Dataset Selection

3.1. **RGB Dataset – Rice Leaf Disease Detection**

In the first phase of the study, the main focus was on the classification of rice leaf diseases using standard RGB images. For this, a curated dataset consisting of high-resolution images of rice leaves, each carefully labeled with both the plant type and the disease category was utilized. The dataset includes a wide range of common rice leaf diseases such as bacterial leaf blight, brown spot, leaf blast, leaf scald, and narrow brown spot, along with healthy leaves and some samples marked as "other disease." This diversity ensures that the model learns to distinguish between multiple disease conditions, not just binary health vs. disease.

Figure 1: Sample images from the RGB dataset

What makes this dataset especially valuable is its structure-it supports multi-output classification, meaning the model learns not only to recognize the disease but also to identify the plant type (i.e., rice vs. non-rice). This dual classification is particularly useful in mixed farming environments, where distinguishing between crops and their respective ailments can help optimize targeted treatments.

The images were sourced from publicly available, research-oriented datasets widely used in plant pathology and AI communities. A robust evaluation pipeline by splitting the data into three distinct sets was ensured: 80% for training, 20% for validation, and a separate test set held out entirely during training. This split allows the model to learn effectively while also providing a reliable way to assess its generalization to new, unseen samples.

By combining clear labeling, disease variety, and high-resolution visual data, this RGB dataset forms a strong foundation for building and evaluating a deep learning model tailored to rice disease detection.

3.2 Hyperspectral Dataset – Tomato Bacterial Wilt

In the second phase of the research, the focus shifted to a more advanced sensing technique-hyperspectral imaging (HSI)-to detect bacterial wilt in tomato plants. Unlike RGB imaging, which captures just three-color channels, hyperspectral imaging collects data across hundreds of narrow spectral bands, including wavelengths far beyond what the human eye can see. This makes it incredibly powerful for detecting subtle physiological changes in plant tissues that might indicate early signs of disease.

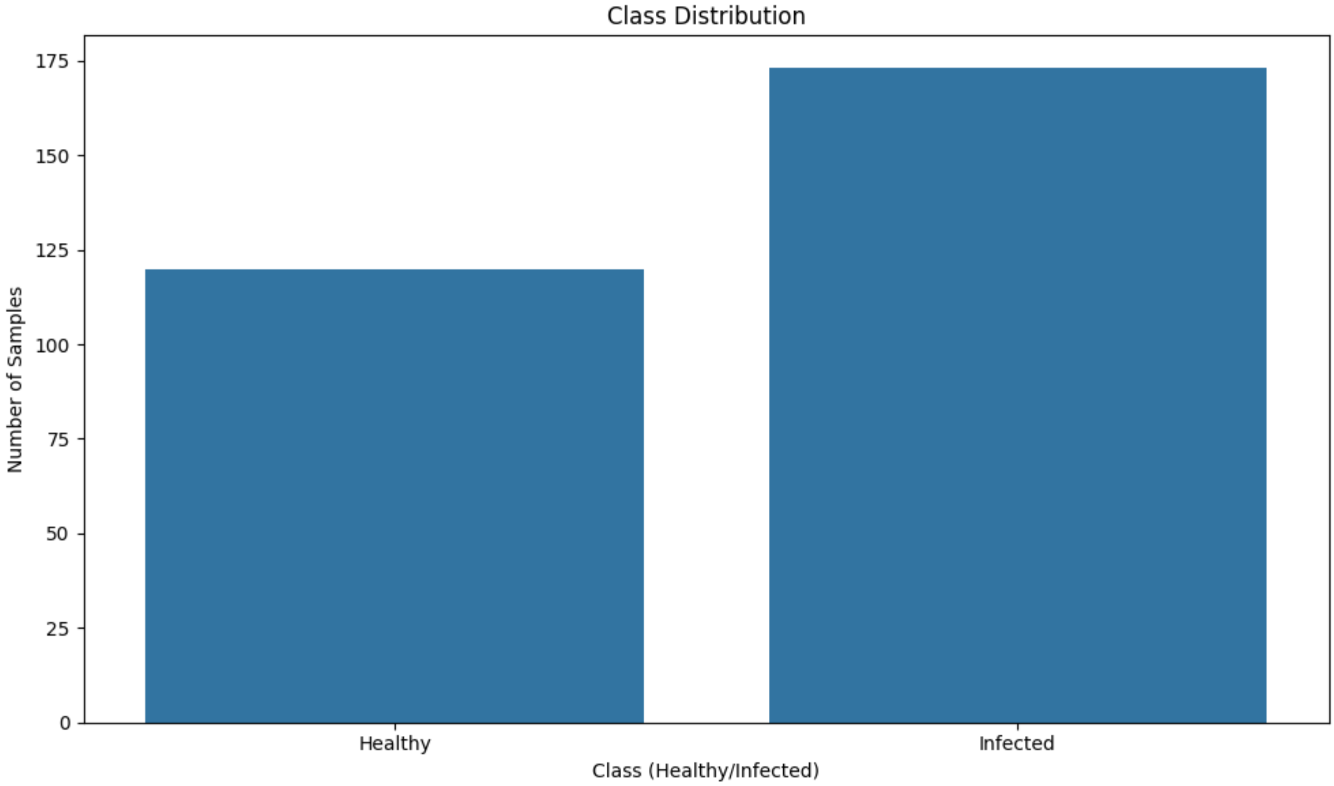

Figure 2: Hyperspectral Dataset Split

The dataset used consists of reflectance spectra extracted from hyperspectral scans of tomato leaves at various stages of infection. Each spectral sample reflects the plant's biochemical makeup-things like chlorophyll content, moisture levels, and cellular structure-which are all affected when the plant is under stress due to bacterial wilt. The disease in question, caused by Ralstonia solanacearum, is particularly aggressive and often progresses before any visible symptoms appear. That’s where HSI can make a real difference.

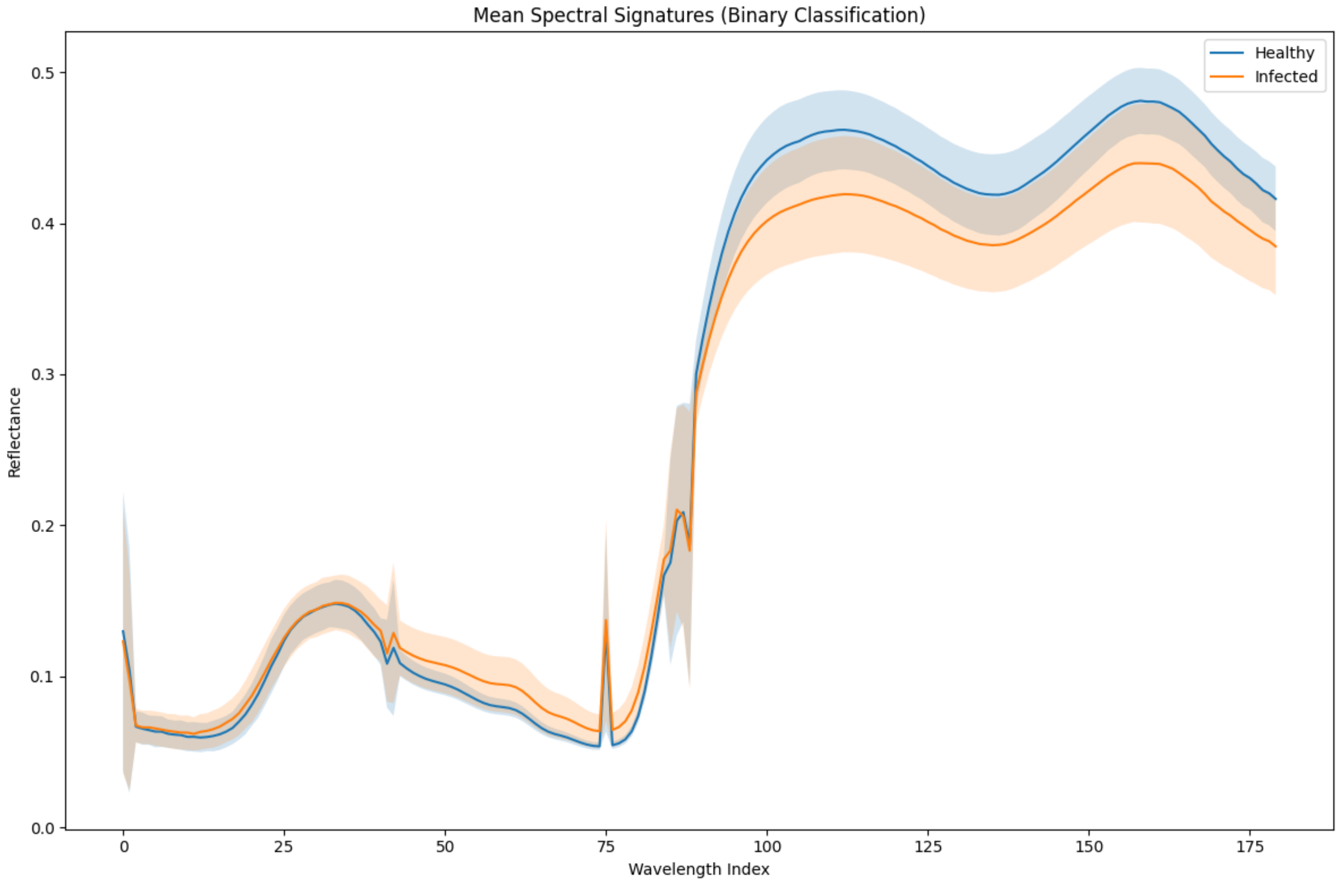

Figure 3: Average Spectral Signatures of healthy vs. infected leaves

While the original dataset labels included multiple severity levels (from healthy to severely infected), the study focused on the practical and impactful task of binary classification-distinguishing between healthy and infected plants. This simplification not only aligns with real-world diagnostic needs (where early detection is key), but also helps streamline the development of a lightweight, deployable model.

