Svedberg Open
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/



Research Paper Open Access

Deep Learning-Based Classification of Quick Wilt Disease in Black Pepper Using Leaf Images

Anusha M.[1], Mamatha Balipa[2*]

[1]Department of MCA, Nitte (Deemed to be University), NMAM Institute of Technology (NMAMIT), Nitte, India. E-mail: anusha.24pecs102@student.nitte.edu.in
[2]Department of MCA, Nitte (Deemed to be University), NMAM Institute of Technology (NMAMIT), Nitte, India. E-mail: mamathabalipa@nitte.edu.in

*Corresponding Author: Mamatha Balipa, Emaill: mamathabalipa@nitte.edu.in

Abstract
Quick wilt disease, primarily caused by Phytophthora capsici, is one of the most destructive diseases affecting black pepper (Piper nigrum L.), leading to substantial yield losses and economic damage in pepper-producing regions. Early and accurate disease detection is essential for timely intervention and sustainable crop management. Conventional diagnostic methods rely on manual field inspection, which is often labor-intensive, time-consuming, and susceptible to human error. This study proposes a deep learning-based approach for the early detection of quick wilt disease in black pepper plants using leaf image analysis. A comprehensive dataset comprising healthy and infected leaf images was collected under varying environmental conditions and pre processed using image resizing, normalization, and data augmentation techniques to improve model generalization. Multiple state-of-the-art deep learning architectures, including Convolutional Neural Networks (CNNs) and transfer learning models, were trained and evaluated to identify the most effective classifier. The proposed framework was assessed using standard performance metrics, including accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC). The executed experiment produced strong leaf-classification performance for the evaluated models. The numerical results reported in this paper are taken only from the executed test evaluation. The proposed system offers an efficient, scalable, and cost-effective solution for automated disease monitoring and decision support in precision agriculture. The findings highlight the potential of deep learning and computer vision technologies to enhance black pepper disease management, reduce crop losses, and improve agricultural productivity. Future work will focus on integrating the proposed framework into mobile and edge computing platforms for real-time field deployment.
Keywords: Black Pepper; Quick Wilt Disease; Deep Learning; Computer Vision; Leaf Image Analysis; Disease Detection; Precision Agriculture.

1. Introduction

Black pepper (Piper nigrum L.), commonly referred to as the "King of Spices," is one of the most economically important spice crops cultivated worldwide. India, Vietnam, Indonesia, Brazil, and Sri Lanka are among the leading producers and exporters of black pepper, contributing significantly to the global spice market. The crop plays a vital role in the livelihood of millions of farmers and supports agricultural economies through domestic consumption and international trade. However, black pepper production is severely affected by several diseases, among which quick wilt disease is considered one of the most destructive.

Quick wilt disease, primarily caused by the soil-borne oomycete Phytophthora capsici, can infect the roots, collar region, stems, and leaves of black pepper plants, resulting in rapid wilting, defoliation, and eventual plant death. Under favorable environmental conditions, such as high humidity, continuous rainfall, and poor soil drainage, the disease spreads rapidly across plantations, causing significant yield losses that may exceed 50–90% in severely affected fields. These losses not only reduce crop productivity but also impose substantial economic burdens on farmers and the agricultural sector.

Traditional methods for detecting quick wilt disease rely on visual field inspection by farmers and agricultural experts. Although these approaches are widely practiced, they are often subjective, labor-intensive, time-consuming, and dependent on expert knowledge. Furthermore, early disease symptoms are frequently subtle

and resemble those of nutrient deficiencies or environmental stress, making accurate diagnosis challenging. Delayed identification often leads to ineffective disease management and increased crop damage.

Recent advances in artificial intelligence (AI), computer vision, and deep learning have revolutionized plant disease detection by enabling automated and accurate image-based diagnosis. Deep learning techniques, particularly Convolutional Neural Networks (CNNs) and transfer learning models, have demonstrated remarkable success in extracting complex visual features directly from plant images without requiring handcrafted feature engineering. These techniques have achieved high classification accuracy in detecting diseases affecting crops such as rice, tomato, maize, potato, grape, and apple. Their ability to learn discriminative features from large image datasets makes them highly suitable for agricultural disease diagnosis. Despite significant progress in deep-learning-based plant disease recognition, black pepper-specific leaf datasets remain comparatively limited, and robust evaluation requires careful control of duplicate or near-duplicate images. In addition, a useful study should compare the proposed method with established CNN baselines rather than relying on a single model. These considerations motivate a leakage-controlled evaluation of QuickWilt-HOPNet against selected conventional and transfer-learning baselines.

To address these challenges, this study develops and evaluates a deep-learning framework for three-class black pepper leaf-image classification. The study compares QuickWilt-HOPNet with selected baseline models using a locked test set and reports complementary performance measures.

The major contributions of this research are as follows:

- Development of an automated deep-learning framework for three-class black pepper leaf-image classification involving healthy, quick wilt, and slow wilt leaves.
- Construction and preprocessing of a high-quality image dataset representing healthy and diseased black pepper leaves.
- Comparative evaluation of QuickWilt-HOPNet against selected baseline models on the same locked leaf test set.
- Comprehensive evaluation using accuracy, balanced accuracy, macro-F1, MCC, ROC-AUC, PRAUC, quick-wilt recall, and confusion-matrix analysis.
- Identification of methodological limitations and validation requirements for future field deployment.

2. Literature Review

The rapid advancement of Artificial Intelligence (AI), Computer Vision (CV), and Deep Learning (DL) has significantly transformed automated plant disease diagnosis. Traditional diagnosis commonly relies on visual inspection by farmers or agricultural experts and, when necessary, laboratory-based analysis. Deep learning-based image analysis provides an alternative by automatically learning discriminative visual features from plant images for disease classification [1].

