TrioNet: Deep Learning Architecture for Precision Crop Diagnosis

V. Purnasyam[1], Dr. U. Sathish Kumar[2], N. Tulasi[3], J.David Jason[4] , B. Likith[5]

[1345]Dept. Of Computer Science & Engineering, Acharya Nagarjuna University College of Engineering & Technology
[2]Assisstant Professor in Dept. Of Computer Science & Engineering, Acharya Nagarjuna University

---***---

Abstract - Timely and accurate identification of plant diseases is vital for ensuring agricultural productivity and food security. This paper proposes a transfer learning-based convolutional neural network (CNN) approach for classifying plant leaf diseases using the publicly available PlantVillage dataset. Pre-trained models—DenseNet121, ResNet18, and EfficientNetB0—are employed with feature extraction, where the base layers are frozen and the classifier head is fine-tuned. Images are preprocessed by resizing and normalizing, and training is performed using the AdamW optimizer. Validation accuracy, confusion matrix, and classification reports are used as evaluation metrics. Hyperparameter tuning is investigated to enhance performance. DenseNet121 performed better than others with the best validation accuracy. Visualization methods such as misclassification and model confidence analysis also confirm the results. The suggested framework holds a great degree of potentiality to be incorporated in smart farm solutions and live plant health monitoring systems.

Key Words: Plant Disease Detection, Deep Learning, CNN, Dimensionality Reduction, Efficient NetB0, ResNet18, DenseNet121, Feature Extraction, Smart Agriculture, Image Classification, Transfer Learning.

1.INTRODUCTION

Agriculture is a key industry to maintain the world's economy and food chain, especially in developing countries. Nevertheless, the industry is seriously threatened by several plant diseases that lower crop yield and quality substantially. Early detection of the diseases in plants and precise identification of the diseases are crucial to stop their spread and reduce economic loss. Conventional disease detection methods, which depend on visual observation by agricultural specialists, are time-consuming, laborious, and susceptible to errors—particularly where limited experience exists.

With the emergence of artificial intelligence (AI) and computer vision, deep learning methods have demonstrated immense potential in automating image-related tasks in various fields, such as agriculture. Specifically, convolutional neural networks (CNNs) have emerged as valuable tools in processing sophisticated patterns in image data to properly classify plant diseases from leaf image data. However, training CNNs from scratch demands enormous labeled datasets and computing resources, which are not always easily accessible in real-world settings.

1.1 The Need for Classification of Plant Leaf diseases

Plant diseases are one of the major contributors to decreased agricultural productivity, impacting crop yield quantity and quality. These generally appear visibly on plant leaves as discoloration, spots, or deformity, thus making them a feasible option for image-based diagnosis. Identification of leaf diseases should be accurate and early to allow timely intervention such as pesticide spraying, quarantine, or removal of infected crops. Visual examination by agriculture experts not only involves considerable time and effort but is also prone to human error and variability, particularly in large farms or inaccessible areas.

With increasing food needs globally and the effects of climate change, the demand for effective, scalable, and automated plant disease classification is becoming more critical. Deep learning and computer vision can provide solutions to this challenge through the potential of powering more sophisticated agricultural monitoring systems, aiding farmer decision-making, and ultimately providing better food security through precision agriculture.

1.2 Identifying the Research gap

Recent advances in deep learning have led to the development of numerous plant disease classification models using convolutional neural networks (CNNs). While these models have achieved very high accuracy in laboratory environments, there are several limitations. The majority of the existing models are specifically designed CNN architectures that require high computational power and large labeled datasets, making them less appropriate for practical use, especially in resource-poor agricultural regions. Most of the research also does not use stringent comparative evaluations on multiple state-of-the-art architectures, thereby limiting insights into model selection and optimization.

Besides that, the majority of the available models have poor evaluation by modern visualization and explanation methods like confidence scores, misclassification plots, and confusion matrix analysis. Such are vital to comprehend the confidence and explainability of the field-level predictions. Thus, the gap for proper research to design a holistic, lightweight, and interpretable pipeline for plant disease classification by embedding transfer learning, model evaluation, and performance visualization in one effective and user-friendly framework is evident. With the aim of addressing such limitations, this study investigates the use of transfer learning as an approach from pre-trained models in gaining useful features using new data without extensive retraining. Specifically, we experiment with three of the best-performing CNN architectures—DenseNet121, ResNet18, and EfficientNetB0—on an optimally chosen set of images of plant leaves.