Before feeding the data into the neural network, several important preprocessing steps were applied to clean and standardize the spectral inputs. First, the spectral data was resampled to 180 uniformly spaced bands to reduce dimensionality while preserving the most informative regions. Then, Savitzky-Golay smoothing was applied, a well-known noise reduction technique that helps clean up minor fluctuations in the spectral signal caused by environmental noise or sensor artifacts. Finally, the data was normalized using standard scaling, ensuring that all wavelengths contribute equally to the model's learning process.

Altogether, this hyperspectral dataset-rich in detail and carefully preprocessed-served as the backbone of the efforts to build an intelligent, non-invasive system capable of detecting tomato bacterial wilt before it becomes visible, helping farmers act early and minimize crop losses.

Despite substantial progress in using RGB and hyperspectral imaging techniques for plant disease detection, many existing approaches rely on either RGB-based models that are reactive in nature—detecting disease only after visible symptoms appear—or computationally intensive hyperspectral systems that are not optimized for early-stage field deployment. Furthermore, while deep learning has shown promise in classification accuracy, few studies have explored lightweight architectures, such as 1D-ANNs, tailored specifically for spectral data in early, asymptomatic stages of infection. Existing works often overlook the integration of efficient preprocessing pipelines, dimensionality reduction, and real-time feasibility for agricultural environments. This study aims to address these limitations by proposing a computationally efficient, early-stage detection framework using hyperspectral reflectance data and a lightweight ANN, optimized for real-time implementation in precision agriculture.

4.Model Design

4.1. RGB-Based CNN Architecture for Rice Leaf Classification

For the first phase of the project-classifying rice leaf diseases using RGB images- a custom deep learning model built on the EfficientNetB0 architecture was designed . The goal wasn't just to detect whether a leaf was diseased or not, but to simultaneously identify the plant type (e.g., rice vs. non-rice) and pinpoint the specific disease affecting it. This required a model capable of handling multi-output classification, where two distinct but related predictions are made from the same image.

EfficientNetB0 was selected due to its well-established balance between performance and computational efficiency. As a lightweight yet powerful convolutional neural network (CNN) pretrained on the ImageNet dataset, it provides a strong foundation for recognizing general image features such as textures, edges, and patterns. This makes it particularly suitable for learning subtle disease-related characteristics in rice leaves without requiring training from scratch.

Here's how the model is structured:

- EfficientNetB0 was used as the base network, with pretrained ImageNet weights loaded and the final classification layers removed to allow task-specific customization.
- A Global Average Pooling layer was added to condense the rich feature maps from the convolutional layers into a compact representation.
- This was followed by a fully connected dense layer with 512 units and a ReLU activation function to capture complex non-linear feature relationships.
- Two separate output branches were then attached, each terminating in a softmax layer:
  - One predicts the plant type (rice or other)
  - The other classifies the disease type across seven categories, including "healthy."

The study also fine-tuned the training strategy to get the best performance. Since disease detection is more critical for the application than plant identification, it was assigned higher importance (weight = 0.9) to the disease classification task, while giving lesser weight (0.1) to the plant type prediction. This ensures that the model learns to be especially accurate when diagnosing leaf diseases.

For training, the Adam optimizer was used, which adapts the learning rate during training for faster and smoother convergence. The study also implemented early stopping to halt training when the validation loss stopped improving, and model checkpointing to save the best-performing version of the model. This helped us avoid overfitting and ensured the final model could generalize well to new, unseen data.

Overall, this architecture allowed us to build a fast, accurate, and efficient model capable of tackling two agricultural challenges at once-identifying the plant and diagnosing its health-making it a practical candidate for real-world deployment in smart farming systems.

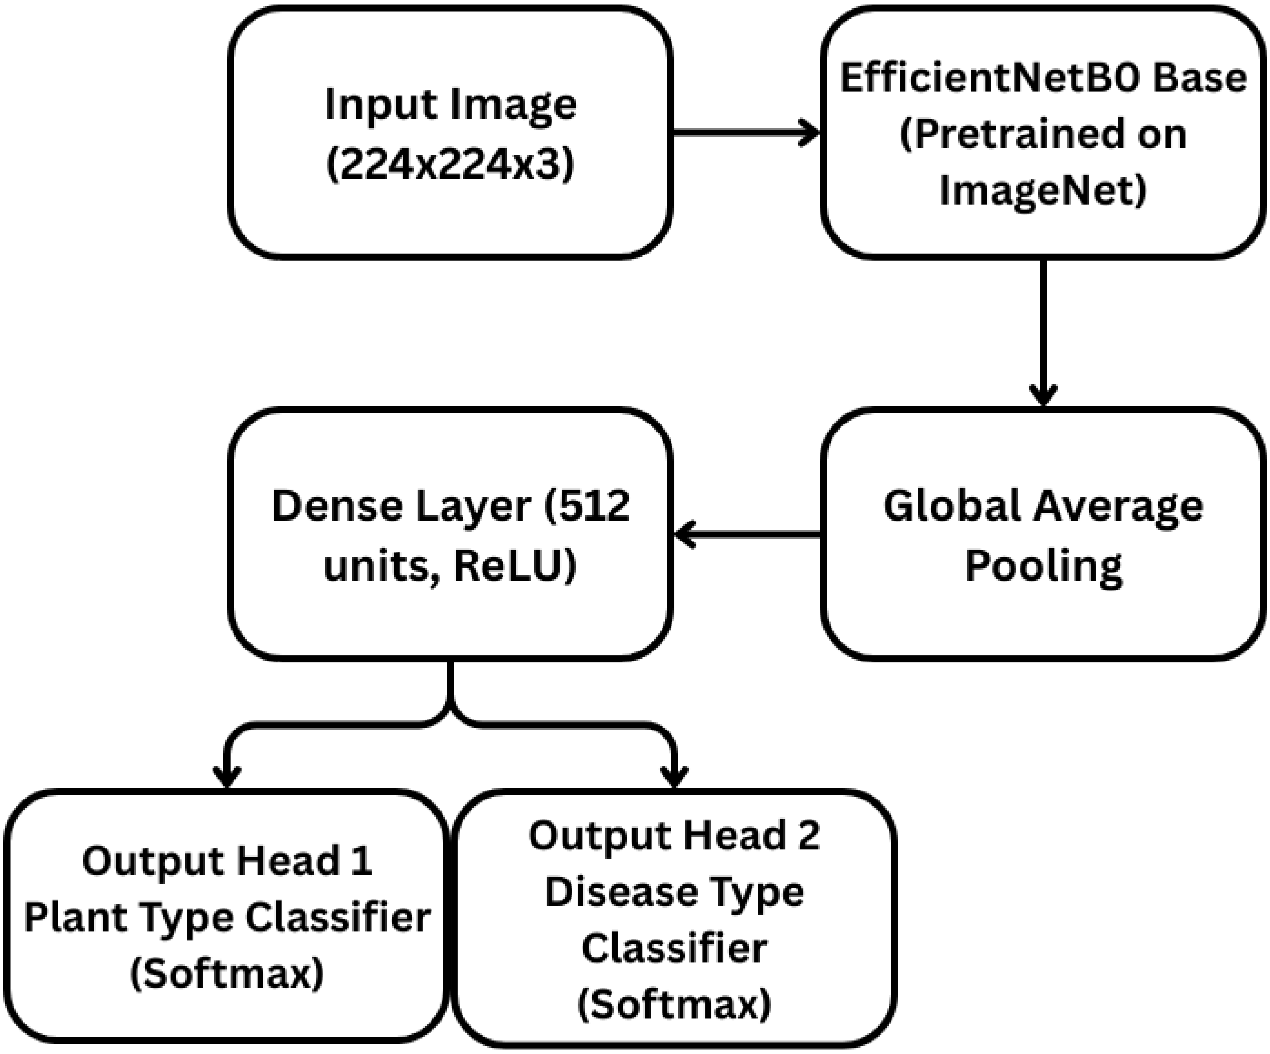

Figure 4: CNN model architecture (EfficientNetB0)

