

Mango Leaf Disease Detection Using Transfer Learning with EfficientNet-B0

Jodell R. Bulaclac¹; Joseph R. Del Carmen²

^{1,2} College of Information and Communication Technology, Nueva Ecija University of Science and Technology, Cabanatuan City, Nueva Ecija, Philippines

Publication Date: 2026/07/25

Abstract: Phytopathological threats to mango (*Mangifera indica* L.) cultivation cause severe global agricultural yield losses, necessitating rapid and accurate diagnostic frameworks. This study presents a highly optimized, end-to-end deep learning pipeline leveraging transfer learning via the EfficientNet-B0 architecture for multi-class mango leaf disease classification. The model was trained and rigorously evaluated on the MangoLeafBD dataset, comprising 4,000 balanced images across eight categorical states (one healthy control and seven distinct fungal and bacterial pathologies) captured under heterogeneous orchard conditions. Utilizing a stratified 70-15-15 data split, the network achieved exceptional convergence within 16 epochs. Empirical evaluation on an isolated test set of 600 images yielded an absolute classification accuracy of 100 percent, with precision, recall, and F1-scores of 1.0 across all classes and zero off-diagonal misclassifications. Furthermore, the pipeline demonstrated near-zero latency inference on standard edge-computing hardware. These findings validate the deployment viability of lightweight convolutional neural networks in resource-constrained agricultural environments, establishing a robust and computationally efficient baseline for automated precision pathology.

Keywords: Computer Vision, Deep Learning, EfficientNet-B0, Phytopathology, Precision Agriculture, Transfer Learning.

How to Cite: Jodell R. Bulaclac; Joseph R. Del Carmen (2026) Mango Leaf Disease Detection Using Transfer Learning with EfficientNet-B0. *International Journal of Innovative Science and Research Technology*, 11(7), 1065-1071. <https://doi.org/10.38124/ijisrt/26jul655>

I. INTRODUCTION

Mango (*Mangifera indica* L.) is one of the most culturally and economically significant tropical fruits cultivated globally. However, mango production is continuously threatened by a wide array of foliar diseases, including Anthracnose, Bacterial Canker, Powdery Mildew, and Sooty Mould. If left undetected and untreated, these pathogens can rapidly spread through orchards, leading to severe defoliation, diminished fruit quality, and catastrophic yield losses [1]. Traditionally, the identification and management of these diseases have relied heavily on direct visual inspection by agricultural experts and extension workers. This manual diagnostic process is not only labor-intensive and time-consuming but is also highly subjective and prone to human error, particularly in the early stages of infection when symptoms are visually ambiguous [2]. Furthermore, the acute shortage of trained phytopathologists in remote agricultural regions exacerbates the delay in applying targeted interventions.

In recent years, the advent of artificial intelligence, particularly deep learning and convolutional neural networks (CNNs), has revolutionized automated plant pathology and precision agriculture [3]. Deep learning models have demonstrated exceptional capability in extracting complex hierarchical features from image data, often matching or

surpassing human expert performance in classifying plant diseases under controlled conditions [4]. Early applications in agricultural computer vision heavily relied on deep, parameter-heavy architectures such as VGG16 and ResNet [5]. While these models achieved high classification accuracy, their significant computational overhead and memory requirements made them impractical for real-time deployment on mobile devices or edge-computing platforms. These lightweight platforms are essential for on-site diagnostic applications in remote orchards.

To bridge the gap between high accuracy and computational feasibility, researchers have increasingly turned to optimized architectures designed for mobile deployment. EfficientNet, introduced by Tan and Le, represents a major breakthrough in this domain by utilizing a compound scaling method that uniformly scales network width, depth, and resolution [6]. The baseline model, EfficientNet-B0, delivers state-of-the-art accuracy while requiring significantly fewer parameters and floating-point operations (FLOPs) compared to legacy architectures.

This research proposes an automated, end-to-end classification pipeline leveraging transfer learning with the EfficientNet-B0 architecture to categorize mango leaves into eight distinct states: Anthracnose, Bacterial Canker, Cutting Weevil, Die Back, Gall Midge, Healthy, Powdery Mildew,

and Sooty Mould. By utilizing the comprehensive MangoLeafBD dataset [7], which features images captured directly from diverse orchard environments via mobile phone cameras, this study aims to validate the efficacy of lightweight CNNs in distinguishing both healthy leaves from diseased ones, as well as accurately differentiating among complex, multi-class fungal and bacterial infections.