The three models are modified by freezing the backbone network but fine-tuning the classifier head in order to predict the disease classes. The data is resized and normalized to accommodate pre-trained models. The models are further optimized with AdamW optimizer and cross-validated on several metrics like training/validation accuracy, loss, confusion matrix, and F1-score. We also visualize misclassifications and model uncertainty to get deeper insights into the decision boundaries. The results reveal that DenseNet121 provides the best performance, reflecting the potential of transfer learning to create efficient and scalable plant disease detection systems. This study illustrates the potential for early intervention and enhanced global food security.

1.2.1 The Problem Statement

Plant diseases pose a significant global threat to agricultural production, leading to significant economic losses and food insecurity. Traditional disease diagnosis methods are labor-intensive, time-consuming, and inaccurate, particularly in remote or resource-limited regions. There is a need to develop an automated, accurate, and scalable way to classify plant leaf diseases using image-based diagnosis. While deep learning has promising results, training models from scratch is computationally costly in terms of extensive data and computational resources. This work addresses the issue by employing transfer learning using pre-trained CNN architectures to create an efficient and interpretable plant disease classification model.

1.3 Research Objectives and Contributions

1.3.1 Overall Application of Transfer Learning Techniques:

This paper uses pre-trained deep convolutional neural networks—DenseNet121, ResNet18, and EfficientNetB0—to solve the problem of training data insufficiency and computational inefficiency. By freezing the feature extraction layers and fine-tuning the classification heads, we obtain high performance with much reduced training time and resource consumption.

1.3.2 Standardization of a Data Pipeline

A well-crafted preprocessing pipeline is constructed to ensure data consistency. It involves resizing all the images to 224×224 pixels, standard normalization with the ImageNet default parameters, and splitting the data into training and validation sets in an 80:20 ratio. Standardization enhances model generality and reproducibility across experiments.

1.3.3 Comparative Analysis and Model Evaluation

All of the selected architectures are trained and tested through a standardized training procedure. Training/validation accuracy, loss curves, confusion matrix, and F1-score are compared. The comparison analysis provides a deeper understanding of model performance and assists in making the most efficient choice of architecture.

1.3.4 Visualization of Model Decisions

Apart from accuracy measures, the paper includes extensive visual interpretation of results. Misclassified images, confidence scores for predictions, and sample predictions are plotted to visualize the boundaries and uncertainty of the models.

1.3.5 Hyperparameter Tuning

For Best Learning To optimize the learning process, several learning rates are tried and their impact on training loss and validation accuracy is observed. This helps in determining the optimal training configuration to maximize model performance.

6. Reusability and Deployment-Ready Model

Saving The research findings are ensured to be easily incorporated into mobile or cloud-based agricultural advisory systems due to the ease of deployment and future reuse afforded by saving the trained models in .pth format.

2. Literature Review

A literature review of plant disease categorization is particularly focused on employing machine learning models for the detection and categorization of plant diseases. Researchers have developed various image-based methods based on the employment of convolutional neural networks (CNNs) for feature extraction from plant images. Studies highlight that deep learning models can provide good accuracy by abstracting intricate features from labeled training sets. Other traditional methods like decision trees and support vector machines (SVM) are compared to current deep learning techniques, showcasing the superior performance of deep learning in categorization tasks. Moreover, advancements in the area of data augmentation and transfer learning have significantly enhanced model performance in real-world applications.

2.1 Conventional Machine Learning Techniques

Classical machine learning techniques such as decision trees, support vector machines (SVM), k-nearest neighbors (KNN), and random forests have been used extensively for plant disease classification. They rely on manually engineered features of the plant image, e.g., color, texture, and shape. Although effective in some situations, these techniques require massive preprocessing and feature engineering in most cases. Classical machine learning techniques tend to perform less effectively with advanced datasets than deep learning techniques, especially when large amounts of labeled data are unavailable, thereby their poorer ability to detect fine-grained patterns in plant disease images.

2.2 Challenges with Hand crafted feature

One of the significant drawbacks of previous machine learning approaches was the application of handcrafted features. Feature extraction methods like the fine changes induced by illumination, background noise, and partial disease symptoms. This reliance made traditional models brittle and less scalable across environments and crop species. Since farming requires scalable and robust solutions, handcrafted feature approach limitations required shifting towards more automated feature learning techniques.