4.2 Hyperspectral Model Design – 1D Artificial Neural Network (ANN)

In the second phase of the study, the focus moved from analyzing images to something a bit more abstract-spectral data. Instead of looking at pictures of leaves, we were now looking at how they reflect light across hundreds of wavelengths, some of which the human eye can't even perceive. The goal was to figure out, based solely on this spectral information, whether a tomato plant is healthy or infected with bacterial wilt-even before the disease shows any visible signs.
To handle this type of data, a 1D Artificial Neural Network (ANN)-a lightweight but powerful model specifically designed to work with numerical sequences like spectral vectors was built. But before feeding the data into the network, it went through a few important preprocessing steps. First, Principal Component Analysis (PCA) was done to reduce the complexity of the spectral data. It compressed the data into fewer dimensions while keeping the most important information. This not only helps the model train faster but also reduces the chances of it getting confused by irrelevant noise.

Once the spectral data was preprocessed and reduced using Principal Component Analysis (PCA), it was input into a custom-designed one-dimensional Artificial Neural Network (1D-ANN) architecture. The model was structured to efficiently process and classify the spectral signatures of tomato leaves as either healthy or infected with bacterial wilt.

The network begins with an input layer, configured to match the dimensionality of the PCA-reduced spectral vectors. This layer serves as the interface between the preprocessed hyperspectral data and the network's internal learning mechanism.

Following the input layer, the network incorporates a series of hidden layers utilizing the Rectified Linear Unit (ReLU) activation function. ReLU was selected due to its computational simplicity and its effectiveness in introducing non-linearity, thereby enabling the network to learn complex patterns within the spectral data that may correspond to physiological changes in infected leaves.

To improve training stability and convergence, batch normalization layers were integrated after selected hidden layers. These layers standardize the input to each activation function, helping to accelerate training and reduce sensitivity to weight initialization.

Additionally, dropout layers were applied to prevent overfitting by randomly deactivating a subset of neurons during each training iteration. This regularization technique encourages the network to learn more robust, generalizable features.

The architecture concludes with a single neuron output layer, equipped with a sigmoid activation function. This configuration outputs a probability value between 0 and 1, indicating the likelihood that a given input spectrum corresponds to an infected leaf. A threshold (typically 0.5) is applied to assign a binary classification label.

The model was trained using the binary cross-entropy loss function, which is well-suited for two-class classification problems. Optimization was performed using the Adam optimizer, chosen for its adaptive learning rate and efficient gradient-based updates, which contribute to faster and more stable convergence.

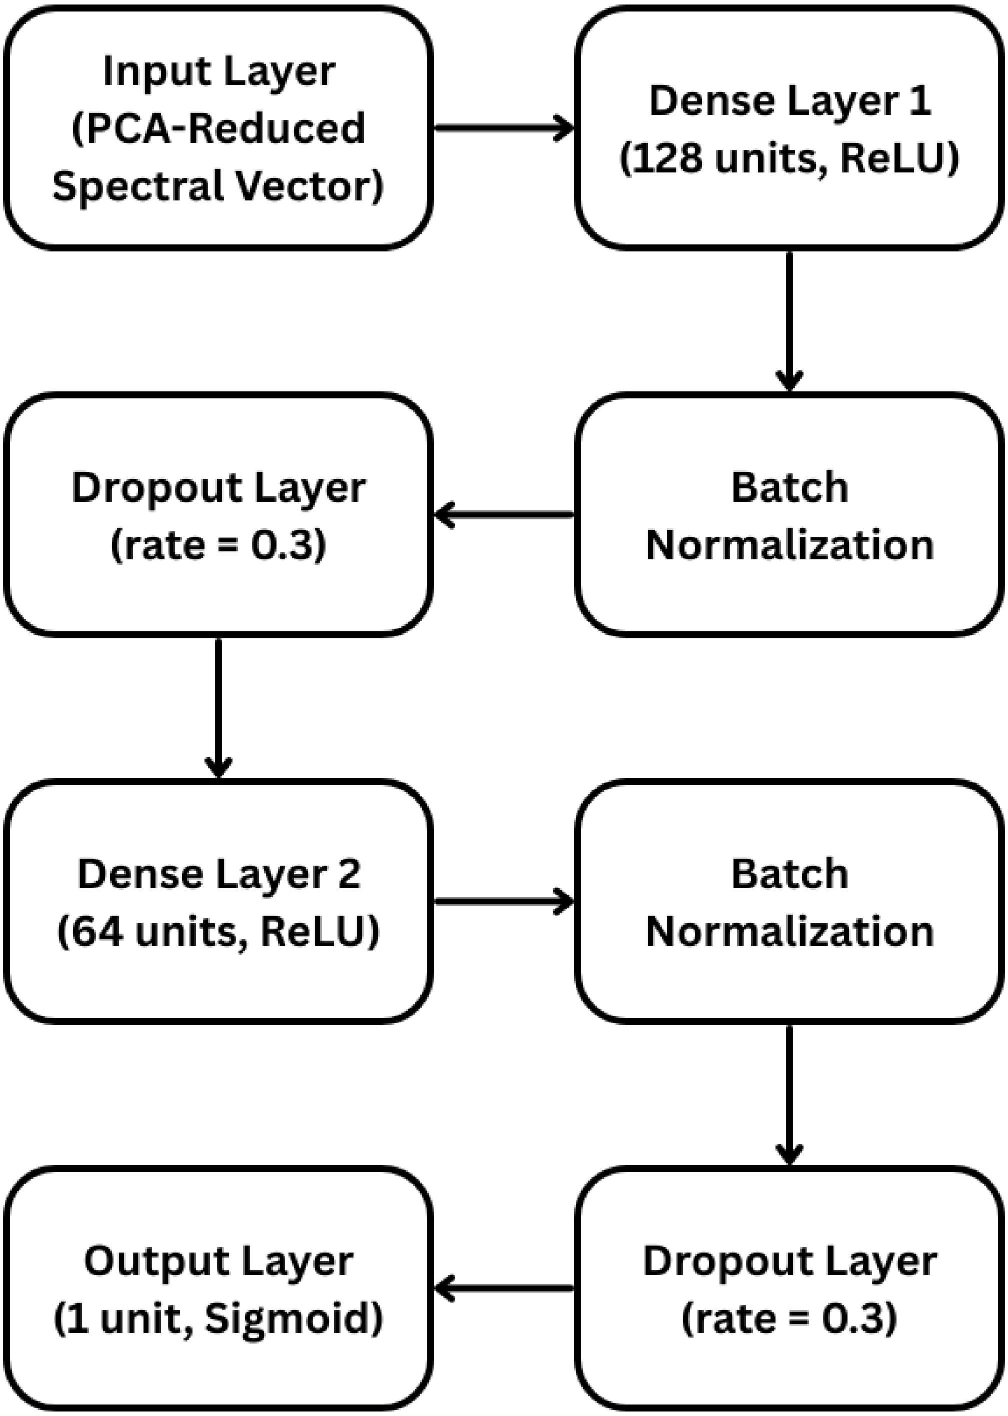

Figure 5: 1D ANN model Architecture

To comprehensively evaluate the model's performance, a suite of standard classification metrics, including Accuracy, Precision, Recall, etc was employed.
This ANN design achieved strong classification results while maintaining a relatively lightweight structure. Its computational efficiency and low input requirements make it a promising candidate for

deployment in real-time, resource-constrained agricultural environments, such as portable field scanners or embedded systems in UAVs.

5.Data Augmentation

To enhance model robustness and generalization, data augmentation techniques were applied to both RGB and hyperspectral data.

5.1. **RGB Image Classification**

To enhance the generalization capability of the convolutional neural network and mitigate overfitting during training, several image augmentation techniques were applied to the RGB rice leaf dataset. These augmentations artificially increased the size and diversity of the training set by introducing controlled variations to the original images, allowing the model to become more robust to environmental noise, camera inconsistencies, and variability in leaf orientation or lighting conditions.

