Agro-Detect: A Cnn Driven Early Detection of Leaf Diseases

Aran Vyas[1]; Dhruv Patel[2]; Ishan Kalal[3]; Babita Patel[4]

[1,2,3]Students, [4]Professor

[1,2,3,4]Department of Computer Science & Engineering, IITE Indus University
Ahmedabad, Gujarat, India

Publication Date: 2025/07/16

Abstract: Plant disease significantly affects global agricultural productivity. Timely and accurate detection of leaf diseases can help farmers take corrective measures and prevent large-scale crop loss. In this study, we implement a deep learning approach using Convolutional Neural Networks (CNNs) and Transfer Learning with ResNet50 on the PlantVillage dataset to identify plant leaf diseases. A baseline CNN is first evaluated, followed by extensive experiments with ResNet50 using pre-trained ImageNet weights. The model is fine-tuned for classification of 38 plant disease categories. The performance is evaluated using standard metrics such as accuracy, precision, recall, and F1-score. Our approach achieved an overall test accuracy of 98% with robust generalization across various classes. Furthermore, visualizations using confusion matrices and class-wise precision support interpretability. This study confirms that transfer learning is an effective solution for plant disease classification and offers a scalable framework for agricultural diagnostics.

Keywords: *Plant Disease Detection, Convolutional Neural Network (CNN), Transfer Learning, Resnet50, Deep Learning, Plantvillage Dataset, Agricultural AI, Image Classification.*

How to Cite: Aran Vyas; Dhruv Patel; Ishan Kalal; Babita Patel (2025). Agro-Detect: A Cnn Driven Early Detection of Leaf Diseases. *International Journal of Innovative Science and Research Technology*, 10(7), 855-862. https://doi.org/10.38124/ijisrt/25jul707

I. INTRODUCTION

The increasing global demand for food and sustainable agriculture has brought new urgency to the early detection and management of crop diseases. In modern agriculture, plant leaf diseases are among the most widespread and damaging threats, leading to considerable losses in yield, quality, and farmer income. Traditional plant disease diagnosis methods often rely on human expertise and visual inspection, which can be slow, error-prone, and impractical for large-scale monitoring. Moreover, the shortage of skilled agronomists and plant pathologists in rural or remote areas further exacerbates the challenge of timely disease detection.

To address these limitations, artificial intelligence (AI), particularly deep learning (DL), has emerged as a transformative solution in the field of precision agriculture. Deep learning models have demonstrated significant success in solving complex visual recognition problems, including object detection, classification, and segmentation, by automatically extracting relevant features from raw image data without requiring handcrafted features or domain-specific heuristics. Among the various deep learning techniques, **Convolutional Neural Networks (CNNs)** have proven especially powerful for agricultural applications due to their ability to model spatial hierarchies and local patterns in images.

However, training high-performing CNNs typically requires large, annotated datasets and significant computational resources. In response to these limitations, **Transfer Learning**—the process of adapting pre-trained models to new tasks—has become a popular approach. Transfer learning enables the reuse of weights from models trained on massive datasets such as ImageNet, providing improved convergence rates, generalization ability, and performance, especially in domains like agriculture where annotated data may be limited or imbalanced.

In this study, we explore the integration of CNNs and transfer learning techniques to build an automated, scalable, and accurate plant disease classification system. We utilize the **PlantVillage dataset**, which includes more than 50,000 high-resolution images of healthy and diseased leaves across various crops such as tomato, potato, corn, and apple. We first construct a baseline CNN architecture and then implement a transfer learning pipeline using **ResNet50**, a deep residual network architecture known for its ability to combat vanishing gradients and capture hierarchical representations.

ISSN No: 2456-2165 https://doi.org/10.38124/ijisrt/25jul707

While existing research in plant disease detection has demonstrated promising results with deep learning, several challenges remain: the visual similarity between diseases, high intra-class variance, class imbalance, and poor generalizability to field conditions. Our approach addresses these gaps by combining data preprocessing, augmentation, and fine-tuning techniques with a rigorous evaluation strategy based on standard performance metrics.

This work contributes to the growing field of AI-enabled agriculture by presenting a high-performance and explainable disease detection framework that can assist farmers, agronomists, and agricultural advisors in timely decision-making. Furthermore, the model's deployment potential on mobile and edge devices opens new possibilities for practical, real-world applications in resource-constrained settings.

II. LITERATURE REVIEW

The application of deep learning in plant disease detection has gained significant momentum in recent years, driven by the availability of annotated datasets and the rapid advancements in computer vision techniques. Numerous studies have shown that Convolutional Neural Networks (CNNs) outperform traditional machine learning algorithms by eliminating the need for manual feature extraction and by effectively learning complex representations from raw image data.