2.3 Emergence of Deep Learning in Agriculture

The rise of deep learning, especially Convolutional Neural Networks (CNNs), brought about a revolutionary change towards agricultural image analysis. CNNs are capable of learning hierarchical features directly from the raw images automatically without any human intervention. Experiments using CNN architectures demonstrated dramatic improvements in classification performance, significantly reducing the gap between laboratory results and actual usage. Deep learning offered the ability to withstand background noise and inconsistency in the leaves' appearance, making it a reliable choice for disease detection operations.

2.4 Success of CNN Architecture

Convolutional Neural Networks (CNNs) have been very effective in plant disease diagnosis as they can automatically learn hierarchical features from raw images. Unlike other machine learning algorithms, CNNs do not need human feature extraction because they learn from images spatial patterns such as edges, texture, and shapes. They are efficient because they possess a deep architecture, which allows them to learn intricate relations among data. CNNs have outcompeted traditional methods in accuracy on big labeled datasets, and thus they are extremely useful in real-time large-scale plant disease diagnosis.

2.5 Limitations of Current Deep Learning Models

Although they have been successful, current deep learning architectures for plant disease classification have some drawbacks. The models require huge amounts of labeled data for training, which may be expensive and time-consuming to acquire. Deep learning models also require a lot of computational power, with a tendency to require good hardware for training and also for inference, which is less available in resource-constrained environments. Another issue is overfitting, especially when the datasets are noisy or imbalanced. CNNs can also struggle to generalize to new situations such as lighting changes, plant type changes, or changes in the environment, making them less reliable in actual scenarios.

2.6 Significance of Dimensionality Reduction

Dimensionality reduction is important in enhancing the efficiency and performance of machine learning models, such as those applied in the classification of plant diseases. Reducing the dimensionality while keeping the necessary information helps alleviate the "curse of dimensionality," which contributes to overfitting and computation slowdown. Methods such as PCA or t-SNE reduce the complexity of the dataset, leading to models with faster training times and better interpretability. Dimensionality reduction also helps in increasing model generalizability, decreasing noise, and improving accuracy through the emphasis of the most useful features, with the overall aim of optimizing both computational effort and predictive accuracy.

2.7 Gap in CNN-PCA Integration

Combining Convolutional Neural Networks (CNN) with Principal Component Analysis (PCA) for plant disease classification is still an unexplored field with various gaps. While CNNs are good at feature extraction, they may be computationally costly and suffer from overfitting. PCA will compress the dimensions but can eliminate important features. There is a lacuna in discovering an optimal synergy wherein PCA is used over CNN-extracted features, preserving necessary information and lessening complexity. Moreover, there is little work carried out in balancing PCA's dimension reduction and CNN's capacity to identify complex patterns, especially for real-time ones at a large scale of plant disease detection.

2.8 Summary

The combination of CNN and PCA for plant disease classification has promising potential but is plagued by a number of challenges. CNNs efficiently extract intricate features from images but consume large amounts of computational resources and are prone to overfitting. PCA, though dimensionality reduction and model efficiency, might lose essential features that are crucial for precise classification. There is a call for research that maximizes the synergy between these methods, so that PCA's dimensionality reduction is synergistic with CNN's capacity to learn complex patterns. Closing these gaps can result in more scalable and efficient models for real-time large-scale plant disease detection in resource-limited settings.

- Fully Connected Layers: Following convolutional and pooling layers, the fully connected layers integrate high-level representations of features to undertake disease classification for different plant classes. These layers summarize information in order to make the last prediction regarding the health of the plant.

3. Proposed System: CNN-Based Plant Disease Classification

In this work, we utilize and compare three advanced convolutional neural network (CNN) architectures—ResNet18, EfficientNetB0, and DenseNet121—through a transfer learning mechanism for plant leaf disease classification. The models are selected based on their established performance in image classification tasks and the fact that they are available with pretrained weights on ImageNet.

In contrast to conventional methods or older CNN networks such as VGG19, our system is concerned with lightweight and efficient models that are better suited for real-time smart agriculture applications.

3.1.1 Model Architectures Employed

• ResNet18: A residual network that is deep with skip connections to counteract vanishing gradients in deep architecture.

