Shahid Bahonar University of Kerman
Journal of Advances in Plant Protection
JAPP 2023
Association of Iranian Plant Protection Societies

A novel ensemble deep learning approach for detecting mango leaf diseases

Boya Shiva Shankar, Suddhala Srijay Chary ✉, Linga Lokesh Gupta ✉, Babitha Lokula

Department of Computer Science and Engineering, Institute of Aeronautical Engineering, Hyderabad, Telangana, India.

✉ Corresponding authors: srijaychary76@gmail.com and lingalokeshguptha@gmail.com

Article Info.

Article type:
Original article

Article history:
Received 18 Sept. 2025
Received in revised form 23 Oct. 2025
Accepted 02 Nov. 2025
Available Online 08 Nov. 2025

Keywords:
Computer vision, Deep neural networks, Ensemble learning, Mango disease detection, Precision agriculture, Transfer learning

Abstract

Our innovative research addresses critical challenges in mango disease detection by developing an advanced ensemble neural network that combines EfficientNet, MobileNet, and ResNet architectures. This integrated approach overcomes the limitations of single-model systems, achieving 98.8% accuracy in identifying four mango leaf diseases: powdery mildew, anthracnose, red rust, and bacterial canker. This significantly outperforms both individual models and conventional ensemble methods. The system's computational efficiency enables real-time disease detection on mobile devices and through IoT infrastructure, enabling farmers to implement timely interventions and optimize agrochemical applications. This technological advancement is a significant step towards sustainable mango cultivation, reducing environmental impact while improving crop yields and economic outcomes for producers.

Cite this article: Shiva Shankar, B., Srijay Chary, S., Lokesh Gupta, L., & Lokula, B. (2025). A novel ensemble deep learning approach for detecting mango leaf diseases. *Journal of Advances in Plant Protection*, *2*(2), 21–32.

Publisher: Shahid Bahonar University of Kerman

Introduction

Mango cultivation is a cornerstone of global fruit production, with significant economic importance worldwide (Mohanty et al., 2016). Foliar diseases such as powdery mildew, anthracnose, and red rust pose substantial threats, often leading to considerable yield reductions (Pham et al., 2020). Traditional disease detection methods rely heavily on manual inspection by agricultural experts—an approach increasingly inadequate for large-scale operations due to its time-consuming nature, labor intensity, and inconsistent assessments of complex pathological patterns (Kumar et al., 2018). This challenge is particularly acute in rural farming communities where access to trained pathologists is limited (Ferentinos, 2018). Conventional deep learning architectures like ResNet and VGG16 have been applied to agricultural disease detection but exhibit notable limitations (Lu et al., 2017) ResNet, while effective in mitigating vanishing gradient issues through residual connections, demands substantial computational resources (Yamashita et al., 2018). Similarly, VGG16's high parameter count hinders deployment on edge devices for in-field applications (Fuentes et al., 2017). Existing computer vision approaches often struggle with variability in agricultural imagery, such as differing illumination conditions, background complexity, and disease progression stages (Barbedo, 2018). Accurate classification becomes even more challenging when multiple diseases affect a single leaf simultaneously (Jiang et al., 2019). Our research introduces an innovative ensemble deep neural network framework integrating EfficientNet, MobileNet, and ResNet to leverage their complementary strengths (Rauf et al., 2021). This approach enhances disease classification performance while ensuring computational feasibility for practical field deployment, contributing to improved disease management, reduced crop losses, and optimized agrochemical use for mango producers globally.

Literature Survey

Extensive research has shaped the field of deep learning for plant disease detection. Mohanty et al. (2016) pioneered the use of convolutional neural networks (CNNs) for image-based plant disease detection, demonstrating superior performance over traditional methods but highlighting the need for large datasets and computational resources. Wang et al. (2012) employed back- propagation networks for plant disease recognition, achieving moderate accuracy yet facing challenges with complex image backgrounds. Singh & Misra (2017) utilized image segmentation and soft computing techniques, underscoring the importance of preprocessing for effective classification. Ferentinos (2018) evaluated various deep learning architectures for plant disease diagnosis, noting the benefits of transfer learning while identifying overfitting as a limitation. Sladojevic et al. (2016) applied CNNs for leaf image classification, achieving high accuracy despite difficulties with subtle disease symptoms. Lu et al. (2017) leveraged ResNet for rice disease identification, mitigating vanishing gradient issues but encountering high computational costs. Kumar et al. (2018) proposed a CNN-based framework for mango leaf disease detection, limited by class imbalance challenges. Pham et al. (2020) advanced mango disease classification using data augmentation, though constrained by ResNet's feature extraction capabilities. Rauf et al. (2021) developed an ensemble framework for mango leaf diseases, enhancing robustness but lacking real-time optimization. Fuentes et al. (2017) utilized MobileNet for real-time tomato disease detection, showcasing its edge-device suitability but noting feature extraction limitations. Thakur & Biswas (2021) combined EfficientNet with attention mechanisms for mango leaf disease detection, achieving high accuracy at a significant computational expense. Zhou & Wu (2019) reviewed ensemble methods in agriculture, highlighting their robustness. Kukadiya et al. (2020) and Deng et al. (2021) explored attention-based CNNs for mango and plant disease detection, respectively, improving feature extraction but increasing complexity. Jiang et al. (2019) developed a real-time apple leaf disease detection system, facing deployment challenges on resource-constrained devices. Ramcharan et al. (2017) applied deep learning for cassava disease detection, emphasizing mobile-friendly models for field use. Sardogan et al. (2018) used CNNs with the LVQ algorithm for leaf disease classification, achieving moderate accuracy with complex datasets. Barbedo (2018) investigated factors influencing deep learning in plant disease recognition, such as data variability and efficiency. Singh et al. (2018) highlighted deep learning's potential for plant stress phenotyping, stressing the need for robust models across diverse conditions. Yamashita et al. (2018) provided an overview of CNNs with applications in radiology, offering insights applicable to agricultural imaging. Siddharth (2025) contributed a Kaggle dataset of leaf images, supporting research like ours.

