International Journal of Advances in Applied Computational Intelligence

♦ **ISSN:** 2833-5600 **Vol:** 08, No. 01, 2026 ♦ PP. 13–19
Article DOI: https://doi.org/10.54216/IJAACI.080103

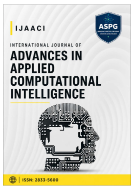

FULL LENGTH ARTICLE

A Hybrid CNN Bi-LSTM Framework for Multi-Class Plant Disease Detection and Health Value Estimation

Janani J.[1,*] **Gautham Praveen R.**[2] **Suguna C.**[2] **Ponni Vinothini K.**[3]

[1] Student, Department of Computer Science and Engineering, Syed Ammal Engineering College, Ramanathapuram-623501, India
[2] Assistant Professor, Department of Computer Science and Business System, Syed Ammal Engineering College, Ramanathapuram-623501, India
[3] Assistant Professor, Department of Computer Science Engineering, Syed Ammal Engineering College, Ramanathapuram-623501, India

Emails: jjananij147@gmail.com · gauthampraveenr@syedengg.ac.in · sugunaranjith@syedengg.ac.in · ponnicse2174@syedengg.ac.in

Received: January 30, 2026 Revised: February 27 2026 Accepted: March 29, 2026 ⋆ Corresponding author

ABSTRACT

Accurate and early identification of plant diseases is essential for ensuring sustainable agriculture and maximizing crop productivity. This paper presents a hybrid deep learning framework integrating Convolutional Neural Networks (CNN) and Bidirectional Long Short-Term Memory (BiLSTM) networks for multi-class plant disease detection, classification, and Plant Health Value (PHV) estimation. The proposed framework begins with a comprehensive data preprocessing pipeline involving image resizing, normalization, and augmentation to improve model robustness. The CNN module extracts critical spatial and visual features such as lesion shape, leaf texture, and color intensity, while the BiLSTM model captures temporal and sequential feature correlations to accurately learn disease progression patterns. A Decision Support System (DSS) is incorporated to compute the Plant Health Value (PHV), where PHV ranges from 0% (Healthy) to 100% (Severely Unhealthy), indicating the severity of disease infection. Additionally, the DSS generates actionable recommendations to assist in early intervention and treatment planning. Experimental results on a multi-species plant dataset demonstrate that the proposed CNN–BiLSTM hybrid model significantly improves accuracy, interpretability, and early disease prediction compared to conventional CNN based methods, offering a robust and intelligent framework for automated plant health monitoring.

Keywords: Plant Disease Detection ▪ CNN ▪ BiLSTM ▪ Deep Learning ▪ Decision Support System ▪ Plant Health Value ▪ Precision Agriculture

Accurate and early identification of plant diseases is essential for ensuring sustainable agriculture and maximizing crop productivity. This paper presents a hybrid deep learning framework integrating Convolutional Neural Networks (CNN) and Bidirectional Long Short-Term Memory (BiLSTM) networks for multi-class plant disease detection, classification, and Plant Health Value (PHV) estimation. The proposed framework begins with a comprehensive data preprocessing pipeline involving image resizing, normalization, and augmentation to improve model robustness. The CNN module extracts critical spatial and visual features such as lesion shape, leaf texture, and color intensity, while the BiLSTM model captures temporal and sequential feature correlations to accurately learn disease progression patterns. A

Decision Support System (DSS) is incorporated to compute the Plant Health Value (PHV), where PHV ranges from 0% (Healthy) to 100% (Severely Unhealthy), indicating the severity of disease infection. Additionally, the DSS generates actionable recommendations to assist in early intervention and treatment planning. Experimental results on a multi-species plant dataset demonstrate that the proposed CNN–BiLSTM hybrid model significantly improves accuracy, interpretability, and early disease prediction compared to conventional CNN based methods, offering a robust and intelligent framework for automated plant health monitoring.