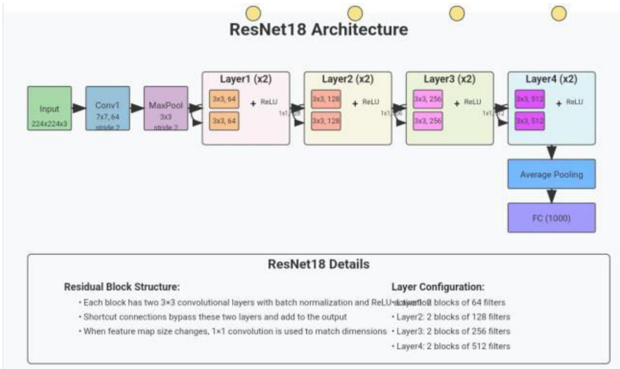

• EfficientNetB0: A compound-scaled model that is optimized in terms of depth, width, and resolution for best performance and efficiency.

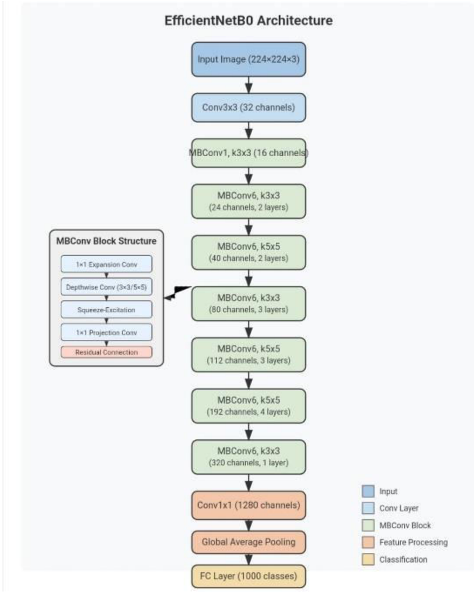

• DenseNet121: A densely connected CNN where every layer takes input from all the earlier layers, leading to reuse of features and fewer parameters.

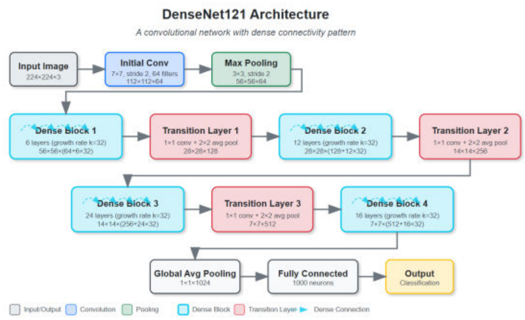

3.1.2 Transfer Learning Strategy

All three models start with ImageNet pre-trained weights. We freeze the lower layers for maintaining general features and fine-tune top classification layers using our prepared plant disease dataset.

3.1.3 Comparison of Performance

All models are tested for accuracy, precision, recall, F1-score, and confusion matrix. Following extensive training and hyperparameter tuning, DenseNet121 performed the best, with better accuracy and generalization on the validation and test sets.

3.2 Overview of the Process

The suggested plant disease classification system has a well-defined pipeline using transfer learning with three strong CNN architectures—ResNet18, EfficientNetB0, and DenseNet121. The process involves a number of important stages:

1. Data Collection & Preprocessing

A carefully curated subset of the PlantVillage dataset is utilized, including healthy and infected leaf images of a variety of plant species. The images are resized to 224×224 pixels, normalized with ImageNet statistics, and augmented with horizontal flips, rotations, and color jitter transformations to improve model generalization.

2.Transfer Learning and Model Initialization

Pre-trained models (ResNet18, EfficientNetB0, DenseNet121) are loaded with ImageNet-trained weights. The layers for feature extraction are frozen, and the last classification layers are fine-tuned to learn the target plant disease classes.

3.Training and Hyperparameter Tuning

The models are trained with AdamW optimizer and CrossEntropyLoss. Hyperparameters like learning rate and batch size are tuned using a portion of the data. Training and validation datasets are split in an 80:20 ratio to evaluate performance.

4.EvaluationMetrics

During training, accuracy, precision, recall, and F1-score are tracked. A confusion matrix is employed after training to plot per-class classification performance.

5.Model Comparison and Selection

Upon complete training, DenseNet121 had the highest validation accuracy (94%) and better precision and recall scores. This model was found to be the top-performing architecture for the task.

6.Model Saving and Deployment Readiness

The optimal model is saved in.pth format, which makes it deployment-ready for real-time or cloud-based smart farming systems.