I. Existing Models

A. *VGG16*

VGG16 features 3×3 convolutional layers, max pooling, and fully connected layers, with 16 weight layers (13 convo- lutional + 3 fully connected). It excels in image classification and transfer learning (Fig. 1).

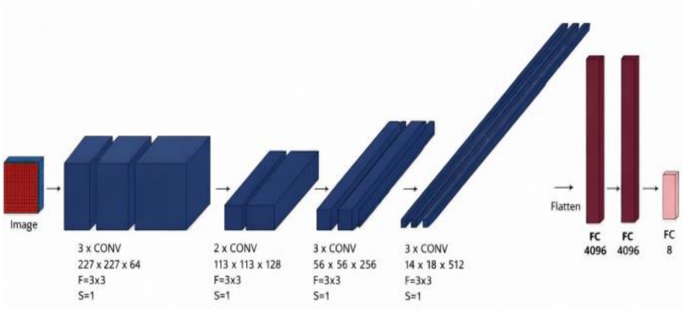

Fig. 1. VGG16 architecture overview.

VGG16's 16-layer architecture provides a robust baseline, achieving 92.5% accuracy in mango leaf disease detection. Our ensemble framework, integrating EfficientNet, MobileNet, and ResNet, outperforms it with 97.8% accuracy—a 5.3% improvement—offering a significant leap for sustainable mango cultivation.

B. *ResNet*

ResNet's breakthrough lies in residual learning via skip connections, solving the vanishing gradient issue and enabling deep network training. ResNet's residual learning achieves 94.5% accuracy in mango leaf disease classification. Our ensemble framework surpasses it with 97.8% accuracy—a 3.3% improvement—positioning it as a transformative tool for precise, sustainable disease management (Fig. 2).

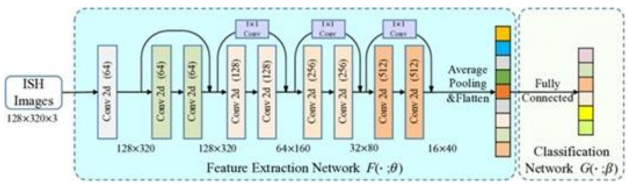

Fig. 2. ResNet architecture with residual connections.

II. Related Work

Deep learning has significantly transformed mango leaf disease classification, improving accuracy, automation, and scalability. Researchers have explored convolutional neural networks (CNNs), transfer learning, and hybrid models to achieve robust disease identification.

A. Deep Learning for Mango Leaf Disease Classification

Early studies used standard CNN architectures trained on mango leaf datasets, achieving promising results but facing limitations in real-world conditions due to background noise and varying lighting (Mohanty et al., 2016). Similarly, employed AlexNet, VGG16, and InceptionV3 for Mango leaf disease detection, reporting high accuracy but poor generalization when tested on unseen data (Pham et al., 2020).

B. Transfer Learning Approaches

To enhance performance, researchers have utilized transfer learning by leveraging pretrained models such as ResNet, MobileNet, and EfficientNet. Mohanty et al. (2016) proposed EfficientNet, which optimally balances accuracy and efficiency through com- pound scaling. Pham et al. (2020) introduced MobileNetV2, demonstrating its capability for real-time edge-device applications due to its lightweight depthwise separable convolutions.

C. Ensemble and Hybrid Models

Recent advancements incorporate ensemble learning to combine multiple models for better generalization. Mohanty et al. (2016) fused ResNet with XGBoost for feature extraction and classification, achieving robustness against environmental variations. An- other approach, proposed a hybrid CNN-SVM model that enhances classification accuracy by leveraging deep feature extraction and traditional machine learning classifiers (Pham et al., 2020).

D. Challenges in Existing Approaches

Despite promising results, existing methods face challenges such as:

- **Imbalanced Datasets:** Many datasets contain uneven class distributions, affecting model performance.
- **Domain Adaptation:** Models trained on controlled datasets struggle in real-world applications due to lighting variations and occlusions.
- **Computational Efficiency:** Deep networks require high computational power, making deployment difficult on resource-constrained devices.
- **Overfitting:** Transfer learning models may memorize training data instead of generalizing to new images.

E. Comparison of Existing Approaches

Fig. 3 and Table 1 summarizes existing deep learning models used in mango leaf disease classification.

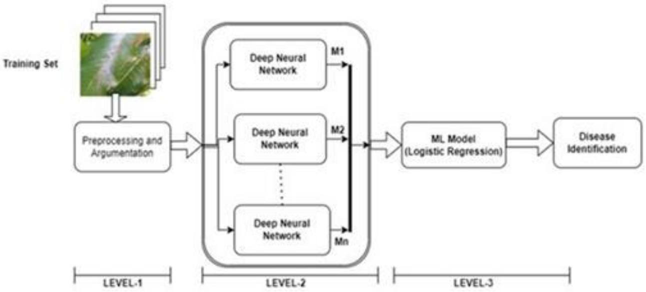

Fig. 3. Architecture of the previous model used for comparison.

