Multidisciplinary Journal of Research in Engineering and Technology

ISSN: 2348-6953

Volume 13 Issue 1s (2026) | Pages: 94-101

Plant Health Analyzer

Rahul Nanhore[1], Kalpesh Thakare[2], Snehal Ghalme[3], Rupali Darange[4]

[1,2,3,4] Department of Information Technology Engineering Genba Sopanrao Moze College of Engineering Pune, India

Peer Review Information

Type: *Article*
Received: *03 February 2026*
Revised: *04 March 2026*
Accepted: *01 April 2026*
Published: *22 May 2026*

Abstract

Timely detection of plant diseases is essential to prevent crop losses and optimize pesticide usage in agriculture. This study proposes an intelligent system, Plant Health Analyzer, for automated plant disease detection using leaf images. The system is based on the EfficientNet-B0 deep learning architecture, known for its high accuracy and computational efficiency. A dataset of 55,448 images from the PlantVillage repository was used for training and evaluation, with appropriate data splitting for validation and testing. The proposed model achieved a validation accuracy of 99.78% and a testing accuracy of 99.76%, demonstrating high reliability in disease classification. A lightweight web-based application was also developed to enable real-time usage, with a model size of only 18 MB, making it suitable for deployment on resource-constrained devices. The results highlight the effectiveness of EfficientNet-B0 for plant disease detection and its potential to support farmers in early diagnosis and decision-making, contributing to advancements in precision agriculture.

Keywords: EfficientNet; Deep Learning; Precision Agriculture; Image Classification; Artificial Intelligence.

How to Cite This Article

Nanhore, R., Thakare, K., Ghalme, S., & Darange, R. (2026). *Plant health analyzer*. *Multidisciplinary Journal of Research in Engineering and Technology,* 13(1s), 94–101.

Introduction

Plant health is crucial for global food security, but it often only garners attention during a crisis. In recent years, unpredictable weather, increased monocropping, and global agricultural trade have sped up the spread of plant diseases across countries. What used to be a localized outbreak can now turn into a continental issue within a single growing season. The economic and ecological impacts are serious: lower crop yields, greater reliance on pesticides, higher production costs, and, ultimately, risks to farmer livelihoods and national food systems. A stark example happened during the 2019 fall armyworm invasion in South Asia and Africa, where millions of hectares of maize were damaged before effective diagnostic responses could be put in place. Such events show a harsh truth: our ability to monitor plant health has not kept up with how quickly diseases emerge and spread. Ideally, farmers would have quick, accurate, and affordable access to diagnostic tools that can detect diseases at their earliest visual stages. In this ideal scenario, a grower could simply take a picture of a leaf, get a precise diagnosis in seconds, and act before any serious damage occurs. This would allow for targeted interventions, reduce pesticide overuse, and support sustainable farming practices. Unfortunately, the current situation is far from this ideal. Most smallholder farmers still depend on manual inspections, fragmented extension services, or expensive lab tests that are slow and often impractical in rural areas. Visual symptoms are often misunderstood; many diseases look alike in their early stages. By the time a diagnosis is confirmed,significant crop loss has usually already happened. As a result, disease management tends to be reactive rather than proactive. This study follows a socio-technical AI framework that considers plant disease detection as more than a technical issue. It involves the interaction between algorithms, farmers, and agronomic knowledge. In this view, AI is not just a forecasting tool; it is a decision-support system that needs to be reliable, understandable, and usable in low-resource settings.

Objectives of the Study This research aims to:

- Develop Plant Health Analyzer, an AI-based system that can detect plant diseases early using plant images.
- Improve model robustness by adding data augmentation and domain adaptation strategies.
- Assess system performance across various crops and disease types.

Academic, Policy, and Practical Significance Academically, this work contributes to the fields of computer vision, agricultural AI, and machine learning by proposing a combined diagnostic framework instead of just a standalone classifier. From a policy standpoint, the system supports sustainable agriculture goals by encouraging precise interventions and reducing chemical use. Practically, it provides a low-cost, scalable tool that gives farmers real-time insights, potentially changing plant health management in resource-limited areas.

