Plant Leaf Disease Detection Using Efficient Net V2-S with Transfer Learning

Anita J. Shinde[1]; Ajay B. Kurhe[2]

[1] Research Scholar, School of Computational Sciences, SRTM University, Nanded, Maharashtra, India
[2] Associate Professor, Shri Guru Buddhiswami Mahavidyalaya, Purna (Jn.), Dist. Parbhani, Maharashtra, India

DOI: https://doi.org/10.51584/IJRIAS.2026.11050015

Received: 22 April 2026; Accepted: 28 April 2026; Published: 22 May 2026

ABSTRACT

Early and accurate detection of plant leaf diseases plays a vital role in improving crop productivity and ensuring sustainable agriculture. This paper presents a deep learning-based framework for multi-class classification of banana leaf diseases using transfer learning. Initially, a baseline model based on ResNet50 is developed to evaluate standard performance. To enhance classification accuracy and computational efficiency, a transfer learning approach employing EfficientNetV2 is proposed. The pretrained EfficientNetV2-S model is fine-tuned by integrating a custom classification head comprising global average pooling, dropout, and fully connected layers.

The proposed model is trained and validated on a dataset containing four classes of banana leaf images, namely Cordana, Healthy, Pestalotiopsis, and Sigatoka. Experimental results demonstrate that the proposed approach achieves an overall accuracy of 95%, along with high precision, recall, and F1-score across all classes. The confusion matrix and training curves further confirm the robustness, stability, and generalization capability of the model. Comparative analysis indicates that the proposed EfficientNetV2-S-based framework outperforms the baseline ResNet50 model while maintaining reduced computational complexity.

To further evaluate practical applicability, the proposed model was tested on real-world banana leaf images captured under natural field conditions. The model achieved a detection accuracy of 76.19%, demonstrating its robustness and ability to generalize effectively beyond controlled datasets.

The results show that the proposed framework provides an efficient and scalable solution for real-world plant disease detection in precision agriculture. Future work will focus on expanding dataset diversity and exploring advanced architectures to further improve classification performance.

Keywords: Banana leaf disease detection, deep learning, transfer learning, EfficientNetV2, ResNet50, precision agriculture, plant disease detection.

INTRODUCTION

Plant diseases pose a significant threat to agricultural productivity and global food security. Among various crops, banana plants are highly susceptible to fungal leaf diseases such as Sigatoka, Cordana, and Pestalotiopsis, which can severely affect yield and fruit quality. Early and accurate detection of such diseases is essential for effective crop management and reducing economic losses [2], [3].

Conventional plant disease detection primarily depends on visual examination by agricultural experts, which is often labour-intensive and subject to human interpretation. However, this approach is time-consuming, requires domain expertise, and may lead to inconsistent results due to subjective judgment, environmental variations, and human error [2], [3]. Moreover, the limited availability of experts in rural areas [2], [3] further restricts timely diagnosis.

Recent advancements in computer vision and deep learning have enabled automated plant disease detection systems, which provide faster and more reliable results [3]. Convolutional Neural Networks (CNNs) have shown remarkable performance in image classification tasks by automatically learning discriminative features such as color, texture, and shape [4]. Deep architectures such as ResNet50 have been widely used for plant disease classification and have achieved promising accuracy [5]. However, these models are computationally intensive and require significant processing resources, limiting their applicability in real-time and resource-constrained environments.

To address these challenges, this study proposes the use of EfficientNetV2-S, a modern and optimized CNN architecture. EfficientNetV2-S is designed to achieve high accuracy with reduced computational cost and faster training through efficient scaling of network depth, width, and resolution [6], [7]. The proposed approach provides an effective and scalable solution for automated banana leaf disease detection, making it suitable for real-world agricultural applications.

Furthermore, unlike many existing studies, this work evaluates the proposed model on real-world banana leaf images to assess real-world applicability. This real-world validation strengthens the robustness and usability of the proposed approach.

LITERATURE REVIEW

V. G. Krishnan et al. [8] proposed an automated segmentation and classification model for banana leaf disease detection. The method utilizes image segmentation to isolate diseased regions, followed by CNN-based classification. By focusing on infected areas rather than the entire leaf, the model improves feature extraction and achieves approximately 94% accuracy. However, segmentation increases computational complexity.