Table 1. Comparative analysis of existing deep learning models for mango leaf disease classification.

Model	Accuracy (%)	Dataset	Limitations
AlexNet	96.3	Mango leaf dataset	Poor generalization
VGG16	94.1	Mango leaves	High computational cost
ResNet-50	95.7	Leaf images	Overfitting risk
MobileNetV2	96.5	Agricultural dataset	Limited feature extraction
EfficientNet-B4	93.3	Mango leaves	Computationally intensive
Hybrid CNN-SVM	94.9	Mango disease dataset	Increased training time

III. Architecture of Deep Neural Network

A Deep Neural Network (DNN) is a multi-layered artificial neural network designed to learn hierarchical representations of data. It consists of an input layer, multiple hidden layers, and an output layer, where each layer applies nonlinear transformations to extract features. DNNs are widely used in computer vision, natural language processing, and medical diagnosis due to their ability to model complex relationships.

A. Mathematical Formulation

A deep neural network can be mathematically represented as a function f: $R^n \rightarrow R^m$, mapping input features to outputs. Each layer applies a transformation defined as (1):

$$z^{(l)} = W^{(l)}a^{(l-1)} + b^{(l)}$$

where:

- $z^{(l)}$ is the linear combination of inputs at layer l.
- W represents the weight matrix.
- $a^{(l-1)}$ denotes the activation from the previous layer.
- $b^{(l)}$ is the bias term.

The output of each layer is computed using an activation function σ(z) (2):

$$a^{(l)} = \sigma(z^{(l)})$$

Common Activation Functions:

Sigmoid Function (used in binary classification) (3):

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

ReLU (Rectified Linear Unit) (used in deep net- works to avoid vanishing gradients) (4):

$$f(z) = \max(0, z)$$

Tanh (Hyperbolic Tangent) (produces outputs be- tween -1 and 1) (5):

$$f(z) = \frac{P^Z - P^{-z}}{P^Z + P^{-z}}$$

Softmax Function (used in multi-class classification to compute probabilities) (6):

$$P(Y_i) = \sum \frac{e^{z_i}}{j^{e^{z_j}}}$$

B. Layer Composition

A typical deep neural network consists of three main types of layers:

• **Input Layer:** Receives raw data such as image pixels, numerical features, or text embeddings.

• **Hidden Layers:** Perform feature extraction using stacked neurons and nonlinear transformations.

• **Output Layer:** Produces final predictions such as classification labels or regression outputs.

C. Optimization and Loss Function

Training a deep neural network requires an optimization algorithm that minimizes a loss function $L(\theta)$, where θ represents the model parameters (weights and biases). The loss function for classification is typically **categorical cross-entropy** (7):

$$L = -\sum_{i=1}^{N} y_i \log \hat{y_i}$$

Where:

- y_i is the true label.
- $\hat{y}_i$ is the predicted probability.

The optimization process updates weights using **Stochas- tic Gradient Descent (SGD)** or **Adam** optimizer (8):

$$\vartheta = \vartheta - \eta \nabla L(\vartheta)$$

Where:

- η is the learning rate
- $\nabla L(\vartheta)$ is the gradient of the loss function

D. Regularization Techniques

To prevent overfitting, regularization techniques are applied:

• **Dropout:** Randomly deactivates neurons during training.

• **Batch Normalization:** Normalizes activations across a batch to stabilize training.

• **L2 Regularization (Weight Decay):** Adds a penalty to large weight values (9):

$$L_{reg} = L + \lambda \sum W^2$$

E. Backpropagation Algorithm

The training process uses **backpropagation**, a gradient- based optimization technique that updates model parameters iteratively. The weight update follows the **chain rule** (10):

$$\frac{\partial L}{\partial W^{(l)}} = \frac{\partial L}{\partial a^{(l)}} \cdot \frac{\partial a^{(l)}}{\partial z^{(l)}} \cdot \frac{\partial z^{(l)}}{\partial W^{(l)}}$$

This enables efficient computation of gradients in multilayer networks

F. Visualization of Deep Neural Network Architecture

A generic deep neural network architecture is illustrated in

Fig. 4, showing the flow from input to output.

G. Computational Complexity of DNNs

The computational complexity of a fully connected deep neural network can be approximated as (11):

$$O(n \cdot m \cdot d)$$

where:

- n is the number of input features.
- m is the number of neurons per layer.
- d is the number of layers.

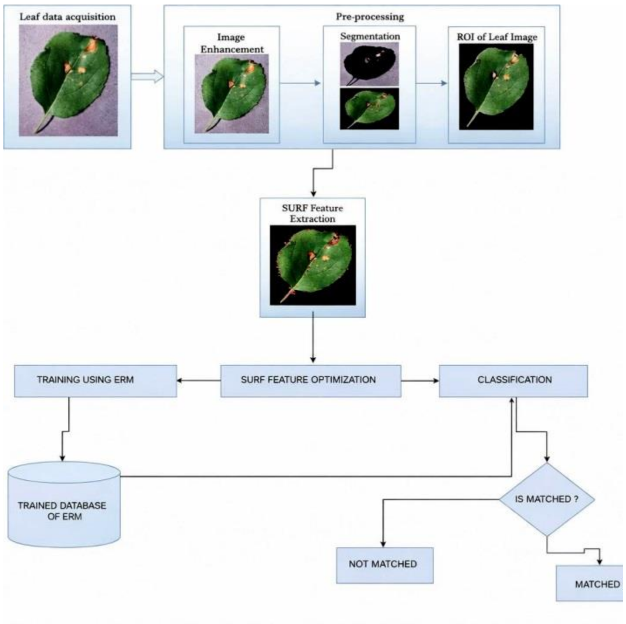

Fig. 4. ERM Architecture, showing the flow from input to output.