The specific transformations applied are as follows:

- Random Horizontal Flipping: This technique flips images along the vertical axis with a set probability, effectively simulating the natural variability in how leaves might appear when captured from different perspectives in the field. This helps the model learn that the spatial orientation of a leaf does not alter its underlying disease characteristics.
- Brightness and Contrast Adjustment: Images were randomly adjusted for brightness and contrast within defined ranges. Since field conditions often vary due to time of day, weather, and camera exposure settings, this augmentation trains the model to ignore such inconsistencies and focus on relevant disease features, such as lesion color or texture patterns.
- Random Zooming: A zoom-in or zoom-out transformation was applied to replicate changes in the distance between the camera and the leaf surface. This helps the model learn to recognize disease symptoms at different scales and discourages overfitting to specific image sizes or framing conditions.
- Rotation and Flipping: Random rotations and additional flips (beyond horizontal) were used to simulate various leaf positions. This ensures the model does not become biased toward a particular orientation of leaf features (e.g., lesions or vein structure) and can accurately classify images taken from diverse angles.

These augmentations were implemented using TensorFlow's built-in image processing pipelines, which allowed for efficient, on-the-fly transformation of image batches during training. This approach reduced memory overhead and ensured that each training epoch was exposed to new image variations, improving model generalization without requiring manual expansion of the dataset.

Overall, these augmentation strategies played a crucial role in improving the robustness and accuracy of the model, particularly when evaluated on unseen test data. They also serve as a practical substitute for large-scale, real-world data collection efforts, which are often resource-intensive in agricultural research contexts.

5.2. **Hyperspectral Data Augmentation**

Hyperspectral imaging (HSI), while rich in information, presents unique challenges due to its high dimensionality and sensitivity to environmental and sensor-related noise. Each sample in a hyperspectral dataset typically contains reflectance values across hundreds of contiguous wavelength bands, many of which are highly correlated or prone to distortion due to factors such as lighting variation, sensor calibration drift, or ambient conditions during data acquisition. As a result, augmenting hyperspectral data requires methods that preserve the physical and biological integrity of spectral signals while introducing enough variability to prevent overfitting.

To address these concerns and improve the generalization ability of the 1D-ANN model, the following spectral-domain augmentation techniques were applied:

- Gaussian Noise Injection: Random Gaussian noise was added to the spectral vectors to simulate electronic and environmental noise that naturally occurs during sensor readings. This augmentation helps the model learn to focus on underlying spectral patterns relevant to plant health, rather than overfitting to noise artifacts that might be present only in the training data. The standard deviation of the noise was carefully controlled to avoid corrupting meaningful signal components.
- Spectral Stretching: This technique involves linearly scaling the reflectance values within a spectral signature, either globally or selectively within certain wavelength bands. Spectral stretching mimics the physiological variation that can occur due to factors like chlorophyll concentration, hydration levels, or leaf thickness-features which may vary naturally even among healthy plants. By exposing the model to stretched spectra, its ability to classify samples under biologically plausible variations was improved.
- Wavelength Shifting: In this augmentation, spectral data is slightly shifted along the wavelength axis to emulate minor misalignments that might occur due to sensor calibration errors, temperature-induced drift, or inconsistent illumination during data collection. The shift is typically within a few spectral bands and is often implemented using interpolation methods. This strategy trains the model to remain invariant to small spectral distortions while still recognizing the core features of healthy or diseased plant tissue.
- Random Spectral Flipping: To introduce further diversity into the training data, horizontal flipping of spectral signatures was applied. This technique mirrors the spectrum about its center, effectively creating a new but structurally similar signal. While not a physically accurate transformation in all cases, spectral flipping has been shown in prior studies to enhance robustness, especially when paired with other augmentations. It forces the model to learn more generalized relationships between spectral patterns and plant health labels.

Together, these spectral augmentation strategies play a critical role in enhancing the robustness of the model. Unlike traditional image augmentation, hyperspectral augmentations must balance randomness with biological plausibility. The transformations applied in this study were carefully designed to preserve the spectral integrity of meaningful features-such as disease-related absorption peaks or reflectance troughs-while introducing sufficient variability to improve model resilience against unseen field conditions.

By simulating real-world spectral distortions and biological variation, this augmentation pipeline enabled the ANN model to generalize more effectively, particularly when classifying individual spectral samples collected under different lighting, sensor configurations, or crop environments.

6.Dataset Loading and Optimization

Efficient data handling is a critical component of any machine learning pipeline, especially when working with high-dimensional inputs such as RGB image datasets and hyperspectral image datasets. In this study, two separate pipelines were implemented-one for RGB image-based disease classification, and another for spectral data-based classification using hyperspectral reflectance signatures. Each pipeline was carefully optimized to ensure fast, reliable data loading, transformation, and model training.

6.1. **RGB Image Dataset Pipeline**

The RGB dataset, used for classifying rice leaf diseases, was organized in a directory-based structure where each subfolder represented a specific combination of plant type and disease label. The dataset was loaded using TensorFlow's tf.data API, which supports parallel data loading, caching, and prefetching-techniques that are essential for maintaining high throughput during training.

Key optimization steps included:

- Image Decoding and Resizing: All images were decoded from JPEG format and resized to a fixed resolution of 224 × 224 pixels to match the input requirements of the EfficientNetB0 model.
- Normalization: Pixel values were scaled to the range [0, 1] using min-max normalization to ensure uniform input distribution and promote stable gradient flow during training.
- Shuffling and Batching: The training data was shuffled with a buffer size proportional to the dataset size and then batched into mini-batches of size 32. This reduces model bias and enhances convergence.
- Caching and Prefetching: To minimize disk I/O bottlenecks, processed image batches were cached in memory and prefetched asynchronously. This ensured that the GPU was never idle, thereby reducing overall training time.
- Label Mapping: Labels were extracted using folder name parsing, and a lookup table was constructed to assign integer codes to both plant types and disease classes for multi-output classification.

6.2. **Hyperspectral Image Dataset Pipeline**

The hyperspectral dataset consisted of one-dimensional spectral vectors, each corresponding to the average reflectance of a tomato leaf across a range of wavelengths. These vectors are stored as NumPy arrays for fast numerical operations.

To optimize the hyperspectral data pipeline:

- Spectral Resampling: Spectral data was uniformly interpolated to a fixed length of 180 bands to ensure consistency across samples and compatibility with the ANN model.
- Normalization: Each spectral vector was standardized using Z-score normalization (mean = 0, standard deviation = 1). This prevents any individual wavelength band from disproportionately influencing the model's learning process.
- Dimensionality Reduction: Principal Component Analysis (PCA) was applied to reduce redundancy and compress the high-dimensional spectral data while preserving essential variance. The number of principal components was determined based on the cumulative explained variance, with a target threshold of 95% retained variance.
- Batching and Shuffling: The processed spectral data was batched into smaller sets for efficient training, and shuffling was applied to ensure randomness in data presentation across epochs.

- Caching and Prefetching: Similar to the image pipeline, spectral batches were cached in memory and preloaded during training to maintain a consistent training loop and eliminate CPU-GPU communication delays.

These optimizations significantly reduced preprocessing overhead and enabled real-time data feeding during training sessions. They also allowed for seamless integration with GPU acceleration, ensuring that both the CNN and ANN models could be trained efficiently on modern hardware configurations, including TPUs and GPUs available via Google Colab.

By designing separate but parallel data pipelines tailored to the specific requirements of image and spectral data, this study demonstrates a flexible and scalable approach to agricultural disease detection using multimodal deep learning systems.

7.Data Analysis

A thorough analysis of the dataset is essential to understand the underlying distribution of features, evaluate data quality, identify class imbalances, and extract meaningful patterns that inform model design. In this study, separate data analysis procedures were applied to the RGB image dataset and the hyperspectral dataset, each aligned with the characteristics of their respective modalities.

7.1. **RGB Image Dataset Analysis**

The RGB dataset used for rice leaf disease classification comprises high-resolution images, each labeled with two attributes: plant type (rice or other) and disease category (e.g., bacterial leaf blight, leaf blast, brown spot, healthy, etc.). To ensure balanced model training and effective evaluation, the distribution of samples across each label was analyzed.

- Class Distribution: In the dataset, while some classes such as "brown spot" and "leaf blast" were well-represented, others like "leaf scald" and "other disease" had relatively fewer samples, indicating a potential class imbalance. This was addressed through data augmentation and loss-weighted training.
- Image Quality Assessment: Images were examined for resolution consistency, brightness variation, and presence of noise or artifacts. While most images were captured under controlled conditions, minor variations in lighting and background were present, justifying the inclusion of brightness and contrast augmentation during preprocessing.
- Visual Feature Consistency: A qualitative visual inspection of each disease category revealed that many symptoms-such as spots, discoloration, and lesions-exhibit similar morphological features, especially under varied lighting. This validated the use of a deep convolutional architecture capable of capturing fine-grained visual patterns across multiple layers.