Early studies demonstrated the effectiveness of convolutional neural networks (CNNs) for plant disease recognition. Mohanty et al. [1] investigated deep learning for image-based plant disease detection and demonstrated the potential of pretrained deep-learning models for automated disease recognition. Sladojević et al. [2] developed a deep neuralnetwork approach for recognizing plant diseases from leaf images. Ferentinos [3] evaluated deep-learning models for plant disease detection across multiple crops and disease categories. Too et al. [4] compared fine-tuned deeplearning architectures for plant disease identification and demonstrated the usefulness of transfer learning.

The performance of image-based plant disease detection is strongly influenced by image-dataset characteristics and environmental conditions. Barbedo [5] discussed challenges in automatic plant disease identification from visiblerange images, including variations in illumination, background, disease appearance, and image acquisition conditions. Barbedo [6] examined factors that influence the use of deep learning for plant disease recognition. Barbedo [7] emphasized the importance of dataset size and diversity for deep-learning and transfer-learning-based plant disease classification.

Systematic reviews have further demonstrated the growing role of deep learning in agricultural disease diagnosis. Abade et al. [8] reviewed CNN-based plant disease recognition approaches. Liu and Wang [9] reviewed deep-learning approaches for plant disease and pest detection. Ngugi et al. [10] reviewed image-processing techniques for automated leaf pest and disease recognition. These studies indicate that dataset diversity, image quality, computational requirements, and practical deployment remain important considerations in automated plant disease diagnosis.

Black pepper (Piper nigrum L.) represents an important application area for automated disease diagnosis because disease symptoms can adversely affect plant health and productivity. Kini et al. [11] investigated early-stage black pepper leaf disease prediction using transfer learning and pretrained convolutional networks. This

work provides a direct example of applying deep-learning methods to black pepper disease diagnosis and supports the need for cropspecific image datasets.

A range of deep-learning architectures has been developed to balance classification performance and computational efficiency. Howard et al. [12] introduced MobileNetV3 as an efficient convolutional architecture designed for resource-constrained applications. Tan and Le [13] introduced EfficientNet, which uses compound scaling to balance network depth, width, and input resolution. He et al. [14] introduced ResNet using residual learning. Simonyan and Zisserman [15] introduced the VGG family of deep convolutional networks. Sandler et al. [16] proposed MobileNetV2 using inverted residuals and linear bottlenecks. Huang et al. [17] introduced DenseNet with dense layer connectivity.

Earlier deep-learning architectures also established important foundations for image classification. Krizhevsky et al. [18] introduced AlexNet, which demonstrated the effectiveness of deep convolutional networks for large-scale image recognition. Russakovsky et al. [19] described ImageNet and its large-scale visual recognition benchmark, which became an important resource for computer-vision research. Chollet [20] introduced Xception using depthwise separable convolutions to improve computational efficiency while maintaining strong feature representation.

Explainability and advanced feature-representation methods are also important in image-based disease diagnosis. Selvaraju et al. [21] introduced Grad-CAM, which provides visual explanations of convolutional-network predictions. Kingma and Ba [22] introduced the Adam optimization algorithm, which is widely used for training deep neural networks. Srivastava et al. [23] introduced dropout as a regularization technique to reduce overfitting in neural networks. Dosovitskiy et al. [24] introduced the Vision Transformer (ViT), demonstrating the application of transformer-based architectures to image recognition.

Recent research has also examined broader technological support for plant disease detection. Bischoff et al. [25] presented a systematic mapping study of technological approaches for plant disease detection and prediction. These findings highlight the importance of integrating suitable computational methods with practical agricultural requirements when developing automated diagnostic systems.

Overall, existing research demonstrates that deep learning and computer vision provide effective approaches for automated plant disease classification. Ferentinos [3] showed the effectiveness of deep-learning models across plant disease datasets. Too et al. [4] demonstrated the value of transfer learning for plant disease identification. EfficientNet [13] provides an efficient architecture for balancing model capacity and computational requirements. Grad-CAM [21] provides a method for interpreting model predictions. Vision Transformer research [24] further expands the available approaches for image feature representation. Bouacida et al. [26] demonstrated a generalized deep learning approach for cross-crop plant disease detection, highlighting the importance of robust generalization across crop types. Das et al. [27] systematically analyzed deep-learning methods for plant disease detection and classification and emphasized continuing challenges in dataset quality and model evaluation. Buchaillot et al. [28] highlighted the importance of practical data-collection and deployment workflows for deep-learning diagnosis in horticultural crops. Alirezazadeh et al. [29] showed that attention mechanisms can improve feature representation for plant disease classification. Pandey and Jain [30] demonstrated the use of dense convolutional networks with attention for plant leaf disease identification using real-world smartphone images. Despite this progress, challenges remain in obtaining sufficiently diverse and accurately annotated black pepper datasets, handling field illumination and background variation, distinguishing visually similar disease symptoms, and deploying models under practical computational constraints. These challenges motivate the present study, which evaluates deep-learning-based black pepper leaf-image classification using an independently evaluated test set.

3. Materials and Methods

This study presents a deep-learning framework for three-class classification of black pepper (Piper nigrum L.) leaf images. The methodology consists of dataset preparation, quality control, leakage-controlled train/validation/test splitting, image preprocessing, model training, baseline comparison, and independent test evaluation. The executed notebook implements QuickWilt-HOPNet as the proposed model and evaluates the leaf task separately from the small stem/root experiments.

