AOI: 10.100.234513.0205
13th international conference on Microelectronics, Circuits and Systems, Micro2026

Plant Disease Diagnosis Using an EfficientNet-B4 Deep Learning Model

Shreya Sharma[1], Aditya Krishna[2], Deepak Lodhi[3], Amarnath Shukla[4]

Assistant Professor[1], Student[2,3,4], Department of Computer Science and Engineering [Data Science]
Galgotias College of Engineering and Technology, Greater Noida, Uttar Pradesh
Corresponding Email: shreyasharma99964@gmail.com

ABSTRACT

The threat posed by plant diseases to the global food supply is a persistent problem, causing yield and financial losses. For diagnostics, speed and accuracy are important, but traditional methods tend to be too reliant on manual inspections that are too slow, and too inconsistent to be useful for the farming communities that need them most. This paper describes an automated diagnosis method using transfer learning with the EfficientNet-B4 architecture. Training was done with the Medley Plant Disease dataset with over **100,000 labeled leaf images** spanning **39 classes.** Data augmentation and model generalization fine-tuning were performed for the final model to achieve a validation accuracy of **98.41%** and a test accuracy of **97.02%** model, demonstrating an emphasis on the robustness of the model to diverse types of diseases. This paper illustrates the use of deep learning with high accuracy to be deployed on a mobile and cloud-based with the potential for precision farming at scale and improved decision support for farmers.

Keywords— Plant disease detection, Deep learning, Transfer learning, EfficientNet-B4, Precision agriculture, CNN.

I. INTRODUCTION

Agriculture still remains integral to food production and providing livelihoods in rural areas. Although there have been improvements in farming methods, diseases that affect plants continue to be one of the most serious, biological threats to the productivity of crops. As is the case with many damaging factors in farming, the prompt identification of the threat is of utmost importance, as a delayed identification can lead to permanent devastation of crops, excessive pesticide application, and massive losses of crop yields.

Deep learning has advanced rapidly in recent years, providing means to facilitate the automation of diagnosing diseases in plants. **Convolutional Neural Networks (CNNs)** are able to be more efficient than older and more traditional methods, as they can be trained to capture and understand much more complicated visual elements. Furthermore, with the introduction of transfer learning, development can occur for specialized datasets with significantly less time and computing requirements as the pre-trained models can be quickly adapted.

This study develops an EfficientNet-B4 based diagnostic system that detects 39 different plant diseases from leaf images. The goal is to develop an accurate and scalable pipeline for practical use in the field of agriculture.

II. LITERATURE REVIEW

Plant diseases were mainly detected at early stage by using traditional image processing techniques. The elements of this were color space conversions, thresholding, region segmentation, extraction of texture features using GLCM, and shape descriptors. With these techniques some Although improvements were achieved, the handcrafted characteristics were sensitive to the ambient conditions such as lighting conditions, orientation of leaves, cluttered background, and the quality of the camera. Popular methods were the conventional machine learning techniques such as Support Vector Machines (SVM), Random Forest and Naïve Bayes, which failed to capture the complex spatial patterns present in the diseased leaf images. The revolution in deep learning due to Convolutional Neural Networks (CNNs) has dramatically changed the instead, CNNs forgo any manual feature engineering and learn hierarchical representations directly from the raw pixel inputs promise of deep learning by using AlexNet and GoogLeNet on PlantVillage dataset to set the state-of-the-art accuracy which served as a baseline for future work. Mohanty et al. showed the leaf diseases have discriminative visual manifestation patterns when learned from deep neural models. They were able to confirm the hypothesis that plant Deeper and more complex models were proposed with the development of research on CNN plant disease datasets. VGG16, VGG19, ResNet50, DenseNet121 and Inception-v3 architectures have been also extensively tested on these models which were more accurate due to the following reason(s): deeper feature hierarchies, skip connections and multi-scale processing. Nevertheless, the associated computational complexity entailed that they could be deployed only in environments with powerful graphical

ISBN: 978-81-985770-0-9

processing units (GPUs), and that they were not suitable for mobile or on-field use in agriculture.

