Optimized Deep Learning Model for Plant Disease Detection Using Leaf Images

By

Md Towhidul Islam
19101159
Ahsanul Haque Kowshik
19301062
Syed Ahnaf Wadud Sami
18201127
Raisul Islam Opu
19101583
Mahmud Hassan
22241160

A thesis submitted to the Department of Computer Science and Engineering
in partial fulfillment of the requirements for the degree of
B.Sc. in Computer Science

Department of Computer Science and Engineering
Brac University
October 2024

Declaration

It is hereby declared that

1. The thesis submitted is my/our own original work while completing degree at Brac University.

2. The thesis does not contain material previously published or written by a third party, except where this is appropriately cited through full and accurate referencing.

3. The thesis does not contain material which has been accepted, or submitted, for any other degree or diploma at a university or other institution.

4. We have acknowledged all main sources of help.

Student's Full Name & Signature:

Md Towhidul Islam
19101159

Ahsanul Haque Kowshik
19301062

Syed Ahnaf Wadud Sami
18201127

Raisul Islam Opu
19101583

Mahmud Hassan
22241160

Approval

The thesis/project titled "Optimized Deep Learning Model for Plant Disease Detection Using Leaf Images " submitted by

1. Md Towhidul Islam (19101159)
2. Ahsanul Haque Kowshik (19301062)
3. Syed Ahnaf Wadud Sami (18201127)
4. Raisul Islam Opu (19101583)
5. Mahmud Hassan (22241160)

Of Summer, 2024 has been accepted as satisfactory in partial fulfillment of the requirement for the degree of B.Sc. in Computer Science on Oct 20, 2024.

Examining Committee:

Supervisor:
(Member)

Dr. Md. Ashraful Alam
Associate Professor
Department of Computer Science and Engineering
BRAC University

Co-Supervisor:
(Member)

Md. Tanzim Reza
Senior Lecturer
Department of Computer Science and Engineering
BRAC University

Thesis Coordinator::
(Member)

Dr. Md. Golam Rabiul Alam
Professor
Department of Computer Science and Engineering
BRAC University

Head of Department:
(Chair)

Dr. Sadia Hamid Kazi
Chairperson and Associate Professor
Department of Computer Science and Engineering
BRAC University

Ethics Statement

This study is practiced to the best ethical standard from every side of planning and execution. Raw sequence data used in the study were downloaded from public databases or obtained under proper regulations and permissions by relevant authorities. The research involved no sensitive or private data, and there were neither any human nor animal participants. The analysis and findings reported in this work have been checked carefully for objectivity. Finally, all of the model training software and resources utilized in our study are properly licensed and procured legally, and we follow data usage laws as well as intellectual property legislation. The points of view and measurements presented are only related to the statistically orientated area in plant pathology detection applications. All applications of the results must respect ethics and regional farming practices.

Abstract

Accurate and timely plant disease detection is very much essential in crop health and agricultural yield. The work below describes an improved deep learning model for the classification of plant diseases through images. In doing so, it considers a dataset of 124,636 raw images taken against 37 categories involving healthy and diseased plants. All of the models that were based on VGG16, VGG19, InceptionV3, MobileNetV2, ResNet50, and a hybrid model that integrated ResNet152 with DenseNet201 have been evaluated. Among the tested models, the proposed hybrid model attained the best result in terms of accuracy and robustness. Also utilising LRA (Learning Rate A) reduced the number of iterations without affecting the accuracy. This approach has significantly enhanced plant disease classification. Therefore, this approach is efficient and extendable for automatic disease detection. Our work depicts how DL can prominently support precision agriculture by enabling early detection and prevention of agricultural damage.

Keywords: Plant disease, deep learning, leaf image, VGG16, VGG19, InceptionV3, MobileNetV2, ResNet50, ResNet152, DenseNet201, hybrid model.

Acknowledgement

First and foremost, we would want to thank Allah for his bounties, which have allowed us to finish our thesis without any major setbacks. Moreover, we are highly grateful to our research supervisor, Dr. Md. Ashraful Alam, and co-supervisor Md.Tanzim Reza for their enlightening comments, insightful criticism, and constant support throughout this project work They were quite useful in everything that led to the completion of this study. Lastly, we would like to thank our families, who always pray and support us. They have played an essential role in helping us reach the graduation threshold.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Precise and timely diagnosis of plant diseases is important in agriculture to sustain the healthiness of the crops and reduce yield losses. The accurate identification of plant diseases is quite a big challenge because there is huge diversity among them. Convolutional Neural Networks, a type of deep learning, have proved very effective and promising in recent times for the automation of image-based plant disease classification. In this work, deep learning models will be adopted for classifying images of healthy and sick leaves representing 32 different types of diseases and healthy plant leaves. Performance is analyzed for CNN architectures such as VGG16, ResNet50, DenseNet201, MobileNetV2, and InceptionV3 for plant disease detection. This paper proposes a hybrid model using ResNet152 and DenseNet201 that can improve the performance of classification. The model implements slightly modified LRA functions that helped reduce epochs and hence were more efficient. [23] Such systems are trained on pre-processed images of leaves and implement transfer learning for precise disease classification. The project uses standard metrics like the F1-score, accuracy, precision, and recall for the evaluation of every model. Saliency maps and confusion matrices represent the classification results by highlighting interesting regions in the images to improve understanding and explainability in the model's prediction. It will be a reliable and efficient system for farmers and agricultural experts, which helps at the earliest opportunity intervene to protect crop health and hence early detection of plant diseases.

1.1 Background and Significance

So far, plant disease detection has been done physically; this is cumbersome and time-consuming, with increased chances of errors, especially where diseases that may bear similar symptoms need to be differentiated. Deep learning and artificial intelligence are among the recent transforming forces for many industries, which may also include agriculture by offering automation techniques for disease detection based on image analysis. Originated from the popularity of Convolutional Neural Networks for image classification, big potentials were recorded in plant disease diagnosis by utilizing images of leaves. Therefore, this paper presents a review in classifying 32 plant diseases utilizing deep learning architectures namely: VGG16, ResNet50, DenseNet201, MobileNetV2, InceptionV3, and hybrid model DenseNet201 combined with ResNet152. The methodology is targeted at improving the accuracy and swiftness of the identification of diseases using CNN architecture and transfer learning, benefiting farmers by leading to better decision-making and timely intervention. This study is significant as the potential it holds toward contributing to sustainable agriculture by reducing crop loss, enhancing productivity, and offering an automated and efficient way of early detection of disease. This increases the model's predictability by providing saliency maps and confusion matrices, serving as a very useful proxy for practical agricultural applications.

1.2 The Emergence of Advanced Image Processing Techniques:

The last couple of decades have seen great enhancements, especially with deep learning in image-processing techniques concerning plant disease identification and classification. These enhancements have enabled the automatic processing of images of leaves for disease diagnosis by making use of huge databases and complex algorithms. [13] In fact, CNNs, such as VGG16, ResNet50, and DenseNet, form the backbone of many plant disease diagnoses due to correct pattern identifications within an image, including texture, color variation, and lesion structures. [7] Transfer learning has been integrated such that the pre-trained networks, with little training, can adapt to specific datasets and significantly improve performance. Other techniques involve saliency mapping, augmentation of data, and segmentation of images that help in the detection process, thus making it trustworthy and understandable. These advanced image-processing algorithms allow for early disease diagnosis and provide scalable agricultural solutions since they minimize the interaction of humans and enhance efficiency regarding crop health monitoring.

1.3 The Role of Deep Learning

Deep learning has emerged as a transformative technology in plant disease identification, significantly enhancing the accuracy and effectiveness of diagnosing plant health issues. Deep learning algorithms utilizing convolutional neural networks (CNNs) can autonomously extract intricate information from images of plant leaves [18]. This facilitates the identification of nuanced visual patterns that may indicate various disorders. Deep learning automated feature selection, unlike conventional image processing techniques that predominantly depend on human assessment. This reduces human mistakes and analytical duration. Deep learning algorithms are especially advantageous in practical applications due to their ability to analyze extensive datasets and derive insights from diverse examples, hence enhancing generalization to novel data [23]. This capability is crucial for addressing the challenges posed by the diverse disease manifestations among various plant species and environmental conditions. Moreover, advanced techniques such as transfer learning facilitate the adaptation of pre-trained models to novel datasets, hence optimizing performance even with limited labeling [19]. Deep learning enhances the accuracy of plant disease identification and optimizes agricultural operations by facilitating timely treatments. This ultimately enhances food security and promotes sustainable agricultural practices.

1.4 The Research Objective

Convolutional neural networks represent a form of deep learning that finds great applications and utility across a wide range of domains, including agriculture. Convolutional Neural Networks can classify images with higher capabilities. Based on the above, our project looks to make use of deep learning in early diagnosis of plant diseases through the analysis of leaf images for timely intervention so as to reduce crop losses. The following goals have been set for our research:

Integrated CNN Framework:

The idea behind designing a robust deep learning framework with various CNN architectures was to design an efficient model that could classify leaf images at a higher resolution rate into 32 different categories of plant diseases. Accordingly, the framework in question will leverage several pre-trained CNN architectures, including but not limited to VGG16, ResNet50, MobileNetV2, and InceptionV3 by integrating their functionalities regarding the capture of various features of image characteristics, whether in texture, color, or shape. An integrated model will enhance the generalization of the model across categories of disease in such a way that it would be resistant against changes in leaf appearances because of environmental or disease stages [14]. This framework will also merge powers of complementary architecture of individual CNN models, such as ResNet to handle vanishing gradients or Inception for capturing multi-scale features [15]. The main objectives are to achieve a high level of classification accuracy, speed up the training time, and hence improve plant disease detection even in challenging conditions by providing a practical tool for farmers and agricultural experts for early detection and efficient preventive measures.

Comparative Analysis of CNN Architectures:

It will find the most appropriate and efficient method for plant disease diagnosis by comparing different CNN models such as VGG16, VGG19, and ResNet50 with a hybrid model that embeds strengths from DenseNet201 and ResNet152. Each of the CNN architectures has some special strengths: VGG models are known to be pretty simple and effective for classification tasks, while ResNet is efficient in handling deep networks due to its skip connections, which reduce the vanishing gradient problem. The paper will further conduct a wide comparison among them by checking their efficacies through various arrays of accuracy, precision, recall, the F1-score, and computational efficiency. The hybrid model with DenseNet201 and ResNet152 should, in fact, outperform the individual models by reusing effective features from DenseNet and the robustness of training deep layers from ResNet. It shall further investigate which of the two or their combination provides the best tradeoff between accuracy and computational cost for plant disease detection applications.

