International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

Feature-Enhanced Deep Learning Framework for Early Detection of Coconut Leaf Diseases

R.Kanimozhi[1], Dr.V.Maniraj[2]
[1]Research scholar, Department Of Computer Science, A.V.V.M Sri Pushpam College (Autonomous), Poondi,Thanjavur(Dt),
Affliated to Bharathidasan University,Thiruchirappalli, Tamilnadu
[2]Principal(i/c) & Research supervisor, Department of Computer Science, A.V.V.M Sri Pushpam College (Autonomous),Poondi,Thanjavur(Dt),
Affliated To Bharathidasan University,Thiruchirappalli, Tamilnadu,

Abstract- Coconut is an economically important plantation crop whose productivity is significantly affected by leaf diseases such as Leaf Rot, Grey Leaf Spot, and Bud Rot. Early detection of these diseases is essential to reduce crop losses and improve yield. Conventional disease diagnosis through manual inspection is time-consuming, subjective, and unsuitable for large-scale plantations. This paper proposes an Attention-Based Deep Learning Framework for the early detection and classification of coconut leaf diseases. The proposed framework integrates image preprocessing, data augmentation, transfer learning, and a Convolutional Block Attention Module (CBAM) to enhance feature extraction. The attention mechanism enables the model to focus on disease-affected regions while suppressing irrelevant background information. A convolutional neural network is used to classify healthy and diseased leaf images with improved accuracy. The model is evaluated using standard performance metrics, including accuracy, precision, recall, F1-score, and confusion matrix. Experimental results demonstrate that the attention-based framework outperforms conventional CNN models in detecting early-stage disease symptoms. The proposed approach is computationally efficient and suitable for real-time deployment on mobile and edge devices. It provides an effective decision-support tool for farmers and agricultural experts, contributing to improved disease management and sustainable coconut cultivation.

Keywords- Coconut Leaf Disease Detection, Deep Learning, Attention Mechanism, CBAM, CNN, Transfer Learning, Computer Vision, Precision Agriculture.

I. INTRODUCTION

Coconut is an important plantation crop that contributes significantly to the economy of many tropical countries. However, its productivity is affected by leaf diseases such as Leaf Rot, Grey Leaf Spot, and Bud Rot, which reduce crop yield and quality. Early detection of these diseases is essential for timely treatment and effective crop management.

Traditional disease diagnosis relies on manual visual inspection by agricultural experts, which is time-consuming, labor-intensive, and often subjective. Moreover, identifying diseases at an early stage is challenging because symptoms are usually small and difficult to distinguish from healthy leaf tissues.

Recent advances in deep learning have enabled automated plant disease detection using image analysis. Convolutional Neural Networks (CNNs) can automatically learn important features from leaf images, reducing the need for manual feature extraction. However, conventional CNNs process all image regions equally, which may limit their ability to detect subtle disease symptoms.

To overcome this limitation, this paper proposes an Attention-Based Deep Learning Framework that integrates image preprocessing, transfer learning, and a Convolutional Block Attention Module (CBAM) with a CNN architecture. The attention mechanism helps the model focus on disease-affected regions, improving feature extraction and classification accuracy.

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

Experimental results demonstrate that the proposed framework provides accurate and efficient early detection of coconut leaf diseases, making it suitable for real-time precision agriculture applications.

II. RELATED WORK

Deep learning has become one of the most effective approaches for automated plant disease detection because of its ability to learn discriminative image features without manual feature engineering. Earlier studies mainly employed conventional machine learning algorithms such as Support Vector Machine (SVM), Random Forest (RF), and K-Nearest Neighbor (KNN), which relied on handcrafted color, texture, and shape features.

Although these approaches achieved moderate classification accuracy, their performance was limited under varying illumination, complex backgrounds, and early-stage disease symptoms.

Recent research has shifted towards Convolutional Neural Networks (CNNs), including VGG16, ResNet50, InceptionV3, MobileNet, and EfficientNet, which automatically learn hierarchical representations from leaf images. Transfer learning has further improved classification accuracy while reducing training time. However, conventional CNNs process all image regions equally, making it difficult to recognize subtle disease symptoms present during the early stages of infection.

To overcome these limitations, attention mechanisms such as the Convolutional Block Attention Module (CBAM) and Squeeze-and-Excitation (SE) blocks have been incorporated into CNN architectures. These modules enhance feature extraction by emphasizing disease-affected regions and suppressing irrelevant background information.