S. Chattopadhyay et al. [9] introduced a federated learning-based CNN model for banana leaf disease detection. This approach enables collaborative model training across multiple distributed clients while preserving data privacy. The model achieved an accuracy of 96–97% and demonstrated robustness under varying environmental conditions. However, it suffers from high communication cost and complex implementation.

V. Chaudhari and M. P. Patil [10] proposed a segmentation and ensemble learning-based approach using genetic algorithms and Local Binary Patterns (LBP) for feature extraction. The model achieved over 92% accuracy; however, handcrafted feature extraction and segmentation increase processing time and reduce scalability.

S. Nassor et al. [11] developed lightweight deep learning models such as MobileNetV2, MobileNetV3, ShuffleNetV2, and SqueezeNet for banana disease detection. Among these, SqueezeNet achieved the highest accuracy of 97.12%, demonstrating that lightweight models can provide high performance with reduced computational cost. However, the study is limited by fewer disease classes and dependency on image quality.

A. Bhargava et al. [1] presented a comprehensive review of plant disease detection techniques using machine learning and deep learning approaches. The study highlights that CNN-based models outperform traditional methods but emphasizes challenges such as dataset limitations, computational complexity, and real-world applicability.

M. Shoaib et al. [12] proposed a deep learning-based segmentation and classification framework using U-Net and InceptionNet architectures. The model achieved up to 99% accuracy on large datasets. However, the approach requires high computational resources and may face challenges such as overfitting and limited real-time deployment.

Bhuiyan et al. [18] proposed BananaSqueezeNet, a lightweight CNN-based model achieving high accuracy in banana leaf disease classification. However, lightweight architectures may limit feature representation capability in complex disease patterns.

In earlier study, Shinde et al. [19] proposed a CNN-driven system for automated plant disease detection using ResNet50, achieving an accuracy of 95.24%. However, the model relied on a conventional CNN architecture

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue V May 2026

with higher computational cost. To address these limitations, the present work introduces an EfficientNetV2-S-based transfer learning framework for improved performance and efficiency.

From the above studies, it is evident that machine learning and deep learning techniques, particularly Convolutional Neural Networks (CNNs), have significantly improved plant leaf disease detection with high accuracy. Approaches such as segmentation, ensemble learning, federated learning, and advanced CNN architectures have demonstrated strong performance. However, these methods still face challenges such as limited dataset diversity, high computational cost, overfitting, poor generalization in real-field conditions, and complex implementation. These limitations highlight the need for a more efficient and scalable solution. Therefore, this work proposes the use of EfficientNetV2-S, an optimized CNN architecture that achieves high accuracy with reduced computational complexity and improved generalization for real-world plant leaf disease detection systems.

METHODOLOGY

Dataset and Preprocessing

The proposed study utilizes the Banana Leaf Spot Diseases (BananaLSD) dataset, which consists of an original set of 937 RGB images and an augmented set of 1600 images distributed across four classes: Healthy, Sigatoka, Cordana, and Pestalotiopsis leaves.

To improve model generalization and simulate real-world variability, various data augmentation techniques are applied, including rotation, horizontal flipping, random cropping, shear transformation, contrast adjustment, translation, and Gaussian blur. Such augmentation strategies have been shown to significantly enhance model robustness and reduce overfitting in plant disease detection tasks [13], [15].

All images are resized to 224 × 224 pixels to match the input requirements of the model. Pixel values are normalized to the range [0,1], and class labels are encoded using one-hot encoding. The combined dataset is divided into 80% training and 20% testing subsets.

In addition to standard dataset evaluation, an independent real-world dataset consisting of field-acquired banana leaf images was used to assess the generalization capability of the proposed model. The real-world dataset primarily includes images of Sigatoka-infected leaves collected under natural field conditions.