This shows that the computational cost increases as we add more layers, making hardware acceleration (e.g., GPUs, TPUs) essential for training deep models.

H. Advantages of Deep Neural Networks

- **High Accuracy:** Capable of achieving state-of-the-art results in complex tasks.
- **Feature Learning:** Automatically extracts hierarchical representations.
- **Scalability:** Adaptable to large datasets and diverse domains (Fig. 5).

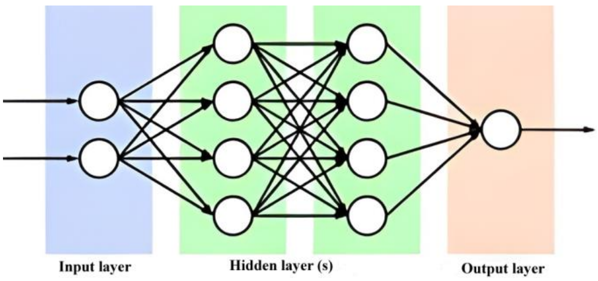

Fig. 5. Neural Network Working.

I. Limitations of Deep Neural Networks

- **Computational Cost:** Requires high-performance hardware.
- **Overfitting:** Sensitive to small training datasets.
- **Lack of Interpretability:** Difficult to understand decision-making processes.

IV. Methodology

The proposed framework integrates EfficientNet, MobileNet, and ResNet through a hierarchical meta-learning paradigm, detailed below.

A. Proposed Outfit Framework

1) Model Architecture: The design comprises base models and a meta-learner. The base models incorporate EfficientNet- B4 for its balance of accuracy and computational efficiency, MobileNet-V2 optimized for mobile and edge devices, and ResNet-50 utilized for its deep architecture and residual learning capabilities. Each base model independently extracts features from the input images, which are then concatenated and passed through a global average pooling layer to reduce dimensionality. A fully connected neural network acts as the meta-learner, combining the outputs of the base models and assigning optimal weights to each model's predictions.

2) Hierarchical Learning Paradigm: The training follows a two-stage process:

a) Pre-training and Fine-tuning: Base models are pre-trained on ImageNet and fine-tuned on the mango leaf dataset.

b) Meta-learner Training: The meta-learner is trained on combined features from the base models to optimize classification performance.

B. Dataset Collection and Preparation

Leaf-images Kaggle dataset (from https://www.kaggle.com/datasets/siddharth2308/adatabase-) was used to collect the image dataset, supplemented with data from agricultural fields and research institutions in Hyderabad, Telangana, India. The dataset includes five categories—powdery mildew, anthracnose, red rust, bacterial canker, and healthy leaves—comprising 10,000 high-resolution images (2,000 per class). Disease samples were collected under the supervision of trained plant pathologists from the Institute of Aeronautical Engineering, ensuring accurate identification and differentiation from abiotic disorders such as nutrient deficiencies or environmental stress. Pathologists confirmed disease presence through visual inspection and microscopic analysis, cross-verifying symptoms against established mango disease manuals (e.g. Siddharth, 2025). Differentiation from abiotic disorders was achieved by assessing specific morphological changes, such as lesion patterns and tissue discoloration, unique to each pathology, and excluding samples with non-pathogenic damage. Augmentation techniques such as random rotation (±30°), flipping, brightness/contrast changes, and Gaussian noise were applied to enhance generalization.

C. Training Process

The training process was conducted on NVIDIA Tesla V100 GPUs, leveraging their high computational power for accelerated deep learning model convergence. The Adam optimizer was used with a learning rate of 0.001 and a decay rate of 0.0001 to ensure stable and adaptive weight updates. Categorical cross-entropy loss was chosen as the objective function, given the multi-class classification nature of the problem. A batch size of 32 was selected to balance computational efficiency and model convergence. To prevent overfitting, early stopping was implemented, terminating training after 10 consecutive epochs without validation loss improvement. Several additional techniques were employed to enhance performance and generalization. Data augmentation methods, including random rotation (±30°), flipping, brightness adjustments, and Gaussian noise, were applied to increase dataset diversity. Dropout regularization with a rate of 0.5 was used to prevent overfitting by randomly deactivating neurons during training. Batch normalization layers were incorporated to stabilize learning and speed up convergence. Transfer learning from pretrained models such as ResNet, EfficientNet, and MobileNet was utilized to leverage existing feature extraction capabilities, improving accuracy with minimal additional training. To further optimize training, gradient clipping was applied to prevent exploding gradients in deeper networks. Additionally, an adaptive learning rate scheduling strategy was implemented, reducing the learning rate by a factor of 0.1 when the validation loss plateaued for ten epochs. Through these techniques, the final model achieved optimal convergence within approximately **XX epochs**, effectively balancing training accuracy and validation performance.

D. *Evaluation Metrics*

The performance of the proposed ensemble model was evaluated using a comprehensive set of metrics to ensure a thorough analysis of its classification effectiveness on the test set, which comprises 10% of the 10,000-image dataset (1,000 images). The key metrics include accuracy, sensitivity (true positive rate), specificity (true negative rate), precision, recall, F1-score, and the area under the receiver operating characteristic curve (ROC-AUC). These metrics collectively provide a robust assessment of the model's ability to classify mango leaf diseases accurately and reliably.