Overcoming Data Limitations:

Another important point to discuss is that one of the major challenges in generalizing deep learning approaches concerning the area of plant disease detection is due to small accessibility of data. In relation to this, the proposed study would create a repository of leaf images with annotations from a wide representation of plant species and diseases. A diversified and well-annotated dataset creates for the models a wider representative training resource so that they increase substantially their generalization ability about the different plant species and disease types. In this manner, it would increase results even on model performance for real-world agricultural applications, where great variation in species and diseases usually occurs.

Model Complication:

Also, the challenge of heavy resource-hungry models and imbalanced datasets has encouraged the research to find more suitable alterations and approaches. These goals shall, therefore, enable us to make valuable contributions in the field of agricultural technology with a workable deep learning solution for early plant disease diagnosis.

1.5 Problem Statement

In the production level, traditional agriculture is severely affected by plant diseases, which can lead to significant crop losses and other inconveniences if not identified and addressed early. Traditional methods of plant disease detection depend on manual inspection by experts or experienced farmers, which is time-consuming, labor-intensive, and often subjective to interpretation. Global demand for sustainable agriculture is exponentially rising along with the need to ensure yield security. There is an urgent need for an automated, accurate, and scalable method to identify plant diseases at an early stage.

In the last decade, automating many manual tasks has been proven efficient and reduced the complexity of our lives. Thus, our proposed model is an approach to solve this problem or find a newer solution.

Many researchers have employed Deep CNN approach to detect plant diseases by analyzing images. There remains a challenge in the proper categorization of the diseases that are affecting plants, especially in those situations when multiple diseases show similar symptoms [8]. Traditional models often cannot deal with distinguishing these diseases, perhaps classifying them wrongly, which has very negative effects on agriculture. The issue at hand requires a more advanced deep learning model incorporating state-of-the-art architectures and techniques to improve performance. Word of mouth: taking the state-of-the-art of advanced CNN architecture combined with new training methods for better feature differentiation capability associated with different diseases.

One of the key challenges as far as detection is concerned pertains to imbalanced datasets, which provide many samples toward some diseases and a few toward other diseases. Not all the diseases had same number of reference pictures, which could have led to bias. The overall model might yield a high accuracy, yet fail to detect rare diseases which don't have enough data to feed the model. These kinds of disparities could culminate in models that do well in regular ones but bad on the rare diseases. Neglecting these challenges might, in turn, reduce generalization and finally the quality of the model developed. Some traditional models might work efficiency on other datasets, yet proven less useful for certain datasets. Many techniques can be used to solve this problem during the image preprocessing stage. Some of the ways used in enhancing dataset balance include data augmentation, generation of synthetic samples, resizing, reduction of bad images, and resampling. Several methods were implemented in this study to enable the model to derive valuable information from the underrepresented classes. A hybrid of the Resnet and Densenet model was tuned using LRA, Learning Rate Adjustment. For the deep learning model to be relied upon reliably over all 32 categories of illness, its general robustness has to be enhanced.

Real-time detection and scalability go hand in hand in agricultural applications. Many preexisting deep learning models are also excessively computationally intensive to be deployed on resource-constrained devices, such as either mobile phones or high-configured computing systems. It is very important that a simplified lightweight deep learning model be developed to enable real-time diagnosis of diseases with high accuracy [3]. This may be achieved by the use of optimized architecture for mobile applications, quantization, or model pruning.

Additionally, detection methods become complex as noise and unpredictability occur in leaf images, and handling it becomes necessary. Fluctuations in illumination, backdrops, and perspectives may mask critical attributes, hindering models from optimized performance [15]. In this regard, improving picture preprocessing techniques by approaching solutions to these inconveniences can increase the input data quality. These picture preprocessing techniques, when combined with a well-organized model architecture, can greatly enhance the feature extraction performance and overall efficiency in detection.

The underlying problem remains that these models are biased and do not perform well when leaf images from different species of plants are tested. The preparation of a model for diagnosis of diseases across a large number of plant species requires training the system with a wide dataset of species and their respective diseases. Such a model is more practical in farming situations; it will definitely give a greater degree of precision across a wide number of plant species.

Finally, and most importantly, the most relevant domain of study toward the improvement of detection accuracy is feature extraction and optimization, especially in the subtle or initial stages of an infection. Distinguishing characteristics that would differentiate ill leaves from healthy ones could appear to be hardly identifiable, especially when symptoms remain modest. Advanced methods for feature extraction, such as transfer learning or attention mechanisms, integrated within deep learning frameworks, could raise the sensitivity of the model to early signs of illness. The project will contribute to the improvement of agricultural sustainability and productivity by enhancing plant disease detection technology, addressing three identified critical challenges.

Chapter 2

Literature Review

[6] In this work, the authors have conducted an extensive comparison study of various CNNs for plant disease classification using the PlantVillage dataset, which includes fourteen plant species and twenty-six diseases. They considered its validation performance, F1-score, and number of training epochs required to rank 18 different deep learning architectures, including but not limited to modified and cascaded/hybrid CNNs, and more commonly used ones. The article also elaborates on various deep learning optimizers, such as Adam, RMSProp, and Adagrad, for the performance of the best models. The Xception model trained with the Adam optimizer yielded the best performance in classification, with a validation accuracy of 99.81% and an F1-score of 0.9978, better compared to other methods. This work serves as a very good starting point for applications of deep learning toward plant disease classification in agriculture. [6].

[23] Chavda investigated the performance of two deep learning architectures, namely, DenseNet121 and VGG16, on identifying and classifying the type of plant disease. Both models were trained and fine-tuned using the PlantVillage dataset of fifteen different types of plants, including but not limited to bell peppers, tomatoes, and potatoes. This paper deals with performance metrics such as accuracy, precision, recall, and F1-score. For instance, DenseNet121 generalizes better in the case of rare disease class recognition compared to VGG16. This is evidenced by its much higher accuracy while using much lower processing power. Conclusively, this work provides evidence that the use of DenseNet121 enhances the efficiency and scalability of a system for the detection of plant diseases, which also points well toward its practical applicability. [23].

[20] Joseph, propose a real-time dataset on plant disease identification for cereal grains such as wheat, rice, and maize. They solve issues such as, but not limited to, a general lack of enough real-world data for determination of plant diseases by creating databases for common diseases that attack the said crops at various stages such as early, progressing, and severe. Using this new dataset, this work studies eight CNN models trained using Xception, MobileNet, and InceptionV3. For rice disease classification, impressive results were shown by the Xception, while for wheat and maize, MobileNet outperformed other models. Additionally, the authors introduce a new model in CNN-MRW-CNN. Its performance was sufficient for fine-tuning to ensure its possible utility. real-world agricultural practices. The authors' work

highlights the importance of more diversified and realistic image datasets, which would serve to improve the performance of plant disease identification systems for real-world applications.[20]

[5] This paper has proposed a combination of the FCM-KM and Faster R-CNN algorithms in fast diagnosis related to rice diseases. The proposed fusion in the paper is actually developed to handle problems such as noise, complicated backgrounds, and poor detection accuracy in rice disease image analysis. For noise reduction, the authors propose a multistep approach using both a 2D filtering mask and an optimized K-Means clustering algorithm. Then, FCM-KM optimizes the bounding boxes and Faster R-CNN is applied on the disease identification problem. This gives an accuracy of 96.71% for rice blast, 97.53% for bacterial blight, and 98.26% for sheath blight when tested on 3010 photos. The detection time is also reduced with the fusion method. The work is going to prove that Faster R-CNN outperforms the previous approaches in plant disease diagnosis, bringing further improvement in efficiency and accuracy.
According to Jiang,[3] this paper proposes a CNN-based deep learning model for detecting, in real time, various diseases affecting apple leaves. These include Alternaria leaf spot, brown spot, mosaic, gray spot, and rust, which are some of the most common diseases that normally attack apple leaves. Among them, the authors have collected 26,377 images of these spots, which they use to create a dataset. The authors have proposed an improved detection performance with a model called INAR-SSD. It uses an improved feature extraction for objects of multiple sizes based on the GoogLeNet Inception module with Rainbow concatenation. In the experiment, it can be found out that the detection speed is 23.13 FPS and mAP is 78.80%. As compared to previous methods, the approach in the paper was more robust and highly efficient in real time for the detection of diseases on apple leaves. This activity is very important in the early detection of diseases in plants related to agricultural purposes.[3].

In their study Adnan, [13] presents a hybrid approach for improving the state-of-the-art performance of the identification of plant diseases across domains by leveraging two popular domain adaptation methods: DANN and CORAL. More exactly, the proposed DANN-CORAL model aims at making deep learning models flexible against changes in light, surroundings, and camera settings that occur during agricultural imaging. Meanwhile, adversarial training in DANN captures domain-invariant features, while CORAL helps in aligning the second-order statistics across source and target domains. Upon testing the model using the PlantVillage dataset, the model performed much better in comparison to the performance of the baseline models by achieving an accuracy of 91.39%, precision of 93.36%, and recall of 88.9%. This work shows tremendous promise for generalization in plant disease detection, allowing models trained in controlled settings to excel in real-world agricultural situations. [13].

[22] It proposes a deep learning model that incorporates DANN and CORAL as a hybrid technique in order to do a better job at plant disease identification across

domains. The proposed DANN-CORAL model targets the facilitation of deep learning models to be more stable and flexible while facing changes in illumination, surroundings, and camera settings occurring in agricultural imaging. DANN features domain-invariant features by adversarial training, and CORAL aligns the second-order statistics between source and target domains. The performance had outperformed some baseline models on the PlantVillage dataset with an accuracy of 91.39%, precision of 93.36%, and recall of 88.9%. This will presumably make a great shift in generalization for plant disease detection, as models trained in controlled environments are performing well in natural agricultural environments. [22].