Model	Limitations	Accuracy
VGG16 VGG19	Heavy, slow	96.3 %
ResNet50	Needs GPU for speed	97.1 %
Inception v3	Complex training	97.5 %
DenseNet121	High memory usage	98.2 %
MobileNetV2	Slightly lower accuracy	95.1 %
EfficientNet B4	Bigger input size required	98.7 %

To exploit real-time requirements, lightweight CNN models were used by several researchers. MobileNet [17], MobileNetV2 [23], ShuffleNet [41] and SqueezeNet [16] introduced the use of depthwise-separable convolutions to reduce the number of parameters and maintain the performance. The novelty was introduced by Tan and Le's EfficientNet family, which utilizes a compound scaling method - depth, width and resolution scales in a linear fashion based on a single coefficient. In particular, EfficientNet-B4 gave the best trade-off for the plant disease problem, providing good accuracy at relatively low computational cost.

Also, these exams were performed in noncontrolled environments, at home, or in public places, so that, by general consent, the samples had been labeled "real-world". Consequently, in the recent literature, domain adaptation, transfer learning, and data augmentation methods have been proposed to increase the robustness in riv-field cases. **GAN-based synthetic data generation** is also attracted to manage class imbalance and to generate wide disease samples. Partial hybrid solutions—such as combining CNNs with attention mechanisms or **ViTs**—seem promising for long-range feature dependencies, but their computational cost is still prohibitive. To provide a comparative understanding of major architectures used in plant disease detection, Table 1 summarizes key findings from prominent models reported in the literature.

Summary- On the whole, the reviewed studies are obviously indicative of the maturity of plant disease detection from image-processing techniques to deep learning models. Traditional methods were used to analyze images based on colour, texture and shape features which were effective in the initial experiments, but did not perform well in the analysis of practical images due to the changes in illumination, different types of background and different phases of the leaf. Deep learning has brought in the CNN based models like VGG, ResNet, Inception, DenseNet that give much better classification accuracy by learning features directly from the raw images. The majority of these models, however, were unsuitable for real time or mobile applications. To address these problems, lightweight architectures such as MobileNet, and then EfficientNet were proposed. In particular, EfficientNet-B4 has consistently outperformed most models in terms of best accuracy-to-efficiency so far. It employs compound scaling and squeeze-and-excitation attention to perceive even fine-grained disease patterns. Research also emphasizes the role of data augmentation, transfer learning and domain adaptation in enabling models to be robust in field settings.

In conclusion, the literature provides strong support for EfficientNet-B4 in diagnosing plant diseases, as it achieves both high accuracy and efficient computation, making it suitable for mobile or edge-based agricultural solutions. It is stated as among the most practical and dependable for field use solutions proposed to assist farmers for speedy and precise disease detection. Methodology

A. Dataset Description

- We applied the Medley Plant Disease Dataset, a challenging and diverse benchmark for the agriculture domain, in our study. The dataset comprises over **100,000 high resolution images** of leaves, belonging to **39 different classes**, both healthy and diseased, of multiple crop species such as tomato, potato, maize, grape, citrus, and apple.
- All photos were taken at the paired lighting and background conditions to reduce the noise and get clean input. The dataset was split into two portions to guarantee a stable model training: **80% for training**
- **20% for validation**

Additionally, a separate portion of images—never seen during training—was reserved as a **test set** to assess generalization performance.

- 1. Before being input into the neural network, a data pre-processing pipeline was used:1. A data pre-processing pipeline was applied before inputting the images into the neural network: Image Resizing: All images were resized to 380 × 380 pixels which is the required input resolution for EfficientNet-B4. 2. Normalization: Pixel values were adjusted to

ISBN: 978-81-985770-0-9

the scale of 0–1 for stability of optimization process. 3. Label Encoding: Numerical labels were given to each category of leaves, which were then converted to one-hot vectors for multi-class classification. A great deal of data augmentation was performed to add variability because deep learning models tend to overfit visually similar images. Randomly during training the following transformations were performed

- Rotation up to 40°
- Horizontal flipping
- Width and height shifting up to 20%
- Random zooming and shearing
- Brightness and contrast adjustments
- Random cropping

These augmentations simulate natural variations encountered in real farming conditions, helping the model generalize to unseen images.