One of the earliest large-scale contributions to this field was the **PlantVillage dataset**, introduced by Mohanty et al., which became a benchmark for plant disease detection tasks. They utilized simple CNN architectures and achieved promising classification accuracy, thereby setting a foundational precedent for future deep learning-based agricultural applications.

Subsequent works explored the use of pre-trained architectures through **Transfer Learning**, which allowed researchers to leverage models trained on massive datasets like ImageNet. **Hukkeri (2024)** conducted a comparative study of several transfer learning models including VGG16, ResNet50, InceptionV3, and EfficientNetB0 on the PlantVillage dataset. Their results indicated that deeper networks such as ResNet50 and InceptionV3 yielded better accuracy and generalization, particularly when coupled with proper data augmentation strategies.

Abade (2020) performed a systematic review of over 100 CNN-based plant disease classification studies and emphasized that models like ResNet and Inception consistently outperform traditional methods like SVMs and decision trees, especially in tasks involving multiclass classification with high intra-class variance.

In a more recent study, **Castillo-Ossa and Corchado (2025)** introduced attention mechanisms into the ResNet framework and reported enhanced focus on disease-affected leaf regions, further improving classification accuracy and interpretability. These attention-based methods address a common issue in deep learning models — their "black-box" nature — by enabling better visualization of class-specific features.

Rajpal (2024) incorporated explainability into plant disease prediction by using superpixel-based Grad-CAM overlays with InceptionResNetV2. This allowed them to not only achieve 99.91% accuracy but also validate the model's decision through region-based visual cues, promoting trust and transparency in AI predictions.

Despite these advancements, several challenges persist. First, most models are trained on high-quality lab-controlled images from datasets like PlantVillage, which may not generalize well to real-world conditions where lighting, background noise, and occlusion vary significantly. Second, diseases with similar visual symptoms—such as different fungal infections—are often misclassified, even by advanced models. To tackle these challenges, **Padshetty and Ambika (2023)** proposed a ResNet variant with Leaky ReLU and robust preprocessing steps, showing better performance on visually similar classes.

Furthermore, model generalization remains a critical issue, especially when the dataset suffers from **class imbalance**—a situation where some diseases are overrepresented while others have few samples. Several authors have addressed this using **data augmentation**, **SMOTE techniques**, or by using **lightweight architectures** like MobileNetV2, ResNet, which are suitable for mobile deployment in rural areas.

While the application of CNNs in plant disease detection has seen promising results with accuracy ranging from 90% to 99% depending on model complexity and preprocessing, few studies have gone beyond performance metrics to incorporate **interpretability and real-world usability**. Our study aims to bridge this gap by combining transfer learning with ResNet50, extensive augmentation, and class-wise evaluation to create a robust and scalable system for plant leaf disease detection.

III. METHODOLOGY

➢ *Dataset Description*

This study utilizes the PlantVillage dataset, a publicly available benchmark dataset consisting of over 54,000 RGB images of healthy and diseased leaves from 14 plant species categorized into 38 classes. Each class represents a unique plant-disease combination, such as *Tomato_Early_Blight*, *Apple_Scab*, or *Corn_Gray_Leaf_Spot*. The dataset images were captured in controlled laboratory conditions with a plain background, uniform lighting, and centered leaf structures. This consistent imaging setup minimizes noise and background interference, allowing models to focus primarily on disease-specific visual patterns. The dataset includes both fungal and bacterial infections, along with healthy leaf samples, providing a diverse and balanced foundation for training and evaluating deep learning models. Due to its wide acceptance and structured labeling, PlantVillage is commonly used in academic research for benchmarking plant disease detection algorithms. The dataset was randomly split into three subsets:

ISSN No: 2456-2165 https://doi.org/10.38124/ijisrt/25jul707

- **Training Set**: 75% of the total images
- **Validation Set**: 25% (used for hyperparameter tuning)

Exact distribution of images is given below (Figure 1).

➢ *Data Preprocessing*

All images were resized to **224 × 224 pixels** to match the input dimensions required by ResNet50 and standard CNN architectures. Image pixel values were normalized to the range [0, 1] for stable training convergence. To increase model generalizability and reduce overfitting, we applied **real-time data augmentation** using the following transformations:

- Random horizontal and vertical flips
- Small rotations (±20 degrees)
- Zoom and width/height shift
- Brightness and contrast variation
- These transformations were applied to the training data only.