This study [18] presents the application of a YOLOv4 deep learning model to leaf image data for plant disease detection and classification. The study utilized an expanded PlantVillage dataset containing more than 50,000 images of healthy and diseased leaves from 14 different plant species. Data augmentation techniques like histogram equalization and horizontal flipping were applied to enhance the diagnosis capability of the YOLOv4 model. The proposed model, YOLOv4, is 99.99% accurate and proves to be the best model in plant disease identification compared to other models, DenseNet and AlexNet. This paper shows how a proposed model works effectively for real-time detection, hence bringing more development in precision agriculture. The proposed method will allow researchers and farmers to detect and effectively diagnose plant illnesses by consuming a minimum amount of time. The paper reviews how YOLOv4 can become a game-changing technique for agricultural crop management as well as plant pathology. [18]

[9]Plant Disease Detection and Classification by Deep Learning—A Review" is a paper that gives an intense review of traditional and advanced deep learning approaches, which have been applied for detecting diseases in plants. It categorizes the methods in deep learning systematically as supervised methods with specific focus on CNN-based models like AlexNet, GoogLeNet, and ResNet, which indeed performed marvelously in diagnosing diseases from plant leaf images. It discusses the limitations imposed by scarce datasets and focuses on transfer learning and GANs for data augmentation, among other strategies. Visualization techniques are also explored, such as saliency maps that would contribute to model explainability, underlining some important gaps with respect to the need for large and diversified datasets and early disease detection in real agricultural applications. This paper will be of great help to the academics in building an enhanced and more capable plant disease detection model by the use of deep learning. [9]

[1] The paper "Plant Disease Detection Using Image Processing" by Suresha et al. gives an overview of image processing applied for the detection and classification of diseases on plant leaves. It systematically describes the steps used in the process of identifying diseases in plants; the method will range from acquisition of the picture, pre-processing, segmentation, features extraction to classification. The segmentation and features extraction methods include k-means clustering, Otsu thresholding, and color co-occurrence methods. ANN, or Artificial Neural Networks, and SVM, or Support Vector Machine, are various approaches that may be used in the classification of diseases. This study has sketched the strength of image processing in automating plant disease identification, which reduces manual involvement in terms

of skill and investment of time. It thus is a good reference for the application of image processing in agriculture regarding disease surveillance to enhance crop yield.[1]

Below is the research [2] entitled "Plant Disease Detection Using Machine Learning" that applied Random Forest to classify healthy and diseased plant leaves targeting papaya leaves. Image processing techniques are used with feature extraction by HOG, Hu Moments, Haralick Texture, and Color Histograms for extracting shape, texture, and color features from the images. The model achieved an accuracy of 70% with a dataset size of 160 images. While the Random Forest performed the best among the three classifiers using SVM and k-Nearest Neighbors, the paper points out various limitations, including the relatively small size of their dataset and that better results may have been possible by employing more sophisticated image feature extraction methods. This research underlines the application potential of machine learning for plant disease diagnosis and points out several limitations that will need to be considered: real-world unpredictability and the need for larger, more diverse datasets in order to attain better accuracy and robustness.[2]

[16]The article "Plant Disease Detection and Classification Using Machine Learning and Deep Learning Techniques: Current Trends and Challenges" critically examines contemporary AI-based approaches for the detection and classification of plant diseases. It further divides these methodologies into machine learning and deep learning approaches, reviewing the pros and cons of each. Machine learning methodologies design the design elements manually, like shape, color, and texture, while they use models like Support Vector Machines and Random Forest extensively. Deep learning methodologies, on the other hand, will extract features automatically such as CNN-based models offering better accuracy and robustness. The paper emphasizes that PlantVillage, among other datasets, has been used for training models, and indicates some key limitations such as limited diversity in these data sets, complexities in segmenting images, and inability to identify multiple diseases or those with identical symptoms. Further, the paper discusses some performance metrics including accuracy, precision, and F1-score. The report concludes with the identification of the direction of further research: mitigation of dataset limits by data augmentation and enhancement of multi-disease detection. This review, therefore, gives a significant contribution towards the development of technologies for detecting plant diseases. [16]

[8]The paper "Plant Disease Detection Using Deep Learning" discusses several deep learning models for the detection of plant diseases, comparing mainly Convolutional Neural Networks with Transformer networks. This work investigates models such as InceptionV3, custom CNN architecture design, and modern Visual Transformers. The dataset used is an improved version of PlantVillage, containing 87.9k images spread over 37 classes of plant-disease pairs. The study underlines that standard CNNs are very powerful in image classification tasks, but the Visual Transformers outperform CNNs, with the best validation accuracy of 97.98% obtained with the Large Transformer Network. The research points out that transformer models are efficient in parameters and accuracy and therefore represent a very promising model

for resource-limited settings. It realizes challenges such as model overfitting and further study required for transformer topologies on vision tasks, placing transformers at a prospective future direction in plant disease detection.[8]

[11]The model is a 14-layer deep convolutional neural network which was developed utilizing 147,750 images, at the horizon of 59 categories. This model was trained for 1000 epochs by utilizing a very high configured computation unit. As expected, the model's outcome was an exceptionally high 99.97% accuracy, surpassing all the similar models that emphasize this problem. In this paper, the researchers tried to improve the diversity of a given dataset by augmentation and eliminating the requirement of new data collection. Various techniques including GAN, flipping, cropping and rotation have been proven effective in successful training, especially when in conjunction. The research observes the effectiveness of BIM, GAN and NST approaches for the categorization of plant leaf diseases. The hyperparameter tuning scores indicate that the dataset size, augmentation techniques and parameter selections are very important for creating effective disease detection models. [11]

[15]This study proposes a standard procedure for identifying the diseases of olive leaves by using a hybrid model. This was made up of the dataset, which contained 4,138 images in three categories of disease and one healthy category. Development of 30 DHMs based on six previously trained CNNs and combined with five existing machine learning classifiers was done in the research. The integration of EfficientNetB0 with logistic regression gave a very high accuracy of 96.14%. It has been confirmed that this is an effective and sophisticated machine learning methodology, enhancing the detection of diseases within the olive cultivation culture and offering an efficient and much more accurate solution than methods at present. [15]

[10]This study addresses the problem of holistic identification of grape leaf diseases like leaf blight, black rot, stable, and black measles. The researchers, using transfer learning in the EfficientNet B7 architecture, have been able to do this task with an accuracy of 98.7%. This model follows a four-step approach: data pre-processing, model training, feature reduction, and classification. The features were further optimized using a logistic regression method in an attempt to improve the accuracy of classification through the reduction of dimensions and classification. The systematic methodology applied in this research is the Hybrid CNN model. The discussed approach lays a benchmark for future progress in the field of automated plant disease identification with artificial intelligence.[10]

[4]This research implies utilizing the classifiers for diagnosing diseases of plants and identifying disease-affected areas to improve understanding and management. The research carefully evaluated more than 30 research articles on deep learning algorithms for plant leaf disease diagnosis, focusing on farmers' challenges diagnosing pathogens including bacteria and fungi. It observes the dataset size, preprocessing, and classifier performance of the models. It studies the transfer learning deep convolutional neural network (CNN) models' data augmentation, image pre-processing, and performance measures of other researchers. The researchers find out that adding classifiers for diagnosing plant diseases and identifying disease-affected zones for bet-

ter comprehension and management of diseases can be very effective. By utilizing convolution neural networks, this research was able to detect and recognize 16 different plant diseases. The experimentations were performed on a comparatively low configured computer, yet the accuracy was remarkable.The training of the deep learning model had 50 epochs and attained an accuracy rate of 98% during training. During the test, this model was capable of detecting diseases at an exceptional rate of 99% accuracy. [4]
[14]This research paper uses one of the latest developed advanced deep-learning models to analyze a method for the detection of diseases found in olive leaves. Moreover, it combines WOA with ANN to achieve higher accuracy in the detection of diseases of olive leaves. The program is set to accurately detect diseases such as Aculus olearius and olive peacock spots through training; over a collection of images of olive leaves. This technique goes beyond what was previously used and designed; therefore, it is considered an efficient tool in the early diagnosis of diseases in agriculture.[14]

[24]The "LeafSpotNet" research article proposes a sophisticated deep learning methodology, MobileNetV3, intended for the detection of leaf spot diseases in jasmine plants. The proposed model can classify with an accuracy of 97%, using Conditional GAN for data augmentation and Particle Swarm Optimization for the optimal selection of features. It will work appropriately on real-time agricultural applications meant for quick diagnosis of the diseases and crop management, since it will be designed to work effectively in the variations of camera angular positions and lighting conditions.[24]

[17]The paper describes a deep learning-based methodology that avails the usage of the Vision Transformers for disease identification in mango leaves and has been able to obtain an accuracy as high as 99.75%. The developed model outperformed traditional CNNs and was further optimized for real-time detection. Further, a mobile application was developed which assisted in on-site testing and showcased the performance of the technology in agriculture while revealing its potential to enhance disease management protocols.[17]

[12]This research paper discusses the effectiveness of various deep learning approaches to plant leaf disease detection, an essential challenge to agriculture. It is basically a project on deep learning-based AI development which would be quite effective and efficient for replacing manual disease detections, which is burdensome and error-prone. They have made a comparison between eight deep learning architectures such as GoogleNet, ResNet, MobileNet, and AlexNet. So far as the dataset benchmark was concerned, they had considered a benchmark plant dataset called PlantVillage consisting of 20,640 images of tomato, pepper, and potato plants. It calibrated several key hyperparameters in an attempt to optimize performance for models considering epochs, learning rates, and mini batch sizes. Some of the most accurate models were ResNet50 and ResNet101; their detection rate was almost flawless. Among these optimizers, Adam proved better than others, such as SGDM, at avoiding local minima and enhancing the process of training. The results have underlined deep learning's potential in the revolution of plant disease detection by providing quicker and more precise agricultural diagnostics.[12]

[19]Early diagnosis of rice leaf diseases is crucial in order to avoid severe damage to the food supply, and the article provides a critical analysis in respect of deep learning methods for the automation of disease detection by four most prominent strategies: ensemble learning, transfer learning, hybrid methods, and customized CNNs. While transfer learning involves the use of pre-trained models such as ResNet and VGG that are initialized in order to improve performance where information is scarce, custom CNNs are built from scratch to detect patterns related to disease from images. Ensemble learning employs several models to enhance performance, whereas hybrid techniques involve the integration of deep learning with machine learning methods for more complicated detection challenges. The main issues that were detected from this study are that the differentiation of various diseases that present similar symptoms is rather difficult, and no richly and extensively annotated dataset is available, which can enable the good generalization performance of the models to practical scenarios. This report further discusses the development of comprehensive datasets, hyperparameter optimization, and other advanced techniques to handle the mentioned challenges, such as Explainable AI. This combination of drones and IoT devices will facilitate real-time disease monitoring, enhance hybrid model-based detection accuracy, and improve learning approaches in this area for further research.[19]

[21]Olive Leaf Disease Detection via Wavelet Transform and Feature Fusion of Pre-Trained Deep Learning Models," proposes an enhanced method of olive leaf disease detection. The model uses wavelet transformations and pre-trained deep learning models such as ResNet-152-V2, DenseNet-201, and EfficientNet-B7 to enhance feature extraction. The dataset has been created right at the commencement with the use of augmentation methods such as resizing, rescaling, flipping, and so on. The wavelet transform decomposes the images into a number of frequency bands. Features are extracted from the third-level subband by pretrained algorithms to form an integrated complete disease classification image. The proposed model outperforms others with the best result on the accuracy of classification of 99.72% in a dataset containing healthy and unhealthy leaves of olives. This result outperforms earlier models and justifies the fact that this very technique is effective enough for early detection of diseases in olive plants. According to a study, this approach in agriculture can automate disease identification and reduce crop loss. The issues of overfitting and complexity were overcome by applying data augmentation and fusing models carefully. Wavelet transformations enhance deep learning model-based image classification performance in disease identification; this is one of the vital enhancements in agricultural technology.[21]