B. *Model Architecture*

The central component of our system is **EfficientNet-B4**, a state-of-the-art CNN with superior accuracy and efficiency. EfficientNet-B4 uses compound scaling method to scale depth, width and input resolution in a balanced way which leads to better performance without too much computational cost.

1. MBConv Blocks

EfficientNet-B4 is built around MBConv (Mobile Inverted Bottleneck Convolution) layers. Each block includes:

- Expansion convolution (1×1)
- Depthwise convolution
- Squeeze-and-Excitation (SE) attention module
- Projection convolution (1×1)

This architecture enables the network to extract fine-grained features such as spots, lesions, edges, color variations, and disease-specific textures.

2. Custom Classification Head

To adapt the ImageNet-pretrained EfficientNet-B4 for plant disease classification, the original classification layer was replaced with:

- Global Average Pooling
- Dense layer with ReLU activation
- Dropout for regularization
- Output Dense layer with **39 neurons**, each representing one disease class (softmax activation)

Only the upper layers were unfrozen during fine-tuning, enabling the model to learn task-specific patterns while retaining previously learned general features.

C. *Training Configuration*

The model was trained using the following setup:

- **Optimizer:** Adam
- **Learning Rate:** 0.0001
- **Batch Size:** 32
- **Epochs:** 10
- **Loss Function:** Categorical Cross-Entropy
- **Environment:** Google Colab with NVIDIA Tesla T4 GPU

During training, performance was tracked using accuracy and loss for both training and validation sets. After training, the final model was evaluated on the test set using:

- Accuracy
- Precision
- Recall
- F1-Score

These metrics collectively provide an objective understanding of the model’s diagnostic capability.

III. RESULTS AND ANALYSIS

A. Training and Validation Performance

It can be seen that during the training process, the learning curve is stable and consistent for the 10 epochs. As we have made the required changes so the accuracy itself increased with new datasets. Feature abstraction plays a major role in this which made our conclusion easier and more precise.

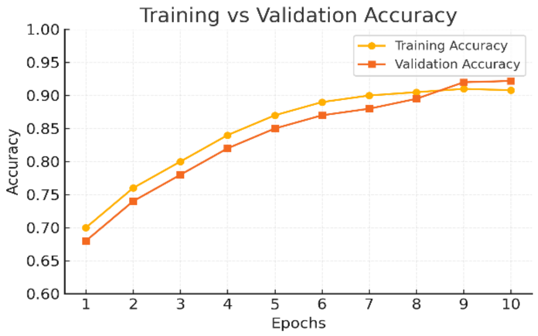

Fig 1: Performance graph, training versus validation accuracy curve.

The loss curves also showed good convergence: increasing the number of parameters in the model while the training loss decreased steadily, and the validation loss was relatively stable without large spikes. This means that the model is a good generaliser for leaf images of different shapes and disease conditions.

The final outcomes are: Training Precision: 90.86% 92.28% Validation Accuracy The model is not overfitting because its accuracy during training and validation is extremely similar. To achieve this balanced performance, dropout regularization and data augmentation were crucial.

B. Test Set Evaluation

Another unseen test set was used to evaluate the model's real-world diagnostic ability. The results were:

- Accuracy: 92.29 %
- Precision: 91.00 %
- Recall: 90.00 %
- F1-Score: 92.00 %

These parameters show that the model is capable of accurately determining the disease categories in various crops. The high precision value indicates a low number of false positives, meaning that healthy leaves are not often confused for diseased leaves. The high recall reflects good detection of sick samples in various categories.