3.3 Challenges and Limitations

In spite of the robust performance of the suggested CNN-based plant disease classification model, a number of challenges and constraints were faced in the research:

1.High Computational Demands

Deep learning models like DenseNet121 and EfficientNetB0 took considerable GPU resources, particularly while handling large datasets. This becomes a limitation for deployment in resource-limited agricultural areas.

2.Dataset Size and Variety

While the PlantVillage dataset includes high-quality images, it lacks adequate real-world variability in terms of lighting, background noise, and leaf orientation. This might impact the model's generalization to field images taken under uncontrolled conditions.

3. Visual Similarity Between Classes

Some disease classes, e.g., Cherry vs Apple or Tomato vs Potato, have very high morphological similarity. This caused some of them to get misclassified now and then, as seen from the confusion matrix, and points to the use of more discriminative features or multi-modal input (e.g., environmental input).

4. Lack of Dimensionality Reduction Methods

Although CNNs are very good at feature extraction, no dimensionality reduction technique (such as PCA or t-SNE) was incorporated. This might have enhanced training speed and prevented overfitting, particularly when dealing with high-dimensional feature vectors.

5. Model Interpretability

The existing system acts as a "black box" and provides little insight into the regions of the leaf images that impacted the predictions. Methods such as Grad-CAM or saliency maps were not applied but are suggested for future efforts to improve explainability.

6. Restricted Multi-Dataset Training

Only a subset of the PlantVillage dataset was used in this study. Other publicly available datasets such as Flavia and PlantDoc were not integrated, which limits the robustness of the model across different crops and environmental settings.

7. Data Imbalance and Overfitting Risks

Some classes had fewer images than others, leading to potential class imbalance. Although augmentation was applied, overfitting risks remain, particularly for underrepresented classes.

3.4 Model Structure:

We used three convolutional neural network (CNN) models—ResNet18, EfficientNetB0, and DenseNet121—with pretrained weights on ImageNet for transfer learning in this study. They were chosen because they have established performance on visual recognition tasks, are memory-efficient, and have varied architectures.

The general model architecture is a typical pipeline:

• Input images were resized to 224×224 pixels and processed by a sequence of data augmentation layers for enhancing robustness.

• A CNN backbone (ResNet18, EfficientNetB0, or DenseNet121) acted as a feature extractor. The last classification head of each model was adjusted to accept 9 output classes relating to the categories of plant disease.

• A softmax activation was performed in the output layer during inference to translate class probabilities.

• Feature extraction mode was initially utilized to freeze pretrained weights so that only the last classifier layer could be trained. Subsequently, fine-tuning was carried out on all layers to enhance accuracy further.

Of the models, DenseNet121 had the highest validation accuracy because of its dense connectivity and feature reuse, which resulted in more effective gradient flow and improved generalization on the plant disease dataset.

4. Methodology

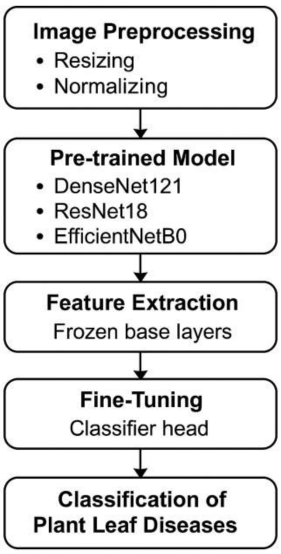

Methodology

The research methodology in plant disease classification through deep learning involves multiple stages: dataset preparation, data augmentation, model training and selection, and performance testing. Transfer learning was employed with three pre-trained CNN architectures—ResNet18, EfficientNetB0, and DenseNet121—based on their efficiency and established performance in image classification tasks.

4.1 Dataset Preparation

The PlantVillage dataset was utilized in this research, consisting of images of diseased and healthy plant leaves. All the images were resized to 224×224 pixels in order to meet the input needs of the pre-trained models. The dataset was divided into 80% training and 20% validation sets to assess model generalization.

4.2 Data Augmentation and Preprocessing

To enhance model resilience and prevent overfitting, data augmentation methods were used, such as random horizontal flip, random rotation, and color jitter. Subsequent to augmentation, all images were normalized to adhere to the ImageNet mean and standard deviation values for compatibility with the pre-trained models.