1. INTRODUCTION

Agriculture remains one of the most important sectors supporting the global economy, and the health of crops plays a critical role in ensuring food production and security. When crops are affected by diseases, productivity can decline significantly, creating challenges for farmers and agricultural systems. Plant diseases caused by bacteria, viruses, and fungi are among the major threats to crop sustainability. In many cases, these diseases spread quickly across fields, reducing both the quality and quantity of agricultural yield.

Traditionally, identifying plant diseases has relied heavily on manual inspection by agricultural experts. Although this approach can be effective, it requires time, experience, and continuous monitoring of crops. In large farming areas, manual observation becomes difficult to maintain consistently and may lead to delayed detection or human error. Because of these limitations, researchers are increasingly exploring automated techniques that can detect and classify plant diseases accurately and at an early stage.

Recent developments in deep learning (DL) have shown promising results in tasks involving image recognition and pattern analysis. These capabilities make deep learning techniques well suited for agricultural applications, particularly for analysing plant leaf images. Convolutional Neural Networks (CNNs) are widely used in image processing because they can automatically learn spatial features such as leaf colour variations, texture differences, and visible disease spots or lesions. By identifying these visual characteristics, CNN models can help distinguish between healthy and diseased plants.

While CNNs are effective at extracting spatial information, understanding complex feature relationships can further improve detection accuracy. For this reason, Recurrent Neural Networks (RNNs), particularly Bidirectional Long Short-Term Memory (BiLSTM) models, are useful. BiLSTM networks process information in both forward and backward directions, allowing the model to capture deeper relationships within the extracted features. This bidirectional learning helps the system recognize patterns that might otherwise remain unnoticed.

The LSTM structure includes three main components known as gates: the forget gate, input gate, and output gate. These gates control how information flows through the network. The forget gate removes unnecessary data, the input gate introduces new relevant information, and the output gate transfers useful information to the next stage of processing. Through this mechanism, LSTM models are able to retain important patterns over time and reduce issues such as the vanishing gradient problem that often occurs in traditional RNN architectures.

In this work, a hybrid deep learning framework that combines CNN and BiLSTM models is proposed for multi-class plant disease detection and classification, along with Plant Health Value (PHV) estimation. The CNN component extracts important visual features from leaf images, while the BiLSTM module analyzes the relationships among these features to improve classification performance. The final output is evaluated by a Decision Support System (DSS), which applies rule-based logic and threshold analysis to estimate the Plant Health Value. In the proposed framework, higher PHV values indicate greater disease severity, where 0% represents a healthy plant and 100% represents severe infection. Based on this value, the system can also suggest early control measures for better crop management.

The overall workflow begins with a pre-processing stage where leaf images are resized, normalized, and augmented to improve model reliability and reduce noise. These processed images are then passed through the CNN module to extract spatial features such as colour differences, lesion patterns, and texture variations. The extracted features are subsequently analyzed by the BiLSTM module, which captures sequential relationships and supports a better understanding of disease progression.

By combining spatial feature extraction, sequential learning, and decision support, the proposed CNN–BiLSTM framework offers an effective approach for early and accurate plant disease detection. Experimental results using a multi-species dataset show that the proposed model performs more reliably than traditional CNN-based methods. As a result, the system can serve as a useful tool for farmers, researchers, and agricultural monitoring systems working toward sustainable and precision-based farming practices.

2. RELATED WORK

Recent advancements in deep learning have significantly improved the performance of plant disease detection systems. Earlier approaches relied on traditional machine learning techniques such as Support Vector Machines (SVM) and k-Nearest Neighbors (KNN), where feature extraction was performed manually using color, texture, and shape descriptors. However, these methods were highly sensitive to variations in lighting, background noise, and image quality, which limited their effectiveness in real-world agricultural environments.