7.2. **Hyperspectral Image Data Analysis**

The hyperspectral dataset, used for classifying tomato plant health status, contains reflectance spectra representing hundreds of wavelengths per sample. Each spectrum is labeled as either healthy or infected based on the presence of bacterial wilt.

- Spectral Profile Comparison: Initial spectral analysis involved plotting the average reflectance curves for healthy and infected leaves. Distinct patterns were observed-infected leaves exhibited subtle shifts in reflectance intensity, particularly in the red-edge and near-infrared (NIR) regions, which are known to correlate with chlorophyll content and leaf water stress.

- Normalized Difference Spectral Index (NDSI): To quantify differences between classes, the Normalized Difference Spectral Index (NDSI) was calculated using specific band pairs sensitive to plant stress.
- Principal Component Analysis (PCA): PCA was applied not only for dimensionality reduction but also as an analytical tool to examine variance distribution across bands. The number of retained components was chosen based on the cumulative explained variance, ensuring that over 95% of the total variance was preserved. The first few principal components captured a significant portion of the variance (over 95%), and visualizing the first two components showed clear clustering between healthy and infected classes. This indicated strong separability in the spectral domain and justified the use of ANN for binary classification.
- Outlier Detection: PCA and Mahalanobis distance were used to identify outlier spectra. A small subset of samples was flagged and manually reviewed. Most anomalies were due to partial occlusion, soil contamination, or sensor noise during data acquisition. These were removed to maintain dataset integrity.

International Journal on Science and Technology (IJSAT)
E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

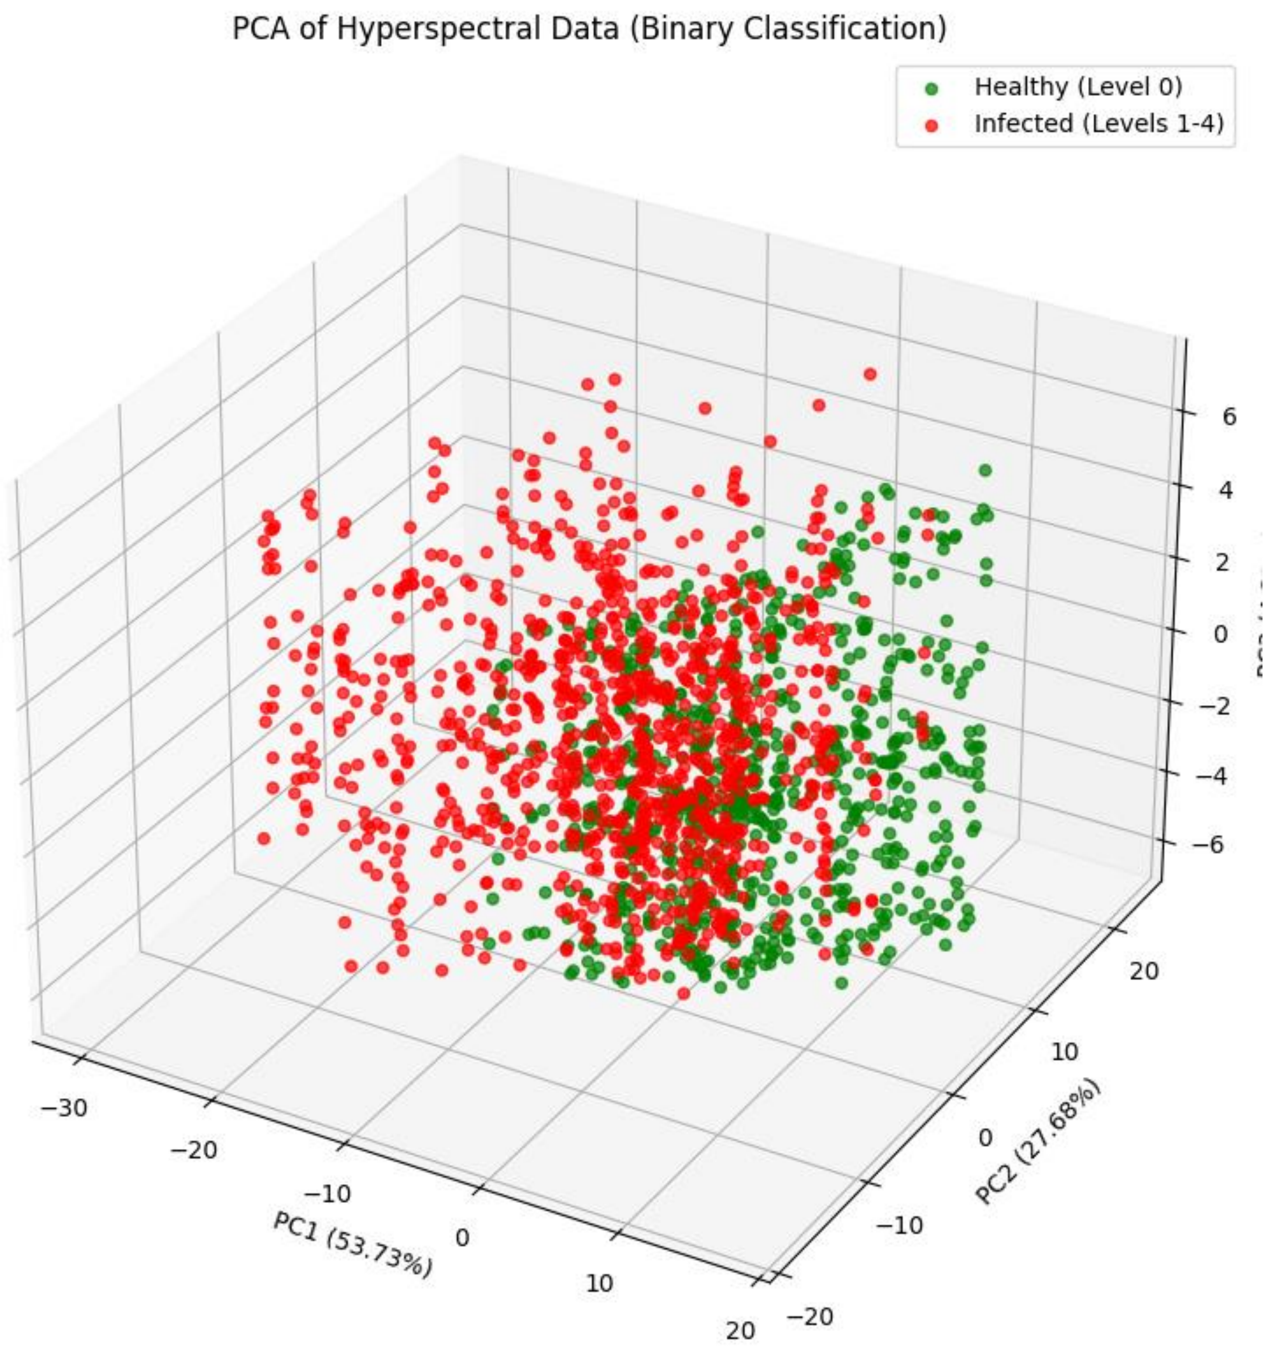

Figure 6: PCA of Hyperspectral Data

By conducting this dual-modality data analysis, it was confirmed that both the RGB and hyperspectral datasets provided sufficient, discriminative features for machine learning models. The RGB dataset benefited from visual inspection and class balancing techniques, while the hyperspectral dataset exhibited strong physiological indicators of disease, suitable for early-stage classification. These insights guided the choice of architecture and preprocessing techniques in subsequent stages of the study.

8.Performance Evaluation

The performance of the proposed models was rigorously evaluated through a series of quantitative metrics and controlled experiments. This section details the evaluation criteria, presents the results obtained for both RGB and hyperspectral models, and discusses the impact of optimization strategies on overall performance.

8.1. Experimental Results

Phase 1: RGB Image Classification (Rice Plant Disease Detection)

The EfficientNetB0-based multi-output CNN model demonstrated strong performance across both classification tasks:

- Plant Type Classification Accuracy: 100%
- Disease Classification Accuracy: 87.5% on the test set
- ROC - AUC: Area Under the Curve (AUC) values ranged from 0.98 to 1.00, with leaf scald achieving perfect separation. This highlights the model's high sensitivity and specificity across all classes.