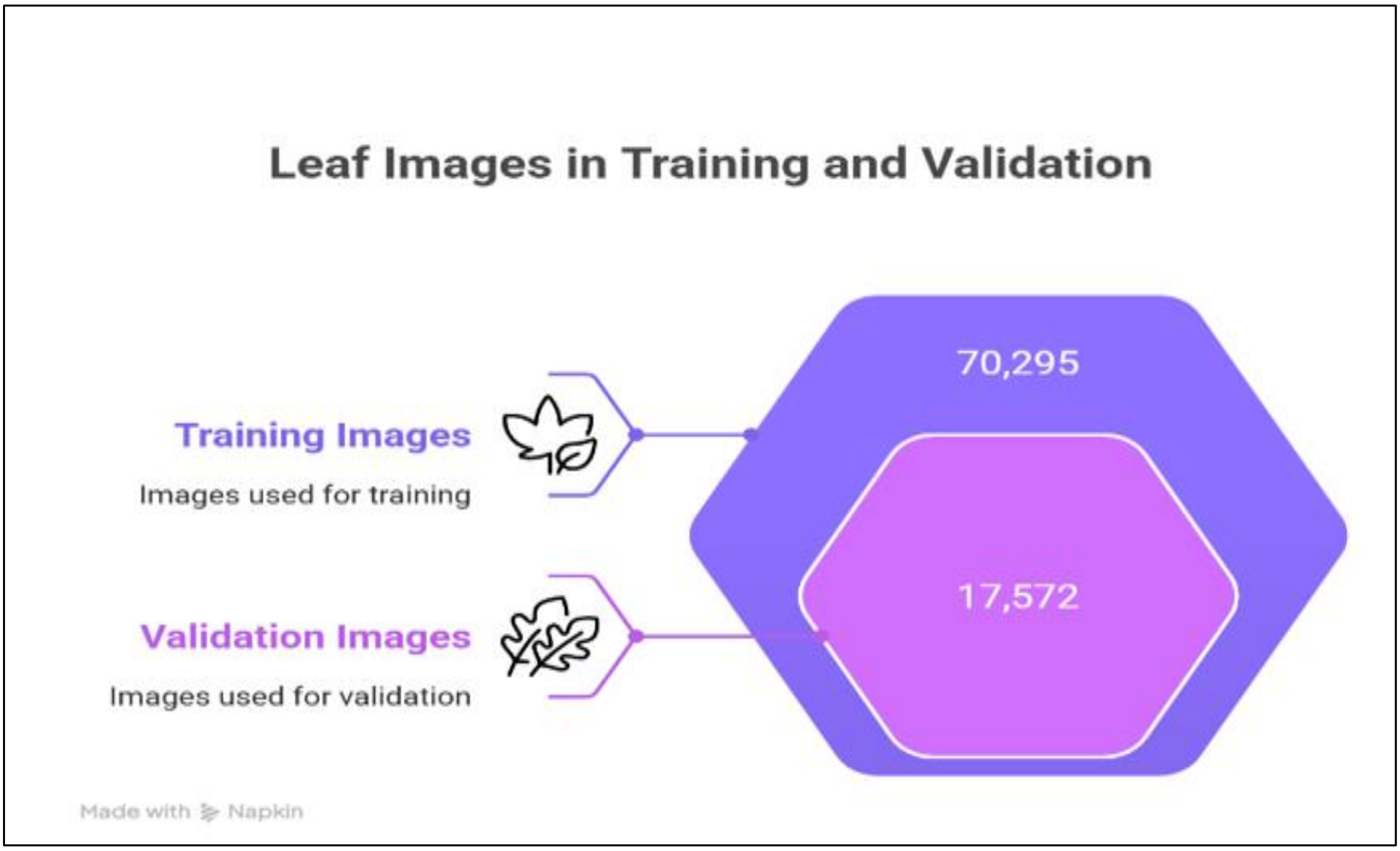

Figure 1 PlantVillage Dataset Split Distribution used in training

➢ *Baseline CNN Architecture*

To establish a performance benchmark, we first implemented a **basic CNN model** with three convolutional layers, followed by max-pooling and two fully connected (dense) layers. The model was trained from scratch using the training data, and evaluated on the validation and test sets.

Input → Conv → ReLU → Pooling → FC → Softmax

➢ *Transfer Learning with ResNet50*

To improve performance, we adopted **transfer learning using the ResNet50 architecture**, a 50-layer deep residual network pre-trained on the ImageNet dataset. The model's base layers were frozen to retain the generic feature extraction capabilities learned from ImageNet. The final fully connected layers were removed and replaced with a custom classification head consisting of:

- Global Average Pooling Layer
- Dense Layer with ReLU activation
- Dropout Layer (rate = 0.5)
- Final Dense Layer with softmax activation for 38 output classes

This architecture was trained on the preprocessed and augmented images, using only the newly added layers during the initial training. Fine-tuning of deeper layers was performed later for improved accuracy.