With the introduction of Convolutional Neural Networks (CNNs), automatic feature extraction became possible, enabling models to learn discriminative patterns directly from raw images. Studies such as Mohanty et al. demonstrated that CNN-based models trained on large-scale datasets like PlantVillage can achieve high classification accuracy. Despite this success, CNN models primarily focus on spatial feature extraction and do not explicitly capture relationships between different feature regions.

To address this limitation, hybrid approaches combining CNN with sequence learning models have been explored. CNN-LSTM models have shown improved performance by learning dependencies among extracted features. However, conventional LSTM processes information in a single direction, which restricts its ability to fully capture contextual relationships.

Bidirectional Long Short-Term Memory (BiLSTM) networks overcome this limitation by processing feature sequences in both forward and backward directions. This enables a more comprehensive understanding of feature dependencies, which is particularly useful for distinguishing visually similar plant diseases. Although existing studies have achieved promising classification accuracy, most of them focus only on disease identification and do not provide meaningful interpretation of plant health status. Therefore, this work proposes a hybrid CNN–BiLSTM framework integrated with a Decision Support System (DSS) to perform multi-class classification along with Plant Health Value (PHV) estimation for better interpretability and practical usability.

3. PROPOSED METHODOLOGY

This study presents a hybrid deep learning framework that integrates Convolutional Neural Networks (CNN) and Bidirectional Long Short-Term Memory (BiLSTM) networks for multi-class plant disease detection and Plant Health Value (PHV) estimation. The proposed system is designed as a multi-stage pipeline that processes input images, extracts meaningful features, models contextual relationships, and generates interpretable outputs.

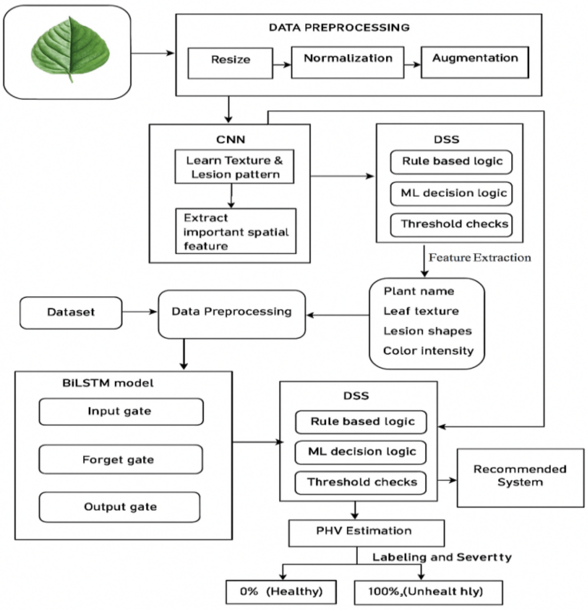

Figure 1. Proposed System Architecture Diagram

3.1 Model Architecture Details

The proposed model combines Convolutional Neural Networks (CNN) and Bidirectional Long Short-Term Memory (BiLSTM) to capture both spatial and contextual features from plant leaf images.

The input images are resized to 224 x 224 x 3 and passed through a CNN consisting of four convolutional layers with 32, 64, 128, and 256 filters respectively. Each convolution layer uses a kernel size of 3 x 3 and is followed by a Rectified Linear Unit (ReLU) activation function. Max pooling layers with a pool size of 2 x 2 are applied to reduce spatial dimensions and computational complexity.

Batch normalization is incorporated to stabilize training, and dropout with a rate of 0.5 is used to prevent overfitting. The resulting feature maps are flattened and reshaped into a sequential format.

The BiLSTM layer, consisting of 128 units, processes the feature sequence in both forward and backward directions to capture dependencies among extracted features. The output is then passed through a fully connected dense layer with 64 neurons, followed by a Softmax layer that produces probability values for each of the 38 classes.

The model is trained using the Adam optimizer with a learning rate of 0.001 and categorical cross-entropy as the loss function.

