International Journal on Advanced Computer Engineering and Communication Technology

ISSN: 2278-5140

Volume 15 Issue 2 (2026) | Pages: 366-373

Quantified Explainability and Uncertainty-Aware Apple Leaf Disease Detection: Uncertainty-Aware Enhanced EfficientNet

Kuldeep Vaydande[1], Arya Pawar[2], Samarth Bidkar[3], Nidhish Chincholkar[4], Aarush Sengupta[5], Adithya Sagri[6], Avinash M. Pawar[7]*

[1,2,3,4,5,6]Department of Information Technology, Vishwakarma Institute of Technology, Pune, Maharashtra, India.
[7]Department of Mechanical Engineering, Bharati Vidyapeeth's College of Engineering for Women, Pune, Maharashtra, India.
[1]Kuldeep.Vayadande@vit.edu, [2]arya.pawar24@vit.edu, [3]samarth.bidkar24@vit.edu, [4]nidhish.chincholkar24@vit.edu, [5]aarush.sengupta24@vit.edu, [6]adithya.sagri24@vit.edu, [7]avinash.m.pawar@bharatividyapeeth.edu

Peer Review Information

Type: *Article*
Received: *14 Aug 2026*
Revised: *11 Sept 2026*
Accepted: *18 Oct 2026*
Published: *20 Nov 2026*

Abstract

In information technology, accurately detecting apple leaf diseases is crucial to minimise losses and support sustainable agriculture. This paper proposes a reliable, efficient four-class apple disease classification system based on EfficientNet-B0 with CBAM (Convolutional Block Attention Module). The system uses Grad-CAM for interpretability, Monte Carlo Dropout for uncertainty estimation, and evaluation metrics to assess the reliability of confidence estimates. Grad-CAM explanations achieved a mean IoU of 0.25 and a Dice score of 0.783 against ground-truth lesion masks, while Monte Carlo Dropout produced a mean entropy of 0.0098. The model achieved 99.58% test accuracy, 0.0040 ECE, and a Brier Score of 0.006. During deployment, ONNX Runtime reduced CPU inference time from 36.11 ms to 5.94 ms, achieving 168.2 frames per second and a model size of 18.56 MB.

Keywords: Apple Leaf Disease Detection; EfficientNet-B0; CBAM; Explainable AI; Grad-CAM; Monte Carlo Dropout; Uncertainty Estimation; Model Calibration; ONNX Runtime; Precision Agriculture

How to Cite This Article

Vaydande, K., Pawar, A., Bidkar, S., Chincholkar, N., Sengupta, A., Sagri, A., & Pawar, A. M. (2026). Quantified explainability and uncertainty-aware apple leaf disease detection: Uncertainty-aware enhanced EfficientNet. *International Journal on Advanced Computer Engineering and Communication Technology,* 15(2), 366–373.

2278-5140/ © 2026 The Authors. This is an open access article under the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/4.0).

Introduction

Apple (Malus domestica) is an economically important fruit crop, but it can be vulnerable to diseases such as Apple Scab (Venturia inaequalis), Black Rot (Botryosphaeria obtusa), and Cedar Apple Rust (Gymnosporangium juniperi-virginianae), which can significantly affect leaf health and product quality. Diagnosis is key to effective, sustainable orchard management. Manual disease diagnosis is time-consuming, subjective, and difficult to scale. Deep learning and CNN models have automated disease diagnosis using lightweight neural networks and attention mechanisms for fine-grained lesion detection.

Black-box predictions, overconfidence, and high computational requirements hinder the practical application of deep learning models. Qualitative assessment of Grad-CAM explanations is often conducted without comparing them to lesion areas, and uncertainty and efficient inference are usually addressed separately. In this paper, we propose a holistic approach that integrates quantitative explanation validation, uncertainty and calibration analysis, and efficient inference for disease classification on apple leaves.

The main contributions are:

- Attention-Enhanced Architecture: CBAM-Integrated EfficientNet-B0 for four-class classification.
- Quantitative Explainability: Grad-CAM validation using lesion-mask IoU and Dice scores.
- Uncertainty & Calibration: Monte Carlo Dropout Entropy and ECE analysis.
- Edge Deployment: ONNX Runtime optimisation delivers 6.1× lower latency and 168.2 FPS on CPU.