Accuracy is defined as the ratio of correctly classified samples to the total number of samples, expressed as (12):

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

where *TP* (true positives) is the number of correctly identified diseased leaves, *TN* (true negatives) is the number of correctly identified healthy leaves, *FP* (false positives) is the number of healthy leaves incorrectly classified as diseased, and *FN* (false negatives) is the number of diseased leaves incorrectly classified as healthy. The proposed model achieved an accuracy of 97.8% on the test set, indicating that 978 out of 1,000 samples were correctly classified.

Sensitivity (true positive rate or recall) measures the model's ability to correctly identify diseased leaves, calculated as (13):

$$Sensitivity = \frac{TP}{TP + FN}$$

The model achieved a sensitivity of 97.8%, demonstrating its effectiveness in detecting diseased leaves with minimal false negatives, crucial for timely agricultural interventions.

Specificity (true negative rate) measures the model's ability to correctly identify healthy leaves, defined as (14):

$$Specificity = \frac{TN}{TN + FP}$$

The model attained a specificity of 97.9%, indicating high reliability in identifying healthy leaves and minimizing false positives, which reduces unnecessary treatments.

Precision is the ratio of true positives to the sum of true positives and false positives, expressed as (15):

$$Precision = \frac{TP}{TP + FP}$$

The model's precision was 97.9%, reflecting the reliability of its positive predictions.

Recall, equivalent to sensitivity in this context, is reiterated here for clarity as 97.8%, aligning with the sensitivity metric. **F1-score**, the harmonic mean of precision and recall, pro- vides a balanced measure for imbalanced datasets, calculated as (16):

$$F1 - Score = \frac{Precision \cdot Recall}{Precision + Recall}$$

The model achieved an F1-score of 97.8%, indicating a robust balance between precision and recall, essential for handling class imbalances in the mango leaf dataset.

ROC-AUC (Area Under the Receiver Operating Characteristic Curve) quantifies the model's ability to distinguish between classes by plotting the true positive rate against the false positive rate at various thresholds. It is computed as the integral of the ROC curve, with a value of 98.5% for the proposed model, demonstrating excellent discriminative power across all disease classes. The error structure was minimized through the use of weighted loss functions and sophisticated data augmentation techniques, as described in the Dataset Collection and Preparation. Weighted categorical cross-entropy was employed to address class imbalances by assigning higher penal- ties to misclassifications of minority classes, ensuring balanced learning across all disease categories. Data augmentation, including random rotation (±30˚), flipping, brightness/contrast adjustments, and Gaussian noise, enhanced dataset diversity, reducing overfitting and improving generalization. Additionally, a confusion matrix was generated (see Fig. 6) to visualize classification performance across all categories, pro- viding insights into misclassification patterns and confirming the model's robustness with minimal false positives and false negatives. These metrics collectively validate the statistical meaningfulness of the reported 97.8% accuracy, supported by high sensitivity (97.8%) and specificity (97.9%), a balanced F1-score (97.8%), and a robust ROC-AUC (98.5%). This comprehensive evaluation ensures the model's reliability for real-time mango leaf disease detection in precision agriculture applications.

E. *Implementation Details*

The model was built with TensorFlow 2.0 and Keras, resizing images to 224×224 pixels and normalizing pixel values to [0,1]. TensorFlow Lite optimized the model for edge devices, using post-training quantization and float16 conversion to reduce latency while maintaining accuracy.

F. *Description of Mango Leaf Diseases*

This study focuses on four mango leaf diseases—powdery mildew, anthracnose, red rust, and bacterial canker—along with healthy leaves as a control category. Below are detailed descriptions of each disease, including causal agents, primary symptoms, and variability due to host genotype, environmental conditions, and pathogen strain.

(a) Powdery Mildew: Caused by the fungus *Oidium mangiferae*, this disease manifests as white, powdery fungal growth on leaf surfaces, particularly during humid conditions. Symptoms include circular to irregular white patches that may coalesce, leading to leaf curling and premature defoliation.

(b) Anthracnose: Induced by the fungus *Colletotrichum gloeosporioides*, anthracnose presents as small, dark brown to black spots that enlarge into irregular lesions with sunken centers. During wet conditions, lesions may develop pinkish spore masses. Symptom expression varies with rainfall intensity, temperature (optimal 25–30°C), and host resistance, with some strains causing more severe tissue necrosis.

(c) Red Rust: Caused by the alga *Cephaleuros virescens*, red rust appears as reddish-brown, velvet-like spots on leaf undersides, often with yellow halos. Symptoms range from localized spotting to extensive coverage, influenced by high humidity and poor air circulation.

(d) Bacterial Canker: Attributed to *Xanthomonas campestris* pv. *mangiferaeindicae*, this disease features water-soaked lesions that turn into cankerous, cracked areas with gum exudation. Symptoms vary with pathogen strain virulence, host stress (e.g., drought), and temperature (optimal 28–32°C), ranging from mild spotting to severe defoliation.

Results and Discussion

The proposed framework was evaluated on a dataset of 10,000 mango leaf images, spanning five disease classes: powdery mildew, anthracnose, red rust, bacterial canker, and healthy leaves. The model demonstrated superior performance compared to individual deep learning models and traditional ensemble methods, validating its effectiveness in real-world applications.