3.2 Dataset Description

The dataset used in this study is derived from the publicly available PlantVillage dataset, which is widely used for plant disease classification research. The dataset contains a total of 54,306 labeled images of plant leaves collected under controlled conditions.

The dataset includes 38 distinct classes representing both healthy and diseased plant leaves across multiple crop species such as tomato, potato, apple, and corn. Each class corresponds to a specific disease type or a healthy category.

All images are resized to a fixed dimension of 224 x 224 x 3 to ensure consistency in input size for the deep learning model. The dataset is divided into three subsets for model development:

Training set: 70% (38,014 images)

Validation set: 15% (8,146 images)

Testing set: 15% (8,146 images)

This split ensures that the model is trained effectively while also being evaluated on unseen data to measure its generalization performance.

3.3 Data Acquisition Module

The Data Acquisition Module is responsible for gathering and organizing plant leaf images that are required to build the proposed plant disease detection system. The effectiveness of the deep learning model largely depends on the quality, diversity, and quantity of the collected images, as these factors influence how well the model learns patterns and generalizes to new data.

In this research, leaf images are collected from publicly available agricultural datasets that include both healthy and diseased plant samples from different plant species. The dataset covers several disease categories, including fungal, bacterial, and viral infections, allowing the system to perform multi-class disease identification.

To improve the reliability of the model in practical agricultural scenarios, the dataset contains images captured under different real-world conditions. These variations include differences in lighting, leaf orientation, background environments, and levels of disease severity. Including such variations helps the model learn realistic disease patterns and improves its ability to perform effectively outside controlled laboratory conditions.

Each image in the dataset is assigned a label that represents its corresponding disease category. The dataset is carefully reviewed to verify the correctness of these labels. During this process, incorrectly labelled or unclear samples are removed, and duplicate or poor-quality images are filtered out to ensure that the dataset remains accurate and consistent.

The final output of the Data Acquisition Module is a clean, diverse, and properly labelled collection of plant leaf images. This dataset forms the basis for the subsequent stages of the system, including pre-processing, feature extraction, and model training for plant disease detection.

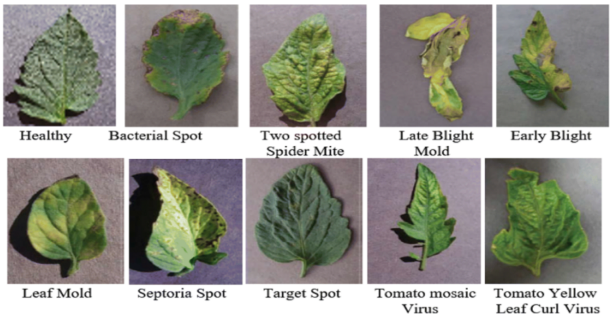

Figure 2. Multi-Class Plant Disease Dataset Samples

3.4 Data Preprocessing Module

The Data Reprocessing Module prepares raw leaf images for effective deep learning analysis by ensuring uniformity and robustness across all samples. Since images collected from different sources vary in size, illumination, and orientation, preprocessing standardizes them before feature extraction. This module performs three main operations: image resizing, normalization, and data augmentation.

Leaf images collected from different datasets often have varying resolutions and dimensions. However, Convolutional Neural Networks (CNNs) require input images to have a fixed size so that convolution operations and tensor calculations can be performed consistently. To satisfy this requirement, all input images are resized to a standard resolution of 224 x 224 pixels.

Resizing ensures that every image maintains the same spatial dimensions when it passes through the network. This uniformity helps reduce computational overhead and memory consumption while still preserving important disease-related visual features such as lesion edges, discoloration patterns, and texture differences on the leaf surface. The chosen resolution provides a good balance between computational efficiency and the retention of relevant visual information, making it well suited for CNN-based models.

Using a consistent input size also enables efficient batch processing during both training and inference. This improves the stability of the training process and allows the model to process images faster.