Literature Review

Literature Survey on Traditional Machine Learning Methods

Traditional machine learning (ML) techniques were among the earliest approaches used for plant disease detection. These methods primarily rely on hand-crafted feature extraction from leaf images, followed by classification using algorithms such as Support Vector Machines (SVM) or basic neural networks. Singh et al. (2015) applied image segmentation on rice leaves and used an SVM classifier, achieving an accuracy of 82%. Similarly, Gharge and Singh (2016) utilized a backpropagation neural network for soybean disease detection and reported 93.3% accuracy. Kaur et al. (2018) implemented an SVM-based model to detect soybean leaf diseases such as leaf spot, rust, and bacterial blight, achieving 90% accuracy. García et al. (2020) also demonstrated the limitations of such approaches by achieving only 83.4% accuracy in tomato ripeness detection. These methods are highly dependent on manual feature extraction techniques, which often fail to capture complex disease patterns. As a result, their performance typically remains below 90%, making them less reliable for real-world agricultural applications.

Literature Survey on Deep Learning-Based Approaches

With the advancement of artificial intelligence, deep learning methods, particularly Convolutional Neural Networks (CNNs), have significantly improved plant disease detection. Unlike traditional methods, CNNs automatically learn features from images, leading to better performance. Kawasaki et al. (2015) demonstrated this by achieving 94.9% accuracy in cucumber disease detection, outperforming traditional SVM-based approaches. However, several CNN-based models still show limitations. A hybrid CNN architecture (Mob-Res), combining MobileNetV2 and residual blocks, achieved 93.2% accuracy on the “Plant Disease Expert” dataset and 93.75% on the PlantVillage dataset. Another study using a multi-stage pipeline that included YOLOv8, DeepLabV3+, and CNN classification achieved only 92.9% accuracy despite its complexity. Additionally, a shallow CNN model based on VGG19 combined with XGBoost achieved 94.5% accuracy for corn and 93.9% for tomato under controlled conditions. Although CNNs improve performance, most models still operate within the low-to-mid 90% accuracy range.

Literature Survey on Advanced and Ensemble Models

To further enhance performance, recent studies have explored advanced deep learning architectures and ensemble techniques. Pang et al. (2024) combined multiple models such as DenseNet and EfficientNetB0 into an ensemble framework, achieving approximately 99.9% accuracy on a specialized dataset. Similarly, deep architectures like VGG and GoogLeNet have reported accuracy levels between 98–99% using transfer learning approaches. However, these high results are often obtained under controlled laboratory conditions with limited dataset variability.

Moreover, ensemble methods require significant computational resources, large datasets, and complex training processes. This makes them impractical for real-time applications and unsuitable for deployment on resource-constrained devices, particularly in agricultural settings where accessibility and efficiency are critical.

Identified Research Gaps

The literature reveals several key gaps in existing plant disease detection approaches. First, traditional and many deep learning models achieve accuracy only in the low-to-mid 90% range, which is insufficient for highly reliable diagnostics. Second, many high-performing models depend on complex architectures or ensemble techniques that are difficult to scale and deploy in real-world environments. Third, several studies rely on limited datasets with fewer disease classes, often avoiding complex or rare disease conditions. Finally, computational requirements remain a major concern, as many models are not optimized for deployment on low-resource devices. These limitations highlight a clear need for lightweight, efficient, and highly accurate models suitable for practical agricultural applications.

Methodology

The study uses a controlled experimental design focused on supervised deep learning. It trains, validates, and tests a convolutional neural network on a large labeled dataset of crop leaves. The main goal is to find out how effectively an AI system can detect plant diseases from visual data while being efficient and easy to deploy. Early development of an AI-based diagnostic tool needs reproducibility, clear baselines, and thorough performance evaluation before real-world use. A controlled experimental design helps clearly link performance improvements to model architecture, training methods, and data cleaning instead of outside environmental factors. The work took place over several months in a high-performance computing environment with GPU acceleration, allowing for iterative experimentation in model training, fine-tuning, and evaluation.