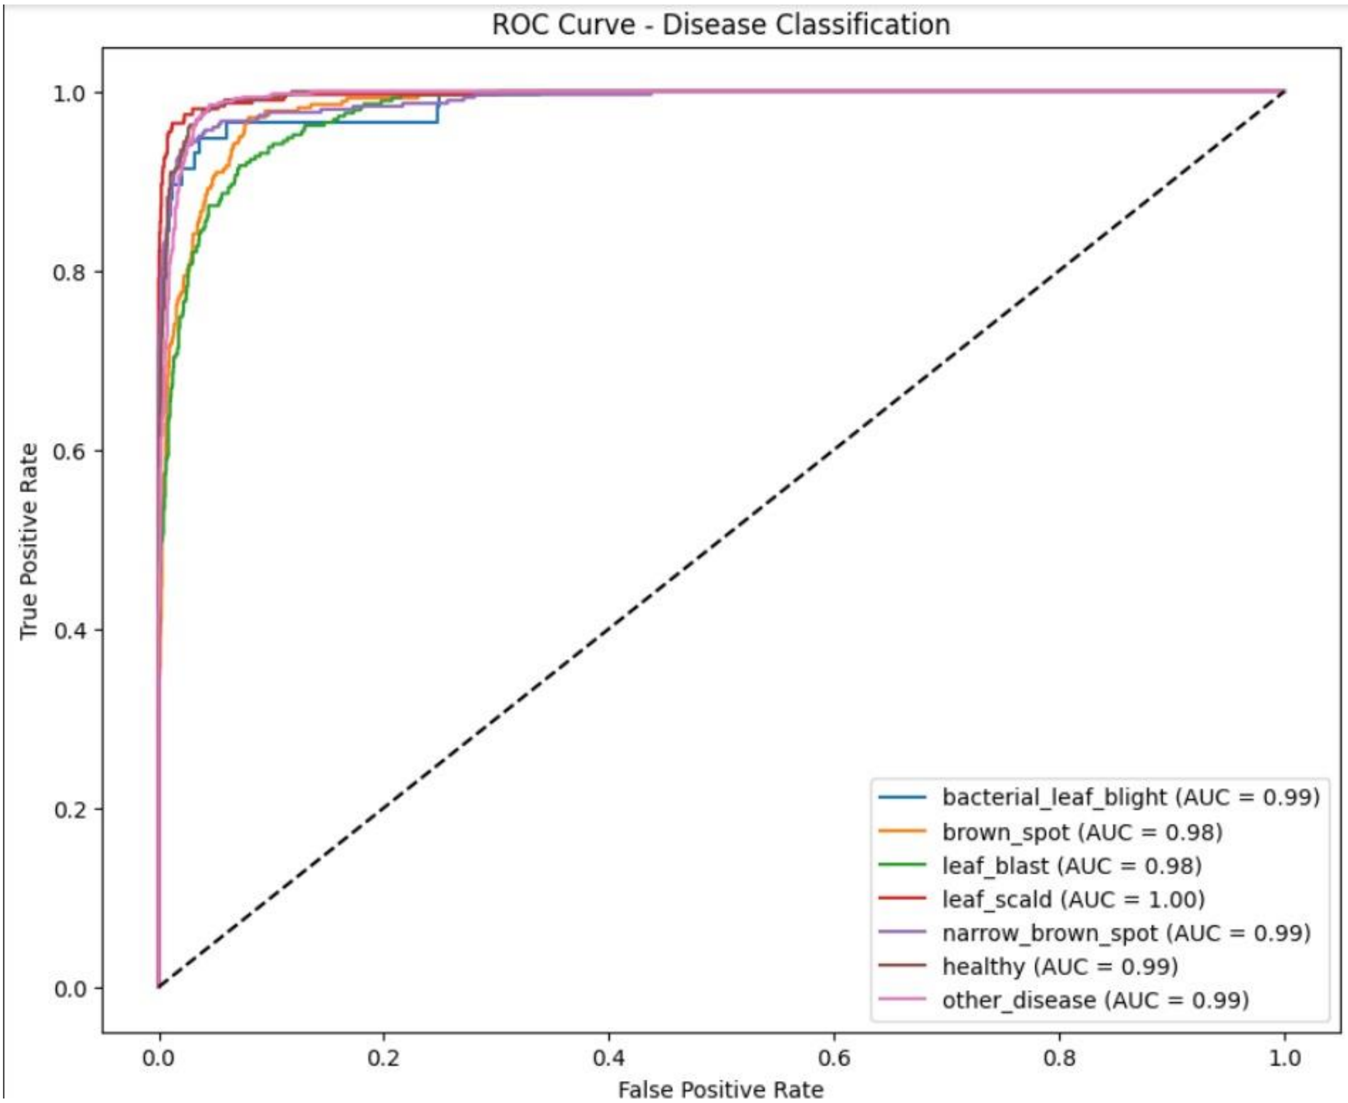

Figure 7: ROC Curve showing True Positive Rate vs. False Positive Rate for each disease class

- Confusion Matrix Analysis: Most confusion occurred between visually similar classes, such as brown spot and narrow brown spot, while classes like healthy and leaf scald were predicted with high precision.

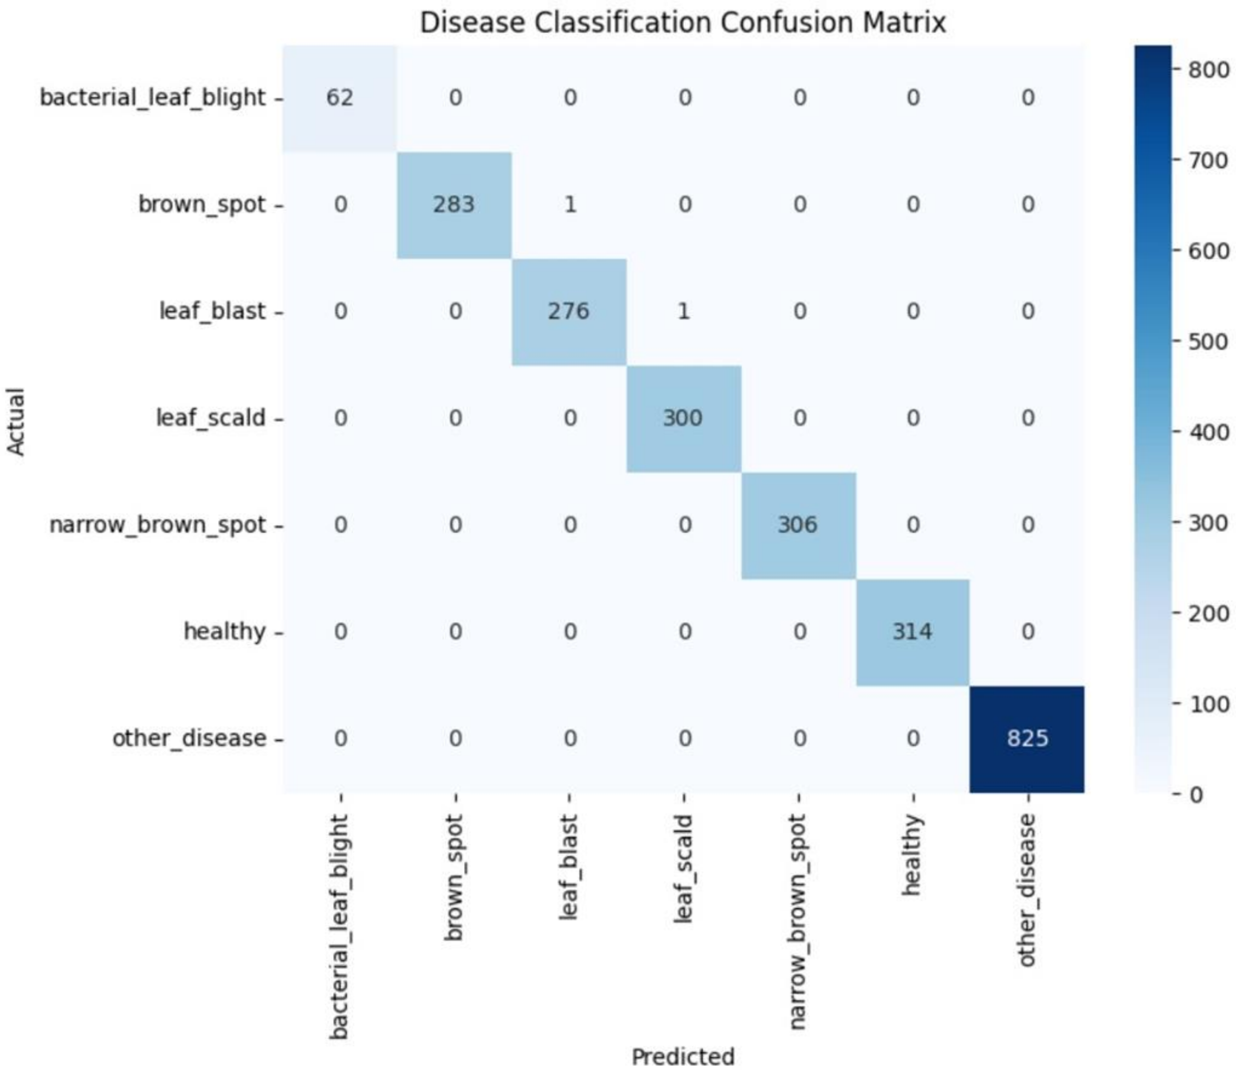

Figure 8: RGB Disease Classification Confusion Matrix

These results indicate that the model successfully learned to distinguish between nuanced disease patterns while maintaining generalization to unseen data.