In digital images, pixel intensity values usually fall within the range of 0 to 255. Large pixel values can negatively affect the training process by causing unstable gradient updates and slower convergence during optimization.

The pixel intensity values are normalized to the range [0,1] using:

Where

I = original pixel intensity value

I norm = normalized pixel value

This scaling process stabilizes gradient descent and improves numerical efficiency during back propagation. It ensures that all input features contribute more evenly during model training, preventing higher intensity values from dominating the learning process. As a result, normalization helps accelerate convergence, improves training stability, and enhances the overall performance of the classification model.

Data augmentation is applied to increase the diversity of the training dataset and enhance the model's ability to generalize to unseen data. In real agricultural environments, leaf images may appear under different lighting conditions, orientations, scales, and background settings. If the model is trained only on limited or uniform images, its performance may decline when used in real-field scenarios. To address this challenge, augmentation techniques are applied to the training images while maintaining their original disease labels.

In this work, several transformation techniques are used to simulate natural variations. Small-angle rotations allow the model to recognize diseases even when leaves appear in different orientations. Horizontal and vertical flipping help the system learn orientation invariance. Zoom operations introduce scale variability, enabling the model to detect disease patterns at different sizes. Brightness adjustments simulate lighting changes caused by sunlight or shadows. Additionally, shearing transformations introduce slight geometric distortions, improving the model's tolerance to perspective variations. These transformations increase the effective size of the training dataset without modifying the original class labels.

By exposing the model to a wider range of visual conditions during training, data augmentation helps reduce over fitting. Instead of memorizing specific samples, the model learns general disease-related features that remain consistent across different environments. This improves the robustness and accuracy of the system when applied to real-world plant disease detection tasks.

3.5 Feature Extraction Module (CNN)

The Feature Extraction Module uses a Convolutional Neural Network (CNN) to automatically identify hierarchical spatial patterns from plant leaf images. The CNN architecture is composed of several layers including convolutional layers, activation functions, batch normalization layers, pooling layers, and dropout regularization. Together, these components enable the network to learn meaningful visual features that represent plant disease characteristics.

During convolution, learnable filters slide across the input image to generate feature maps that capture important patterns such as lesion boundaries, irregular textures, color changes, and the distribution of disease spots. The convolution operation can be represented mathematically as

The convolution operation used in the CNN layer can be mathematically expressed as:

$$F(x,y) = \sum_{i,j} I(i,j)\,K(x-i,y-j). \tag{1}$$

where $I(i, j)$ represents the input image pixel values and K denotes the convolution kernel. This operation extracts important spatial features such as edges, textures, and disease patterns from the image.

After convolution, the resulting feature maps pass through nonlinear activation functions, which increase the representational power of the network by introducing nonlinearity. Pooling layers are then applied to reduce the spatial dimensions of the feature maps while preserving the most relevant information. This dimensionality reduction decreases computational requirements and helps the model focus on the most informative features.

In contrast to conventional machine learning methods that depend on manually designed features, CNNs automatically learn discriminative patterns directly from raw image data. This hierarchical learning process allows the model to detect complex visual characteristics associated with plant diseases, improving classification accuracy while minimizing human bias in feature design.

3.6 Feature Sequence Learning Module (BiLSTM)

While the CNN module effectively extracts spatial features such as texture, color variations, and lesion patterns from plant leaf images, it does not explicitly capture relationships between different regions of the extracted feature maps. To model these dependencies, a Bidirectional Long Short-Term Memory (BiLSTM) network is incorporated after the CNN feature extraction stage.

The output of the CNN is a three-dimensional feature map of size $H \times W \times C$, where H and W represent spatial dimensions and C denotes the number of feature channels. This feature map is reshaped into a sequential format of length $T = H \times W$, where each time step corresponds to a feature vector of dimension C. Thus, the input to the BiLSTM can be represented as:

$$\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \ldots, \mathbf{x}_T\}. \tag{2}$$

where each $\mathbf{x}_t \in \mathbb{R}^C$ represents the feature vector at position t.