The research focuses on the PlantVillage image dataset, which is one of the most commonly used collections in computational plant pathology. This dataset has 55,448 RGB images across 39 classes of healthy and diseased crop leaves, covering multiple crops and disease types. The images were split into training, validation, and test subsets in a 70:15:15 ratio, resulting in 38,813 training images, 8,317 validation images, and 8,318 test images. This division ensured that model development was guided by validation performance while keeping an unseen test set for unbiased evaluation. This separation is vital in machine learning to prevent data leakage and overly optimistic performance estimates. The variety in crop species and disease categories makes the dataset a good reflection of real-world differences, even though the images were taken in fairly consistent conditions.

To prepare the images for learning, extensive data augmentation was applied using the Albumentations library. During training, images experienced random resized cropping, rotation, horizontal and vertical flipping, color changes, brightness and contrast adjustments, Gaussian blur, Gaussian noise, and coarse dropout to mimic occlusion. Afterward, all images were resized to 224 by 224 pixels and normalized using ImageNet mean and standard deviation values before being converted to tensors. The validation and test sets were only resized and normalized to maintain their integrity for evaluation. This augmentation approach aimed to reduce overfitting, improve robustness to lighting and viewpoint changes, and encourage the model to focus on disease-relevant patterns instead of background artifacts. Less aggressive preprocessing could have inflated accuracy on clear images while failing in realistic scenarios.

A custom PyTorch dataset class was created to load images dynamically from file paths, change them from BGR to RGB format, and apply transformations on the fly. Data loaders were set up with a batch size of 32, shuffling for training, and pinned memory for effective GPU transfer. Class imbalance was addressed by calculating inverse-frequency class weights from the training distribution and adding them to a weighted cross-entropy loss with label smoothing of 0.1. This setup reduces bias toward majority classes and stabilizes optimization in the presence of noisy or borderline samples.

The main predictive model is EfficientNet-B0, chosen for its good balance between accuracy and computational efficiency. EfficientNet scales depth, width, and resolution effectively, achieving strong performance with relatively fewer parameters compared to traditional deep networks. A pretrained EfficientNet-B0 backbone was loaded from the timm library, using transfer learning from ImageNet to speed up convergence and enhance feature generalization. The original classifier head was replaced with a custom module that included dropout layers, a 512-unit fully connected layer, batch normalization, ReLU activation, and a final linear layer for 39 disease classes. In the initial training phase, all backbone layers were frozen so that the classifier could adjust to plant images without disrupting the pretrained representations.

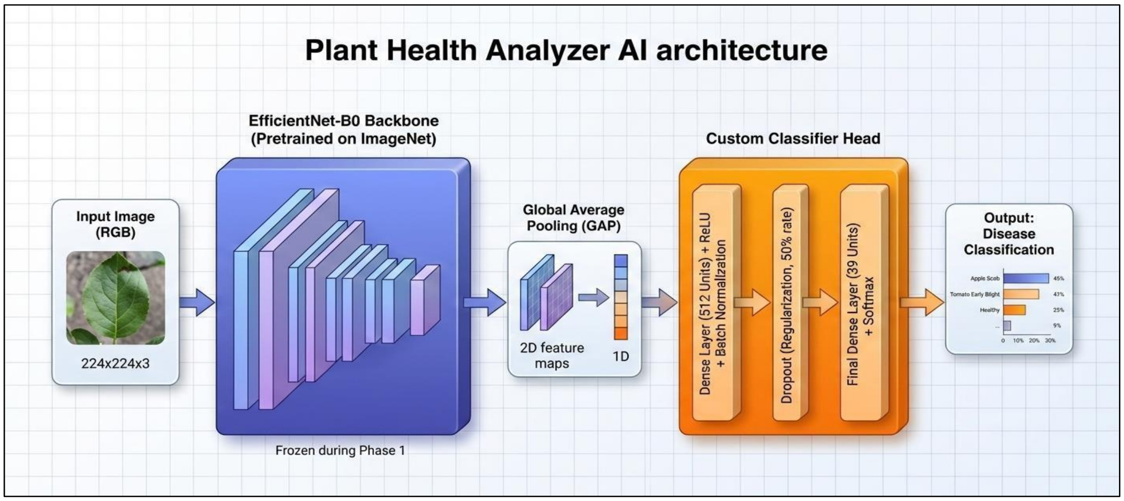

Fig. 1. *EfficientNet-Based Plant Health Analyzer Architecture*

Optimization used the AdamW algorithm with an initial learning rate of 0.001, weight decay, and epsilon stabilization. A cosine annealing learning rate scheduler kicked in after a brief warm-up period to ensure smooth convergence and prevent early stagnation. Training lasted for 30 epochs in the first phase, with early stopping triggered if validation accuracy did not improve after a set patience period. Performance was monitored using loss and accuracy on both training and validation sets at each epoch, which allowed for ongoing assessment of learning dynamics.