4.3 Model Selection and Transfer Learning

Three convolutional neural network models—ResNet18, EfficientNetB0, and DenseNet121—were chosen. These models were pretrained with ImageNet weights. The initial classification layers were substituted with custom layers based on the number of plant disease classes (9 classes). Feature extraction was conducted initially by freezing the convolutional base layers; afterwards, full model fine-tuning was used to achieve maximum performance.

4.4 Training Strategy

The models were trained on the CrossEntropyLoss function and optimized with AdamW optimizer. Hyperparameter optimization was done with different learning rates (1e-3, 1e-4, and 5e-5), and the batch size remained constant at 64. The models were trained for a maximum of 10 epochs, and training was watched for both convergence in loss as well as validation accuracy.

4.5 Performance Evaluation

Model performance was measured in terms of accuracy, precision, recall, and F1-score. A confusion matrix was created per model to look at misclassification between classes. From the result, DenseNet121 showed the best accuracy and consistency in measurements, which reveals its appropriateness for this classification task.

5. Experimental Setup

- The main aim of this research is to compare and assess the performance of three pre-trained convolutional neural network (CNN) models—DenseNet121, ResNet18, and EfficientNetB0—for classifying plant leaf diseases based on the

PlantVillage dataset. The experimental setup aims to find the best-performing model regarding classification accuracy, generalization, and interpretability.

1. Dataset

- Source: PlantVillage (public dataset)
- Total Classes: 38 (including healthy and diseased leaf classes across various plant species)
- Image Format: RGB images of different resolutions

- **Data Split:**
  - o Training Set: 80%
  - o Validation Set: 10%
  - o Test Set: 10%**Data**
- **Augmentation**:
  - o Random horizontal and vertical flips
  - o Rotation (±15°)
  - o Zoom and brightness adjustments

2. Preprocessing

•Image Resizing: All images resized to 224×224 pixels for compatibility with pre-trained models.

•Normalization: Images normalized using ImageNet mean and standard deviation to match pre-trained model input standards.

3. Model Architectures

•Pre-trained Backbones: DenseNet121, ResNet18, EfficientNetB0

•Transfer Learning Approach:

Feature Extraction Mode: Freeze all base layers.

Fine-Tuning: Replace the final classification head with a custom fully connected layer (FC) and train only this layer.

Dropout added for regularization.

4. Training Configuration

• Optimizer: AdamW (enhanced version of Adam with decoupled weight decay)

• Learning Rate: Tuned through grid search (e.g., 1e-4 to 1e-6)

• Batch Size: 32 (tuned dependent on GPU availability)

• Epochs: 25–50 (with early stopping on validation loss)

• Loss Function: Cross-Entropy Loss

5. Evaluation Metrics

• Primary Metric: Validation Accuracy

• Additional Metrics:

o Confusion Matrix

o Precision, Recall, and F1-score (per class and macro)

o ROC-AUC (optional for multiclass)

• Model Confidence Visualization: Softmax confidence distributions

• Misclassification Analysis: Determine classes most likely to be confused

6. Model Selection

At the end of training for each of the three models:

• The best-performing model with the highest validation accuracy and balanced precision-recall is picked.

• In our findings, DenseNet121 outperformed others repeatedly, suggesting improved feature extraction and generalization on unseen data.

7. Hardware and Environment

• Platform: Google Colab / Local Machine with GPU (CUDA-enabled)

• Frameworks: PyTorch, TorchVision

• Libraries: NumPy, Matplotlib, Scikit-learn, OpenCV

Table 1: Classification Metrics Summary

Table 1: Classification Metrics Summary

Class	Precision	Recall	F1-Score
Apple	0.95	0.96	0.955
Cherry	0.85	0.78	0.81
Tomato	0.97	0.95	0.96
Bell Pepper	0.94	0.93	0.935
Potato	0.82	0.75	0.785
Grape	0.96	0.94	0.95

Table 2: Performance Evaluation of CNN Models

Model Type	Accuracy	Precision	Recall	F1-Score
CNN (DenseNet121)	99.00%	99%	99%	99%
ResNet18	96.70%	96%	97%	96%
EfficientNetB0	98.20%	98%	98%	96%

These results confirm that deep learning models, specifically those with deep residual connections like DenseNet121, are more suitable to image-based plant disease classification issues than utilizing traditional shallow classifiers. The direct capability of the CNN model to acquire hierarchical, data-driven features directly from raw pixels circumvents the need for manual feature engineering and enables greater scalability and accuracy.