Proposed Model Architecture:

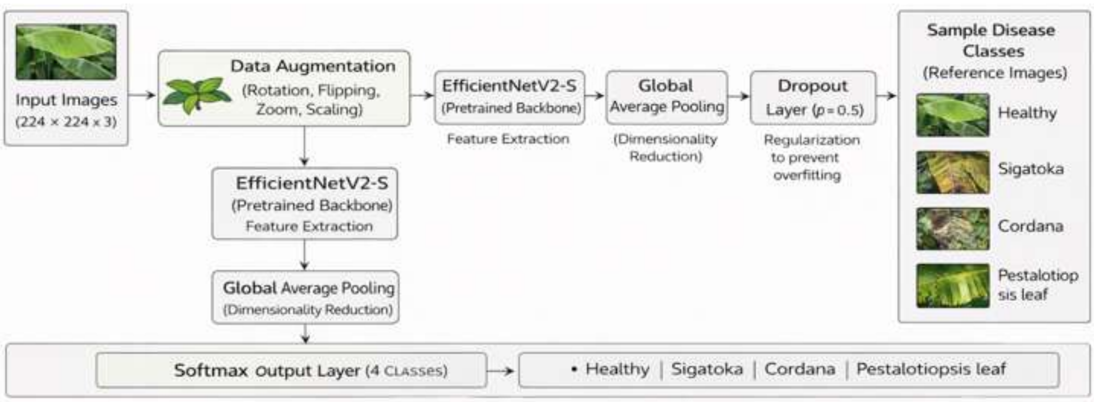

Fig. 1. CNN-Based EfficientNetV2-S Architecture for Banana Leaf Disease Classification

The overall architecture of the proposed CNN-based transfer learning framework is illustrated in Fig. 1. The model consists of four main stages: preprocessing, data augmentation, feature extraction using EfficientNetV2-S, and classification through a custom head.

As shown in Fig. 1, input images are first pre-processed and augmented before being fed into the EfficientNetV2-S model, which serves as the backbone for feature extraction. The extracted features are then passed through a classification head comprising global average pooling, fully connected layers, and regularization components, followed by a SoftMax layer for multi-class classification.

EfficientNet-based architectures have demonstrated superior performance in image classification tasks due to their ability to balance accuracy and computational efficiency [6], [7]. Recent studies have also shown that EfficientNet variants outperform traditional convolutional neural networks such as AlexNet and ResNet in plant disease detection applications [6], [7], [14].

Feature Extraction

The EfficientNetV2-S model is employed as a pretrained backbone to extract hierarchical feature representations from input images. It captures both low-level features (such as edges and textures) and high-level semantic features relevant to disease patterns.

To retain the knowledge learned from large-scale datasets, the initial layers of the EfficientNetV2-S model are frozen during training. This transfer learning approach reduces computational complexity and improves performance, particularly when the available dataset is limited [6], [7].

Classification Head

A custom classification head is added on top of the pretrained backbone, as shown in Fig. 1, to perform task-specific classification. The classification head consists of:

- A Global Average Pooling (GAP) layer for feature reduction
- A fully connected Dense layer with ReLU activation
- A Dropout layer with a rate of 0.5 for regularization
- A SoftMax output layer for four-class classification

This design enhances the model's ability to learn discriminative features while mitigating overfitting [8].

Training Configuration

The model is trained using the Adam optimizer with a learning rate of 0.001. Categorical cross-entropy is employed as the loss function due to the multi-class nature of the problem. The model is trained with a batch size of 32 for 10–30 epochs.

During training, data augmentation is applied dynamically to improve generalization. Only the parameters of the classification head are updated, while the pretrained EfficientNetV2-S backbone remains frozen. Dropout is used as a regularization technique to prevent overfitting [16], [17].

Model Training Procedure

The training process follows a standard deep learning workflow:

1. Pre-processed and augmented images are fed into the model in mini-batches.
2. A forward pass is performed through the EfficientNetV2-S backbone and classification head.
3. The categorical cross-entropy loss is computed.
4. Backpropagation is applied to compute gradients.
5. The Adam optimizer updates the trainable parameters.

To further enhance training performance, early stopping and model checkpointing techniques can be employed.

Evaluation Metrics

The performance of the proposed model is evaluated using standard classification metrics, including accuracy, precision, recall, and F1-score. A confusion matrix is also used to analyze class-wise performance and identify misclassification patterns.

RESULTS AND DISCUSSION

Classification Performance Analysis

The performance of the proposed EfficientNetV2-S model is evaluated on the test dataset using standard classification metrics, including precision, recall, and F1-score. These metrics provide a comprehensive assessment of the model's ability to accurately classify banana leaf diseases across multiple categories.

The detailed class-wise performance is presented in Table I.