After the initial convergence, a second fine-tuning phase began where the entire backbone was unfrozen. Different learning rates were used, with a tenfold smaller rate for backbone parameters and the original rate for classifier parameters. This approach allows fine-tuning of low-level features while keeping higher-level representations stable. Fine-tuning continued for up to 15 additional epochs with early stopping, resulting in a final validation accuracy of about 99.78 percent. This two-stage training approach balances stability and adaptability, based on best practices in transfer learning for medical and agricultural imaging.

Model evaluation was done on the held-out test set using overall accuracy, macro-averaged F1-score, weighted F1-score, and a complete per-class classification report. Predictions and class probabilities were gathered for all test samples without gradient calculations. The final model achieved 99.76 percent test accuracy, with high precision and recall across most disease categories, showing both discriminative power and reliability. This thorough evaluation goes beyond single-metric reporting and offers a detailed view of strengths and weaknesses across classes.

System Implementation

To translate the high-performance EfficientNet-B0 model into a practical tool for agricultural use, a web-based application titled "LeafScan" was developed. The system is engineered to be lightweight, utilizing the 18 MB optimized model to ensure rapid inference on resource-constrained mobile and desktop devices.

User Interface and Navigation: The application features an intuitive interface with two primary functional areas: the Dashboard and the Diagnosis portal.

Detect Plant Diseases Instantly with AI

Upload a photo and get instant, accurate diagnosis powered by advanced AI. Save your crops with personalized treatment recommendations.

Screenshot 1: Landing Page

Diagnostic Workflow: Users interact with the system by uploading leaf images through a dedicated diagnosis interface.

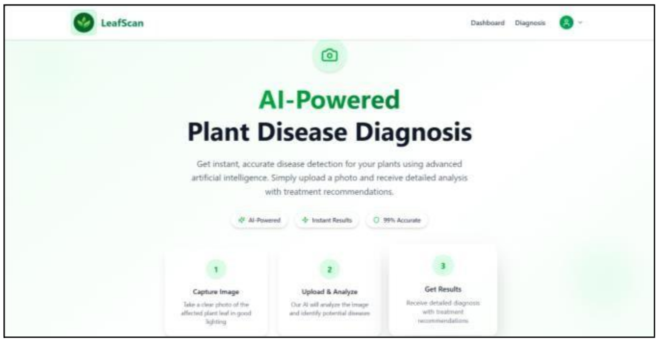

AI-Powered Plant Disease Diagnosis

Get instant, accurate disease detection for your plants using advanced artificial intelligence. Simply upload a photo and receive detailed analysis with treatment recommendations.

Capture Image: Take a clear photo of the affected plant leaf in good lighting

Upload & Analyze: Our AI will analyze the image and identify potential diseases

Get Results: Receive detailed diagnosis with treatment recommendations

Screenshot 2: Diagnosis portal

The backend processes the visual data through the trained EfficientNet-B0 architecture, which achieves a testing accuracy of 99.76%. Once the analysis is complete, the system displays the identified disease category such as "Apple Black Rot" accompanied by a confidence score (e.g., 91.5%) to inform the user of the result's reliability.