II. DATASET AND PREPROCESSING

➤ Data Acquisition and Characteristics

The study utilizes the MangoLeafBD dataset, which consists of 4,000 mango leaf images. All images were captured directly from mango trees using mobile phone cameras and are standardized to a resolution of 240x320 pixels in JPG format. To ensure environmental and lighting diversity, the data collection spanned four distinct mango orchards in Bangladesh: Sher-e-Bangla Agricultural University orchard, Jahangir Nagar University orchard, Udaypur village mango orchard, and Itakhola village mango orchard. Within the total corpus of 4,000 images, approximately 1,800 represent distinct, unique leaves. The remaining images were synthetically generated to enhance the

dataset's volume and robustness by applying transformations such as zooming and rotating where deemed necessary.

➤ Class Distribution and Applicability

The dataset categorizes the leaves into eight distinct classes. This includes one healthy category and seven disease categories: Anthracnose, Bacterial Canker, Cutting Weevil, Die Back, Gall Midge, Powdery Mildew, and Sooty Mould. The dataset is perfectly balanced, ensuring that each of the eight categories contains exactly 500 image instances. This uniform distribution and comprehensive categorization make the data highly suitable for both binary classification (distinguishing healthy from diseased leaves) and complex, multi-class disease prediction tasks.

➤ Data Splitting

To ensure a rigorous and unbiased evaluation of the EfficientNet-B0 model, the dataset was partitioned using a stratified split. The split configuration allocated 70% of the data (350 images per class) for training, 15% (75 images per class) for validation during the training process, and the remaining 15% (75 images per class) for final testing. A fixed random seed was utilized during this process to guarantee complete reproducibility of the subsets.




Fig 1 Test Support Per Class.

The Bar Chart Demonstrates a Uniform Distribution of Exactly 75 Test Samples for Each of the 8 Classes, Confirming the Integrity of the Stratified Test Split.

III. METHODOLOGY

➤ Research Design

This study employs a quantitative, experimental research design focused on evaluating the performance of a lightweight deep learning architecture for multi-class image classification. The core framework relies on transfer learning,

a technique that leverages the generalized feature-extraction capabilities of a model pre-trained on a massive dataset, such as ImageNet, and fine-tunes it for a specific, narrower domain [8]. This approach significantly reduces the computational burden and training time required compared to training a network from scratch, while also mitigating the risk of overfitting on localized datasets [9]. The design

incorporates a strictly controlled training pipeline where the independent variables are the image classes (seven diseases and one healthy state) and the dependent variables are the evaluation metrics (loss, accuracy, precision, recall, and F1-score) calculated across the predefined stratified splits.

➤ Method and Architecture

The chosen core architecture for this pipeline is EfficientNet-B0. Unlike traditional convolutional neural networks that scale dimensions (width, depth, or resolution) arbitrarily, EfficientNet utilizes a compound scaling method that uniformly scales all three dimensions using a fixed set of scaling coefficients [6]. This ensures maximum accuracy with a minimal number of parameters, making it highly suitable for agricultural edge-computing applications.

To adapt the model for this specific task, the pre-trained EfficientNet-B0 backbone was utilized as a feature extractor. The final fully connected layer of the original network was removed and replaced with a custom classification head, which was specifically designed to output probability distributions for the eight target mango leaf categories.

➤ Experimental Setup and Training Configuration

The model was developed and trained using the PyTorch deep learning framework (version 2.11.0+cu128) [10]. To ensure high-speed tensor operations, the training pipeline was accelerated using a dedicated NVIDIA GeForce RTX 3060 GPU. The training parameters were explicitly tuned to this hardware profile to prevent memory saturation while maintaining steady model convergence. The final hyperparameters used were:

- Epochs: 16
- Batch Size: 24
- Data Loader Workers: 2

During the training phase, the model's weights were iteratively updated to minimize classification error. An explicit checkpointing system was integrated into the training loop to monitor the validation accuracy at the end of every epoch. This system ensured that only the "best model" weights, representing the lowest validation loss and highest validation accuracy, were retained and saved to disk for the final evaluation phase [11].

IV. RESULTS AND DISCUSSION

➤ Training Dynamics and Convergence

The training process of the EfficientNet-B0 model was monitored over 16 epochs to evaluate its learning stability and convergence behavior. The performance metrics, specifically categorical cross-entropy loss and classification accuracy, were recorded for both the training and validation subsets at the end of each epoch to track the model's progression.

The model exhibited rapid and highly stable convergence. As illustrated in the loss and accuracy tracking, the training loss decreased sharply during the initial three epochs. This rapid decline was mirrored by a corresponding steep increase in classification accuracy. Importantly, the validation metrics tracked almost identically with the training metrics. By the fifth epoch, both training and validation accuracies stabilized near 100 percent, and the loss values approached zero. This accelerated learning capacity highlights the high efficiency of the pre-trained EfficientNet-B0 weights when fine-tuned for this specific agricultural domain. The pre-existing feature extraction capabilities of the model required minimal adjustment to accurately map the visual patterns of the mango leaf diseases.