[7]Optimizing Deep Learning Model for Plant Disease Detection Using Particle Swarm Optimizer: This research is related to plant disease detection using deep learning. The authors have used a pre-trained AlexNet CNN with classifications of plant diseases in five crops into 25 categories. Particle Swarm Optimization selects the most relevant attributes from the deep learning model so that the model is optimized by reducing the complexity and improving the accuracy of the model. Overfitting has been checked by modifying the AlexNet architecture with the addition of dropout and weight regularization. To increase the robustness of the dataset, data

augmentation techniques, such as image rotation and flipping, were applied. The revised model achieved a very good accuracy of 98.83%, with sensitivity, specificity, precision, and F-score all greater than 98%. Deep learning goes very well in concert with optimization algorithms regarding PSO for the classification of plant diseases; this would indicate, in effect, that ensemble models may assist in improvements of other crops. [7]

Chapter 3

Methodology

3.1 Work Plan Diagram

This workflow involves some major steps: data collection, splitting of the dataset into 124,636 images of leaves representing 37 categories of plant diseases into training, validation, and test sets. The dataset needed further cleaning to remove corrupted or mislabeled images, and preprocessing included resizing of images and normalization. In this experiment, six CNN models were selected: VGG16, VGG19, ResNet50, MobileNetV2, InceptionV3, and ensembled DenseNet201 and ResNet152. The performance was analyzed using training vs. validation loss/accuracy graphs, confusion matrices, and classification reports to choose the most accurate and simultaneously the most efficient network that performs plant disease classification. Cross-validation was also used in verifying model stability and its generalization capability on random splits.

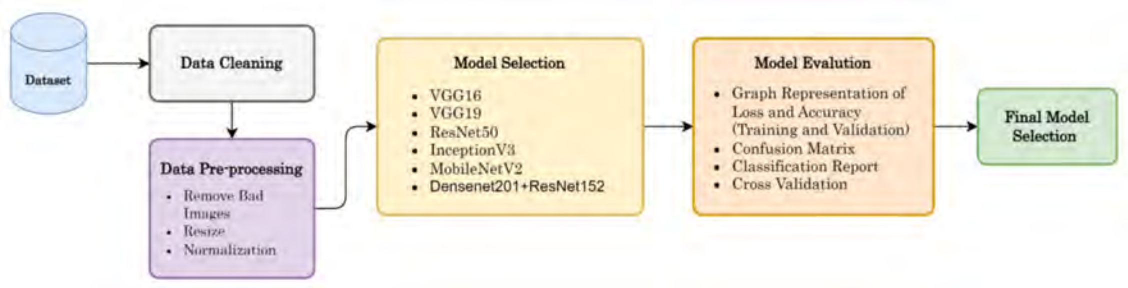

Figure 3.1: Workflow Diagram

3.2 Data Collection

The dataset for plant disease identification comprises 124,636 annotated photos of leaves, including 37 unique classes of plant illnesses and healthy leaf categories. These labels include a broad spectrum of plants and illnesses, guaranteeing a varied dataset that allows train models to generalize efficiently across numerous plant species and situations.
The dataset's class distribution comprises a substantial quantity of images for each disease, including Tomato Yellow Leaf Curl Virus with 7,877 images and Citrus greening with 5,787 images, alongside healthy leaf instances, such as Apple Healthy

Figure 3.2: Bell Paper Bacterial

Figure 3.3: Apple Black Rot

Figure 3.4: Strawberry Leaf Scorch

Figure 3.5: Tomato Healthy

Figure 3.6: Dataset Images

(4,234 images) and Tomato Healthy (4,053 images). This variety guarantees that models will encounter both sick and healthy leaf pictures, facilitating accurate categorization. The dataset contains classes with markedly fewer photos, such as Citrus Healthy (87 images) and Citrus Black Spot (312 images), which may provide issues owing to uneven class representation.

The whole dataset is partitioned into three subsets for model training and assessment: training set (99,708 pictures), validation set (12,464 images), and test set (12,464 images). The picture has dimensions of 224 by 224 pixels and has three color channels (RGB). Keras' ImageDataGenerator is used for data augmentation and preprocessing, implementing pixel scaling to a range of -1 to 1, hence facilitating accelerated convergence during model training.

The dataset's richness, including many plant species and disease circumstances, renders it optimal for training deep learning models. The meticulous allocation of training, validation, and test sets guarantees that models are assessed under diverse settings, mitigating overfitting and enhancing generalization. Moreover, issues like

class imbalance, especially in underrepresented categories such as Citrus Healthy and Citrus Black spot, need the use of approaches like class-weight balancing or data augmentation to enhance model performance across all categories.
This dataset, including a broad spectrum of illnesses and healthy circumstances, is essential for training an optimal deep learning model proficient in properly identifying plant diseases across various species.

3.2.1 Cleaning and Preprocessing

The cleaning and pre-processing of the dataset are essential steps in preparing the leaf images for training the deep learning model. The dataset is first examined for faulty or corrupted images, which are eliminated to maintain data integrity. Subsequently, photos are downsized to a standardized dimension of 224x224 pixels to provide consistency throughout the dataset, which is crucial for effective model training. To improve the model's learning capacity, the pixel values of the images are normalized by a scaling function, converting the original values to a range of -1 to 1. This normalization process accelerates convergence during training by ensuring that all input features possess a comparable scale. Also, for further improvement on the models performance, we implemented procedures like removing and resizing deficient images .The dataset is divided into training, validation, and testing subsets to facilitate effective evaluation of model performance, ensuring the model adaptability from a varied array of examples while retaining the capacity to assess its predicted accuracy on novel data.

Figure 3.7: Dataset Preprocessing

3.2.2 Class Dropping

Sometimes, the deep learning model's extracted features, especially CNNs, are so robust that removing some of the less important or redundant attributes does not change the performance; for example, edges, textures, and patterns from images.

Given that the dropped attributes are either redundant or much overlapped with other features, performance changes for this model since the important attributes remain there and can be used by the model in classification. Secondly, if the dropped attributes are not crucial in distinguishing the classes, a powerful CNN like ResNet or DenseNet, well known for extracting hierarchical features, will generalize well. In the case of class imbalance, when the dropped attributes are from the under-represented classes, the model may stay biased toward the majority classes because dropping such attributes has a minimal effect on the overall performance of the model. However, if those dropped attributes have a high importance in the model for classifying specific categories, like the "Tomato Yellow Leaf Curl Virus" class in our dataset, this may degrade their accuracy of detection and perform negatively for such categories.

To act on this, our procedure also includes the dropping classes, which will be shown in Result and Analysis part of our paper. The classes dropped includes the following classes: Citrus Black Spot, Citrus Healthy and Citrus Canker. They were removed due to having insignificant numbers of data within them as they would make the procedure of the model's performance much more complex; not to mention our dataset has 37 classes in total. This procedure results in model to work with fewer classes and work more profoundly. Primarily all, including our proposed model, had utilized 37 attributes with amazing outcome. As expected, we achieved good overall performance, but the model was underperforming on the low-density classes. As our proposed model achieved exceptional performance, we dropped 3 less dense attributes and we reran the whole procedure; amazingly, the result stayed unchanged.

```
labels
Tomato Yellow Leaf Curl Virus    7877
Citrus greening                  5787
Peach Bacterial spot             4594
Tomato Bacterial spot            4355
Tomato Late blight               4323
Apple Healthy                    4234
Tomato Septoria leaf spot        4091
Tomato Healthy                   4053
Bell pepper Healthy              4014
Tomato Spider mites              3854
Grape Black Measles              3783
Tomato Target Spot               3688
Corn Common rust                 3682
Grape Black rot                  3596
Corn Northern Leaf Blight        3548
Potato Early blight              3532
Potato Late blight               3521
Corn Healthy                     3486
Tomato Early blight              3478
Bell pepper Bacterial spot       3449
Tomato Leaf Mold                 3389
Strawberry Leaf scorch           3327
Apple Scab                       3232
Grape Isariopsis Leaf Spot       3228
Cherry Healthy                   3182
Cherry Powdery mildew            3156
Apple Black rot                  3105
Strawberry Healthy               2819
Tomato Mosaic virus              2655
Corn Gray leaf spot              2629
Peach Healthy                    2623
Grape Healthy                    2594
Cedar apple rust                 2553
Potato Healthy                   2432
```

Figure 3.8: Visualization of Dataset After Attribute Dropping

3.2.3 Model Selection

We decided to use the synergies between different state-of-the-art deep learning architectures known for their excellence in image classification tasks, in order to choose models to be used for our project. Both VGG16 and VGG19 were evaluated since they are both layered architectures, a form which has been very effective in representations for extracting complex image features. Among these, ResNet50 was selected because its residual learning architecture makes the training of deeper networks easier and reduces the problem of a vanishing gradient. MobileNetV2 was included for its lightweight architecture, hence making it suitable for resource-constrained environments without sacrificing performance necessarily. InceptionV3 was selected to take advantage of its multi-scale processing, thus giving the model the capability of capturing features at higher resolutions. We proposed a hybrid model that combined DenseNet201 and ResNet152 in order to utilize the advantages of both architectures, hence optimizing feature propagation and reuse, which eventually improves detection performance across a broad dataset of 37 plant disease categories. This combination of models allows for the systematic study of singular and collective model efficacies in accurately diagnosing plant diseases from leaf images.