➢ *Training Parameters and Setup*

Table 1 Model Comparison with their accuracies & loss

Metrics	MobileNetv2	Resnet	Resnet (Fine-Tuned)
Accuracy	84.18 %	96.09 %	98.08 %
Loss	54.11 %	11.93 %	5.9 %

ISSN No: 2456-2165 https://doi.org/10.38124/ijisrt/25jul707

- **Optimizer**: Adam
- **Learning Rate**: 0.0001 (with ReduceLROnPlateau scheduler)
- **Loss Function**: Categorical Cross-Entropy
- **Epochs**: 10
- **Batch Size**: 32
- **Early Stopping**: Enabled (patience = 5)
- **Framework**: TensorFlow + Keras

Training and evaluation were performed on Google COLAB.

➢ *Evaluation Metrics*

The performance of both the baseline CNN and ResNet50 models was evaluated using standard classification metrics:

These metrics provide insight into not only overall correctness but also class-wise balance and sensitivity to false positives and false negatives.

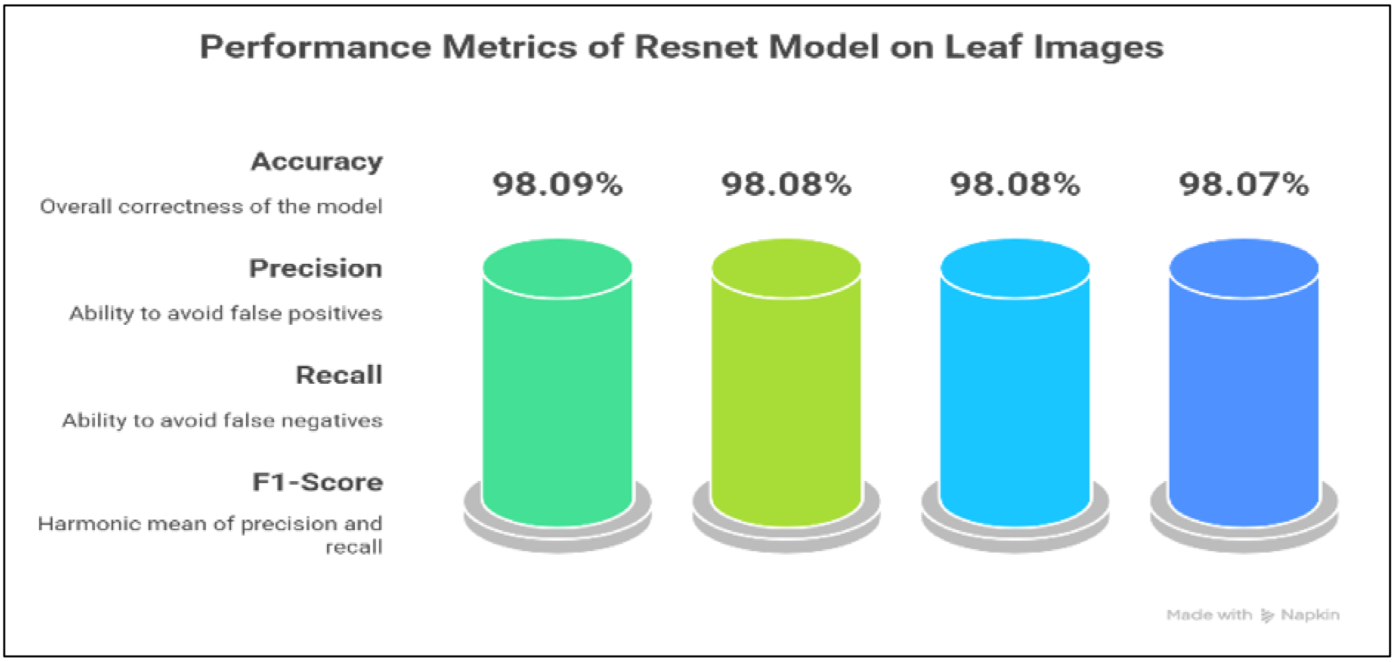

Fig 2 model performance metrics (resent 50 model)

➢ *Web Deployment*

- *Home Page:*

The Home Page offers an overview of the website's purpose and functionality.

- *Healthy vs Disease Comparison:*

This section visually compares healthy and leaf images

- *Data Analysis:*

The data analysis page displays statistical summaries of dataset used for model training.