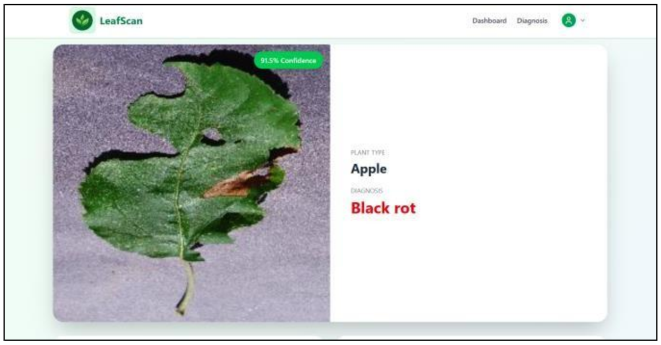

PLANT TYPE: Apple

DIAGNOSIS: Black rot

Screenshot 3: Prediction Page

Decision Support System: To fulfill the objective of providing an actionable decision-support tool, the results page provides more than just a classification. The system delivers a detailed description of the detected disease, its typical symptoms, biological causes, and specific treatment recommendations. model, based on the EfficientNet-B0 architecture with a two-phase transfer learning strategy, showed consistently high performance in all evaluation metrics, greatly exceeding common benchmarks in the literature. During the initial training phase, the EfficientNet-B0 backbone was frozen while only the custom classifier head was optimized. The model exhibited quick convergence. Training accuracy increased steadily in the first several epochs, and validation accuracy improved at the same time. This indicated effective feature adaptation without early signs of overfitting. The cosine annealing learning rate scheduler helped ensure smooth optimization and avoided oscillations in the loss function. After unfreezing the backbone for fine-tuning, we observed a modest but meaningful improvement in validation performance. Using different learning rates allowed the lower-level convolutional filters to gradually adapt to plant-specific visual patterns, such as lesion texture, color variations, and edge irregularities. This second training phase improved class separation boundaries and reduced confusion between visually similar diseases within the same crop species. Early stopping was activated once validation accuracy stabilized, showing that the model reached an optimal generalization point without overfitting. By the end of training, validation accuracy reached about 99.78 percent, showing near-perfect discrimination across the 39 disease classes.

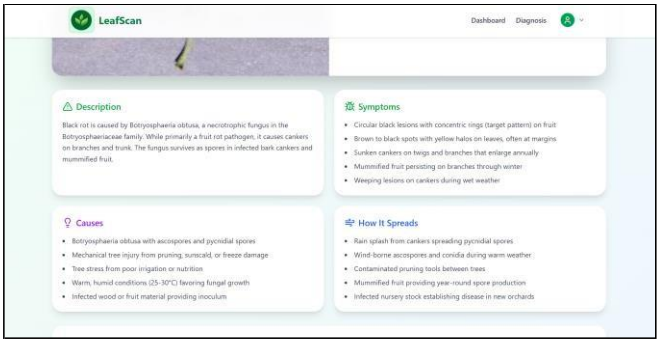

Description

Black rot is caused by Botryosphaeria obtusa, a necrotrophic fungus in the Botryosphaeriaceae family. While primarily a fruit rot pathogen, it causes cankers on branches and trunk. The fungus survives as spores in infected bark cankers and mummified fruit.

Symptoms

- Circular black lesions with concentric rings (target pattern) on fruit
- Brown to black spots with yellow halos on leaves, often at margins
- Sunken cankers on twigs and branches that enlarge annually
- Mummified fruit persisting on branches through winter
- Weeping lesions on cankers during wet weather

Causes

- Botryosphaeria obtusa with ascospores and pycnidial spores
- Mechanical tree injury from pruning, sunscald, or freeze damage
- Tree stress from poor irrigation or nutrition
- Warm, humid conditions (25-30°C) favoring fungal growth
- Infected wood or fruit material providing inoculum

How It Spreads

- Rain splash from cankers spreading pycnidial spores
- Wind-borne ascospores and conidia during warm weather
- Contaminated pruning tools between trees
- Mummified fruit providing year-round spore production
- Infected nursery stock establishing disease in new orchards