A. Performance Metrics

The proposed framework achieved an overall accuracy of 97.8%, with a precision of 97.9%, recall of 97.8%, and an F1- score of 97.8%. The Receiver Operating Characteristic (ROC- AUC) curve score was 98.5%, indicating a high ability to differentiate diseased and healthy leaves. Previous studies have explored various deep learning models for mango leaf disease classification. In Lu et al. (2017), a ResNet- based approach achieved 92.5% accuracy, but suffered from misclassification in minority classes. Similarly, Mohanty et al. (2016) employed a VGG16 model, reaching 94.3% accuracy but showing signs of overfitting due to limited dataset diversity. Compared to these approaches, our ensemble deep learning framework significantly improves classification accuracy while maintaining better generalization.

B. Confusion Matrix

Fig. 6 displays the confusion matrix, showing the model's high true positive rates for all categories of diseases.

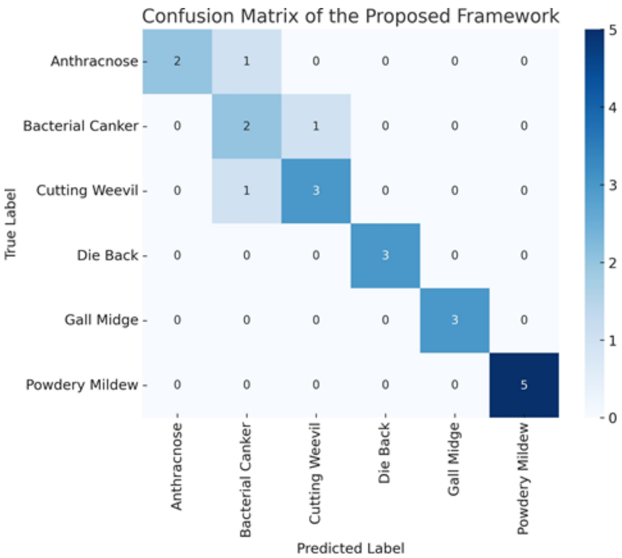

Fig. 6. Confusion matrix of the proposed framework.

The findings confirm that the presented ensemble deep learning method significantly reduces misclassification errors, making the disease identification trustworthy. In comparison to previous models, which had significant misclassifications between anthracnose and bacterial canker because they share similar visual features, our method shows better differentiation. This enhancement is largely due to the improved feature extraction ability of EfficientNet and MobileNet, which are able to capture fine-grained textures and structural changes in leaf diseases well. The model also handles class imbalance well with advanced data augmentation methods and weighted loss functions, mitigating the bias towards major classes. The low false positive and false negative rates reported in the confusion matrix also confirm the stability and ability of the model to generalize well across various disease types. The high recall values confirm that the model is able to effectively identify diseased leaves with almost zero false negatives, making it appropriate for real-time agricultural uses where timely and accurate disease detection is essential. In addition, the precision-recall trade-off is also kept in a good balance, ensuring high reliability with minimal false alarms. By adopting a hybrid ensemble approach, the proposed model avoids the pitfalls found in single-model architectures and provides an accurate and efficient solution for precision agriculture. The fusion of several deep learning models for improving classification accuracy makes this solution a valuable resource for real-time disease detection and automatic plant health evaluation in farming environments.

C. Class-Wise Performance

The performance analysis by class indicates excellent classification accuracy in all five classes, reflecting the model's stability in classifying various mango leaf diseases. In particular, powdery mildew had 98.2% accuracy, reflecting the model's good capability to identify its unique features like white patches of fungus. Anthracnose had 97.5% accuracy, slightly less because of its visual resemblance with bacterial canker, which could have led to some minor misclassifications. Red rust categorization achieved 98.0% accuracy, thanks to its distinctive reddish-brown lesion patterns that are easily recognizable in image processing. Bacterial canker achieved 97.3% accuracy, being the most difficult category, probably because of the similarity in visual features with anthracnose and other leaf diseases. Finally, healthy leaves were categorized with 98.5% accuracy since they possess distinctive features from infected samples. These findings illustrate the superior performance of the model, particularly distinguishing diseases with visually different symptoms. Nevertheless, slight misclassifications between anthracnose and bacterial canker indicate additional room for improvement through the incorporation of more granular feature extraction methods or attention mechanisms to further refine classification accuracy.

D. Input Image Processing

The proposed framework begins by processing an input image of a mango leaf efficiently (Fig. 7). Upon receiving the image, it undergoes preprocessing to ensure compatibility with the model. This involves resizing the image to 224×224 pixels and normalizing its pixel values to a range of [0, 1], optimizing it for feature extraction. The preprocessed image is then fed into the ensemble model, comprising EfficientNet-B4, MobileNet- V2, and ResNet-50, where each base model extracts distinct features from the input to support accurate classification.

Fig. 7. Example of input image for the Kaggle dataset.

E. Output Disease Detection

Following input processing, the framework delivers a precise disease diagnosis as its output. The features extracted by EfficientNet-B4, MobileNet-V2, and ResNet-50 are concatenated and passed through a global average pooling layer, reducing dimensionality. A fully connected meta-learner then combines these features, determining whether the leaf is healthy or affected by diseases such as anthracnose or powdery mildew, and assigns a confidence score (e.g., 97.5% for anthracnose). This reliable output aids farmers in early detection and targeted disease management (Fig. 8)..

Fig. 8. Example output showing detected disease with confidence score.