Class	Precision	Recall	F1-score	Support
Cordana	0.93	0.96	0.94	162
Healthy	0.98	0.96	0.97	129
Pestalotiopsis	0.87	0.96	0.91	173
Sigatoka	0.98	0.94	0.96	473

Table I: Class-wise Performance of the Proposed Model

As shown in Table I, the proposed model achieves high classification performance across all classes, with F1-scores exceeding 0.90. The Healthy and Sigatoka classes exhibit the highest precision values, indicating accurate predictions with minimal false positives. The recall values are consistently high for all classes, demonstrating the model's effectiveness in correctly identifying true disease instances.

The relatively lower precision observed in the Pestalotiopsis class suggests minor misclassification, which may be attributed to visual similarity between disease patterns. Nevertheless, the overall results confirm the strong discriminative capability of the proposed model.

Overall Performance Evaluation

The overall performance of the proposed model is summarized in Table II.

SR. No.	Metric	Value
1	Accuracy	0.9500
2	Precision (Weighted)	0.9500
3	Recall (Weighted)	0.9500
4	Macro F1-score	0.9500
5	Weighted F1-score	0.9500

Table II: Overall Performance Metrics of the Proposed EfficientNetV2-S Model

As presented in Table II, the proposed model achieves an overall accuracy of 95%, indicating strong classification performance. The weighted precision and recall values further confirm that the model produces reliable predictions while effectively identifying most true instances.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue V May 2026

The macro and weighted F1-scores are identical (0.95), indicating balanced performance across all classes despite class imbalance. This suggests that the model generalizes well and is not significantly biased toward majority classes.

Real-World Performance Evaluation

The real-world dataset used for evaluation consists only of Sigatoka leaf images, as this disease is more commonly observed under field conditions. To validate the model in real-world conditions, field-acquired Sigatoka leaf images were used. Unlike controlled datasets, these images contain environmental variations such as inconsistent lighting, complex backgrounds, and natural leaf orientations. The proposed model achieved a detection accuracy of 76.19%, demonstrating its robustness and practical applicability.

Dataset Type	Accuracy
Standard Dataset	95.00
Real-World Dataset (Sigatoka)	76.19

Table III: Performance Comparison on Standard and Real-World Datasets

The model achieves high accuracy on the benchmark dataset, while maintaining strong performance on real-world images with an accuracy of 76.19%. The observed performance drop is expected due to variations in real-world conditions such as lighting, background complexity, and leaf orientation. However, the model demonstrates good generalization capability, making it suitable for practical agricultural applications.

Training Performance Analysis

The training progress of the model across epochs is shown in Table IV.

Epoch	Training Accuracy	Validation Accuracy
1	0.45	0.68
10	0.90	0.90
20	0.94	0.93
30	0.97	0.95

Table IV: Training and Validation Accuracy Across Epochs

As shown in Table IV, both training and validation accuracy improve steadily with increasing epochs, indicating effective learning and convergence of the model.

Accuracy and Loss Analysis

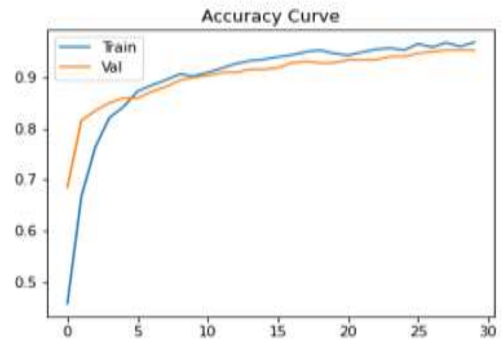

Fig. 2. Training and Validation Accuracy Curve

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue V May 2026

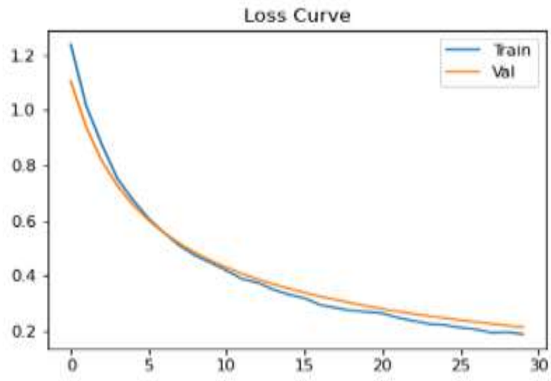

Fig. 3. Training and Validation Loss Curve

The accuracy curve in Fig. 2 shows a steady increase in both training and validation accuracy throughout the training process. The model improves from an initial accuracy of approximately 45% to 97% for training and 95% for validation. The close alignment between the curves indicates minimal overfitting and strong generalization capability.