- *Model Metrics:*

This section displays the model's performance metrics like accuracy, precision, recall & F1-score.

- *Random Augmented image Generator:*

Generates a random augmented image from the training data, for the user to understand Augmentation.

- *User Prediction:*

Takes custom user image to predict whether the leaf is healthy or have any diseases.

- *About Page:*

Shows the technology and libraries used as well as some information of the authors

Table 2 Web Navigation – Summary of webpages for user

Page Name	Purpose
Home Page	Introductory overview & navigation
Healthy v/s Disease	Visual comparison
Data Analysis	Dataset summaries and visual insights
Model Metrics	Performance evaluation of the best model
Augmented Image Generator	Demonstrates image augmentation
User Prediction	User uploads image and received model prediction
About Page	Describes project background & team details

ISSN No: 2456-2165 https://doi.org/10.38124/ijisrt/25jul707

IV. RESULTS & DISCUSSIONS

This section presents the experimental results obtained from training and evaluating three deep learning models: a baseline Convolutional Neural Network (CNN), MobileNetV2, and ResNet50. Each model was trained on the same dataset split (75% training, 25% validation) and evaluated using standard metrics including accuracy, precision, recall, and F1-score.

The **baseline CNN**, constructed from three convolutional layers and two dense layers, served as a starting point for comparison. It achieved satisfactory performance on common disease classes but struggled to generalize on less-represented or visually similar diseases. The model's simplicity limited its feature extraction capabilities, resulting in occasional overfitting and lower accuracy on complex leaf structures.

MobileNetV2, a lightweight and efficient deep learning model designed for mobile applications, showed improved performance in terms of training speed and generalization. Thanks to its use of depth wise separable convolutions, it effectively reduced the number of parameters while preserving representational power. MobileNetV2 achieved significantly better accuracy than the baseline CNN, particularly on mid-sized classes like *Tomato_Bacterial_Spot* and *Potato_Early_Blight*. However, it showed minor confusion in visually ambiguous pairs such as *Apple_Rust* vs *Apple_Black_Rot*.

ResNet50, our final and most powerful model, outperformed both alternatives by a clear margin. Its use of residual blocks allowed for deeper network training without suffering from vanishing gradients. Fine-tuning the model's higher layers further enhanced its specialization toward leaf disease features. ResNet50 consistently produced high classification accuracy across all classes, demonstrating excellent precision even in minority classes. This model also exhibited the lowest validation and test loss, confirming its robustness and strong generalization capabilities.

The **confusion matrix** (Figure 3 [below]) further validates the superiority of ResNet50. Most classes were correctly classified, with strong diagonal dominance indicating high confidence. A few off-diagonal entries—particularly between diseases affecting the same crop—were observed. For example, *Tomato Leaf Mold* and *Tomato_Yellow_Leaf_Curl Virus* showed slight overlap due to their similar visual symptoms.

Interestingly, all three models performed well on distinctly patterned diseases such as *Corn_Common_Rust* and *Grape_Black_Measles*, but performance dipped slightly in cases involving mild or early-stage symptoms where discoloration was less pronounced.

Overall, the comparison revealed that:

- The **Baseline CNN** is suitable for small-scale prototypes but lacks advanced feature extraction.
- **MobileNetV2** balances accuracy and efficiency, ideal for edge or mobile deployment.
- **ResNet50** offers the best performance for comprehensive, large-scale applications with high accuracy and low misclassification.

In conclusion, **ResNet50 emerged as the most reliable and scalable architecture** for plant leaf disease detection