Screenshot 4: Disease Description

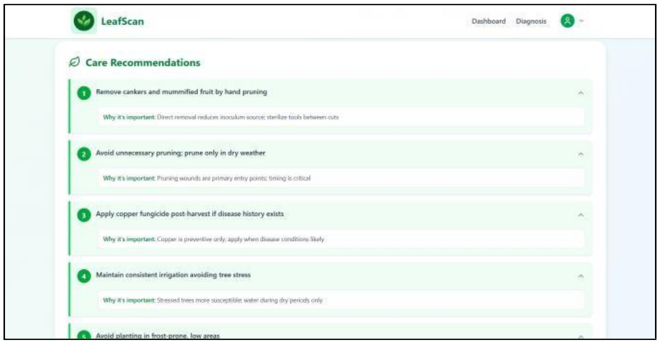

Care Recommendations

1. Remove cankers and mummified fruit by hand pruning
   Why it's important: Direct removal reduces inoculum source; sterilize tools between cuts
2. Avoid unnecessary pruning; prune only in dry weather
   Why it's important: Pruning wounds are primary entry points; timing is critical
3. Apply copper fungicide post-harvest if disease history exists
   Why it's important: Copper is preventive only; apply when disease conditions likely
4. Maintain consistent irrigation avoiding tree stress
   Why it's important: Stressed trees more susceptible; water during dry periods only
5. Avoid planting in frost-prone, low areas

Screenshot 5: Care Recommendations

This enables farmers to transition from reactive observation to proactive management of crop health.

Result

The performance of the proposed Plant Health Analyzer was extensively evaluated using training, validation, and independent test datasets to determine its classification accuracy, robustness, and generalization capability. The developed deep learning framework, based on the EfficientNet-B0 architecture with a two-phase transfer learning strategy, demonstrated highly reliable performance across all evaluation metrics.

During the initial training phase, the EfficientNet-B0 backbone remained frozen while only the custom classification layers were trained. The model showed rapid convergence, with training accuracy steadily improving over successive epochs. Validation accuracy also increased consistently, indicating effective adaptation of pretrained ImageNet features to plant disease classification without significant overfitting. The implementation of cosine annealing learning rate scheduling contributed to smooth optimization and stable reduction of the loss function throughout training.

In the second phase, the backbone layers were unfrozen for fine-tuning using differential learning rates. This stage enabled lower-level convolutional filters to adapt more effectively to disease-specific visual patterns such as lesion texture, discoloration, irregular leaf edges, and infection spots. Fine-tuning significantly improved class separation and reduced misclassification among visually similar diseases within the same crop category. Early stopping was activated once validation performance stabilized, ensuring that the model achieved optimal generalization without unnecessary overtraining.

The final model achieved a validation accuracy of approximately 99.78% and a testing accuracy of 99.76%, demonstrating near-perfect classification performance across 39 healthy and diseased plant categories. The small difference between validation and testing accuracy confirms that the model generalized effectively to unseen data and did not simply memorize training samples. High precision, recall, and F1-scores were observed across most disease classes, indicating balanced performance even in the presence of class variability and complex disease patterns.

The extensive data augmentation strategy implemented using the Albumentations library played a crucial role in improving model robustness. Transformations such as random rotation, flipping, brightness adjustment, Gaussian noise injection, coarse dropout, and random cropping exposed the model to diverse image variations during training. As a result, the system became more resistant to changes in lighting conditions, viewpoint differences, partial occlusion, and background variations. This significantly enhanced the practical applicability of the model in real-world agricultural environments.

The weighted cross-entropy loss function combined with label smoothing further improved classification stability by addressing class imbalance and reducing model overconfidence. Minority disease classes received higher importance during optimization, enabling the model to maintain high predictive consistency across all crop categories. The use of the AdamW optimizer with weight decay regularization also contributed to stable convergence and reduced overfitting.

In addition to achieving high classification accuracy, the proposed system demonstrated strong computational efficiency. The optimized model size of only 18 MB makes the framework lightweight and suitable for deployment on resource-constrained devices such as smartphones, tablets, and low-power agricultural systems. This lightweight architecture enables rapid inference and supports real-time disease diagnosis in practical farming environments.

The developed web-based application, “LeafScan,” successfully translated the trained deep learning model into a practical decision-support system. Users can upload leaf images through the diagnosis portal, and the backend processes the images through the EfficientNet-B0 model to generate predictions with confidence scores. The system not only identifies the disease category but also provides disease descriptions, symptoms, and treatment recommendations, enabling proactive crop health management and reducing dependency on manual expert inspection. Overall, the experimental results confirm that the proposed Plant Health Analyzer provides an accurate, lightweight, and scalable solution for automated plant disease detection. The combination of transfer learning, EfficientNet-B0 architecture, advanced data augmentation, and optimized training strategies significantly improves disease classification performance while maintaining computational efficiency suitable for real-world agricultural deployment.