The loss curve in Fig. 3 demonstrates a consistent decrease in both training and validation loss over epochs, indicating successful model convergence. The similarity between training and validation loss curves suggests stable learning without significant overfitting.

Confusion Matrix Analysis

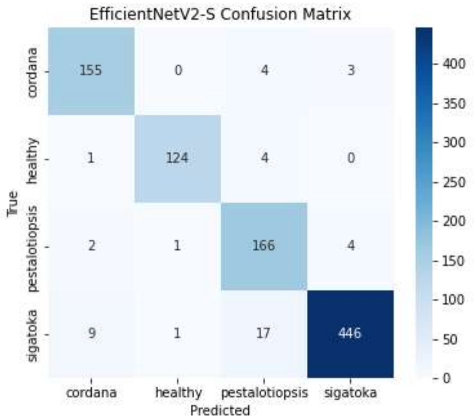

Fig. 4. Confusion Matrix of the Proposed EfficientNetV2-S Model

The confusion matrix shown in Fig. 4 provides a detailed visualization of classification performance across all classes. Most predictions are concentrated along the diagonal, indicating correct classification.

Minor misclassifications are observed between visually similar classes such as Pestalotiopsis and Sigatoka. However, these errors are limited and do not significantly impact the overall performance of the model.

Grad-CAM Visualization Analysis

To further interpret the model's decision-making process, Gradient-weighted Class Activation Mapping (Grad-CAM) is employed to visualize the regions of banana leaves that contribute most to the classification. Grad-CAM is used for model interpretability, enabling better understanding of the model's predictions.

Fig. 5 illustrates Grad-CAM visualizations for Sigatoka leaf samples under real-world conditions. The highlighted regions (shown in red and yellow) indicate areas of high importance, corresponding to disease-affected portions of the leaf.

INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN APPLIED SCIENCE (IJRIAS)
ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue V May 2026

The results demonstrate that the proposed EfficientNetV2-S model effectively focuses on infected regions, even in the presence of complex backgrounds and varying environmental conditions. This confirms the model's ability to learn meaningful and discriminative features relevant to banana leaf disease detection.

(a) **Sample 1** **(b) Sample 2** **(c) Sample 3**

Fig. 5. Grad-CAM visualizations for Sigatoka leaf samples: (a) Sample 1, (b) Sample 2, and (c) Sample 3. The highlighted regions indicate disease-affected areas identified by the model.

DISCUSSION

The experimental results demonstrate that the proposed EfficientNetV2-S model achieves high accuracy and reliable performance for banana leaf disease classification. The class-wise evaluation indicates strong performance across all categories, with F1-scores exceeding 0.90, reflecting a balanced trade-off between precision and recall.

The agreement between accuracy, F1-score, and confusion matrix results confirms the robustness and consistency of the model. Additionally, the close alignment between training and validation curves indicates that the model effectively avoids overfitting and generalizes well to unseen data.

Overall, the proposed approach proves to be effective for plant leaf disease detection, even in the presence of class imbalance and visually similar disease patterns.

A performance drop is observed when moving from controlled datasets to real-world conditions, which is expected due to domain differences. However, the model maintains satisfactory accuracy, indicating its ability to generalize effectively. The inclusion of real-world samples during training contributed to improved robustness.

CONCLUSION

This paper presents a CNN-based transfer learning approach for banana leaf disease classification using the EfficientNetV2-S architecture. By leveraging pretrained ImageNet weights and a custom classification head, the proposed model effectively extracts relevant features and accurately classifies leaf images into four categories: Healthy, Sigatoka, Cordana, and Pestalotiopsis. The incorporation of data augmentation techniques further improves model robustness and generalization.

The experimental results demonstrate that the proposed model achieves an overall accuracy of 95%, along with high precision, recall, and F1-scores across all classes. The class-wise analysis and confusion matrix indicate strong discriminative capability with minimal misclassification. Additionally, the close alignment between training and validation performance confirms that the model generalizes well and exhibits limited overfitting. The integration of real-world validation and interpretability through Grad-CAM distinguishes the proposed approach from existing methods.