The BiLSTM processes this sequence in both forward and backward directions. The forward hidden state is computed as:

$$\overrightarrow{\mathbf{h}}_t = \mathrm{LSTM}_f(\mathbf{x}_t, \overrightarrow{\mathbf{h}}_{t-1}). \tag{3}$$

In the backward pass, the hidden state is computed as:

$$\overleftarrow{\mathbf{h}}_t = \mathrm{LSTM}_b(\mathbf{x}_t, \overleftarrow{\mathbf{h}}_{t+1}). \tag{4}$$

The final contextual representation at time step t is obtained by concatenating both hidden states:

$$\mathbf{h}_t = [\overrightarrow{\mathbf{h}}_t; \overleftarrow{\mathbf{h}}_t]. \tag{5}$$

where $\mathbf{x}_t$ is the input feature vector, $\overrightarrow{\mathbf{h}}_t$ is the forward hidden state, $\overleftarrow{\mathbf{h}}_t$ is the backward hidden state, and $\mathbf{h}_t$ is the concatenated BiLSTM output representation.

This bidirectional processing enables the model to capture relationships between different spatial regions of the leaf image, allowing better discrimination between visually similar disease patterns. Compared to standard LSTM, which processes data in a single direction, BiLSTM provides a more comprehensive understanding of feature dependencies. This improves classification performance, particularly in cases where disease characteristics are distributed across multiple regions of the leaf.

The output of the BiLSTM layer is then passed to a fully connected layer followed by a Softmax classifier for multi-class disease prediction.

Although the input is not inherently temporal, representing spatial features as sequences enables the BiLSTM to learn inter-region dependencies, improving feature representation.

3.7 Decision Support System (DSS) Module

The Plant Health Value (PHV) is computed based on the predicted probability of disease infection. Lower PHV values indicate healthy plants, whereas higher PHV values indicate severe disease infection. The Plant Health Value (PHV) is computed based on the predicted probability of disease as

$$\mathrm{PHV} = P_{\text{disease}} \times 100 \tag{6}$$

where P_{disease} represents the probability of disease

The Plant Health Value (PHV) ranges from 0% (Healthy) to 100% (Severely Unhealthy), indicating the severity of plant disease infection. Unlike conventional health scoring systems, the proposed PHV metric represents disease severity. Therefore, lower PHV values correspond to healthier plants, while higher PHV values indicate severe infection.

Table 1. PHV-Based Plant Disease Severity Classification

PHV (%)	Plant Condition
0-20	Healthy
21-40	Mild Infection
41-70	Moderate Infection
71-100	Severe Infection

Based on predefined threshold ranges, the DSS assigns a severity level to the detected plant condition Higher PHV values represent severe disease conditions, while lower PHV values indicate healthier plants. While lower ranges correspond to mild, moderate, or severe disease stages. These threshold values are determined empirically during experimental evaluation.

In addition to computing a health score, the DSS integrates domain knowledge to improve interpretability. The module considers both the predicted disease category and the calculated PHV when determining the final severity level. This combined analysis ensures that classification confidence is properly reflected in the plant health assessment.

The DSS therefore transforms the system from a simple classification model into a complete plant health monitoring framework. By converting raw probability outputs into an intuitive percentage-based health metric, the system becomes easier to interpret and more practical for real agricultural use.

Furthermore, the DSS reduces the risk of misinterpretation that may occur when relying only on probability values. For example, if the predicted disease probability is moderately high while the healthy probability is still significant, the PHV provides a balanced representation of plant condition rather than giving only a single categorical output.

F .Recommendation Module

The Recommendation Module enhances the practical applicability of the proposed system by providing actionable guidance based on the predicted disease class and severity level. Unlike traditional classification systems that only identify diseases, this module translates predictions into meaningful agricultural recommendations.