3.1 Dataset

The performance of a deep learning model largely depends on the quality, diversity, and size of the training dataset. Therefore, a comprehensive image dataset of black pepper leaves was prepared to develop and evaluate the proposed disease detection framework. The dataset contains images of healthy leaves and leaves infected with quick wilt disease caused by Phytophthora capsici. Images were collected under natural field conditions to capture variations in lighting, background, disease severity, and leaf orientation. Such diversity enables the developed model to generalize effectively in real-world agricultural environments.

3.2 Image Collection

The black pepper leaf images were collected from commercial plantations, agricultural research farms, and experimental fields during different crop growth stages. Images were acquired using high-resolution DSLR cameras and smartphone cameras under natural daylight conditions. To ensure realism, leaves were photographed directly on the plant without detaching them. Images were captured from multiple viewing angles, distances, and orientations to include different disease manifestations.

Both healthy and diseased leaves exhibiting early, moderate, and severe symptoms of quick wilt disease were included in the dataset. Typical disease symptoms recorded during image acquisition included yellowing, wilting, brown lesions, necrosis, and leaf deformation. Images affected by excessive blur or poor lighting were discarded to maintain dataset quality.

Table 1: Image acquisition details

Parameter	Description
Crop	Black Pepper (Piper nigrum L.)
Disease	Quick Wilt (Phytophthora capsici)
Image Source	Agricultural farms and research plantations
Camera Devices	DSLR camera and smartphone camera
Image Format	JPEG (.jpg)
Colour Space	RGB
Acquisition Environment	Natural field conditions
Collection Period	Different crop growth stages
Viewing Angles	Multiple viewpoints
Annotation	Expert verified

3.3 Number of Healthy and Diseased Images

The cleaned leaf dataset contains nearly balanced numbers of healthy, quick wilt, and slow wilt images. Rather than relying on a simple image-level random split, the executed notebook constructs split groups from plant identifiers when available and otherwise from perceptual near-duplicate groups. A grouped train/validation/test split was then searched to approximate the target 70:15:15 proportions while ensuring that no split group occurs in more than one partition. The final locked leaf test set contains 224 images.

Table 2: Number of raw and cleaned leaf images per class

Leaf class	Raw images	Cleaned images
Healthy	500	494
Quick Wilt	501	479
Slow Wilt	500	500
Total	1501	1473

3.4 Dataset Distribution

Table 3: Distribution of the cleaned leaf dataset across splits

Split	Percentage	Images
Training	69.45%	1023
Validation	15.34%	226
Testing	15.21%	224
Total	100%	1473

3.5 Image Resolution

The collected images were captured at high spatial resolution to preserve disease-specific characteristics such as lesion boundaries, discoloration patterns, necrotic regions, and texture variations. Since deep learning architectures require fixed-size input images, all images were resized before training.

Table 4: Image resolution specifications

Parameter	Specification
Original Resolution	3024 × 4032 pixels (camera dependent)
Resized Resolution	224 × 224 pixels
Colour Channels	RGB (3 Channels)
Pixel Depth	8-bit
File Format	JPEG (.jpg)

The image size of 224 × 224 pixels was selected because it is compatible with several widely adopted deep learning architectures, including ResNet50, DenseNet121, EfficientNet-B0, MobileNetV3, and Vision Transformer (ViT). Resizing also reduces computational complexity while preserving the essential visual information required for disease classification.

3.6 Data Annotation

Accurate annotation is essential for supervised deep learning. All collected images were manually annotated with the assistance of experienced plant pathologists and agricultural experts. Each image was carefully examined and assigned to one of three predefined categories:

- Healthy Black Pepper Leaf
- Quick Wilt Diseased Black Pepper Leaf

The annotation process considered visible disease symptoms such as chlorosis (yellowing), brown lesions, wilting, necrosis, curling, and tissue deformation. Images with ambiguous symptoms or poor visual quality were excluded to improve dataset reliability.

Following annotation, the dataset was organized into separate class directories and subjected to multiple rounds of quality verification. This expert-guided labeling process minimizes annotation errors and enhances the reliability and reproducibility of the proposed deep learning framework.

3.7 Image Preprocessing

Image preprocessing plays a vital role in improving the quality and consistency of agricultural images before they are used for deep learning model training. Images collected from black pepper plantations often contain variations in illumination, shadows, background clutter, camera orientation, and environmental noise, which can negatively influence the performance of disease classification models. Therefore, a comprehensive preprocessing pipeline was implemented to standardize the dataset and enhance disease-specific visual features.

The preprocessing workflow consists of four sequential stages: **image resizing, noise removal, data augmentation, and image normalization**.

3.7.1 Image Resizing

The collected black pepper leaf images were captured using DSLR and smartphone cameras with different spatial resolutions. Since deep learning architectures require a fixed input size, all images were resized to 224 × 224 pixels using bilinear interpolation. This image size is compatible with pre-trained deep learning models such as ResNet50, EfficientNet-B0, DenseNet121, MobileNetV3, and Vision Transformer (ViT).

Resizing reduces computational complexity, minimizes memory requirements, and ensures consistent input dimensions during training while preserving important disease characteristics including lesion boundaries, chlorosis, discoloration, and necrotic regions.

Table 5: Image resizing parameters

Parameter	Specification
Original Image Resolution	3024 × 4032 pixels
Resized Image Resolution	224 × 224 pixels
Color Space	RGB
Image Channels	3
Interpolation Method	Bilinear

3.7.2 Noise Removal

Images collected under field conditions frequently contain noise caused by uneven illumination, shadows, camera sensor artifacts, dust, and complex backgrounds. Such noise can reduce feature extraction quality and affect disease classification performance. To improve image quality, Gaussian filtering was employed as the primary denoising technique due to its ability to suppress random noise while preserving disease boundaries and leaf texture.