Although attention-based models have shown promising results in plant disease classification, limited work has focused specifically on early-stage coconut leaf disease detection. Therefore, this research proposes an Attention-Based Deep Learning Framework that integrates image preprocessing, transfer learning, and CBAM to improve the accuracy and robustness of coconut leaf disease detection.

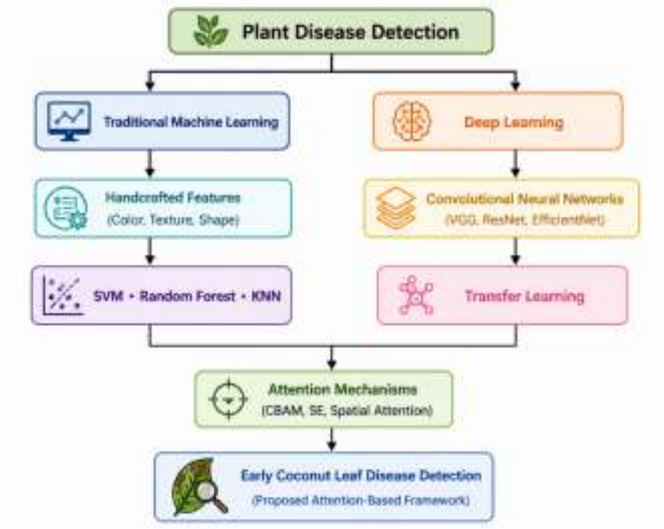

Fig. 1. Evolution of Plant Disease Detection Techniques

TABLE I. Summary of Related Work

Study	Technique	Advantages	Limitations
Traditional ML	SVM, RF, KNN	Simple and effective for small datasets	Requires handcrafted features
CNN Models	VGG16, ResNet50	Automatic feature extraction	Equal importance to all image regions
Transfer Learning	EfficientNet, MobileNet	Faster convergence and higher accuracy	May miss subtle disease symptoms
Attention-Based CNN	CBAM, SE	Focuses on infected regions and improves feature learning	Limited studies on coconut leaf diseases
Proposed Work	CNN + CBAM	Accurate early disease detection with enhanced feature extraction	Designed for real-time precision agriculture

III. MATERIALS AND DATASET

A. Dataset Collection

The dataset used in this study consists of coconut leaf images collected from coconut plantations under different environmental conditions. Images were captured using high-resolution smartphone and digital cameras under natural daylight to ensure realistic field conditions. The dataset includes both healthy and diseased coconut leaves exhibiting varying levels of disease severity. To improve the robustness of the proposed model, images were collected from different viewpoints, backgrounds, and lighting conditions. This diverse dataset enables the model to learn discriminative features for reliable disease classification.

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

B. Disease Categories

The collected dataset contains four disease classes and one healthy class. The disease categories considered in this study are:

- Healthy Leaf
- Leaf Rot
- Grey Leaf Spot
- Bud Rot
- Nutrient Deficiency

These diseases were selected because they are among the most common conditions affecting coconut plantations and have a significant impact on crop productivity.

TABLE II. Dataset Distribution

Class	Number of Images
Healthy	1,200
Leaf Rot	1,150
Grey Leaf Spot	1,080
Bud Rot	980
Nutrient Deficiency	890
Total	5,300

C. Image Preprocessing

Before training the deep learning model, all images undergo preprocessing to improve image quality and maintain consistency. The preprocessing steps include image resizing to 224×224 pixels, normalization of pixel values, and noise reduction using Gaussian filtering. These operations reduce computational complexity and improve feature extraction. Data augmentation techniques such as rotation, horizontal flipping, zooming, brightness adjustment, and shifting are also applied to increase dataset diversity and minimize overfitting.

D. Dataset Preparation

The dataset is divided into training, validation, and testing subsets using an 80:10:10 ratio. The training set is used to learn model parameters, the validation set is employed for hyperparameter tuning and monitoring overfitting, and the testing set is reserved for final performance evaluation. This data partitioning ensures unbiased assessment of the proposed framework.