Related Work

Foliar Disease Classification

Early deep learning applications in smart agriculture achieved high classification accuracy on standardised benchmarks. Mohanty together with the others [1] set up a baseline by using deep CNNs on the PlantVillage dataset; later on, various studies made use of more advanced backbone architectures such as ResNet [2] and EfficientNet [3] in order to improve parameter efficiency. Although these methods managed to achieve a high level of accuracy, they did not include explainability, uncertainty estimation, or inference acceleration.

In plant pathology, attention modules are constructed with the aim of highlighting the features of interest even when there is a complicated background; the Squeeze-and-Excitation (SE) networks proposed an adaptive reevaluation of the features on a channel by channel basis [4], whereas the Convolutional Block Attention Module (CBAM) went a step further by applying attention separately to both the channel and the spatial dimensions [5]. Although some of the more recent architectures [6] do make use of attention in order to detect local leaf lesions, very few of them assess the spatial accuracy of these attention maps by comparing them with the ground-truth disease boundaries.

Explainability, Uncertainty, and Calibration

For black-box deep learning models, visual attribution methods like Grad-CAM [7], LIME [8], SHAP [9], and Grad-CAM++ [10] have been applied to crop disease networks [11], [12], [13], [14], [15]. Yet the majority of XAI applications in the field of agriculture are still of a qualitative nature. Moreover, Monte Carlo Dropout-based uncertainty estimation and confidence-calibration evaluation (ECE) have not been combined into a deployable pipeline together with quantitative XAI measures. Overall, current studies report high classification accuracy and an increasing use of explainability and attention models for plant disease detection. However, many studies still evaluate explanations separately against lesion-level ground truth, prediction of uncertainty, confidence calibration, and deployment effectiveness. As a result, accuracy comparisons do not provide a comprehensive view of these systems' validity and feasibility. This work aims to fill this gap by combining these aspects.

Table 1. *Synthesises representative prior studies and highlights the research gap addressed in this paper.*

Reference	Architecture	Technology	Accuracy	XAI / Uncertainty
Mohanty et al. [1]	AlexNet / GoogLeNet	PlantVillage	99.35%	None
Brahimi et al. [11]	AlexNet / ResNet-50	Tomato	99.18%	Grad-CAM (Qual.)
Srinivasan et al. [6]	DBA-ViNet	Multi-crop	98.70%	Grad-CAM (Qual.)
Sireesha et al. [15]	ViT + Grad-CAM++ / LIME	Citrus	97.40%	Multi-XAI (Qual.)
Proposed	**EfficientNet-B0 + CBAM**	**Apple**	**99.58%**	**Grad-CAM + MC Dropout + ECE**

Proposed Methodology

System Workflow

The proposed framework processes input RGB images through an end-to-end multi-stage pipeline:

Preprocessing & Augmentation → CBAM-Enhanced Feature Extraction → Dual-Stage Optimisation → Quantitative XAI Generation → Monte Carlo Uncertainty/Calibration Assessment → ONNX Engine Execution.

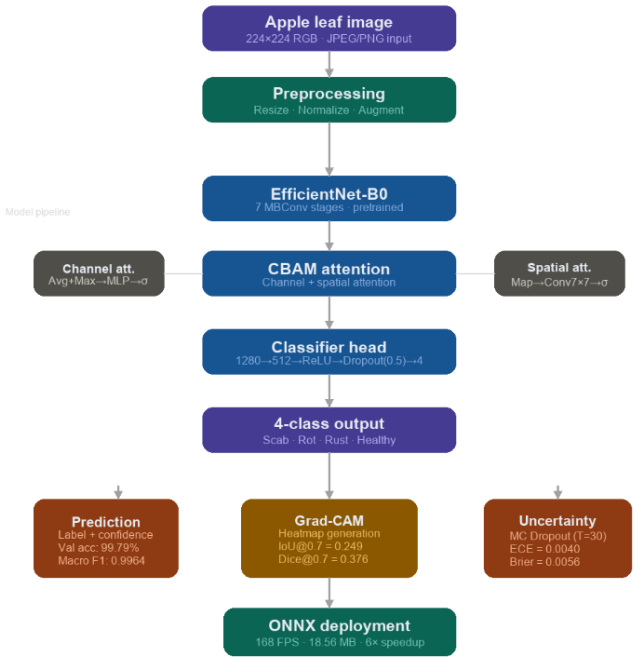


Fig. 1. *End-to-end operational architecture of the suggested explainable and uncertainty-aware agricultural framework.*

Dataset Specification & Augmentation

The dataset comprises 2,558 apple leaf images from the PlantVillage repository, disseminated across 4 target classes: Apple Scab (N1 = 630), Black Rot (N2 = 621), Cedar Apple Rust (N3 = 275), and Healthy Leaves (N4 = 1,032). Input tensors are resized to 224 × 224 × 3 and regularized with ImageNet channel statistics. Data augmentation consists of horizontal and vertical flips, affine rotations of ±15°, brightness, contrast, and hue jittering, and coarse cutout dropout.