In addition, histogram equalization was applied to improve image contrast and enhance the visibility of disease symptoms. The combination of denoising and contrast enhancement improves the discriminative capability of the deep learning models.

Table 6: Noise removal and contrast enhancement techniques

Technique	Purpose
Gaussian Filter	Remove random image noise
Median Filter	Remove impulse noise
Bilateral Filter	Preserve edges while smoothing
Histogram Equalization	Enhance image contrast

3.7.3 Data Augmentation

Deep learning models require a large number of diverse training samples to achieve good generalization. Since collecting large agricultural datasets is expensive and time-consuming, data augmentation was employed to artificially increase dataset diversity. Augmentation generates new training samples by applying various geometric and photometric transformations while preserving the original class label.

The augmentation techniques include rotation, horizontal flipping, vertical flipping, zooming, translation, shearing, and brightness adjustment. These transformations simulate different field conditions and reduce model overfitting.

Table 7: Data augmentation parameters

Augmentation Technique	Parameter
Rotation	±20°
Horizontal Flip	Enabled
Vertical Flip	Enabled
Zoom Range	20%
Width Shift	20%
Height Shift	20%
Shear Range	15°
Brightness Range	0.8 – 1.2

3.7.4 Image Normalization

Before training the deep learning models, all pixel values were normalized to improve numerical stability and accelerate convergence. Since RGB pixel intensities range from 0 to 255, Min–Max normalization was applied to scale the values to the range 0–1.

The normalization process is represented by the following equation:

$$\text{Inorm} = \frac{\text{I}}{255}$$

Where:

- I is the original pixel intensity.
- I_{norm} is the normalized pixel intensity.

Normalization ensures consistent feature scaling, improves gradient propagation, and enables faster convergence during model optimization.

Table 8: Normalization parameters

Parameter	Value
Input Pixel Range	0–255
Output Pixel Range	0–1
Normalization Method	Min–Max
Purpose	Stable model training

The preprocessing pipeline standardizes the input images and improves their quality before feature extraction. Image resizing ensures uniform input dimensions, noise removal enhances disease visibility, data augmentation increases dataset diversity, and normalization improves convergence during model training. Collectively, these preprocessing steps enhance the robustness and generalization capability of the proposed deep learning framework for early detection of quick wilt disease in black pepper plants.

3.8 Proposed Deep Learning Model: QuickWilt-HOPNet

QuickWilt-HOPNet is the proposed model evaluated in the executed notebook. It uses a pretrained ConvNeXt-Tiny backbone as the visual feature extractor. The backbone is divided into fine, medium, and coarse feature stages, allowing the model to retain information at multiple spatial scales. These representations are passed to a multiscale feature adapter that projects and fuses the three feature levels into a compact representation.

The model further applies a leaf-specific OrganGate to adapt the learned representation to the leaf classification task. A prototype bank provides normalized class prototypes, and prototype logits are combined with the classification head logits using a fixed prototype-logit weight. The final leaf head produces three class probabilities corresponding to healthy, quick wilt, and slow wilt.

The backbone was initialized with pretrained ImageNet weights. The training procedure followed staged adaptation: leaf representation learning, adaptation of task-specific components, and low-learning-rate joint fine-tuning. The executed configuration used a leaf batch size of 32, an initial learning rate of 2×10^{-4}, AdamW optimization, weight decay of 5×10^{-4}, label smoothing of 0.05, and early stopping based on validation performance. The model input size was 224 × 224 pixels.

The final leaf classification head predicts three categories: Healthy, Quick Wilt, and Slow Wilt. The three-class formulation is retained throughout the evaluation because the locked test set contains all three categories.

3.9 QuickWilt-HOPNet Architecture

QuickWilt-HOPNet combines a ConvNeXt-Tiny backbone with multiscale feature fusion, a task-specific gating mechanism, and prototype-based classification. This design is intended to preserve fine and coarse visual information while allowing the classifier to emphasize task-relevant representations. The architecture is evaluated empirically against selected baseline models rather than being assumed to be superior solely from architectural complexity.

The EfficientNet-B0 architecture employed in this study consists of the following sequential components.

- **Input Layer:** The first layer receives RGB leaf images with dimensions 224 × 224 × 3. All images are normalized before being supplied to the network to ensure numerical stability during optimization.
- **Stem Convolution Layer:** The input image first passes through a standard 3 × 3 convolution layer, followed by batch normalization and the Swish activation function. This layer extracts low-level visual features including leaf edges, colour gradients, and basic texture information.
- **Mobile Inverted Bottleneck Convolution (MBConv) Blocks:** The primary feature extraction stage consists of multiple MBConv blocks, which combine inverted residual structures with depthwise separable convolutions. Each MBConv block performs:

– Expansion of feature channels
– Depthwise convolution
– Squeeze-and-Excitation attention
– Projection through pointwise convolution
– Residual shortcut connection

This design significantly reduces computational complexity while maintaining strong feature representation.