Fig 2 Loss and Accuracy Curves.

The Training and Validation Metrics Track Closely Together, Indicating Rapid Convergence and Stable Learning without Early Divergence.

To further investigate the model's robustness and check for potential overfitting, the generalization gap was analyzed.

The generalization gap represents the mathematical difference between the training accuracy and the validation

accuracy across the epochs. A widening gap typically indicates that a model is beginning to memorize the training data rather than learning generalized features.

The analysis demonstrated that this gap remained exceptionally narrow throughout the entire 16-epoch training

cycle. As the training concluded, the gap approached near-zero variance. The absence of any significant divergence between the training and validation sets confirms that the model maintained strong generalization capabilities and did not overfit the 70 percent training data allocation.




Fig 3 Generalization Gap by Epoch.

The Gap Approaches Near-Zero as Training Concludes at Epoch 16, Confirming that the Model Successfully Abstracted the Data without Heavily Overfitting the Training Set.

➤ Classification Performance

Following the training phase, the optimized EfficientNet-B0 model was subjected to a rigorous evaluation using the completely isolated test set. The test set comprised a total of 600 images, distributed uniformly with exactly 75 images per class, providing an objective baseline to determine the model's true performance on unseen data.

The evaluation yielded an absolute classification accuracy across all categories. The model successfully classified every single test image into its true respective

category, resulting in zero misclassifications. To verify the class-level performance, a complete confusion matrix was computed. The matrix demonstrates that all 75 instances for each of the eight mango leaf states were correctly mapped to the true labels along the main diagonal. The off-diagonal quadrants, which represent false positives and false negatives, remained completely populated by zeros. This indicates that the network achieved perfect separability among closely resembling diseases, such as Anthracnose and Bacterial Canker, which often present overlapping visual lesions in field conditions.




Fig 4 Mango Leaf Disease Confusion Matrix.

The Diagonal Contains the Maximum Value of 75 for All Classes, with Absolute Zeroes in All Off-Diagonal Cells, Confirming Zero Misclassifications.

To provide a granular assessment of the model's predictive precision, the individual metrics of Precision, Recall, and F1-score were calculated for each class. Precision gauges the exactness of the classifier by measuring the ratio of true positives to total predicted positives, while Recall assesses the completeness by measuring the ratio of true positives to actual positives. The F1-score provides the harmonic mean of the two parameters, offering a single balanced measure of performance [4]. Because the model generated zero false positives and zero false negatives during the test phase, all three metrics achieved a perfect score of 1.0 across every individual category. The model showed no bias toward any particular disease type, demonstrating uniform feature mapping capabilities across both the healthy control samples and the seven distinct pathologies.




Fig 5 Per-Class Metrics.

Precision, Recall, and F1-Score are Consistently Maximized at 1.0 for Every Individual Mango Leaf Category.

Consequently, the aggregate macro-average and weighted-average scores for Precision, Recall, and F1-score also reached a flawless value of 1.0. The macro average treats all classes equally regardless of instance volume, whereas the weighted average scales the metric based on the

class distribution support. Because the MangoLeafBD dataset features a perfectly balanced distribution of test instances, both aggregate evaluation methods converge seamlessly, confirming the total effectiveness of the transfer learning framework under balanced class conditions.




Fig 6 Macro vs. Weighted Aggregate Metrics.

The Uniform Aggregate Scores of 1.0 Across Both Calculation Methods Reflect the Balanced Distribution and Perfect Classification Success on the Test Set.

➤ Discussion and Limitations

While achieving a perfect classification accuracy of 100 percent is statistically ideal, such flawless performance necessitates rigorous academic scrutiny. In computer vision tasks, an absolute 1.0 F1-score can sometimes indicate latent data leakage or a dataset lacking sufficient background and environmental variance [12]. The MangoLeafBD dataset provides a robust baseline; however, the highly controlled nature of the augmented leaf images may not fully represent the chaotic visual conditions of a live orchard.

Furthermore, evaluating computational efficiency is as critical as evaluating classification accuracy. While the initial training phase was accelerated using a dedicated GPU, subsequent inference testing and deployment simulations were conducted utilizing an AMD 5600GT processor. The results demonstrate that the EfficientNet-B0 model maintained exceptional efficiency, processing single-image predictions with near-zero latency on the AMD 5600GT. This confirms that the proposed pipeline does not rely on high-end hardware for execution and can be effectively deployed on standard commercial processors or agricultural edge devices.