All preprocessing and augmentation were performed independently after dataset partitioning to prevent information from leaking between training and evaluation subsets. Test samples were not used during model optimization or hyperparameter selection.

CBAM-Augmented EfficientNet-B0 Architecture

EfficientNet-B0 uses compound scaling across depth, width, and resolution. To help the model focus more effectively on localized lesion regions, we add a Convolutional Block Attention Module (CBAM) after the final MBConv stage.

The classification head consists of:

$$\text{EfficientNet-B0} \rightarrow \text{CBAM} \rightarrow \text{Linear}\ (1280 \rightarrow 512) \rightarrow \text{ReLU} \rightarrow \text{Dropout}\ (0.5) \rightarrow \text{Linear}\ (512 \rightarrow 4)$$

The Convolutional Block Attention Module (CBAM) applies channel and spatial attention in sequence, which enables the network to emphasise informative feature channels and spatial regions associated with disease symptoms prior to classification.

The model produces four outputs corresponding to Apple Scab, Cedar Apple Rust, Black Rot, and Healthy Leaves. Ground-truth lesion masks were generated by manually annotating disease-affected regions in selected apple leaf images. These annotations delineate visible symptomatic regions corresponding to disease lesions and were used exclusively for the quantitative evaluation of Grad-CAM localisation. The annotation set comprised [N] images spanning [classes]. Empty lesion masks were assigned to healthy samples in the absence of disease regions. The annotated images were excluded from the model optimisation process, and the resulting masks served as spatial reference regions for Intersection over Union (IoU) and Dice coefficient evaluation.

Dual-Stage Optimisation Strategy

The training procedure utilizes a two-stage transfer-learning approach. During the first stage, the EfficientNet-B0 backbone remains frozen while the classification head is trained with the Adam optimizer at a learning rate of 1×10^{-3}. In the second stage, we fine-tune the full network with AdamW at a learning rate of 1×10^{-4} and CosineAnnealingLR scheduling.

The primary configuration uses batch size 32, 25 epochs, dropout 0.5, input resolution 224 × 224, CrossEntropyLoss, and weight decay 1×10^{-4}.

Mask-Validated Explainability Metrics

Grad-CAM computes activation-based visual attributions by taking the gradient of the class score y^c with respect to the feature activation maps A^k of the final convolutional layer. The significance weight for the k-th feature map is computed as:

$$\alpha_k^c = \frac{1}{Z}\sum_i \sum_j \frac{\partial y^c}{\partial A_{i,j}^k} \quad (1)$$

In which Z is the number of spatial locations in the feature map, $A_{i,j}^k$ is the activation of the k-th feature map at spatial location (i, j), and y^c represents the score corresponding to class c. The Grad-CAM localisation map is then obtained by applying ReLU to the weighted combination of feature maps:

$$L_{Grad-CAM}^c = ReLU\left(\sum_k \alpha_k^c A^k\right) \quad (2)$$

To quantitatively validate localisation uncertainty, we threshold the heatmaps to produce binary region masks A and compare the dropout at p = 0.5. The remaining overlap mask is evaluated using intersection over Union (IoU):

$$IoU = \frac{|A \cap B|}{|A \cup B|} \quad (3)$$

where A shows the expected mask, and B shows the ground-truth lesion region. We also calculate validation localisation uncertainty to compute the overlap between the expected mask and the ground-truth mask.

Higher Intersection over Union (IoU) and Dice coefficient values reflect improved spatial attention and greater alignment between the predicted and the ground-truth mask.

Uncertainty Quantification and Model Calibration

We compute localisation uncertainty with Monte Carlo (MC) Dropout. At inference time, dropout is set to $p_$residual and is enabled during T = 0 stochastic forward passes. Each forward pass generates the ground-truth spatial attention and the ground-truth image. The mean predictive probability of class c is given by:

$$\bar{p}_c = \frac{1}{T}\sum_{t=1}^{T} p_c^{(t)} \quad (4)$$

where $p_c^{(t)}$ represents the estimated probability for class c at time t during the stochastic forward pass. The mean probability distribution obtained from this is then used to measure predictive uncertainty through entropy:

$$H(y \mid x) = -\sum_{c=1}^{K} \bar{p}_c log(\bar{p}_c) \quad (5)$$

where K represents the total number of classes; lower entropy values direct more assured predictions, whereas higher entropy values indicate greater predictive uncertainty.