VGG16:

The VGG16 is a convolutional neural network architecture that presents an ideal tradeoff between simplicity and excellence in performance for image classification tasks. It was proposed by the Visual Geometry Group of the University of Oxford, which used small 3×3 convolutional filters consecutively to allow deeper networks with minimal increases in computational complexity.[1]

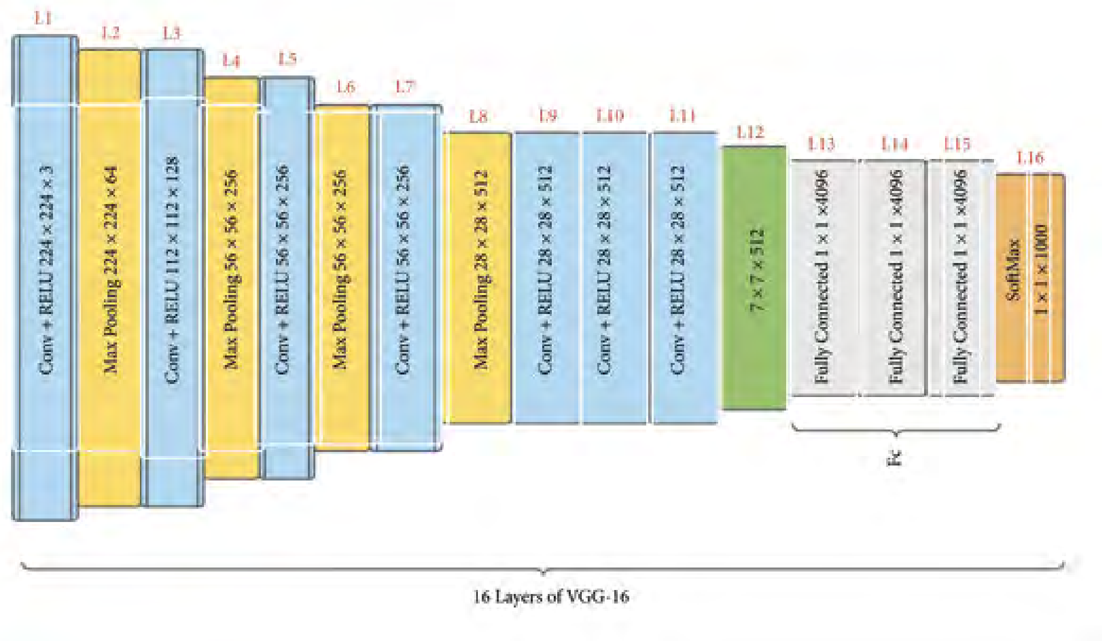

Figure 3.9: Basic Architecture of VGG16 Model

The figure 3.4 contains, of which all the 16 layers of the VGG16 model are divided into 5 different layers (conv, ReLU, max-pooling, softmax, and fully connected). These convolutional layers takes a 224 by 224 RGB value with a fixed length as the source.Then the data gets updated by a bunch of convolutional layers with a 3 ×3 visual field (the minimum size to retain the concepts of rightmost/leftmost, rise/bottom, and core). The max-pooling layer follows the convolutional layer. Convolution and ReLU, combining their working methods, process the data together. In a few of the configurations, it additionally includes 11 filters of convolution, which is usually regarded as a proportional change of the input networks. Every convolution stride is specified as a single pixel, and the spatial padding of such a convolution layer source is set the same as a single pixel for 3 by 3 convolution layers, following convolution, keeping spatial resolution. Instead of having a bunch of hyperparameters, VGG16 concentrated on having a 3 ×3 filter input layer with a stride of 1 and would always use the same padding and max-pool structure. Spatial pooling is done via layers of 5 max-pooling that replicate part of the convolutional layers.

VGG19:

VGG19 is an extended model architecture from VGG16, having a total of 19 layers: 16 convolutional and 5 levels of max-poling and 3 layers connected in a fully connected manner. This architecture is similar to that of VGG16; however, VGG19 has more convolutional layers and thus allows finer feature extraction from an image. Small 3x3 convolutional filters have been used in VGG19 with ReLU as the activation function, enhancing the nonlinear nature of the model [1]. It is mainly known for its fine tuning of image patterns, which makes it very powerful for tasks related to image classification. The architecture also has dropout layers that have been a technique aimed at reducing overfitting by randomly switching off some portions of input units during training. Thus, VGG19 successfully combined depth and simplicity, providing a strong basis for various computer vision tasks.

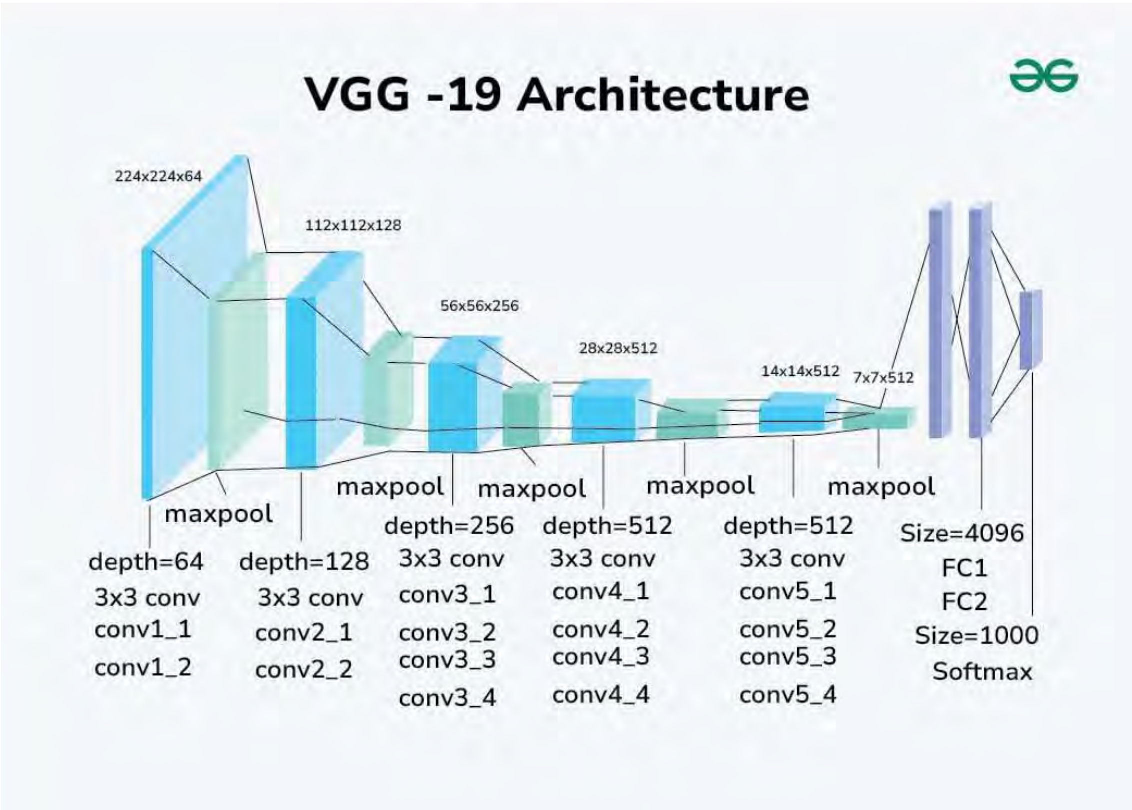

Figure 3.10: Basic Architecture of VGG19 Model

MobileNetV2:

MobileNetV2 is a lightweight deep learning architecture for mobile and edge devices. It emphasizes efficiency and has been duly optimized for accuracy. It was an extension of the original MobileNet, using depthwise separable convolutions that vastly reduced the number of parameters and computational cost. This model uses a slim architecture with an inverted residual structure that allows it to extract features in a very fast manner while retaining a slim model size. The main novelty of MobileNetV2 is the use of linear bottlenecks, which reduce the information loss by maintaining representational capacity while being in lower-dimensional layers [16]. As a result, MobileNetV2 achieves competitive accuracy on many tasks, including image classification and object detection, with very impressive efficiency in terms of speed and resource consumption. It is thus suitable for deployment in resource-constrained environments.

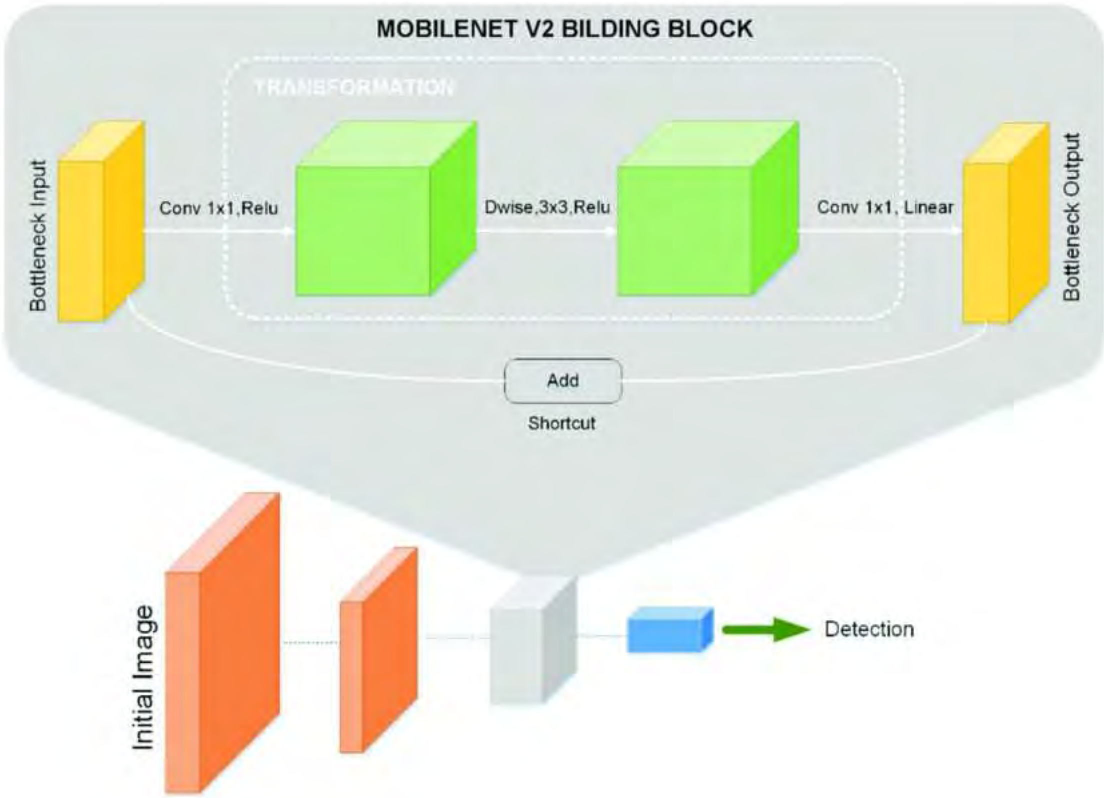

Figure 3.11: Basic Architecture of MobileNet V2 Model

ResNet50

ResNet50 is the residual network comprising 50 layers and 26 million parameters. In the year of 2015, Microsoft introduced a deep convolutional neural network model called residual network [10]. The residual network learned from residuals, that is, the difference between features acquired before its layer's input instead of acquiring features itself. ResNet introduced skip connections to make the data flow between levels easy. The following graph shows how the 50 layers of ResNet50 are interrelated. How ResNet50's architecture is divided into four portions may be shown as in the below picture. It can take an input with a stride of three in the channel and multiples of 32 in the dimension. Let's consider that the size of the filter input is 224 x 224 x 3. In all ResNet models, kernel size is 77% of initial convolution and 33% of max pooling. Thus, the first stage of the network is three residual blocks, each comprising three layers. The size of the kernel for the convolution in each of three layers within the stage 1 unit is 64, 64 and 128, respectively. The bending arrow means the same connection. It contains a residual block with a convolution stride of two, hence reducing the input dimensions by half in both height and width while doubling the channel width, as shown by the dotted arrows between the connections. Coming to the bottom layer of the system, the layer is an average pooling layer.