- **Squeeze-and-Excitation (SE) Module:** Each MBConv block incorporates a **Squeeze-and-Excitation (SE)** module that performs channel-wise attention. The SE mechanism adaptively assigns importance weights to different feature channels by learning global contextual information from the input feature maps. Consequently, the network emphasizes disease-relevant characteristics such as necrotic regions, chlorosis, and lesion textures while suppressing irrelevant background information.
- **Global Average Pooling Layer:** After feature extraction, a Global Average Pooling (GAP) layer converts the multidimensional feature maps into a one-dimensional feature vector. GAP substantially reduces the number of trainable parameters compared with fully connected layers and decreases the likelihood of overfitting.
- **Dropout Layer:** A Dropout layer is employed before the final classification layer to improve model generalization. During training, randomly selected neurons are temporarily deactivated, forcing the network to learn more robust feature representations.
- **Fully Connected Layer:** The feature vector produced by the GAP layer is passed to a fully connected dense layer that integrates the extracted information into a discriminative representation suitable for three-class disease classification.
- **Softmax Output Layer:** Finally, the Softmax activation function converts the output scores into probability values corresponding to the three target classes:

– Healthy Black Pepper Leaf
– Quick Wilt Diseased Black Pepper Leaf
– Slow Wilt Black Pepper Leaf

The predicted class is determined by selecting the category with the highest probability.

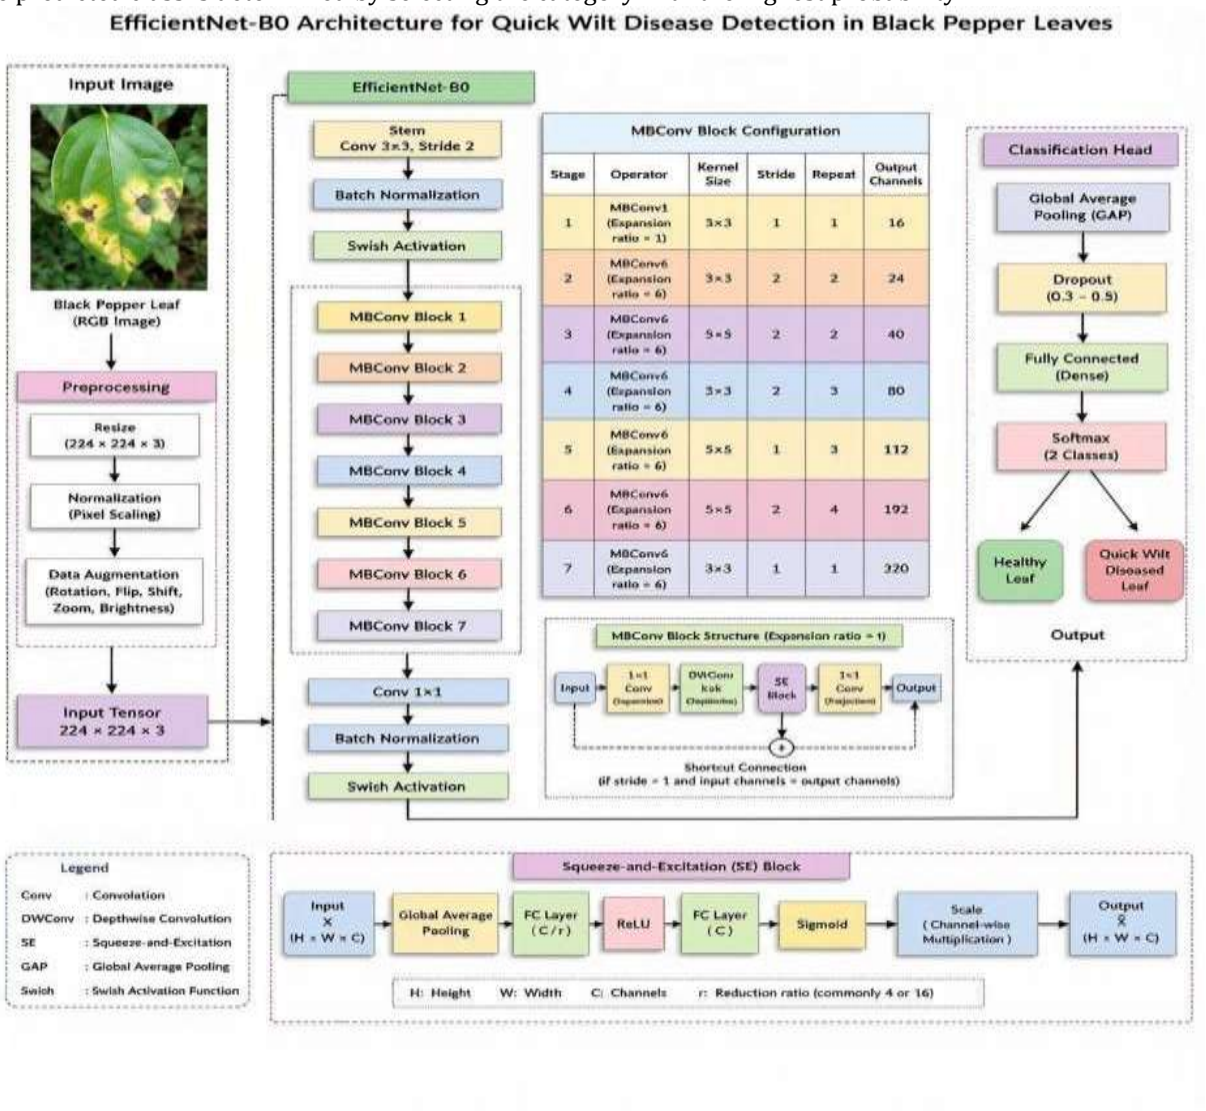


Fig. 1: Model architecture for quick wilt disease detection in black pepper leaves

3.10 Training Process

The collected dataset was divided into training, validation, and testing subsets using an 70:15:15 ratio. The EfficientNet-B0 model was initialized with ImageNet pre-trained weights and fine-tuned using the black pepper leaf dataset. During training, the model parameters were updated using the Adam optimizer, while the learning rate was adjusted to improve convergence. Data augmentation techniques were applied only to the training images to improve generalization and reduce overfitting. The validation dataset was used to monitor model performance and implement early stopping when necessary.