The module operates using a rule-based approach, where each predicted disease category is mapped to a predefined set of treatment strategies. These recommendations are further refined based on the Plant Health Value (PHV), which indicates the severity of infection.

For instance, in cases of mild infection, preventive measures such as removal of affected leaves and monitoring are suggested. For moderate conditions, targeted treatments such as appropriate fungicides or bactericides are recommended. In severe cases, more intensive interventions, including chemical treatments or crop isolation, are advised.

This integration of predictive modeling with domain-specific knowledge ensures that the system not only detects plant diseases but also supports informed decision-making for effective crop management.

3.8 Output & Visualization Module

The Output and Visualization Module presents the final predictions in a clear and interpretable format for end users. The system displays key information including the predicted disease class, classification confidence score, Plant Health Value (PHV), severity level, and recommended treatment actions.

To improve interpretability, the module incorporates visual representations such as probability distribution charts and PHV-based health indicators. These visual elements help users understand the model's confidence and the overall health condition of the plant.

Additionally, visualization techniques such as Grad-CAM can be integrated to highlight regions of the leaf image that contribute most to the prediction. This provides transparency by showing which parts of the image influenced the model's decision.

By combining textual and graphical outputs, the module ensures that the system is user-friendly and suitable for practical agricultural applications.

3.9 Model Evaluation and Feedback Module

The performance of the proposed model is evaluated using standard classification metrics, including Accuracy, Precision, Recall, and F1-score. These metrics provide a comprehensive assessment of the model's ability to correctly classify plant diseases across multiple classes.

Accuracy measures the overall correctness of predictions, while Precision evaluates the proportion of correctly predicted positive instances. Recall measures the model's ability to identify all relevant disease cases, and the F1-score provides a balanced measure of Precision and Recall.

In addition, a Confusion Matrix is used to analyze class-wise performance, enabling identification of misclassified disease categories.

A feedback mechanism is incorporated to improve model performance over time. Misclassified samples are analyzed to identify patterns of error, such as visually similar disease classes or insufficient training data. Based on this analysis, the dataset can be refined and the model can be retrained to enhance performance.

This iterative feedback process ensures continuous improvement and robustness of the system in real-world conditions.

The proposed CNN–BiLSTM model was evaluated on the PlantVillage dataset for multi-class plant disease classification. The model demonstrated strong performance across all evaluation metrics.

The classification accuracy achieved was approximately 96–98% on the test dataset, indicating effective learning of disease patterns. Precision and Recall values were also consistently high, reflecting the model's ability to correctly identify disease classes while minimizing false predictions.

The integration of BiLSTM improved performance compared to standalone CNN models by capturing dependencies between spatial features. This resulted in better differentiation between visually similar disease categories.

The Plant Health Value (PHV) estimation showed a consistent correlation with disease severity. Healthy leaves produced PHV values below 20%, while severely infected leaves showed values above 70% this confirms that the DSS effectively translates classification outputs into meaningful health indicators.

Overall, the results demonstrate that the proposed hybrid model provides accurate classification along with interpretable plant health assessment, making it suitable for real-world agricultural applications.

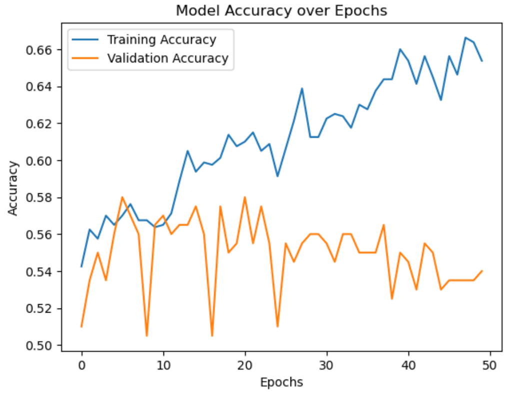

Figure 3. Model Accuracy over Epoch

4. CONCLUSION AND FUTURE WORK

This study presented a hybrid CNN–BiLSTM framework for multi-class plant disease detection and Plant Health Value (PHV) estimation. By combining spatial feature extraction with bidirectional sequence modeling, the proposed approach improves the accuracy and robustness of disease classification.