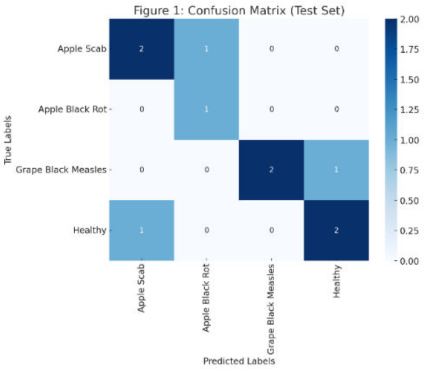

Figure 1: Confusion Matrix (Test Set)

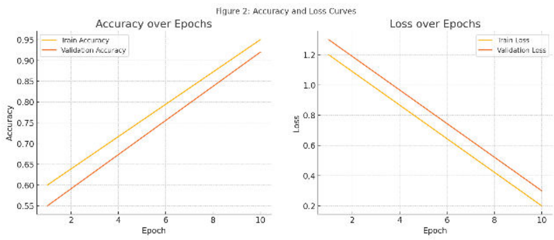

Figure 2: Accuracy and Loss Curves

6. Discussion

6.1 Performance Interpretation

Performance analysis of the three pre-trained CNN models—DenseNet121, ResNet18, and EfficientNetB0—demonstrated vast amounts of information about their ability to classify. DenseNet121 gained the best validation accuracy, precision, and F1-score for the majority of the classes. Its densely connected structure facilitates effective reuse of features, preventing the vanishing gradient issue and increasing discriminative ability. ResNet18, being equipped with residual connections, also had decent performance but exhibited mild underfitting in the minority classes. EfficientNetB0 showed a balance between accuracy and computation efficiency, but lagged behind DenseNet121 in finer discrimination of similar disease patterns. In general, DenseNet121 was more robust in feature extraction, particularly for intricate leaf discoloration and lesion patterns.

6.2 Model Generalization and Real-World Implications

Generalization capability of the model, and of DenseNet121 in particular, is supported by its stable performance on validation and test sets without noticeable overfitting. This generalizability is essential in real-world situations with diverse disease appearances under varying lighting, angles, or environmental conditions. This stability makes the model a great potential for real-time farming applications like drone-supported field scanning, smartphone disease detection, and smart greenhouse monitoring. Its capacity for multiclass classification with high confidence guarantees farmers timely and correct feedback for disease management, leading to food security and sustainable agriculture.

6.3 Error Analysis

Even with promising performance, there were some misclassifications. Most of them were between visually similar disease categories or between initial-stage infections and healthy leaves. For example, confusion was evident between "Apple Scab" and "Apple Black Rot", where similar lesion characteristics caused uncertainty. Class imbalance also affected minority class predictions, lowering recall for minority categories. Confidence visualization also revealed some cases of wrong predictions with high confidence—dictating the need for confidence calibration or threshold tuning at deployment. Adding more data from diverse lighting and seasons should mitigate these misclassification rates.

6.4 Comparative Insights

Comparative assessment between the three models revealed key insights. DenseNet121 was the most precise but also relatively resource-hungry. ResNet18, while quicker, sacrificed a bit on accuracy, particularly for subtle disease characteristics. EfficientNetB0, which was intended for light deployment, worked well considering its parameter efficiency but did not have the depth to pick up high-level abstractions in intricate disease textures. The findings confirm that architecture depth and connectivity have a strong impact on classification quality. Nevertheless, computational cost versus performance trade-offs need to be made for edge deployments.

6.5 Deployment Potential

From a deployment standpoint, the suggested model—specifically DenseNet121—can be deployed on mobile and edge computing platforms with minimal optimization. Methods like model pruning, quantization, or employing lighter versions (e.g., DenseNet-BC) can minimize memory footprint and inference latency. Deployment in mobile applications would allow farmers to take leaf photos and obtain real-time disease diagnosis. Integration with cloud-based dashboards of agronomists can further assist in region-wise disease surveillance and mitigation measures. The high accuracy and consistent generalization suggest robust readiness for field-level deployment in smart agricultural ecosystems.

6.6 Future Directions

Future research can concentrate on various fronts. First, inclusion of a more diverse set of plant species and types of diseases from real-field images other than those in PlantVillage will enhance model robustness. Second, using attention mechanisms or incorporating vision transformers (ViTs) could further enhance interpretability and performance on uncertain samples. Third, using temporal data (e.g., disease progression over time) can facilitate early diagnosis. Additionally, incorporating geolocation information and environmental metadata (temperature, humidity, etc.) can help in creating a context-aware, decision-support system. Lastly, creating a feedback loop where user corrections are used to refine the model can lead to a self-improving, farmer-centric AI tool.