Phase 2: Hyperspectral Classification (Tomato Bacterial Wilt Detection)

The 1D-ANN model achieved high accuracy in detecting bacterial wilt based on spectral reflectance:

- Validation Accuracy: 96%
- Test Accuracy: 89.77%
- Precision (Infected): 100%
- Recall (Infected): 82.69%
- F1-Score (Infected): 90.57%
- Confusion Matrix:

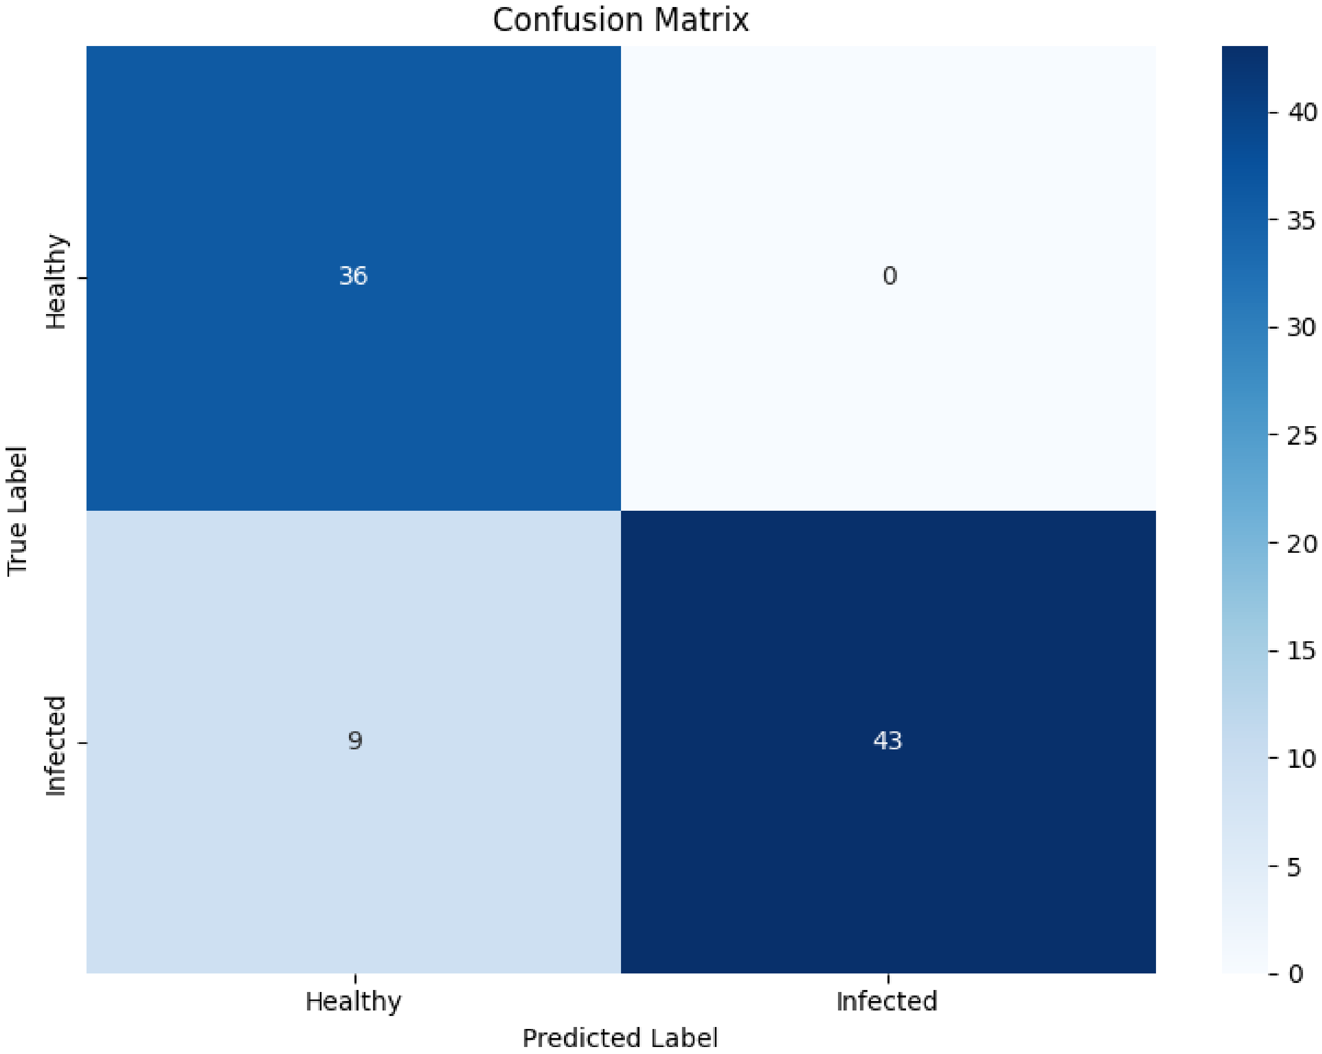

Figure 9: Hyperspectral Dataset Confusion Matrix

These metrics confirm the model's ability to accurately differentiate between healthy and infected leaves, even in the early, asymptomatic stages of the disease. The high precision indicates minimal false positives, while the high recall underscores the model's capacity to detect a majority of infected samples.

8.2. Model Optimization Impact

Multiple optimization techniques were implemented during model development, each contributing significantly to performance and training stability.

For RGB Model

- Loss Weighting: Assigning higher importance to disease classification (weight = 0.9) and lower to plant classification (weight = 0.1) resulted in improved learning focus on the more critical disease detection task.
- Data Augmentation: Application of random flipping, brightness/contrast adjustments, and zooming led to an observable increase in generalization, especially for underrepresented disease classes.

- Early Stopping and Checkpointing: Prevented overfitting by halting training at optimal validation performance and preserving the best model.

For Hyperspectral Model

- Dimensionality Reduction (PCA): Reduced training time and memory consumption while preserving over 95% of spectral variance, significantly improving model efficiency.
- Spectral Data Augmentation: Techniques like Gaussian noise injection and wavelength shifting made the model more resilient to real-world variability and measurement noise.
- Dropout and Batch Normalization: These regularization strategies effectively reduced overfitting and stabilized training, especially in the presence of limited spectral data samples.
- Adaptive Learning Rate (Adam Optimizer): Enabled faster convergence and smoother loss minimization across epochs, contributing to the high final accuracy.

Overall, the performance evaluation validates the reliability, accuracy, and practical applicability of both models in their respective domains. The RGB model offers a scalable solution for visual disease detection, while the hyperspectral model introduces a powerful method for early, non-invasive diagnosis using spectral patterns-demonstrating clear potential for integration into real-time precision agriculture systems.

9.Future Works

While the current study demonstrates promising results in plant disease detection using both RGB imaging and hyperspectral data, there remain several opportunities for further research and development. These future directions aim to enhance model scalability, improve robustness under real-world conditions, and expand the scope of agricultural applications.

9.1. Real-Time Field Deployment

One of the most important next steps is the deployment of these models in real-world agricultural environments. This involves integrating the trained models into edge devices such as mobile phones, drones (UAVs), or portable hyperspectral scanners. Real-time deployment will require further model optimization, including conversion to lightweight formats (e.g., TensorFlow Lite) and reducing inference latency. Field trials should also be conducted to evaluate performance under uncontrolled conditions such as variable lighting, occlusion, and background noise.

9.2. Multimodal Data Fusion

Future systems can benefit significantly from fusing multiple sensing modalities, such as combining hyperspectral data with thermal imaging, RGB video feeds, LiDAR, or fluorescence spectroscopy. By integrating multimodal information, models can capture both surface-level and internal stress indicators, improving the reliability and precision of diagnosis. Multimodal architectures, including attention-based fusion networks, could be explored to enhance feature learning.

9.3. Generalization Across Crops and Conditions

The current models are trained and validated on specific crops (rice and tomato) and diseases. However, crop-specific and condition-specific tuning limits the scalability of such systems. Future work should focus on developing generalized frameworks capable of detecting a wider range of plant diseases across different crops, varieties, soil types, and environmental conditions. This may involve using domain adaptation and transfer learning techniques to enable cross-crop applicability.