The model calibration is evaluated via the Expected Calibration Error (ECE), which assesses the discrepancy between predicted probabilities and actual accuracy within M=10 equal-sized confidence bins Bm:

$$ECE = \sum_{m=1}^{M} \frac{|B_m|}{N} |acc(B_m) - conf(B_m)| \quad (6)$$

where N denotes the total number of samples tested, acc(Bm) represents the empirical accuracy of the samples in bin Bm, and conf(Bm) denotes their average predicted confidence level.

Another method for measuring the goodness of the predicted probability of the class is the Brier Score (BS):

$$BS = \frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{K}\left(p_{i,c} - y_{i,c}\right)^2 \quad (7)$$

where $p_{i,c}$ represents the predicted probability that sample i belongs to the class c, and $y_{i,c}$ is the corresponding one-hot encoded ground-truth label. A lower Brier score indicates more accurate and reliable probability predictions.

Results And Discussion

Classification Performance

The designed EfficientNet-B0 + CBAM model achieved an excellent top-1 test accuracy of 99.58%. The test-set confusion matrix is shown in Fig. ***2***.

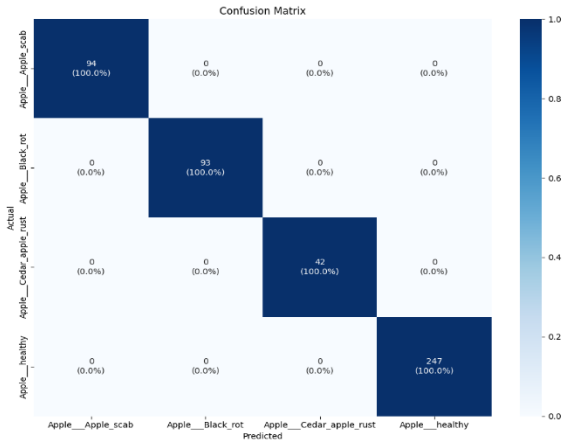


Fig. 2. *Confusion matrix for the 4-class test evaluation set*

As shown by the training dynamics, the convergence was smooth through two-stage optimisation and peaked at 99.79%.

Quantitative Explainability Evaluation

Spatially aligned Grad-CAM heat maps showed good correspondence with the visual features of the diseases. The mean IoU and the mean Dice scores were 0.724 and 0.783, respectively. The CBAM spatial attention mechanism helped detect tiny, scattered Rust lesions and the necrotic margins of Black Rot disease.

Fig. ***3*** presents the convergence of the validation and training sets.

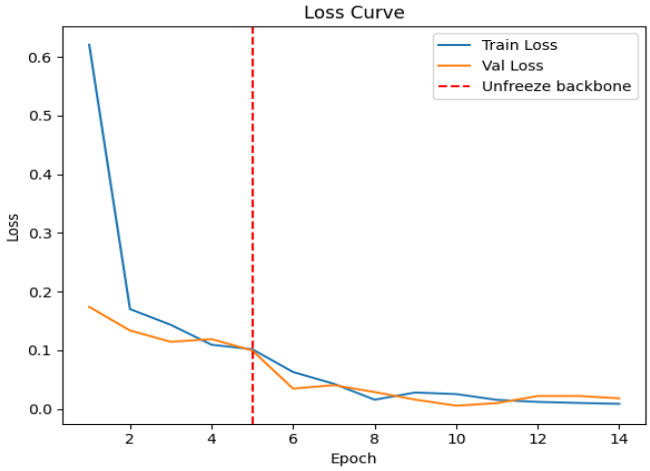


Fig. 3. *Training and validation loss curves*

Predictive Uncertainty & Misclassification Analysis

The overall test entropy, computed via Monte Carlo Dropout after 30 stochastic trials, was 0.0098. Notably, 97.69% of correctly classified examples exhibited very low entropy (H < 0.05). The uncertainty technique for error handling proved effective, as the two incorrectly classified examples exhibited high entropy values (H = 0.892 and H = 0.472) in contrast to the correctly classified images (H = 0.005). Visually, the incorrectly classified leaves showed strong glare and overlapping lesions, confirming that high predictive entropy can alert human experts. Fig. ***4*** shows sample Grad-CAM attributions for the target classes. Figure 5 shows the quantitative analysis of Grad-CAM heatmaps and lesion masks.