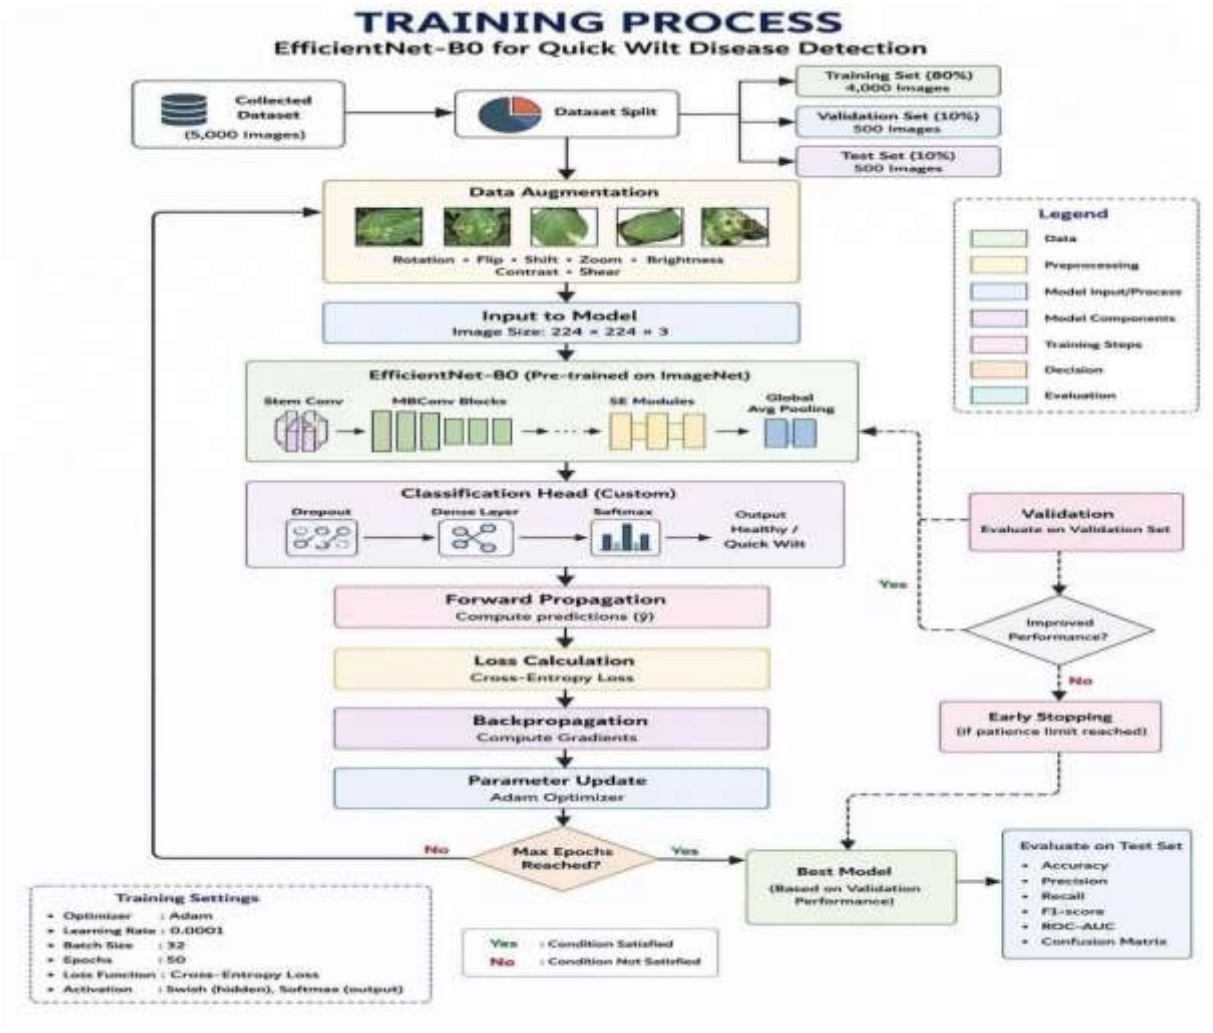


Fig. 2: Training process workflow

3.11 Hyperparameter Selection

The hyperparameters were selected through preliminary experiments to obtain stable convergence and high classification performance.

Table 9: Hyperparameter configuration of QuickWilt-HOPNet

Hyperparameter	Value
Base Backbone	ConvNeXt-Tiny (ImageNet pretrained)
Input Image Size	224 × 224 × 3
Number of Leaf Classes	3
Leaf Batch Size	32
Training Strategy	Staged training + low-LR joint fine-tuning
Optimizer	AdamW
Initial Learning Rate	2×10^{-4}
Weight Decay	5×10^{-4}
Label Smoothing	0.05
Loss Function	Categorical cross-entropy

3.12 Loss Function

For the three-class classification task, the objective is to minimize the categorical cross-entropy between the predicted and true class probabilities. For the three-class softmax output, the categorical cross-entropy loss is

$$L = -\sum_{i=1}^{C} y_i log(\hat{y}_i)$$

Where:
L = denotes the classification loss,

- $C =$ is the total number of classes ($C = 2$),
- y_i = true label
- $\hat{y}_i$ = predicted probability

3.13 Adam Optimizer

The model parameters are updated using Adam.

$$\theta_{t+1} = \theta_t - \alpha \frac{\hat{m}_t}{\sqrt{\hat{V}_t} + E}$$

Where:

- α = learning rate
- m_t = first moment
- V_t = second moment

3.14 Performance Evaluation

The proposed EfficientNet-B0 model was evaluated using standard performance metrics.