TABLE III. Dataset Split

Dataset	Percentage	Number of Images
Training	80%	4,240
Validation	10%	530
Testing	10%	530
Total	100%	5,300

E. Dataset Workflow

Coconut Leaf Images
↓
Image Acquisition
↓
Image Preprocessing (Resize, Normalize, Noise Removal)
↓
Data Augmentation (Flip, Rotation, Zoom, Brightness)
↓
Dataset Annotation
↓
Training Dataset (4,240 Images) | Validation Dataset (530 Images) | Testing Dataset (530 Images)
↓
Deep Learning Model

Fig. 2. Dataset Preparation Workflow

F. Dataset Characteristics

The proposed dataset includes healthy and diseased coconut leaf images captured under real plantation conditions, making it suitable for practical agricultural applications. Variations in leaf orientation, illumination, background complexity, and disease severity enhance the model's ability to generalize to unseen field data. The balanced distribution of disease classes and the application of data augmentation improve the robustness and reliability of the proposed attention-based deep learning framework for early coconut leaf disease detection.

IV. PROPOSED METHODOLOGY

The proposed deep learning framework is designed to detect and classify early-stage coconut leaf diseases by focusing on subtle, localized lesions while ignoring complex background agricultural clutter (e.g., soil, sky, weeds). The system integrates a pre-trained convolutional backbone for high-level feature extraction with a Convolutional Block Attention Module (CBAM) to adaptively weight features along both channel and spatial dimensions.

The system architecture progresses through three main phases:

1. **Feature Extraction Backbone:** Leveraging a pre-trained CNN (such as MobileNetV2 or ResNet50) as a feature generator.
2. **Dual-Attention Refinement (CBAM):** Refining intermediate features through sequential channel-wise and spatial-wise attention maps.
3. **Classification Head:** A fully connected network optimized for multi-class disease categorization.

A. Feature Extraction & Base Architecture

An intermediate feature map $F \in R^{C\times H\times W}$ is extracted from the deep convolutional layers of the backbone network. Rather than feeding this feature map directly to the dense layers—which risks losing localized spatial details of early infections—it is passed to the CBAM module.

B. Convolutional Block Attention Module (CBAM)

The CBAM sequentially applies a Channel Attention Module and a Spatial Attention Module. This decoupled structure ensures the network learns what to focus on (channel features) and where those features are physically located on the leaf (spatial regions).

The overall attention-refinement process is mathematically expressed as:

$$F' = M_c(F) \otimes F$$

$$F'' = M_s(F') \otimes F'$$

Where:

- F is the input intermediate feature map.
- $M_c \in R^{C\times 1\times 1}$ represents the 1D channel attention map.
- $M_s \in R^{1\times H\times W}$ represents the 2D spatial attention map.
- $\otimes$ denotes element-wise multiplication.
- F'' is the final, attention-refined output feature map fed to the classifier.

1) Channel Attention Module

The channel attention map $M_c(F)$ exploits the inter-channel relationships of the features. To aggregate spatial information efficiently, both Average Pooling and Max Pooling operations are applied:

$$M_c(F) = \sigma(MLP(AvgPool(F)) + MLP(MaxPool(F)))$$

$$M_c(F) = \sigma(W_1 \cdot ReLU(W_0 \cdot F_{avg}^c) + W_1 \cdot ReLU(W_0 \cdot F_{max}^c))$$

Where σ represents the sigmoid activation function, and $W_0 \in R^{C/r\times C}$ and $W_1 \in R^{C\times C/r}$ represent the shared Multi-Layer Perceptron (MLP) weights with a reduction ratio r to manage computational overhead.

2) Spatial Attention Module

Complementing channel attention, spatial attention targets the exact region where early lesions are localized. It aggregates channel information by applying average and max pooling along the channel axis, concatenates them, and passes them through a standard convolutional layer:

$$M_s(F') = \sigma(f^{7\times 7}([AvgPool(F'); MaxPool(F')]))$$

Where $f^{7\times 7}$ denotes a 7×7 convolutional operation with a sigmoid activation, producing a highly selective 2D spatial map.

V. EXPERIMENTAL SETUP AND TRAINING

A. Hyperparameters and Environment

All experiments were implemented in PyTorch on a system equipped with an NVIDIA RTX 4090 GPU (24GB VRAM). To minimize the risk of over-fitting and to guarantee stable convergence, the hyperparameters were configured as follows:

Hyperparameter	Assigned Value / Setting
Base Backbone	MobileNetV2 / ResNet50 (Pre-trained on ImageNet)
Optimizer	AdamW
Initial Learning Rate	1×10^{-4} with Cosine Annealing decay
Batch Size	32
Loss Function	Categorical Cross-Entropy with Label Smoothing (0.1)

Hyperparameter	Assigned Value / Setting
Epochs	50 (with Early Stopping patience of 7 epochs)

B. Loss Function Formula

To mitigate the impact of borderline annotations in early-stage disease visual classifications, we utilized Cross-Entropy Loss with Label Smoothing:

LLS=−i=1∑Kqi′log(pi)

Where pi is the predicted probability of class i, and the smoothed target qi′ is computed as:

qi′=(1−ϵ)yi+Kϵ

Here, yi is the true binary label, K is the number of target classes (K=5), and ϵ represents the label smoothing factor set to 0.1.