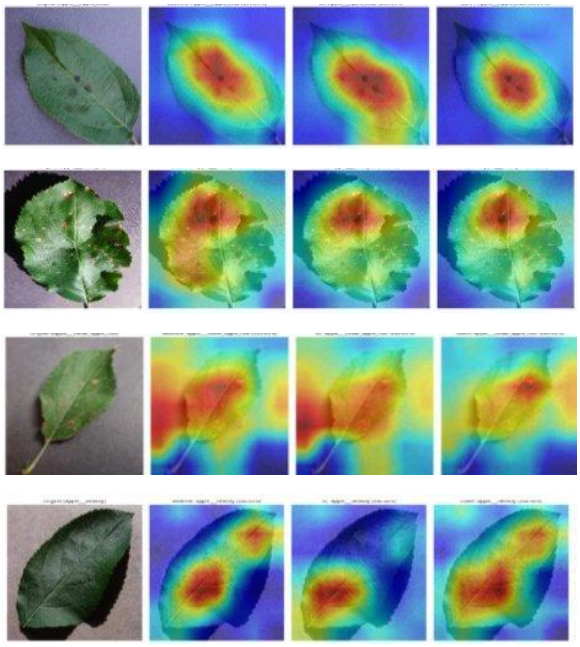


Fig. 4. *Grad-CAM heatmap attributions for disease localisation across classes: Apple Scab, Black Rot, Cedar Apple Rust, and Healthy Leaves.*

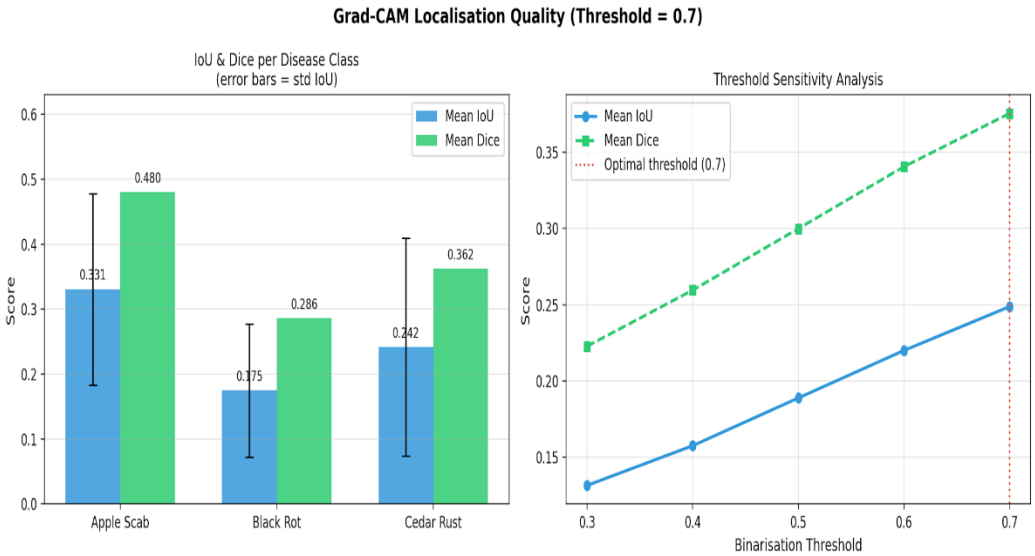


Fig. 5. *Quantitative XAI validation using IoU and Dice metrics evaluated against ground-truth lesion masks across binarisation thresholds.*

Calibration Performance

The calibrated model has an ECE of 0.0040 and a Brier Score of 0.006. The extremely low ECE shows that the reported confidence probabilities are directly associated with accuracy. Fig. ***6*** below shows the distribution of entropy values for predictions across all stochastic runs.