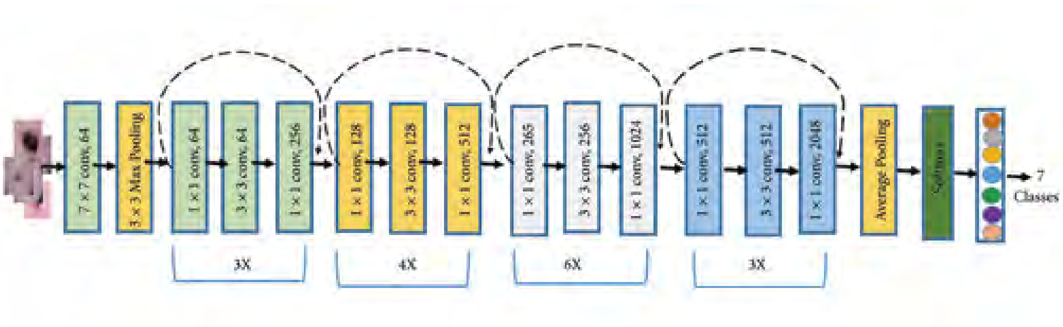

Figure 3.12: Basic Architecture of Resnet50 Model

Inception V3:

This is a deep learning architecture that extends the conventional CNN architecture with a special structure combining several convolutional filters of different sizes all together. This, therefore, enables the model to capture a wide range of features as well as patterns from input images. A very significant novelty of InceptionV3 consists in "inception modules" that let the network develop different levels of abstraction because of the different kernel sizes and pooling within one layer [9]. Also, InceptionV3 implemented batch normalization and factorized convolutions with the aim of lightening computational burdens and achieving improvement in efficiency during training and in model performance. It has achieved state-of-the-art performance in image classification tasks, and this architecture is well noted for its robustness on complex datasets. In fact, most applications, from object identification to image categorization, have found InceptionV3 useful.

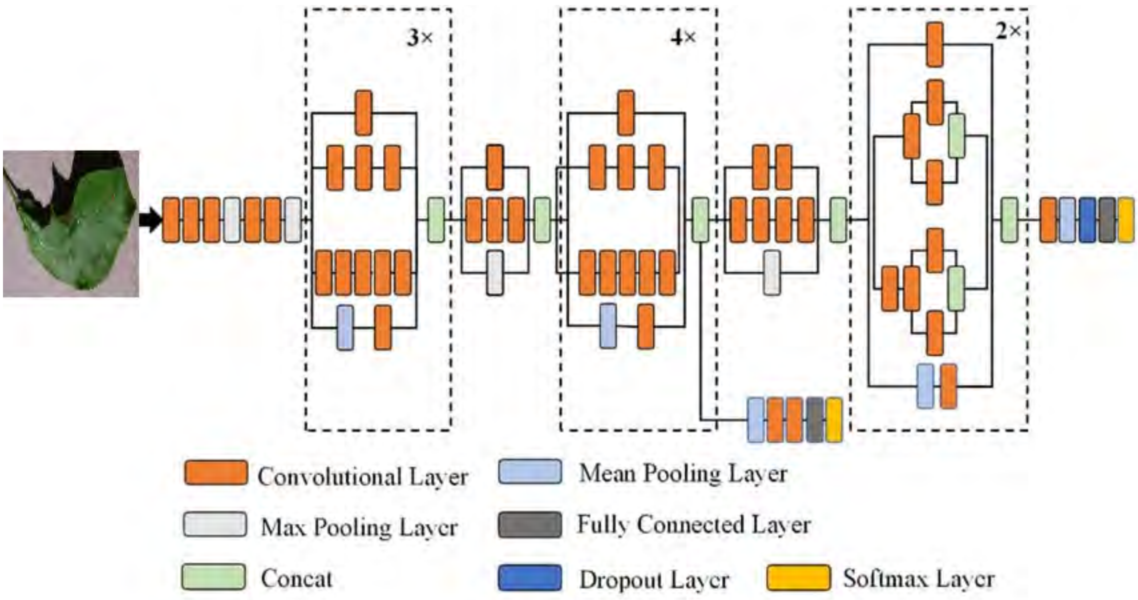

Figure 3.13: Basic Architecture of Inception V3 Model

Hybrid Model (combining DenseNet201 and ResNet152):

The concept of the hybrid deep learning model unites many architectures that could improve its effectiveness in solving such complicated tasks as plant disease detection.

A hybrid deep learning model is designed which combines many CNN architectures like DenseNet201 and ResNet152 for leveraging their complementary strengths: the effective feature reuse in DenseNet and skip connections in ResNet overcome the problem of vanishing gradients. This is because the architecture learns deep features in the leaf images and provides very similar performance for a wide range of settings. A hybrid approach works exceptionally well on multiple categories of plant diseases, improving both accuracy and generalization compared to single models.

Layer (type)	Output Shape	Param #	Connected to
input_layer_5 (InputLayer)	(None, 224, 224, 3)	0	-
densenet201 (Functional)	(None, 7, 7, 1920)	18,321,984	input_layer_5[0][0]
resnet152 (Functional)	(None, 7, 7, 2048)	58,370,944	input_layer_5[0][0]
concatenate_1 (Concatenate)	(None, 7, 7, 3968)	0	densenet201[0][0], resnet152[0][0]
flatten_1 (Flatten)	(None, 194432)	0	concatenate_1[0][0]
dense_3 (Dense)	(None, 1024)	199,099,392	flatten_1[0][0]
dense_4 (Dense)	(None, 512)	524,800	dense_3[0][0]
dense_5 (Dense)	(None, 256)	131,328	dense_4[0][0]
dense_6 (Dense)	(None, 37)	9,509	dense_5[0][0]

Figure 3.14: Layer of Hybrid Model (DenseNet201 and ResNet152)

The chosen hybrid combination of DenseNet201 and ResNet152 is particularly effective for handling imbalanced datasets due to its ability to leverage the complementary strengths of both architectures:

- DenseNet201 utilizes feature reuse, ensuring that important details, especially those relevant to minority classes, are retained throughout the network. This prevents the model from losing critical information about underrepresented classes, making it better at classifying rare examples.
- ResNet152, on the other hand, is well-suited for deeper feature extraction, enabling the model to capture complex, high-level patterns in the data while preserving the flow of gradients. This helps the model remain robust even as the number of layers increases, which is important for learning from a large number of classes in an imbalanced dataset.
- Some classes in the dataset are significantly underrepresented, which can lead to poor generalization for these classes.
- Due to the large number of data points and classes, the model's training time began to increase exponentially.
- To mitigate this issue, two key modifications were introduced:
    1. **Early Stopping:**When two consecutive epochs yield the same performance in terms of validation loss or another metric, the model training is halted early. This technique, called early stopping, prevents the

model from continuing training unnecessarily when no further improvement is observed. It also helps to avoid overfitting by selecting the best-performing model based on validation performance, thus reducing the number of iterations required for training while achieving high accuracy.

2. **Learning Rate Adjustment (LRA):**In addition to early stopping, learning rate scheduling was implemented to further optimize training. Learning Rate Adjustment (LRA) reduces the learning rate when the model's performance plateaus, encouraging better convergence. This adjustment allows the model to escape local minima or saddle points and helps the optimizer make finer adjustments, leading to higher accuracy with fewer epochs.
   This research used a slightly modified LRA implementation that proved effective in achieving superior accuracy with fewer iterations. The learning rate was adjusted dynamically based on the model's performance, which prevented wasted computation and sped up the convergence process.

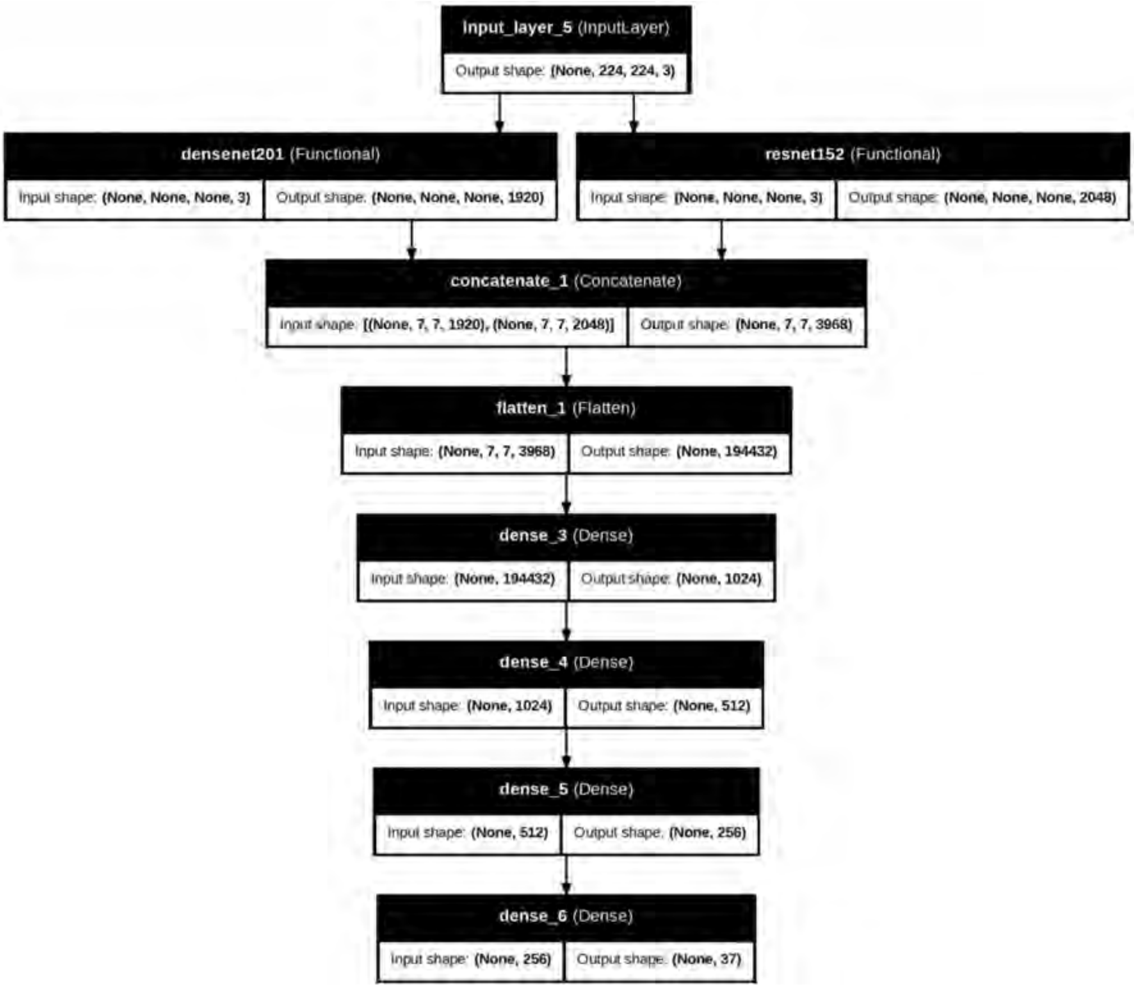

Figure 3.15: Architecture of Hybrid Model (DenseNet201 and ResNet152)