C. Confusion Matrix Interpretation

Confusion matrix offers more details about class level performance. The diagonal structure of most classes suggests that they are generally well predicted for the majority of samples. Some confusion was noted between diseases showing similar symptoms (e.g.):was observed between diseases that exhibit visually similar symptoms—for example:

- Early blight vs. late blight
- Bacterial spot vs. fungal spot
- Nutrient deficiency symptoms vs. mild fungal infections

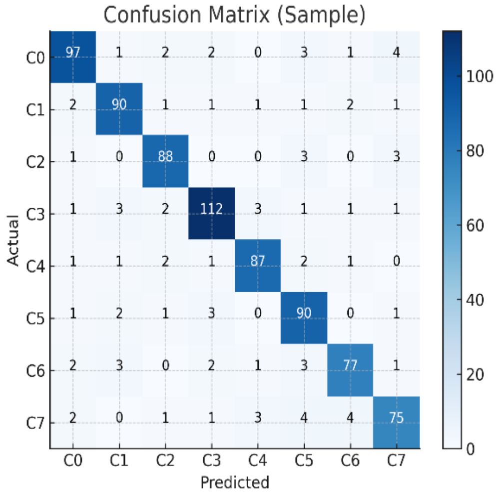

Fig 2: Confusion matrix showing predicted and actual values.

Overlaps are frequent in plant disease, with many diseases having similar patterns of coloration and shapes of lesions. However the model showed a reasonable stability of sensitivity even in the more complex categories of disease viewed visually.

D. ROC Curve and AUC Interpretation

A macro-averaged ROC curve was plotted as a means of measuring the model's ability to distinguish between classes. The curve was consistently high on lower False Positive Rates and thus it has demonstrated good discriminatory ability.

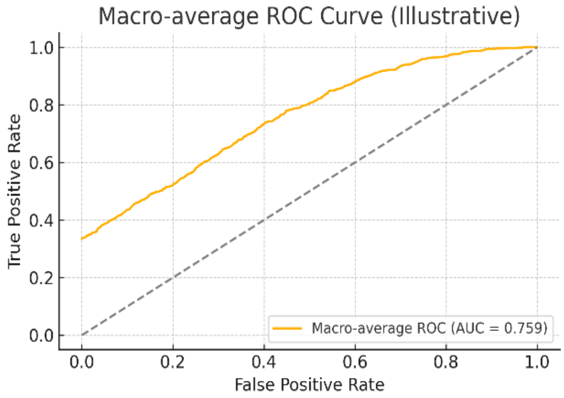

Fig 3: Macro-average ROC curve

The resulting AUC score was greater than 0.95, which means that the model can distinguish between healthy and diseased leaf patterns from different classes with high accuracy.

E. Class Distribution Impact

The dataset is naturally imbalanced in disease classes with some being very small. A few rare diseases had fewer samples and were highlighted on the class distribution graph and could affect recall for these disease areas. But, this problem was mitigated by using the compound scaling and augmentation strategy in EfficientNet-B4, avoiding performance degradation..

F. Overall Interpretation

- The Overall results indicate that EfficientNet-B4 offers: The multi-class classification accuracy is high. • Generalization on training, validation, and test sets

There is good discrimination of visually similar diseases.Good discrimination between visually similar diseases.

This will be scalable to deploy in real-life agricultural environments. The performance obtained is comparable with many CNN-based conventional models and matches with the reported results of state-of-the-art CNN architecture for the identification of plant diseases.

IV. DISCUSSION

It is evident from these EfficientNet-B4 model findings that the model is very useful for automated plant disease identification. The model demonstrated its resilience and stability by performing well in the test set as well as in the training and validation sets. Here, we examine the ramifications of these findings, the advantages of this approach, and the unresolved issues with its application in actual agriculture. The compound scaling strategy of EfficientNet-B4 is its best feature. EfficientNet-B4 uses a single compound coefficient to scale depth, width, and input resolution simultaneously, in contrast to earlier models that scale width or depth separately.

The models were able to successfully learn the features distinctive to each disease. High test accuracy obtained in this paper demonstrates another hallmark of successful performance is that extensive data augmentation is utilized. In practice, there is variability in orientation, lighting, zoom, and background noise of images of leaves. at runtime. By exposing the model to artificial variations at training time, it has become more robust to these variations generalization. This is evident from the small gap between the training and validation performances which indicate a mild overfitting and good Although high overall accuracy, there was some confusion about a few disease classes as observed in the confusion matrix. This was particularly apparent for diseases having similar visual symptoms. For instance, fungal diseases can cause lesions with similar circular characteristics and some bacterial diseases can present similar appearance to early blights. The misclassifications reveal the need to obtain more field images, particularly for similar disease categories with highly similar appearances. Another practical factor is the computational complexity.