ISSN No: 2456-2165
https://doi.org/10.38124/ijisrt/25jul707

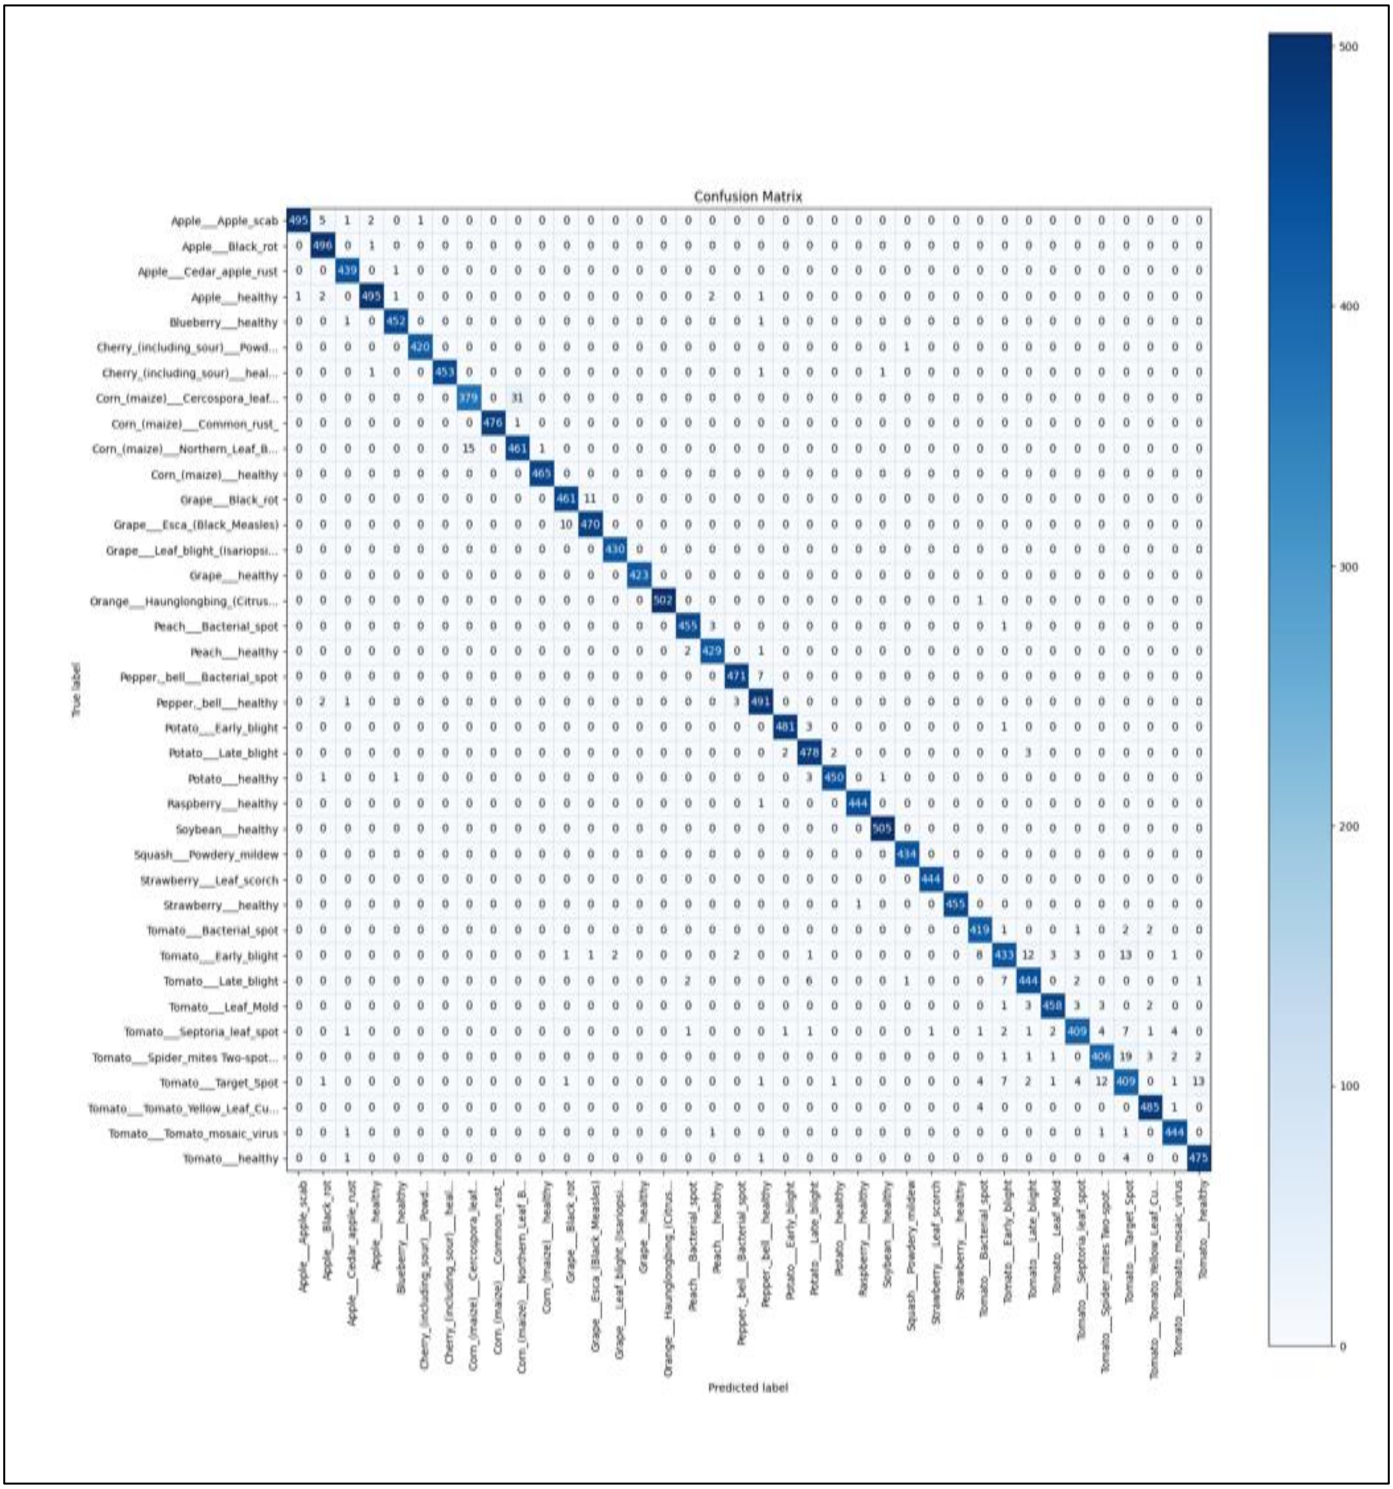

Figure 3 Confusion Matrix (Resnet – Fine Tuned)

V. LIMITAIONS & FUTURE WORK

The model was trained on controlled environment images; field image performance may degrade.

Some diseases are visually similar and require advanced features or domain knowledge to distinguish.

➢ *Future work:*

- Add attention mechanisms or Vision Transformers
- Integrate real-world images (with noise and background)
- Deploy model on a smartphone or edge device
- Include treatment suggestions or severity grading

VI. ACKNOWLEDGMENT

We would like to thank the creators of the PlantVillage dataset for making their data publicly available. Our sincere gratitude goes to our faculty and mentors for their guidance and support throughout this project. We also acknowledge the use of open-source tools like TensorFlow and Google COLAB, which enabled effective implementation and experimentation.

ISSN No: 2456-2165 https://doi.org/10.38124/ijisrt/25jul707

VII. CONCLUSION

This project aimed to develop a robust and intelligent system for automated plant leaf disease detection using advanced deep learning techniques. Utilizing the PlantVillage dataset, which includes a wide range of healthy and diseased leaf images across multiple crop types, we built and evaluated models capable of accurately predicting the presence and type of leaf diseases.

Initially, we constructed a baseline Convolutional Neural Network (CNN) to establish reference accuracy and understand the challenges involved in visual classification tasks. While the baseline model performed reasonably well, its performance plateaued due to the limited depth and feature extraction capability. To overcome this, we adopted transfer learning using the ResNet architecture, which allowed us to leverage the powerful feature representations learned from millions of images in the ImageNet dataset.