3.3 Model Evaluation

The models will be evaluated using various metrics such as precision, accuracy, recall, and F1 score, which collectively offer a comprehensive understanding of each model's performance.
Precision measures the proportion of correctly identified positive instances among all predicted positive instances. It is calculated using the formula:

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive} \tag{3.1}$$

Recall, also referred to as sensitivity or the true positive rate, assesses the proportion of actual positives that are correctly identified. The formula for recall is:

$$Recall = \frac{True\ Positive}{True\ Positive + False\ Negetive} \tag{3.2}$$

The F1 score combines both precision and recall, providing a balance between the two by calculating their harmonic mean:

$$F1 = 2 * \frac{Precision\ and\ Recall}{True\ Positive + False\ Positive} \tag{3.3}$$

These requirements will be utilized to assess the performance of the models for effective comparison. This comparison facilitates the identification of the most effective model for the specific task and offers essential insights for model selection and implementation. Furthermore, confusion matrix and classification reports will be included for further assessment. This rigorous evaluation guarantees that the chosen models obtain high accuracy while preserving a robust balance between precision and recall, therefore augmenting their dependability in practical applications.

Chapter 4

Result and Analysis

Model	Training		Validation	
	Loss	Accuracy	Loss	Accuracy
VGG16	0.6743	0.9138	0.6017	0.9336
VGG19	0.7229	0.9035	0.6469	0.9223
ResNet50	1.5486	0.7097	1.4130	0.7567
MobilenetV2	0.6527	0.9266	0.5979	0.9418
InceptionV3	0.7727	0.9186	0.7080	0.9338
Proposed Model (DenseNet201+ResNet152)	0.0028	0.9991	0.0839	0.9835

Table 4.1: TRAINING AND VALIDATION LOSS/ACCURACY FOR ALL THE MODELS

Model	Precision	Recall	F1 Score	Accuracy
VGG16	0.89	0.90	0.89	0.93
VGG19	0.89	0.88	0.88	0.92
ResNet50	0.72	0.70	0.70	0.76
MobilenetV2	0.90	0.91	0.90	0.94
InceptionV3	0.90	0.89	0.89	0.93
Proposed Model (DenseNet201+ResNet152)	0.97	0.96	0.97	0.98

Table 4.2: PRECISION, RECALL, F1 SCORE and ACCURACY of ALL THE MODELS

4.1 VGG16

- **Training Loss:** 0.6743
- **Training Accuracy:** 91.38%
- **Validation Loss:** 0.6017
- **Validation Accuracy:** 93.36%

- **Analysis:** In terms of accuracy, VGG16 does very well, especially on the validation set (93.36%). Low loss values during training and validation suggest that the model has learned to generalize effectively with minimal overfitting. However, with training and validation losses that are substantially higher, the model may not be as efficient as others.
- **Precision:** 0.93
- **Recall:** 0.92
- **F1 Score:** 0.92
- **Accuracy:** 93.36%
- **Summary:**With a high precision of 0.93, VGG16 accurately and efficiently identifies diseased plants while minimizing false positives. The model does a decent job at identifying most sick plants with a recall of 0.92. The model's balanced F1 score of 0.92 shows that it does a decent job at detecting plants with diseases and avoiding false positives. When compared to other models, VGG16's accuracy of 93.36% is reasonable but not exceptional.

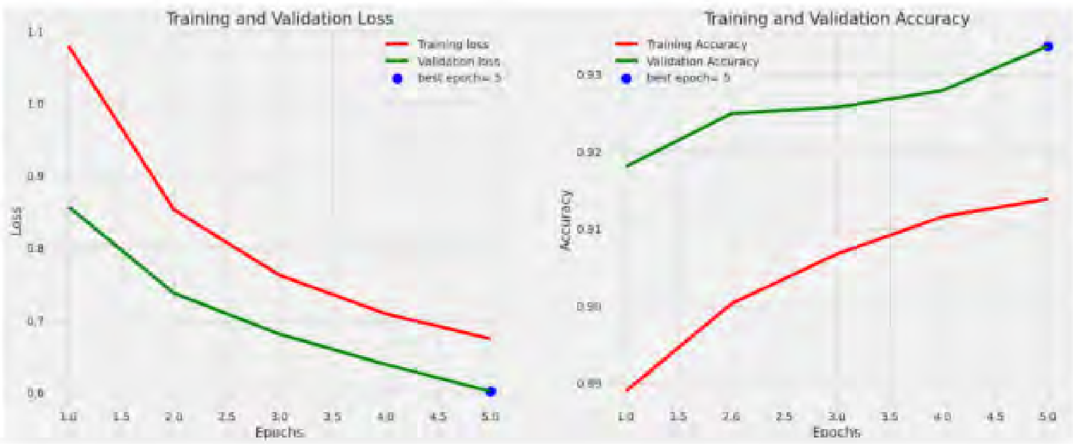

Figure 4.1: Training and Validation Loss & Accuracy graph for VGG16

The graphs illustrate the model's performance regarding loss and accuracy across the training epochs. On the left side, the training loss, depicted by the red line, commences at roughly 1.1 during the initial epoch, signifying incorrect predictions, but progressively reduces to around 0.6 by the fifth epoch as the model enhances its performance. The validation loss, indicated by the green line, commences at approximately 1.0 and progressively decreases, achieving roughly 0.6 by the fifth epoch, demonstrating the model's improved capacity for generalization to novel data. On the right side, the training accuracy, represented by the red line, commences at almost 90% in the initial epoch and marginally increases to surpass 91% by the fifth epoch. The validation accuracy, represented by the green line, commences at approximately 92% and subsequently increases, achieving roughly 93% by the fifth epoch, signifying that the model is working effectively on additional data.

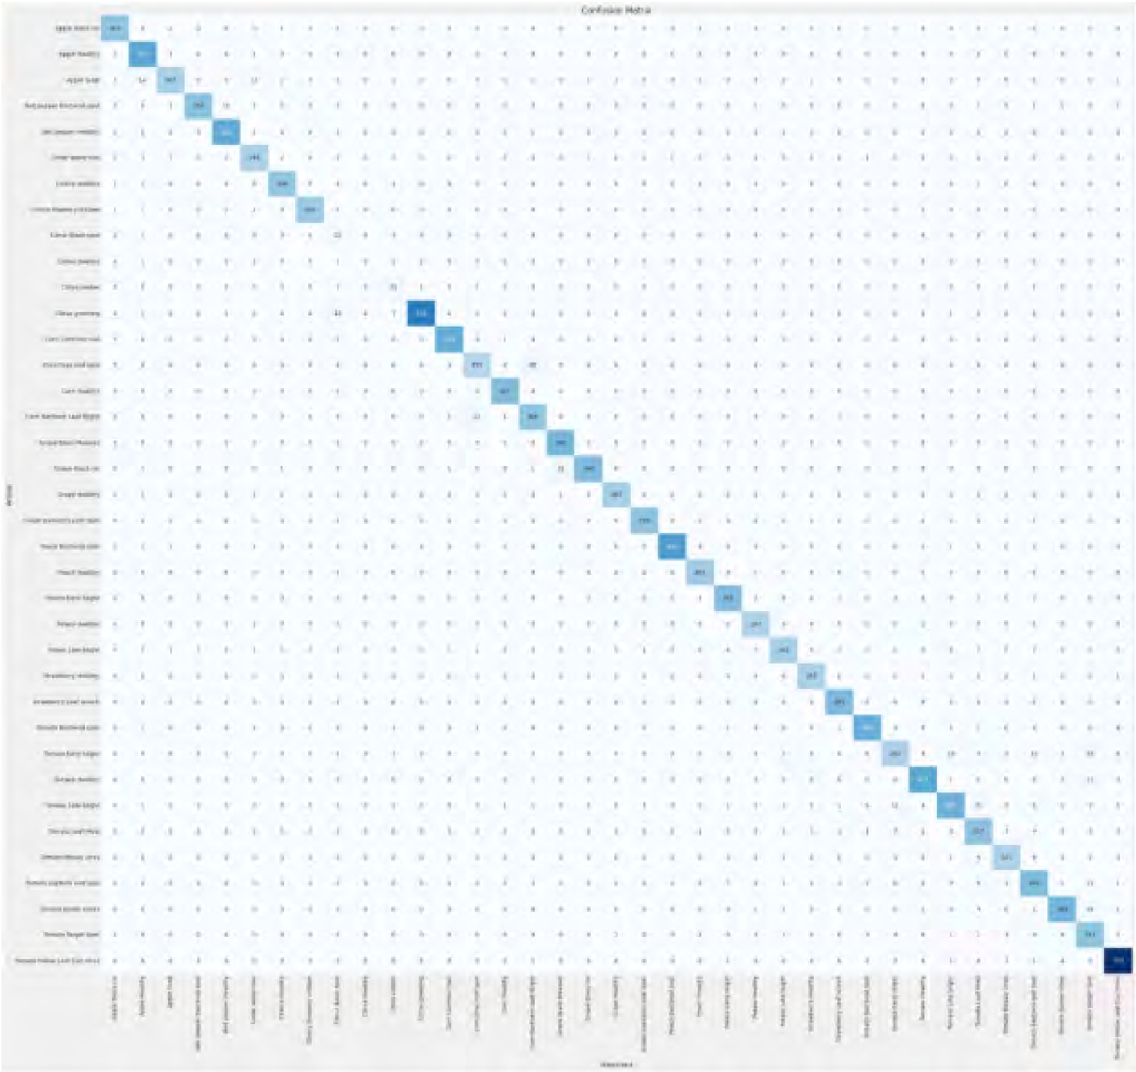

Figure 4.2: CONFUSION MATRIX FOR VGG16

The confusion matrix for the VGG16 model demonstrates its efficacy in categorizing images into 32 plant disease classifications, with a significant diagonal reflecting a high accuracy in predictions. The algorithm accurately diagnoses 346 samples of Grape Black Rot and 380 samples of Tomato Late Blight, demonstrating overall excellent classification accuracy. The misclassification rates are minimal, with categories such as Apple Scab (411 correct) and Tomato Yellow Leaf Curl Virus (773 correct, the highest) exhibiting strong performance. Minor misclassifications arise, such as Tomato Healthy being incorrectly identified as Tomato Mosaic Virus, possibly because to analogous symptoms, and Corn Gray Leaf Spot being mistaken with Corn Common Rust. The errors are negligible. Rare diseases such as Peach Bacterial Spot (446 accurate) demonstrate commendable classification performance, despite the limited number of samples. The VGG16 model exhibits robust performance across several illness categories, with misclassifications primarily occurring among visually analogous disorders.