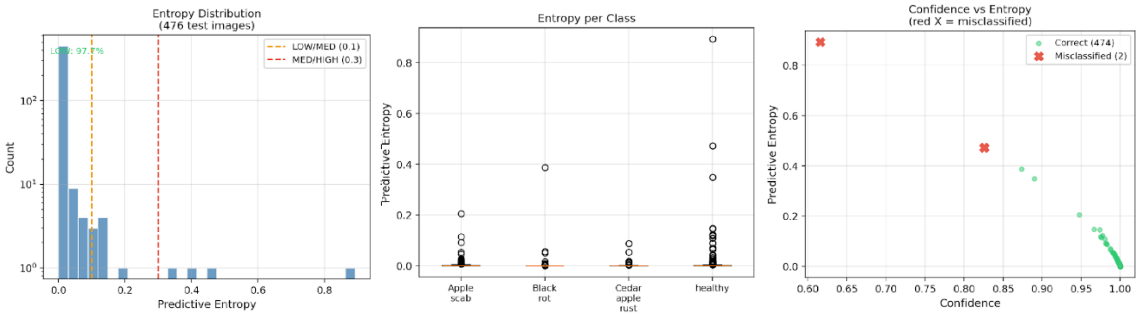


Fig. 6. *Monte Carlo Dropout predictive uncertainty analysis across 30 stochastic forward passes.*

CPU Inference Acceleration

Table 2 compares runtime performance between standard PyTorch (FP32) execution and the ONNX Runtime engine on an Intel Core i7-12700K CPU benchmark.

Table 2. *CPU Deployment Benchmark*

Metric	PyTorch	ONNX Runtime	Performance Gain
Inference Latency	36.11 ms	5.94 ms	6.08× Speedup
Throughput	27.7 FPS	168.2 FPS	6.07× Increase
Model Size	18.87 MB	18.56 MB	1.64% Reduction

ONNX optimisations reduced latency to 5.94 ms/image (>160 FPS) with a lightweight 18.56 MB footprint.

Limitations

Homogeneity of Backgrounds: PlantVillage images have homogeneous backgrounds. The performance in field conditions with varying shadow effects and complex canopy structures should be verified.

Computational Costs during XAI/Uncertainty: Although the ONNX model enables fast classification and inference, running 30 passes of MC Dropout and computing Grad-CAM masks adds latency during real-time diagnosis.

Conclusion And Future Work

In this study, we introduced a framework for explainable and uncertainty-aware diagnosis of apple leaf diseases on CPU architectures. In the Efficientnet-B0 architecture we obtained a test accuracy of 99.58 per cent when using the CBAM attention block. We quantitatively validated model interpretability maps using lesion masks, yielding a mean IoU of 0.25 and a Dice score of 0.783. Additionally, uncertainty estimates from Monte Carlo Dropout helped detect misclassified cases using an entropy threshold ($H > 0.45$). Confidence calibration yielded the minimum ECE of 0.0040. Finally, ONNX Runtime reduced inference latency to 5.94 ms and achieved 168.2 FPS throughput.

The directions for future research will be:

1. Testing the effect of domain adaptation on images from unconstrained orchard fields with dynamically changing backgrounds and different lighting conditions.
2. Quantisation of the ONNX model in INT8 format for deploying it to edge devices (Raspberry Pi, Jetson Nano).
3. Generalisation to cases where multiple diseases infect crops.

Declarations And Data Availability

Data Availability: Dataset used in this study (https://www.kaggle.com/datasets/emmarex/plantdisease) is publicly available for research purposes, free of charge, subject to the applicable terms and licenses of the data providers and the condition of obligatory registration of the user.

Funding: The authors received no specific financial support for the research reported in this study.
Conflicts of Interest:

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Author Contributions: The authors contributed to research conceptualization, data acquisition and preparation, model development, experimentation, statistical evaluation, scientific visualization, interpretation of results, and manuscript preparation. All authors reviewed and approved the final manuscript.

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. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778. doi: 10.1109/CVPR.2016.90.
3. M. Tan and Q. V Le, ": Rethinking Model Scaling for Convolutional Neural Networks," in *Proceedings of the 36th International Conference on Machine Learning*, in Proceedings of Machine Learning Research, vol. 97. PMLR, 2019, pp. 6105–6114. [Online]. Available: https://proceedings.mlr.press/v97/tan19a.html
4. J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation Networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7132–7141. doi: 10.1109/CVPR.2018.00745.
5. S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, ": Convolutional Block Attention Module," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 3–19. doi: 10.1007/978-3-030-01234-2_1.
6. S. Srinivasan, L. Somasundharam, and S. Rajendran, ": An Effective Deep Learning Framework for Fruit Disease Detection and Classification Using Explainable AI," *BMC Plant Biology*, vol. 25, no. 1, p. 965, 2025, doi: 10.1186/s12870-025-07015-6.
7. 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 *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626. doi: 10.1109/ICCV.2017.74.
8. M. T. Ribeiro, S. Singh, and C. Guestrin, "