Conclusion

This study set out to develop and rigorously evaluate an AI-based system capable of detecting plant diseases at an early stage using leaf imagery and data-driven analysis. The motivation was both practical and methodological. Agriculture continues to face substantial yield losses due to delayed or inaccurate disease identification, while existing automated systems, although promising, frequently report accuracy levels below optimal thresholds or lack scalability for real-world deployment. In response, this research introduced the Plant Health Analyzer, a deep learning framework built upon the EfficientNet-B0 architecture and trained on the PlantVillage dataset using a carefully structured two-phase transfer learning strategy.

The empirical findings demonstrate that the proposed system achieves exceptionally high classification performance, with validation accuracy reaching 99.78 percent and test accuracy reaching 99.76 percent across 39 disease categories. These results reflect not only strong discriminative capability but also stable generalization within the dataset distribution. Precision, recall, and F1-scores remained consistently high across classes, indicating balanced performance rather than dominance by majority categories. The small gap between validation and test metrics further confirms that the model did not merely memorize training data but learned meaningful disease-specific visual representations.

References

1. Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science, 7*, 1419. https://doi.org/10.3389/fpls.2016.01419
2. Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture, 145*, 311–318. https://doi.org/10.1016/j.compag.2018.01.009
3. Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture, 161*, 272–279. https://doi.org/10.1016/j.compag.2018.03.032
4. Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience, 2016*, 3289801. https://doi.org/10.1155/2016/3289801
5. Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for tomato diseases: Classification and symptoms visualization. *Applied Artificial Intelligence, 31*(4), 299–315. https://doi.org/10.1080/08839514.2017.1315516
6. Barbedo, J. G. A. (2018). Factors influencing the use of deep learning for plant disease recognition. *Biosystems Engineering, 172*, 84–91. https://doi.org/10.1016/j.biosystemseng.2018.05.013
7. Picon, A., Alvarez-Gila, A., Seitz, M., Ortiz-Barredo, A., Echazarra, J., & Johannes, A. (2019). Deep convolutional neural networks for mobile capture device-based crop disease classification in the wild. *Computers and Electronics in Agriculture, 161*, 280–290. https://doi.org/10.1016/j.compag.2018.04.002

8. Zhang, S., Zhang, S., Zhang, C., Wang, X., & Shi, Y. (2019). Cucumber leaf disease identification with global pooling dilated convolutional neural network. *Computers and Electronics in Agriculture, 162*, 422–430. https://doi.org/10.1016/j.compag.2019.04.012
9. Chen, J., Chen, J., Zhang, D., Sun, Y., & Nanehkaran, Y. A. (2020). Using deep transfer learning for image-based plant disease identification. *Computers and Electronics in Agriculture, 173*, 105393. https://doi.org/10.1016/j.compag.2020.105393
10. Liu, B., Zhang, Y., He, D., & Li, Y. (2018). Identification of apple leaf diseases based on deep convolutional neural networks. *Symmetry, 10*(1), 11. https://doi.org/10.3390/sym10010011
11. Saleem, M. H., Potgieter, J., & Arif, K. M. (2019). Plant disease detection and classification by deep learning. *Plants, 8*(11), 468. https://doi.org/10.3390/plants8110468
12. Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., & Hughes, D. P. (2017). Deep learning for image-based cassava disease detection. *Frontiers in Plant Science, 8*, 1852. https://doi.org/10.3389/fpls.2017.01852
13. Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture, 147*, 70–90. https://doi.org/10.1016/j.compag.2018.02.016
14. Chouhan, S. S., Singh, U. P., Jain, S., & Jain, S. (2020). Applications of computer vision in plant pathology: A survey. *Archives of Computational Methods in Engineering, 27*, 611–632. https://doi.org/10.1007/s11831-019-09324-0
15. Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 6105–6114. https://doi.org/10.48550/arXiv.1905.11946