To ensure the model is entirely robust for commercial field deployment, future iterations of this study must incorporate cross-validation techniques. Testing the model against external, out-of-distribution datasets captured under

varying lighting, weather, and camera conditions will be essential to validate its real-world generalization [13].

V. CONCLUSION

This study successfully implemented a highly accurate and computationally lightweight pipeline for detecting mango leaf diseases. By leveraging transfer learning with the EfficientNet-B0 architecture, the model achieved perfect convergence on the balanced 4,000-image dataset. Evaluation on an isolated 600-image test set resulted in an aggregate accuracy and F1-score of 1.0, with zero misclassifications across all eight categories. The system's ability to maintain high precision while functioning efficiently on standard hardware like the AMD 5600GT proves its viability as a baseline diagnostic tool for precision agriculture. Future research will focus on expanding dataset diversity and integrating explainable artificial intelligence techniques, such as Grad-CAM, to visually interpret the specific leaf lesions driving the network's classifications.

RECOMMENDATION

Based on the empirical findings and computational efficiency demonstrated by the proposed EfficientNet-B0 pipeline, several strategic recommendations are proposed to advance both the academic rigor and practical field viability of automated mango leaf disease detection systems.

➤ Dataset Expansion and Environmental Diversification

While the MangoLeafBD dataset provides an excellent, perfectly balanced baseline, future research must prioritize expanding the data to capture real-world operational challenges. It is highly recommended to collect multi-temporal images across different seasons, varying times of day, and diverse weather conditions to account for shifting illumination, shadows, and wet leaf surfaces [12]. Furthermore, integrating out-of-distribution datasets from geographically distinct mango-growing regions would test the model's true geographic robustness and prevent localization bias. Incorporating complex background noise, such as multi-layered canopy depth, weeds, and soil interference, will also ensure the system can operate seamlessly within messy field settings without requiring manual leaf segmentation.

➤ Advanced Architectural and Analytical Enhancements

To move past the limitations of an absolute 100 percent classification score, future iterations should explore architectural and training upgrades. Implementing k-fold cross-validation is highly recommended to establish a more statistically reliable performance baseline across multiple data permutations. Researchers should also integrate explainable artificial intelligence frameworks, specifically Gradient-weighted Class Activation Mapping (Grad-CAM), into the inference pipeline. This will allow pathologists to visually audit the network by generating heatmaps that pinpoint the exact morphological boundaries and lesion patterns driving the model's classification decisions. Additionally, exploring modern lightweight architectures, such as MobileNetV4 or Vision Transformers (ViTs) optimized for mobile systems, could yield deeper feature extractions while preserving edge-computing feasibility.

➤ Field Deployment and System Integration

Given that the optimized model achieved near-zero latency when running on standard processing hardware like the AMD 5600GT, the pipeline is highly recommended for direct edge deployment. Developers should focus on compiling the PyTorch model into optimized mobile runtimes, such as ONNX or TensorFlow Lite, to build a native mobile application for Android and iOS devices. This would empower local farmers and extension workers to perform real-time, on-site diagnostics without requiring a continuous internet connection. On a larger operational scale, this lightweight deep learning model can be integrated into edge-computing payloads mounted on Unmanned Aerial Vehicles (UAVs) or autonomous orchard-monitoring robots, enabling automated, high-throughput disease surveillance across expansive commercial mango plantations.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the Nueva Ecija University of Science and Technology (NEUST) for providing the necessary approval, resources, and institutional support to successfully conduct and complete this research.

REFERENCES

- [1]. [R. C. Ploetz, "Diseases of mango," in Diseases of Tropical Fruit Crops, CABI, 2003, pp. 327–363.
- [2]. A. Kamilaris and F. X. Prenafeta-Boldu, "Deep learning in agriculture: A survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.
- [3]. 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.
- [4]. J. G. A. Barbedo, "Factors influencing the use of deep learning for plant disease recognition," Biosystems Engineering, vol. 172, pp. 84–91, 2018.
- [5]. 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.
- [6]. M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proceedings of the 36th International Conference on Machine Learning (ICML), 2019, pp. 6105–6114.
- [7]. S. Ali, M. Ibrahim, S. I. Ahmed, M. Nadim, M. R. Mizanur, M. M. Shejuti, and T. Jabid, "MangoLeafBD Dataset," Mendeley Data, V1, 2022. doi: 10.17632/hxsnvwt3r.1.
- [8]. S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359, 2009.
- [9]. K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," Journal of Big Data, vol. 3, no. 1, pp. 1–40, 2016.
- [10]. A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Systems 32, 2019, pp. 8024–8035.
- [11]. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, MIT Press, 2016.
- [12]. 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.
- [13]. D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv preprint arXiv:1511.08060, 2015.