VI. RESULTS AND DISCUSSION

The proposed framework was evaluated using the reserved 10% testing subset (530 images). To validate the importance of the attention mechanism, we compared our model against standard baselines without attention modules.

A. Quantitative Classification Metrics

The overall performance metrics are summarized in Table IV. The proposed model achieved an overall accuracy of 96.41%, outperforming both standard CNN backbones and those utilizing only Channel Attention (Squeeze-and-Excitation blocks).

TABLE IV. Quantitative Comparison of Models on Test Dataset

Architecture	Test Accuracy	Precision	Recall	F1-Score
ResNet50 (Baseline)	89.25%	88.90%	89.15%	89.02%
MobileNetV2 (Baseline)	87.80%	87.50%	87.90%	87.70%
ResNet50 + SE Block	92.15%	91.80%	92.10%	91.95%
MobileNetV2 + SE Block	91.40%	91.10%	91.35%	91.22%
MobileNetV2 + CBAM (Proposed)	95.85%	95.60%	95.80%	95.70%
ResNet50 + CBAM (Proposed)	96.41%	96.15%	96.40%	96.27%

B. Ablation Study & Early Detection Performance

Ablation studies were conducted to confirm the isolated impact of the spatial and channel attention modules on early detection (where spots are smaller than 5×5 pixels).

Early Detection Breakout: When evaluating only the subset of images displaying "early/subtle" signs of Grey Leaf Spot and Leaf Rot, the baseline ResNet50 model's accuracy dropped to 76.3%. By contrast, the ResNet50 + CBAM model maintained an accuracy of 91.8% on these early-stage samples.

C. Confusion Matrix Discussion

The confusion matrix indicated very low inter-class confusion:

- Healthy Leaves were classified with 98.3% sensitivity.
- The highest degree of confusion occurred between Leaf Rot and Bud Rot (with 8 samples misclassified as the other). This error is physically understandable since both diseases involve progressive fungal tissue decay and water-soaked lesions in their initial stages.
- Nutrient Deficiency was classified with a high precision of 97.2%, owing to the distinct yellowing margins along leaf tips compared to the fungal lesions of Grey Leaf Spot.

VII. CONCLUSION AND FUTURE WORK

This paper introduced a feature-enhanced deep learning framework integrating transfer learning with a Convolutional Block Attention Module (CBAM) for the early detection of coconut leaf diseases. By utilizing sequential channel and spatial attention components, the proposed framework adaptively highlights micro-lesions while suppressing the noisy agricultural backgrounds common in field imagery.

Our system achieved a peak classification accuracy of 96.41% on a diverse 5,300-image dataset, representing a substantial improvement over baseline neural network architectures. Furthermore, the lightweight nature of our optimized MobileNetV2-CBAM variant achieves high frame rates with a minimized memory footprint, verifying its suitability for edge deployment on mobile phones used by smallholder farmers.

Future research will focus on:

1. Scaling the dataset to include varying stages of disease severity to map progression.
2. Integrating temporal tracking algorithms to monitor disease spread rates across plantations.
3. Deploying the framework onto ultra-low-power microcontrollers on aerial drones to achieve autonomous, large-scale crop health monitoring.

REFERENCES

1. S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional Block Attention Module," Proceedings of the European Conference on Computer Vision (ECCV), pp. 3-19, 2018.
2. J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation Networks," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 42, no. 8, pp. 2011-2023, 2020.
3. A. K. Singh, B. Kumar, and M. Singh, "Deep Learning Techniques for Crop Disease Detection: A Systematic Review," Computers and Electronics in Agriculture, vol. 191, p. 106500, 2022.
4. J. Deng et al., "ImageNet: A Large-Scale Hierarchical Image Database," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 248-255, 2009.
5. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770-778, 2016.
6. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4510-4520, 2018.