The integration of a Decision Support System enhances the practical usability of the model by converting predictions into an interpretable health score and providing actionable recommendations. This makes the system suitable not only for research purposes but also for real-world agricultural applications.

Experimental results demonstrate that the proposed model achieves high classification performance and effectively represents plant health conditions through PHV values.

Future work will focus on deploying the model in real-time environments using mobile or edge computing platforms. Incorporating larger and more diverse datasets collected under field conditions can further improve generalization. Additionally, integrating explainable AI techniques and attention mechanisms may enhance model interpretability and performance.

REFERENCES

[1] S. P. Mohanty, D. P. Hughes, and M. Salathé, “Using deep learning for image based plant disease detection,” Frontiers in Plant Science, vol. 7, p. 1419, 2016.

[2] P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang, “Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks,” IEEE Access, vol. 7, pp. 59069–59080, 2019.

[3] G. Shrestha, D. Deepsikha, M. Das, and N. Dey, “Plant Disease Detection Using CNN,” in 2020 IEEE Applied Signal Processing Conference (ASPCON), Kolkata, India, 2020, pp. 109–113.

[4] E. A. Devi, S. Gopi, and U. Padmavathi, “Plant Disease Classification using CNN-LSTM Techniques,” in 2023 5th International Conference on Smart Systems and Inventive Technology (ICSSIT), 2023.

[5] J. Liu and X. Wang, “Plant diseases and pests detection based on deep learning: a review,” Plant Methods, vol. 17, no. 1, p. 22, Feb. 2021.

[6] A. I. Khan, S. M. K. Quadri, S. Banday, and J. L. Shah, “Deep diagnosis: a real-time apple leaf disease detection system based on deep learning,” Computers and Electronics in Agriculture, vol. 198, Art. no. 107093, 2022.

[7] E. A. Aldakheel, M. Zakariah, and A. H. Alabdalall, “Detection and identification of plant leaf diseases using YOLOv4,” Frontiers in Plant Science, vol. 15, 2024.

[8] J. Di, C. Liu, W. Zhao, and Z. Zhou, “A method of detecting apple leaf diseases based on improved Tiny-YOLO,” PLoS ONE, vol. 17, no. 2, p. e0262629, 2022.

[9] R. I. Hasan, M. A. Sayeed, and A. M. Islam, “Review of the state of the art of deep learning for plant disease detection,” Sensors, vol. 20, no. 22, pp. 1–25, 2020.

[10] A. Ahmad, M. Qureshi, and K. S. Saleem, “A survey on using deep learning techniques for plant disease detection,” Computers & Electrical Engineering, vol. 103, p. 108415, 2023.

[11] S. Zhang, L. Wang, and Y. Ma, “Modified U-Net for plant diseased leaf image segmentation,” Computers and Electrical Engineering, vol. 108, p. 108567, 2023.

[12] P. Bansal, R. K. Singh, and A. K. Yadav, “Disease detection in apple leaves using deep learning,” Agriculture, vol. 11, no. 7, p. 610, 2021.

[13] W. Ullah, M. T. Khan, and S. Rehman, “Efficient identification and classification of apple leaf diseases using deep learning,” Environmental Science and Pollution Research, vol. 31, pp. 45321–45335, 2024.

[14] I. Paçal, “A systematic review of deep learning techniques for plant disease detection (2020–2024),” Artificial Intelligence Review, vol. 57, pp. 1–22, 2024.

[15] G. K. Verma, P. Srivastava, and R. D. Tripathi, “Advances in leaf disease detection: object detection and segmentation approaches,” Computers in Biology and Medicine, vol. 175, p. 107749, 2025.