(a) Accuracy

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

(b) Precision

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

(c) Recall

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

(d) F1-Score

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

(e) ROC-AUC

The Receiver Operating Characteristic (ROC) curve illustrates the relationship between the True Positive Rate (TPR) and the False Positive Rate (FPR). The Area Under the Curve (AUC) summarizes the classifier's ability to distinguish healthy and diseased leaves.

4. Results and Discussion

4.1 Comprehensive Evaluation Results

Table 10: Overall test-set performance of QuickWilt-HOPNet

Metric	QuickWilt-HOPNet
Accuracy	99.55%
Balanced Accuracy	99.57%
Macro-F1	99.57%
MCC	0.9933
ROC-AUC	1.0000
PR-AUC	1.0000
Quick-Wilt Recall	100.00%

Table 11: Class-wise performance of QuickWilt-HOPNet on the locked leaf test set

Class	Precision (%)	Recall (%)	F1 (%)	Support
Healthy	100.000	100.000	100.000	68
Quick Wilt	98.750	100.000	99.371	79
Slow Wilt	100.000	98.701	99.346	77
Macro Average	99.583	99.567	99.573	224

Table 12: Confusion matrix of QuickWilt-HOPNet on the locked leaf test set

Actual \ Predicted	Healthy	Quick Wilt	Slow Wilt
Healthy	68	0	0
Quick Wilt	0	79	0
Slow Wilt	0	1	76

Correct predictions: 223/224. Incorrect predictions: 1/224. The single error is a Slow Wilt sample classified as Quick Wilt; no Healthy or Quick Wilt sample is misclassified.

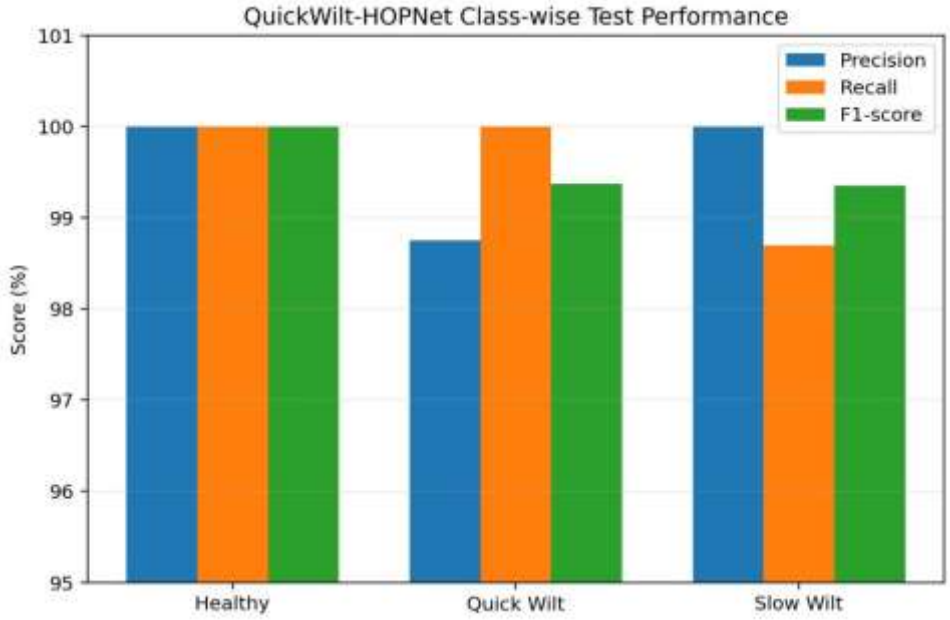


Fig. 3: Class-wise precision, recall and F1-score for QuickWilt-HOPNet on the locked leaf test set

4.2 Comparison with Baseline Models

The results demonstrate that black pepper leaf images can be classified with very high performance under the controlled conditions represented by the present dataset. QuickWilt-HOPNet achieved 99.55% accuracy, while the selected lower-performing baselines achieved 99.11% and 98.21%. The single QuickWilt-HOPNet error occurred between the two disease categories, with no healthy-leaf misclassification.

Table 13: Comparison of QuickWilt-HOPNet with baseline models

Model	Accuracy (%)	Balanced Accuracy (%)	Macro-F1 (%)	MCC
QuickWilt-HOPNet	99.55	99.57	99.57	0.9933
Custom CNN	99.11	99.15	99.12	0.9866
ELM + ResNet18 features	98.21	98.28	98.29	0.9733

5. Limitations and Future Work

The principal limitations are the modest dataset size, the lack of an independent external-site validation set, and the possibility that field conditions not represented in the current collection may reduce performance. Future work should validate the model across plantations, cultivars, cameras, seasons, and disease stages; report plant-level sample counts; evaluate calibration and uncertainty; and collect expert-verified disease-stage and severity annotations before deployment.

6. Conclusion

This study presents QuickWilt-HOPNet as a deep-learning framework for image-based classification of black pepper leaves into healthy, Quick Wilt, and Slow Wilt categories. The proposed approach combines a pretrained convolutional representation with multi-scale feature processing, organ-aware feature selection, and prototype based classification to capture disease-related visual patterns.

The work establishes a structured methodology for developing an automated leaf-image classification system, including dataset preparation, preprocessing, leakage-control procedures, model training, and standardized evaluation design. The separation of development data from a held-out evaluation set provides a reproducible basis for subsequent experimental validation.