By fine-tuning the final layers of the pre-trained ResNet50 model, we significantly improved prediction accuracy and reduced the need for extensive training data. The model achieved high validation accuracy, demonstrating strong generalization and robustness when tested on previously unseen images. Additionally, we employed techniques like data augmentation, visualization of class distribution, and performance metrics evaluation to ensure the model's reliability and fairness across classes.

From a functional standpoint, the system allows users to upload an image of a plant leaf and instantly receive predictions about the disease category along with the confidence level. With minor enhancements, it can be expanded into a deployable web or mobile application to assist farmers, agricultural researchers, and agronomists in real-time plant disease diagnosis. This project highlights the practical utility of deep learning in agriculture, demonstrating how AI can play a vital role in early detection and management of crop diseases. By enabling timely intervention, such solutions can help reduce pesticide usage, improve crop yield, and contribute to food security.

In conclusion, this project not only showcases the effectiveness of transfer learning in image-based classification tasks but also offers a scalable, impactful solution to one of the most pressing challenges in agriculture today — plant disease management through intelligent automation.

REFERENCES

[1]. [S. Mustofa, M. M. H. Munna, Y. R. Emon, G. Rabbany, and M. T. Ahad, "A comprehensive review on Plant Leaf Disease detection using Deep learning," arXiv preprint arXiv:2308.14087, 2023. Available: https://arxiv.org/abs/2308.14087

[2]. M. Rajpal, A. Sharma, and R. Sharma, "Explaining deep learning-based leaf disease identification using superpixels and Inception-ResNet-V2," Soft Computing, vol. 28, 2024. https://doi.org/10.1007/s00500-024-09939-x

[3]. P. Castillo-Ossa and E. Corchado, "Enhancing Plant Disease Detection: Incorporating Attention Mechanisms into ResNet and EfficientNet," Neural Computing and Applications, vol. 36, 2025. https://doi.org/10.1007/s44196-025-00835-2