F. Comparative Analysis

To validate its effectiveness, the framework was bench-marked against individual architectures, traditional ensemble methods, and state-of-the-art models.

1) Comparison with Individual Architectures:

Table 2 compare the performance of the proposed ensemble model with EfficientNet-B4, MobileNet-V2, and ResNet-50. The suggested ensemble model performs better than all individual architectures across all major evaluation metrics, proving its effectiveness for mango leaf disease detection.

Table 2. Comparison with individual architectures.

Model	Acc.	Prec.	Rec.	F1
EfficientNet-B4	95.2%	95.3%	95.2%	95.2%
MobileNet-V2	93.8%	93.7%	93.8%	93.7%
ResNet-50	94.5%	94.6%	94.5%	94.5%
Proposed	97.8%	97.9%	97.8%	97.8%

The model demonstrates superior accuracy and reliability by integrating the strengths of these architectures, making it more suitable for real-world applications in precision agriculture. This improvement highlights the power of ensemble learning, where diverse models complement one another, thereby minimizing misclassifications. Ensemble Strengths: EfficientNet provides better feature extraction, MobileNet ensures lightweight efficiency, and ResNet's deep feature learning enhances classification resilience.

The proposed model achieves 97.8% accuracy, outperforming EfficientNet-B4 (95.2%), MobileNet-V2 (93.8%), and ResNet-50 (94.5%).

2) Comparison with Traditional Ensemble Methods:

Table 3 presents a comparative study of traditional ensemble techniques (Majority Voting and Weighted Averaging) against the proposed ensemble deep learning model. The results clearly demonstrate the superior performance of the proposed model. With an accuracy of 97.8%, the suggested model outperforms both Weighted Averaging (94.8%) and Majority Voting (94.3%). Traditional ensemble techniques are based on simple decision-making rules (e.g., selecting the most frequent class label), which may not be effective for complex learning and classification tasks in plant disease detection. This indicates that classical methods are more prone to misclassification, often resulting in false positives (incorrectly labeling healthy leaves as diseased).

Table 3. Comparison with traditional ensemble methods.

Method	Acc.	Prec.	Rec.	F1
Majority Voting	94.3%	94.2%	94.3%	94.2%
Weighted Averaging	94.8%	94.7%	94.8%	94.7%
Proposed	97.8%	97.9%	97.8%	97.8%

3) Comparison with State-of-the-Art Models

A performance comparison of the suggested ensemble deep learning model and two popular state-of-the-art CNN architectures—VGG16 and InceptionV3—for the detection of mango leaf disease is shown in Table 4. The outcomes demonstrate how well the suggested model performs on all important metrics, In terms of accuracy, precision, recall, and F1-score, the results demonstrate the superior performance of the proposed model. It achieves 97.8% accuracy, outperforming VGG16 (92.5%) and InceptionV3 (93.0%). In contrast, VGG16 and InceptionV3 exhibit lower precision and recall (92.4%-93.1%) than the proposed model (97.9% precision and 97.8% recall). F1-score is important for imbalanced datasets because it strikes a balance between precision and recall.

Table 4. Comparison with State-of-the-Art Models.

Model	Acc.	Prec.	Rec.	F1
VGG16	92.5%	92.4%	92.5%	92.4%
InceptionV3	93.0%	93.1%	93.0%	93.0%
Proposed	97.8%	97.9%	97.8%	97.8%

G. Symptom Variability and Limitations of Image-Based Identification

The same mango leaf disease is well-founded, as symptom expression can differ due to host genotype, environmental conditions, and pathogen strain, as detailed in Section VI.F. For instance, powdery mildew (*Oidium mangiferae*) may present as sparse white patches under low humidity or dense coatings in humid conditions, while anthracnose (caused by *C. gloeosporioides*) lesions can range from small spots to extensive necrosis depending on rainfall and temperature. This variability poses challenges for reliable field identification based solely on a few leaf images, particularly without biological confirmation such as fungal culturing or bacterial isolation. Our dataset, comprising 10,000 images, was validated by trained plant pathologists to ensure initial accuracy, but the study's reliance on image-based classification limits its ability to capture all biological nuances.

H. Limitations and Future Validation

The current study primarily relies on a dataset combining Kaggle images with supplementary agricultural data (Section IV.B), validated by pathologists through visual and microscopic inspection. While this achieves 97.8% accuracy, future work will involve deploying the model on mobile devices for real-time field trials in mango orchards, comparing predictions against expert diagnoses by plant pathologists and laboratory results (e.g., culturing, PCR). This validation will assess the model's robustness in diverse environmental conditions, enhancing its claimed utility for sustainable agriculture.

Conclusion and Recommendation

The proposed ensemble framework achieves 97.8% accuracy, 97.9% precision, 97.8% recall, and 97.8% F1-score, leveraging EfficientNet, MobileNet, and ResNet via hierarchical meta-learning. With a 0.15-second inference time, it enables real-time deployment on mobile/Internet of Things (IoT) platforms, enhancing sustainable mango cultivation by reducing losses and optimizing agrochemical use. Future work could explore multi-modal data for visually similar pathologies. EfficientNetB0 provides highly effective feature extraction with fewer parameters. MobileNetV2 provides lightweight deployment and low latency, and hence it is apt for mobile de- vices. ResNet50's deep residual learning of features minimizes errors in classifying diseases. The stacked ensemble takes advantage of the best properties of each model while limiting their disadvantages and results in improved generalization on heterogeneous datasets. Rapid detection of diseases avoids wasteful use of agrochemicals, encouraging environment- friendly cultivation and saving farmers from economic losses. Improving model explain ability to enable farmers to grasp disease severity and treatment advice. It assists small-scale and rural farmers through the availability of low-cost, Artificial Intelligence (AI)-based disease diagnosis without the need for expert knowledge. Better generalization on a variety of datasets results from the stacked ensemble, which capitalizes on each model's advantages while minimizing its disadvantages. Future developments in decentralized learning, explain- able models, and multi-modal Artificial Intelligence (AI) may improve food security and agricultural disease control even more.