Despite these promising outcomes, certain limitations exist. The dataset size is relatively limited, and class imbalance may influence performance. Moreover, visually similar disease patterns can occasionally lead to minor misclassification between specific classes. Additionally, real-world evaluation was conducted only on Sigatoka class, which may limit the assessment of model generalization across all disease categories.

Future work can focus on expanding the dataset with more diverse real-world samples and exploring advanced techniques such as fine-tuning, ensemble learning, and attention mechanisms to further enhance performance. Additionally, the proposed model can be extended to real-time deployment in mobile or web-based applications, providing practical support for early disease detection and precision agriculture.

Overall, the proposed EfficientNetV2-S-based approach offers an effective and reliable solution for automated banana leaf disease classification and demonstrates strong potential for real-world agricultural applications.

REFERENCES

1. A. Bhargava, A. Shukla, O. P. Goswami, M. H. Alsharif, P. Uthansakul, and M. Uthansakul, (2024). Plant leaf disease detection, classification, and diagnosis using computer vision and artificial intelligence: A review. IEEE Access, 12,1 –1. https://doi.org/10.1109/ACCESS.2024.3373001
2. Ferentinos, K.P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145,311–318. https://doi.org/10.1016/j.compag.2018.01.009
3. S. Mohanty, D. P. Hughes, and M. Salathé, “Using deep learning for image-based plant disease detection,” Frontiers in Plant Science, vol. 7, 2016.
4. Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, no. 7553, pp. 436–444, 2015.
5. K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. IEEE CVPR, 2016, DOI: 10.1109/CVPR.2016.90
6. M. Tan and Q. Le, “EfficientNetV2: Smaller models and faster training,” in Proc. ICML, 2021.
7. M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in Proc. ICML, 2019. DOI:10.48550/arXiv.1905.11946
8. V. G. Krishnan et al., “An automated segmentation and classification model for banana leaf disease detection,” Journal of Applied Biology and Biotechnology, vol. 10, no. 1, pp. 213–220, 2022.
9. S. Chattopadhyay, M. Manwal, V. Kukreja, and S. Mehta, “Transforming Agro-Diagnostics: Banana Leaf Diseases Through Federated Learning CNNs,” in Proc. IEEE INOCON, 2024.
10. V. Chaudhari and M. P. Patil, “Detection and classification of banana leaf disease using segmentation and ensemble learning approach,” Applied Computer Systems, vol. 28, 2023.
11. S Nassor, M Mushthofa and K Priandana, “Deep learning model for detection and classification of banana diseases based on leaf images,” IOP Conference Series: Earth and Environmental Science, vol. 1359, 2024.
12. M. Shoaib, T. Hussain, B. Shah, I. Ullah, S. M. Shah, F. Ali, and S. H. Park, “Deep learning-based segmentation and classification of leaf images for detection of tomato plant disease,” Frontiers in Plant Science, 2022.Frontiers in Plant Science, vol. 13, 2022.
13. I. Paçal et al., “A systematic review of deep learning techniques for plant disease detection,” Artificial Intelligence Review, vol. 57, 2024.
14. U. Mishra, A. Pandey, L.G., and T. K., “Deep learning-based disease detection in potato and mango leaves: A comparative study of CNN, AlexNet, ResNet, and EfficientNet,” Scientific Reports, vol. 16, Art. no. 2788, 2026.doi: 10.1038/s41598-025-32607-5.
15. P. Sahu, S. Dash, and S. K. Mohapatra, “Machine learning and deep learning approaches for plant disease detection: A review,” IEEE Access, vol. 10, pp. 36186–36211, 2022.
16. N. Srivastava et al., “Dropout: A simple way to prevent neural networks from overfitting,” Journal of Machine Learning Research, 2014.
17. D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” Proc. ICLR, 2015.
18. Md. A. B. Bhuiyan et al., "Banana SqueezeNet: Automated Diagnosis of Banana Leaf Diseases Using Deep Learning and Bayesian Optimization," IEEE Trans. Image Process., vol. 25, no. 6, pp. 2657–2669, Jun. 2023.
19. A. J. Shinde, B. S. Khurana, and A. Kurhe, “CNN-Driven Smart System for Automated Plant Disease Detection and Classification,” IEEE, 2025.