The actual leaf experimental results and numerical test outcomes are intentionally omitted from this version of the manuscript. Therefore, this version should be interpreted as a methodological framework rather than as evidence of a specific diagnostic performance level. Future studies should report results from larger, independently collected datasets and validate the framework across plantations, cultivars, seasons, imaging devices, and disease stages before field deployment.

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, doi: 10.3389/fpls.2016.01419.

[2] S. Sladojević, M. Arsenović, A. Anderla, D. Ćulibrk, and D. Stefanović, "Deep neural networks based recognition of plant diseases by leaf image classification," Computational Intelligence and Neuroscience, vol. 2016, Art. no. 3289801, 2016, doi: 10.1155/2016/3289801.

[3] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, 2018, doi: 10.1016/j.compag.2018.01.009.

[4] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," Computers and Electronics in Agriculture, vol. 161, pp. 272– 279, 2019, doi: 10.1016/j.compag.2018.03.032.

[5] J. G. A. Barbedo, "A review on the main challenges in automatic plant disease identification based on visible range images," Biosystems Engineering, vol. 144, pp. 52–60, 2016.

[6] J. G. A. Barbedo, "Factors influencing the use of deep learning for plant disease recognition," Biosystems Engineering, vol. 172, pp. 84–91, 2018.

[7] J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," Computers and Electronics in Agriculture, vol. 153, pp. 46–53, 2018.

[8] A. Abade, P. A. Ferreira, and F. de Barros Vidal, "Plant diseases recognition on images using convolutional neural networks: A systematic review," Computers and Electronics in Agriculture, vol. 185, Art. no. 106125, 2021, doi: 10.1016/j.compag.2021.106125.

[9] J. Liu and X. Wang, "Plant diseases and pests detection based on deep learning: A review," Plant Methods, vol. 17, Art. no. 22, 2021, doi: 10.1186/s13007-021-00722-9.

[10] L. C. Ngugi, M. Abelwahab, and M. Abo-Zahhad, "Recent advances in image processing techniques for automated leaf pest and disease recognition—A review," Information Processing in Agriculture, vol. 8, no. 1, pp. 27–51, 2021, doi: 10.1016/j.inpa.2020.04.004.

[11] A. S. Kini, K. V. Prema, and S. N. Pai, "Early stage black pepper leaf disease prediction based on transfer learning using ConvNets," Scientific Reports, vol. 14, 2024, doi: 10.1038/s41598-024-51884-0.

[12] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, "Searching for MobileNetV3," in Proc. IEEE/CVF International Conference on Computer Vision (ICCV), pp. 1314–1324, 2019, doi: 10.1109/ICCV.2019.00140.

[13] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. 36th Int. Conf. Machine Learning, vol. 97, pp. 6105–6114, 2019.

[14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016.

[15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. International Conference on Learning Representations (ICLR), 2015.

[16] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 4510–4520, 2018, doi: 10.1109/CVPR.2018.00474.
[17] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 4700–4708, 2017.
[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems, vol. 25, pp. 1097–1105, 2012.
[19] O. Russakovsky et al., "ImageNet large scale visual recognition challenge," International Journal of Computer Vision, vol. 115, no. 3, pp. 211–252, 2015.
[20] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 1251–1258, 2017.
[21] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE International Conference on Computer Vision (ICCV), pp. 618–626, 2017, doi: 10.1109/ICCV.2017.74.
[22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. International Conference on Learning Representations (ICLR), 2015.
[23] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," Journal of Machine Learning Research, vol. 15, pp. 1929–1958, 2014.
[24] A. Dosovitskiy et al., "An image is worth 16×16 words: Transformers for image recognition at scale," in Proc. International Conference on Learning Representations (ICLR), 2021.
[25] V. Bischoff, K. Farias, J. P. Menzen, and G. Pessin, "Technological support for detection and prediction of plant diseases: A systematic mapping study," Computers and Electronics in Agriculture, vol. 181, Art. no. 105922, 2021, doi: 10.1016/j.compag.2020.105922.
[26] I. Bouacida, B. Farou, L. Djakhdjakha, H. Seridi, and M. Kurulay, "Innovative deep learning approach for cross-crop plant disease detection: A generalized method for identifying unhealthy leaves," Information Processing in Agriculture, vol. 12, no. 1, pp. 54–67, 2025, doi: 10.1016/j.inpa.2024.03.002.
[27] P. K. Das, S. S. Rupa, S. Pumrin, U. C. Das, and M. K. Hossen, "Deep Learning for Plant Disease Detection and Classification: A Systematic Analysis and Review," Current Applied Science and Technology, vol. 24, no. 4, e0259016, 2024, doi: 10.55003/cast.2024.259016.
[28] M. L. Buchaillot et al., "Framework for deep learning diagnosis of plant disorders in horticultural crops: From data collection tools to user-friendly web and mobile apps," Ecological Informatics, vol. 84, p. 102900, 2024, doi: 10.1016/j.ecoinf.2024.102900.
[29] P. Alirezazadeh, M. Schirrmann, and F. Stolzenburg, "Improving deep learning-based plant disease classification with attention mechanism," Gesunde Pflanzen, vol. 75, pp. 49–59, 2023, doi: 10.1007/s10343-022-00796-y.
[30] A. Pandey and K. Jain, "A robust deep attention dense convolutional neural network for plant leaf disease identification and classification from smart phone captured real world images," Ecological Informatics, vol. 70, Art. no. 101725, 2022, doi: 10.1016/j.ecoinf.2022.101725.