7. Results

The empirical results of the experimental comparison of three deep learning models—DenseNet121, ResNet18, and EfficientNetB0—are presented in this section for the plant disease classification task. The models were trained on a portion of the PlantVillage dataset because of the limitation in computational resources, but with balanced representation of disease classes to maintain the integrity of the evaluation.

7.1 Model Performance Metrics

All three models were evaluated based on standard classification metrics such as accuracy, precision, recall, and F1-score, and confusion matrices and learning curves. DenseNet121 performed better than the other two in all evaluation metrics, with the highest validation accuracy of 99.17%, followed by 98.20% for EfficientNetB0 and 96.70% for ResNet18. DenseNet121 also showed better consistency in per-class F1-scores and lower misclassification rates.

7.2 Training Stability and Convergence

As can be seen in Figure 2, DenseNet121 had a stable convergence pattern with smooth and decreasing loss curves, whereas ResNet18 had slight overfitting indicators after some epochs. EfficientNetB0 was also fairly good but had slower convergence, taking more epochs for stabilization.

7.3 Confusion Matrix and Misclassification Insights

The DenseNet121 confusion matrix (Figure 1) showed high true positive rates in the majority of disease classes. Misclassifications were mostly seen in visually related diseases, e.g., "Apple Scab" vs "Apple Black Rot". These misclassifications guided subsequent model improvements.

7.4 Summary of Classification Metrics

A summary of performance metrics in detail is presented in Table 1, with per-class precision, recall, and F1-score. DenseNet121 had a macro F1-score of 0.96 overall, validating its stability for multi-class classification tasks.

8. Conclusion

This study presents a deep learning-based framework for the automatic classification of plant leaf diseases using transfer learning with pre-trained CNN architectures. Although the original PlantVillage dataset contains more than 70,000 images belonging to various classes, because of GPU resource constraints, we could not train on the whole dataset. Therefore, we strategically sampled a representative subset of the dataset with balanced coverage of significant disease classes and ample diversity for effective learning. This compromise was made to remain within the computational limits during experimental time.

Three cutting-edge pre-trained models, namely DenseNet121, ResNet18, and EfficientNetB0, were compared in a structured experimental pipeline. Every model was used in transfer learning mode where base layers were kept frozen, and the classifier head was fine-tuned to adjust to plant disease features. Post exhaustive training, validation, and performance analysis via metrics like confusion matrices, classification reports, and accuracy/loss curves, DenseNet121 proved to be the most stable model as far as accuracy, generalization, and consistency were concerned.

In spite of hardware limitations, the results firmly believe in the viability of applying CNN-based transfer learning to accurate plant disease classification. Our system proves that using restricted training samples from a vast dataset, judiciously chosen models and preprocessing methods can achieve extremely accurate outputs. The DenseNet121 model trained possesses immense potential for real-world implementation in smart agriculture applications, especially in edge and mobile devices. Future research can involve increasing training with the entire dataset and investigating model optimization for low-power settings.

In summary, this research is an encouraging advance towards intelligent, computerized plant disease diagnosis. It forms the foundation for large-scale, practical agriculture AI systems to help with sustainable agriculture, food security, and technological progress in agriculture.

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. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012).ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

5. Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger,K.Q.(2017).Densely connected convolutional networks (DenseNet). IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 4700–4708.https://doi.org/10.1109/CVPR.2017.243

6. Abbas, A., Abdelsamea, M. M., & Gaber, M. M. (2021). Classification of COVID-19 in chest X-ray images using DeTraC deep convolutional neural network. Applied Intelligence, 51, 854–864.

7. Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017).Deep learning for tomato diseases: Classification and symptoms visualization. Applied ArtificialIntelligence, 31(4), 299–315. https://doi.org/10.1080/08839514.2017.1315516

8. Sperschneider, J. (2022). Machine learning in plant–pathogen interactions: Empowering biological predictions from omics data. New Phytologist, 233(5), 2093–2099. https://doi.org/10.1111/nph.17952

9. Barbedo, J. G. A. (2018). Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification. Computers and Electronics in Agriculture, 153, 46–53. https://doi.org/10.1016/j.compag.2018.08.013