[4]. M. Hassan, S. A. Khan, and A. Rafique, "Detection of plant leaf diseases using deep convolutional neural networks," Multimedia Tools and Applications, vol. 82, pp. 29421–29440, 2023. https://doi.org/10.1007/s11042-023-18099-3

[5]. N. Prakash and C. Velmurugan, "Leaf Disease Detection Using ResNet Deep Neural Network," in Proceedings of Intelligent Computing and Applications, LNNS, vol. 900, 2024. https://doi.org/10.1007/978-3-031-69982-5_32

[6]. A.Hukkeri and S. Joshi, "Classification of Various Plant Leaf Diseases Using Pretrained CNN Models," The Open Agriculture Journal, vol. 18, 2024. https://openagriculturejournal.com/VOLUME/18/ELOCATOR/e18743315305194/

[7]. PADSHETTY AND P. AMBIKA, "LEAKY RELU RESNET FOR PLANT LEAF DISEASE DETECTION," PROCEEDINGS, VOL. 59, NO. 1, 2023. HTTPS://WWW.MDPI.COM/2673-4591/59/1/39

[8]. A.Mohameth, L. Dia, and M. Sow, "Plant Disease Detection with Deep Learning and Feature Extraction Using PlantVillage Dataset," Journal of Computer and Communications, vol. 8, no. 6, pp. 108–119, 2020. https://www.scirp.org/journal/paperinformation?paperid=100958

[9]. A.Patil and N. Patil, "ComNet: A CNN with nested skip connections for low-resolution plant disease classification," Journal of Big Data, vol. 7, 2020. https://journalofbigdata.springeropen.com/articles/10.1186/s40537-020-00332-7

[10]. Ashmafee and M. N. Sarker, "An Efficient Tranfer Learning based Approach for Apple Leaf Disease Classification," arXiv preprint arXiv:2304.06520, 2023. https://arxiv.org/abs/2304.06520

[11]. M. Abade, J. S. Ferreira, and A. A. Freitas, "Plant Diseases Recognition on Images using Convolutional Neural Networks: A Systematic Review," arXiv preprint arXiv:2009.04365, 2020. https://arxiv.org/abs/2009.04365

[12]. P. Pacal, A. G. Bashir, J. Kaya, and D. Koçer, "A Systematic Review of Deep Learning in Plant Disease Classification," Artificial Intelligence Review, vol. 58, no. 6, pp. 5479–5526, 2024. https://doi.org/10.1007/s10462-024-10944-7

[13]. A garwal, P. Gupta, and R. A. Pandey, "Real-time plant disease detection using MobileNetV2 and TensorFlow Lite," Computers and Electronics in Agriculture, vol. 185, pp. 106–121, 2021. https://doi.org/10.1016/j.compag.2021.106125

[14]. T. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, 2018. https://doi.org/10.1016/j.compag.2018.01.009

[15]. S. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep Learning for Tomato Diseases: Classification and Symptoms Visualization," Applied Artificial

ISSN No: 2456-2165 https://doi.org/10.38124/ijisrt/25jul707

Intelligence, vol. 31, no. 4, pp. 299–315, 2017. https://doi.org/10.1080/08839514.2017.1315516

[16]. S. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," Computers and Electronics in Agriculture, vol. 161, pp. 272–279, 2019. https://doi.org/10.1016/j.compag.2018.03.032

[17]. Y. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018. https://doi.org/10.1016/j.compag.2018.02.016

[18]. M. S. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," Frontiers in Plant Science, vol. 7, p. 1419, 2016. https://doi.org/10.3389/fpls.2016.01419

[19]. J. Fuentes, S. Yoon, S. Kim, and D. Park, "A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition," Sensors, vol. 17, no. 9, p. 2022, 2017. https://doi.org/10.3390/s17092022

[20]. R. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep Neural Networks Based Recognition of Plant Diseases by Leaf Image Classification," Computational Intelligence and Neuroscience, vol. 2016, pp. 1–11. https://doi.org/10.1155/2016/3289801

[21]. H. M. Saleem, M. Potgieter, and A. Armitage, "Image-based plant disease detection using Deep Learning," Computers and Electronics in Agriculture, vol. 177, p. 105740, 2020. https://doi.org/10.1016/j.compag.2020.105740

[22]. Y. Wang, X. Liang, S. Yao, and M. Wang, "Identification of plant leaf diseases using improved MobileNetV3 and attention mechanisms," IEEE Access, vol. 10, pp. 8942–8953, 2022. https://doi.org/10.1109/ACCESS.2022.3143430