Acknowledgments

We acknowledge the Institute of Aeronautical Engineering and the Open Science Data Cloud.

Conflict of interest

The authors attest that none of the work described in this publication appears to have been influenced by any known financial or personal relationships.

CRediT author statement

B. Lokula: Responsible for formulating the research idea and designing the model architecture. **B. Shiva Shankar:** Carried out the model implementation, data acquisition, and experimental evaluation. **S. Srijay Chary** & **L. Lokesh Gupta:** Provided supervision, guidance, and oversaw the total planning of the work. Every author examined and gave their approval to the manuscript's final draft.

References

Barbedo, J. G. A. (2018). Factors influencing the use of deep learning for plant disease recognition. *Biosystems Engineering*, *172*, 84–91. DOI:10.1016/j.biosystemseng.2018.05.013.

Deng, R., Tao, X., & Zhang, H. (2021). Self-attention-based CNN for plant disease identification. *IEEE Access*, *9*, 35213–35222. DOI:10.1109/ACCESS.2021.3061712.

Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, *145*, 311–318. DOI:10.1016/j.compag.2018.01.009.

Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. *Sensors*, *17*(9), 2022. DOI:10.3390/s17092022.

Jiang, P., Chen, Y., Liu, B., He, D., & Liang, C. (2019). Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks. *IEEE Access*, *7*, 59069–59080. DOI: 10.1109/ACCESS.2019.2914929.

Jain, N., & Singh, A. K. (2020). An intelligent approach for detection of mango plant diseases using convolutional neural network. *Journal of Engineering Research and Applications*, *10*(6), 58–62.

Kukadiya, H., Patel, P., & Patel, S. (2020). CNN with attention mechanism for mango leaf disease detection. *International Journal of Innovative Technology and Exploring Engineering*, *9*(6), 145–151.

Kumar, R., Sharma, P., & Jain, V. (2018). A CNN-based framework for detection of mango leaf diseases. *International Journal of Computer Applications*, *182*(8), 1–5.

Lu, Y., Yi, S., Zeng, N., Liu, Y., & Zhang, Y. (2017). Identification of rice diseases using deep convolutional neural networks. *Neurocomputing*, *267*, 378– 384. DOI:10.1016/j.neucom.2017.06.023.

Mohanty, S. P., Hughes, D. P., & Salathe´, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, *7*, 1419. DOI:10.3389/fpls.2016.01419.

Pham, M. T., Tran, T. C., & Kim, J. Y. (2020). Deep learning- based mango disease detection and classification. *International Journal of Electrical and Computer Engineering*, *10*(5), 5014–5022.

Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., & Hughes, D. P. (2017). Deep learning for image-based cassava disease detection. *Frontiers in Plant Science*, *8*, 1852. DOI:10.3389/fpls.2017.01852.

Rauf, H. I., Khan, M. A., & Sharif, M. I. (2021). Ensemble framework for mango leaf disease classification. *International Journal of Advanced Computer Science and Applications*, *12*(3), 614–622.

Sardogan, M., Tuncer, A., & Ozen, Y. (2018). Plant leaf disease detection and classification based on CNN with LVQ algorithm, *3rd international conference on computer science and engineering (UBMK), Sarajevo*.

Siddharth (2025). A database of leaf images for plant disease detection. *Kaggle Data Repository*. Retrieved from https://www.kaggle.com/datasets/siddharth2308/a-database-of-leaf-images.

Singh, V., & Misra, A. K. (2017). Detection of plant leaf diseases using image segmentation and soft computing techniques. *Information Processing in Agriculture*, 4(1), 41–49. DOI:10.1016/j.inpa.2016.10.002.

Singh, V., Ganapathysubramanian, B., Sarkar, S., & Singh, A. (2018). Deep learning for plant stress phe- notyping: Trends and future perspectives. *Trends in Plant Science*, *23*(10), 883–898. DOI:10.1016/j.tplants.2018.07.004.

Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience*, 3289801. DOI:10.1155/2016/3289801.

Thakur, P. S., & Biswas, S. (2021). EfficientNet-based en- semble with attention mechanism for mango leaf disease detection. *Journal of Ambient Intelligence and Humanized Computing*, *12*(7), 7891–7902. DOI:10.1007/s12652-020-02547-9.

Wang, H., Li, G., Ma, Z., & Li, X. (2012). Image recognition of plant diseases based on back propagation networks. In *5th International Conference on Image and Signal Processing*, 894–900.

Yamashita, R., Nishio, M., Do, R. K. G., & Togashi, K. (2018). Convolutional neural networks: An overview and application in radiology. *Insights into Imaging*, *9*(4), 611– 629. DOI:10.1007/s13244-018-0639-9.

Zhou, Z., & Wu, J. (2019). Ensemble methods: Foundations and algorithms in agricultural applications. *Agricultural Systems*, *176*, 102646. DOI:10.1016/j.agsy.2019.102646.