9.4. Explainable AI (XAI)

While deep learning models often achieve high accuracy, they are typically treated as "black boxes." For agricultural adoption, especially by domain experts and farmers, interpretability is essential. Implementing explainable AI methods such as SHAP (Shapley Additive Explanations), LIME (Local Interpretable Model-Agnostic Explanations), or Grad-CAM (Gradient-weighted Class Activation Mapping) can provide insights into which spectral bands or image features are influencing predictions, thereby increasing user trust and model transparency.

9.5. Robustness to Environmental Variability

Variations in lighting, temperature, humidity, and plant age can significantly affect spectral and visual data. Developing adaptive preprocessing algorithms that can dynamically adjust to changing field conditions is crucial for ensuring consistent performance. Future studies should also investigate training models using augmented datasets that simulate real-world noise and variability.

9.6. Data Collection and Label Expansion

Improving the diversity and size of the training datasets is another key area of future work. This includes collecting more labeled data from different geographical regions, capturing early-stage infections, and including rare or co-occurring diseases. Additionally, introducing disease severity scoring (e.g., mild, moderate, severe) rather than binary labels could enable more granular diagnosis and better guide treatment recommendations.

9.7. IoT and Smart Farming Integration

To realize the full potential of intelligent plant disease detection, future systems should be integrated with Internet of Things (IoT) infrastructures. Sensors placed in the field can continuously monitor plant health, feeding data to edge or cloud-based models for real-time inference. Combining this with automated irrigation, fertilization, or pesticide delivery systems can create closed-loop precision agriculture systems that respond proactively to disease outbreaks.

In summary, while this study lays a strong foundation for AI-assisted plant disease detection using RGB and hyperspectral imaging, continued research and innovation are required to scale these solutions for real-world deployment. Advancements in model generalization, interpretability, multimodal integration, and smart device deployment will be key to making early, accurate, and automated plant disease diagnosis a standard part of modern agriculture.

10. Conclusion

This study explored the use of deep learning techniques combined with both RGB imaging and hyperspectral sensing for accurate and early-stage plant disease detection in rice and tomato crops. By leveraging multimodal datasets and carefully optimized machine learning models, the research demonstrates the feasibility and effectiveness of integrating artificial intelligence into precision agriculture.

In the first phase, a multi-output Convolutional Neural Network (CNN) based on the EfficientNetB0 architecture was developed to classify rice leaf diseases using RGB images. The model achieved high classification accuracy across both plant type and disease categories, demonstrating robustness in handling visual variability and class imbalance. Data augmentation and efficient training strategies contributed significantly to its performance.

In the second phase, a 1D Artificial Neural Network (ANN) was employed to detect bacterial wilt in tomato plants using hyperspectral reflectance data. Through preprocessing techniques such as Savitzky-Golay filtering and Principal Component Analysis (PCA), the spectral data was refined and reduced for efficient model training. The hyperspectral model achieved strong results in binary classification, with particularly high precision and recall, highlighting its potential for early, non-invasive disease detection.
The integration of both traditional RGB imaging and advanced hyperspectral analysis allowed for a comparative and complementary approach to plant disease identification. While RGB methods are accessible and effective for detecting visible symptoms, hyperspectral imaging offers unique advantages in identifying biochemical changes before symptoms appear-an essential capability for proactive crop management.
Overall, the findings underscore the transformative role of artificial intelligence and spectral sensing in agriculture. These technologies can significantly enhance disease surveillance, reduce crop loss, and support sustainable farming practices. Moving forward, expanding these models to operate in real-world environments, across different crops, and in real-time systems will be critical for achieving their full potential. With continued research and technological advancement, AI-driven plant disease detection can play a vital role in ensuring global food security and promoting smarter, data-informed agricultural practices.

References

1. Zhang, M.; Tang, S.; Lin, C.; Lin, Z.; Zhang, L.; Dong, W.; Zhong, N. Hyperspectral Imaging and Machine Learning for Diagnosing Rice Bacterial Blight Symptoms Caused by Xanthomonas oryzae pv. Oryzae Pantoea ananatis and Enterobacter asburiae. Plants 2025, 14, 733. https://doi.org/10.3390/plants14050733
2. A. Bhargava, A. Shukla, O. P. Goswami, M. H. Alsharif, P. Uthansakul and M. Uthansakul, “Plant Leaf Disease Detection, Classification, and Diagnosis Using Computer Vision and Artificial Intelligence: A Review,” in IEEE Access, vol. 12, pp. 37443-37469, 2024, doi: 10.1109/ACCESS.2024.3373001.
3. W. Shafik, A. Tufail, A. Namoun, L. C. De Silva and R. A. A. H. M. Apong, “A Systematic Literature Review on Plant Disease Detection: Motivations, Classification Techniques, Datasets, Challenges, and Future Trends,” in IEEE Access, vol. 11, pp. 59174-59203, 2023, doi: 10.1109/ACCESS.2023.3284760.
4. R. Jiang et al., “Hyperspectral Detection of Pesticide Residues in Black Vegetable Based on Multi-Classifier Entropy Weight Method,” in IEEE Access, vol. 13, pp. 55701-55711, 2025, doi: 10.1109/ACCESS.2025.3553384.
5. Dolatabadian, A., Neik, T.X., Danilevicz, M.F., Upadhyaya, S.R., Batley, J. & Edwards, D. (2025) Image-based crop disease detection using machine learning. Plant Pathology, 74, 18–38. Available from:https://doi.org/10.1111/ppa.14006
6. L. Li, S. Zhang and B. Wang, “Plant Disease Detection and Classification by Deep Learning-A Review,” in IEEE Access, vol. 9, pp. 56683-56698, 2021, doi: 10.1109/ACCESS.2021.3069646.
7. Md. Manowarul Islam, Md Abdul Ahad Adil, Md. Alamin Talukder, Md. Khabir Uddin Ahamed, Md Ashraf Uddin, Md. Kamran Hasan, Selina Sharmin, Md. Mahbubur Rahman, Sumon Kumar Debnath, DeepCrop: Deep learning-based crop disease prediction with web application, Journal of

Agriculture and Food Research, Volume 14, 2023, 100764, ISSN 2666-1543 https://doi.org/10.1016/j.jafr.2023.100764.
8. Jung, M., Song, J.S., Shin, AY. et al. Construction of deep learning-based disease detection model in plants. Sci Rep 13, 7331 (2023). https://doi.org/10.1038/s41598-023-34549-2
9. A. Chakrabarti and T. Zickler, “Statistics of real-world hyperspectral images,” CVPR 2011, Colorado Springs, CO, USA, 2011, pp. 193-200, doi: 10.1109/CVPR.2011.5995660.
10. G. K. T and K. V. Shashank, “Smart Farming based on AI, Edge Computing and IoT,” 2022 4th International Conference on Inventive Research in Computing Applications (ICIRCA), Coimbatore, India, 2022, pp. 324-327, doi: 10.1109/ICIRCA54612.2022.9985023.
11. S. S and J. J. Nair, “Wheat Disease Detection And Growth Stage Monitoring Using Deep Learning Architectures,” 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT), Delhi, India, 2023, pp. 1-5, doi: 10.1109/ICCCNT56998.2023.10307548.
12. Neeraj S. Kumar, Jyothika Sony, Arathi Premkumar, Meenakshi R, Jyothisha J Nair, Transfer Learning-based Object Detection Models for Improved Diagnosis of Tomato Leaf Disease, Procedia Computer Science, Volume 235, 2024, Pages 3025-3034, ISSN 1877-0509 https://doi.org/10.1016/j.procs.2024.04.286.
13. Hridya Krishna, R., Manoj, A., Vaishnavi, K.P., Nair, J.J. (2023). Deep Learning Approach for Brown Spot Detection and Nitrogen Deficiency Estimation in Rice Crops. In: Tuba, M., Akashe, S., Joshi, A. (eds) ICT Systems and Sustainability. Lecture Notes in Networks and Systems, vol 516. Springer, Singapore. https://doi.org/10.1007/978-981-19-5221-0 65
14. Thirumurthy, B., Parameswaran, L. & Vaiapury, K. A deep learning and transfer learning model for intra-change detection in images. J Supercomput 80, 11640–11660 (2024). https://doi.org/10.1007/s11227023-05878-w
15. Uchimuthu, M., Sonai, V., Chitra, S., Roy, R.V. (2024). Metaheuristics Feature Selection Algorithms for Identification and Classification of Mango Pests Diseases. In: Gopi, E.S., Maheswaran, P. (eds) Proceedings of the International Conference on Machine Learning, Deep Learning and Computational Intelligence for Wireless Communication. MDCWC 2023. Signals and Communication Technology. Springer, Cham. https://doi.org/10.1007/978-3-031-47942-7 3