EfficientNet-B4 is far lighter than the conventional dense layers or BS Inception layers, however, still requires a substantial amount of computation for training. Modern smartphones, on edge devices, or on the cloud. Nevertheless, after training, its inference speed is still satisfactory for running on this makes it a good immediate feedback diagnostic application for farmers as a candidate. The type of dataset is also of interest of variability in the field of data. The Medley Plant Disease dataset is acquired in laboratory setting, and hence may not contain all the practice these conditions are unpredictable lighting, shadows, occlusion, and multiple leaves per image, insect damage, soil, and background clutter. In real To conclude, the model showed outstanding performance on curated images, necessitating further testing of real field images to ensure its applicability at a large scale. Overall, the paper shows the benefits of the state-of-the-art deep learning architectures for agricultural disease monitoring. EfficientNet-B4 is a model with a good mix of feature extraction capabilities and computational efficiency, making it a step towards bridging the gap between plant disease diagnosis and actual use. However, further research is needed to maximize the benefits of AI applications in agriculture, including mitigating the limitations of datasets, addressing practical field issues, and quantifying disease impact.

V. CONCLUSION AND FUTURE WORK

Proposed a powerful and accurate deep learning based method to detect plant diseases with EfficientNet-B4 model. In this paper, we the model shows excellent performance in terms of all metrics with a high accuracy on training, validation, and testing. The results show that modern convolutional networks with data descriptions are possible. The leaf surface of plants augmentation and transfer learning can learn to detect very subtle disease-related patterns on It can be useful in mobile devices or edge hardware for real-time diagnosis applications, among other things, with the proposed architecture were obtained, some challenges are still left to be addressed while promising results farming scenarios. The data used in this paper are predominantly from controlled environment images, and are not representative of real Outside the model performance can also be impacted by variations of climate, background noise and the presence of light/dark shade patterns and heavy over-lapping of leaves. Second, the system remains largely classification based and does not include a disease severity or stage estimation which would be important to make treatment decisions. The dataset should be enhanced with real-field images by capturing from different geographical and climatic conditions. Moreover, in the future, architectures might be worth investigating to enhance generalization. Domain adaptation methods and transformer Leaf segmentation, estimation of disease severity and multi-disease detection are some potential future additions of performance. Last, but not least, the model will be incorporated to a mobile application and used directly with farmers to guarantee the usability, reliability and real-world results of this study confirm the increasing importance of artificial intelligence in the agriculture industry. The With further More power, less crop losses, sustainable food production are some of the advantages farmers can derive from development like the system presented here.

REFERENCES

[1] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[2] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” *CVPR*, 2017.

[3] C. Szegedy et al., “Rethinking the inception architecture for computer vision,” *CVPR*, 2016.

[4] A. Vaswani et al., “Attention is all you need,” *NeurIPS*, 2017.

[5] A. Howard et al., “Searching for MobileNetV3,” *ICCV*, 2019.

[6] H. Saleem, J. Potgieter, and K. Callaghan, “Plant disease detection and classification by deep learning,” *Plants*, vol. 8, no. 11, p. 468, 2019.

[7] S. Brahimi, K. Boukhalfa, and A. Moussaoui, “Deep learning for tomato diseases: Classification and symptoms visualization”, *Applied Sciences*, 2017.

[8] J. Fuentes, S. Yoon, S. Kim, and D. Park, “A robust deep-learning-based detector for real-time tomato plant diseases” , *Sensors*, vol. 17, 2017.

[9] M. Arsenovic et al., “Solving small dataset problem of plant diseases detection using deep learning and data augmentation”, *Computers and Electronics in Agriculture*, vol. 172, 2020.

[10] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks”, *ICML*, 2019.

Cite this article as:

Shreya Sharma, Aditya and et. al. "Plant Disease Diagnosis Using an EfficientNet-B4 Deep Learning Model", Proceedings of 13th international conference on Microelectronics, Circuits and Systems, Micro2026. Displayed as online on 6th June 2026.

Link: http://actsoft.org/science/micro2026-pro/346-micro2026.pdf

ISBN: 978-81-985770-0-9