4.2 VGG19

- **Training Loss:** 0.7229
- **Training Accuracy:** 90.35%
- **Validation Loss:** 0.6469
- **Validation Accuracy:** 92.23%
- **Analysis:**

  There was no discernible improvement in performance due to the additional depth, as VGG19 exhibited somewhat worse training and validation accuracy compared to VGG16. This model seems to have encountered more challenges in learning the dataset, as evidenced by its higher training and validation losses compared to VGG16.

- **Precision:** 0.92
- **Recall:** 0.91
- **F1 Score:** 0.91
- **Accuracy:** 92.23%
- **Summary:**While it's not quite as precise as VGG16 (0.92), it does a decent job of preventing the false diagnosis of healthy plants as unhealthy. When it comes to detecting all sick plants, VGG19 does marginally worse with a recall of 0.91. A model that achieves a balance between recall and precision but falls short of VGG16 (F1 score of 0.91) is considered to be performing adequately. The fact that VGG19's accuracy is somewhat lower than VGG16's (92.23%) suggests that the extra depth of VGG19 may not have substantially enhanced performance.

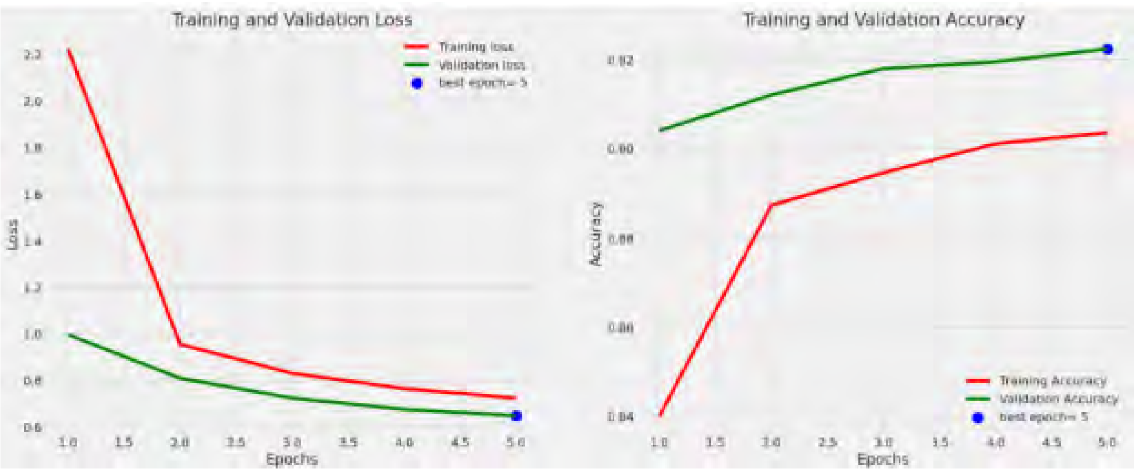

Figure 4.3: Training and Validation Loss & Accuracy graph for VGG19

The graphs illustrate the model's performance regarding loss and accuracy across the training epochs. On the left side, the training loss, indicated by the red line, commences at approximately 2.2, signifying suboptimal first predictions. The model rapidly assimilates the data patterns, resulting in a substantial reduction in loss following the initial epoch, which progressively declines to approximately 0.7 by the fifth epoch. The validation loss, indicated by the green line, begins at approximately 1.0, which is lower than the training loss, implying that the model initially exhibits superior performance on unknown data. The training and validation losses consistently decline, obtaining roughly 0.7 by the fifth epoch, signifying that the model is not overfitting and demonstrates strong generalization capabilities to unique data. The training accuracy, represented by the red line on the right side, commences at 84% and steadily increases, attaining approximately 91% by the fifth epoch as the model enhances its classification of the training data. The validation accuracy, represented by the green line, begins at 91% and rises to approximately 92% by epoch five, with the peak validation accuracy denoted by a blue dot, indicating optimal model performance on unseen data at this point in time.

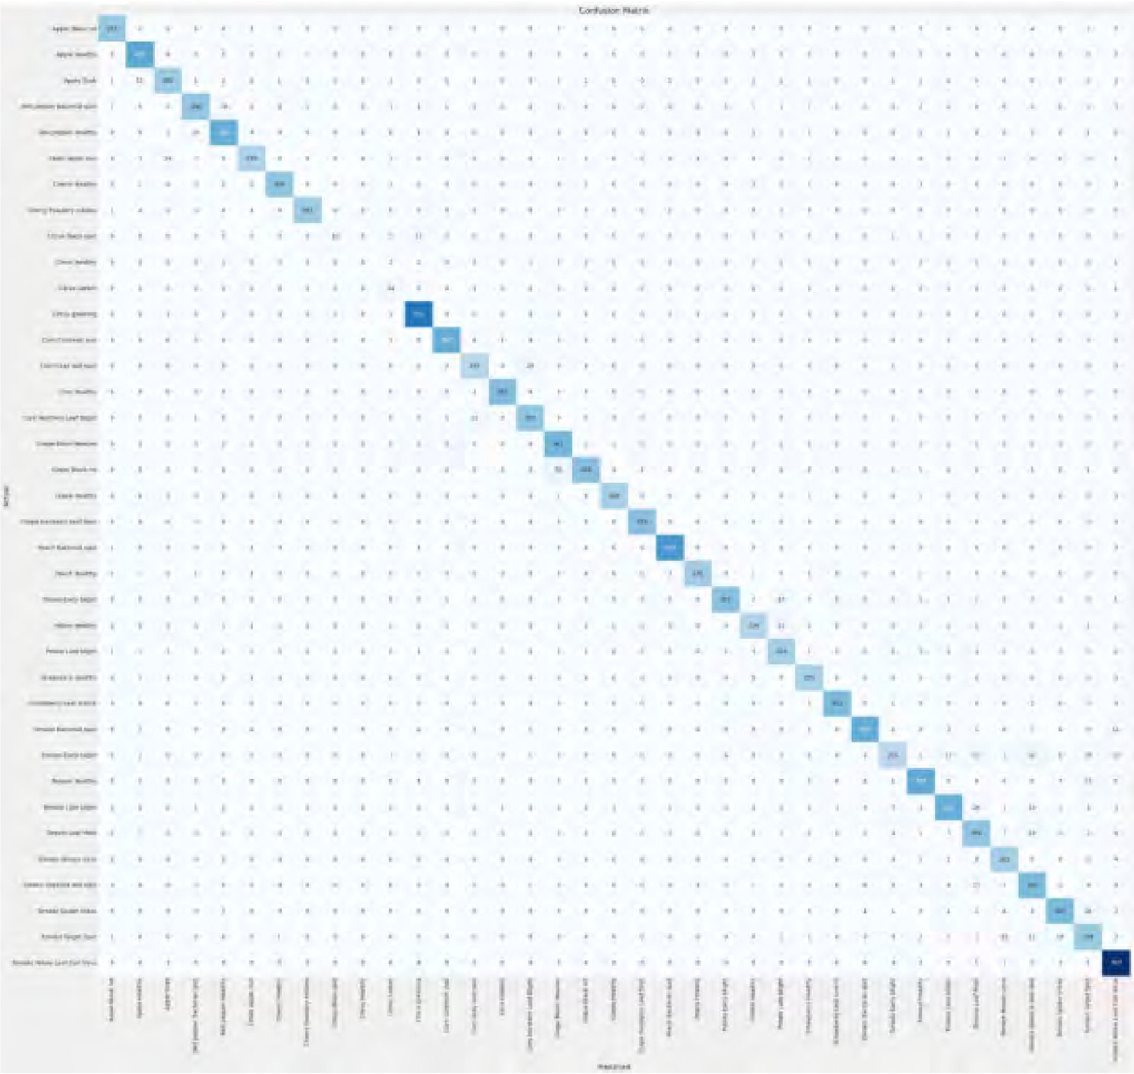

Figure 4.4: CONFUSION MATRIX FOR VGG19

The confusion matrix for the VGG19 model illustrates its performance across 32 plant disease categories, with a significant diagonal that signifies accurate classifications for the majority of samples. For instance, 281 samples of Apple Scab and

767 samples of Tomato Yellow Leaf Curl Virus are accurately diagnosed. Numerous categories have great accuracy, including Grape Black Rot (325 correct) and Tomato Mosaic Virus (263 correct), illustrating the model's efficacy. Nonetheless, certain misclassifications arise, such as between Apple Scab and Apple Black Rot, as well as among corn illnesses like Corn Gray Leaf Spot and Corn Northern Leaf Blight, perhaps attributable to analogous visual symptoms. Certain categories, such as Tomato Late Blight (386 correct) and Cherry Healthy (335 correct), demonstrate nearly flawless classification. The matrix underscores difficulties in differentiating between illnesses exhibiting identical symptoms, such as the merging of Potato Early Blight and Potato Late Blight; yet, overall misclassifications remain negligible. Broader categories, such as Tomato Yellow Leaf Curl Virus, exhibit great accuracy, while even narrower categories like Strawberry Leaf Scorch (342 accurate) have remarkable performance, highlighting the efficacy of the VGG19 model in plant disease categorization.

4.3 ResNet50

- **Training Loss:** 1.5486
- **Training Accuracy:** 70.97%
- **Validation Loss:** 1.4130
- **Validation Accuracy:** 75.67%
- **Analysis:** When contrasted with the other models, ResNet50's accuracy and loss are comparatively low. It fails spectacularly at capturing the dataset's complexity, with training accuracy at 70.97% and validation accuracy at 75.67%. The high loss values during training and validation suggest that the model may have underfitted the data, failing to adequately capture the patterns.
- **Precision:** 0.72
- **Recall:** 0.70
- **F1 Score:** 0.71
- **Accuracy:** 75.67%
- **Summary:**There are more false positives since ResNet50 has a low accuracy of 0.72 when it comes to identifying sick plants compared to healthy ones. It fails to detect a large percentage of diseased plants, with a recall of 0.70. Because of its unsatisfactory precision-to-recall ratio, ResNet50 has the lowest F1 score (0.71) of all the models. Perhaps as a result of underfitting, the model was unable to adequately manage the complexity of the dataset, as seen by its accuracy of 75.67 percent.

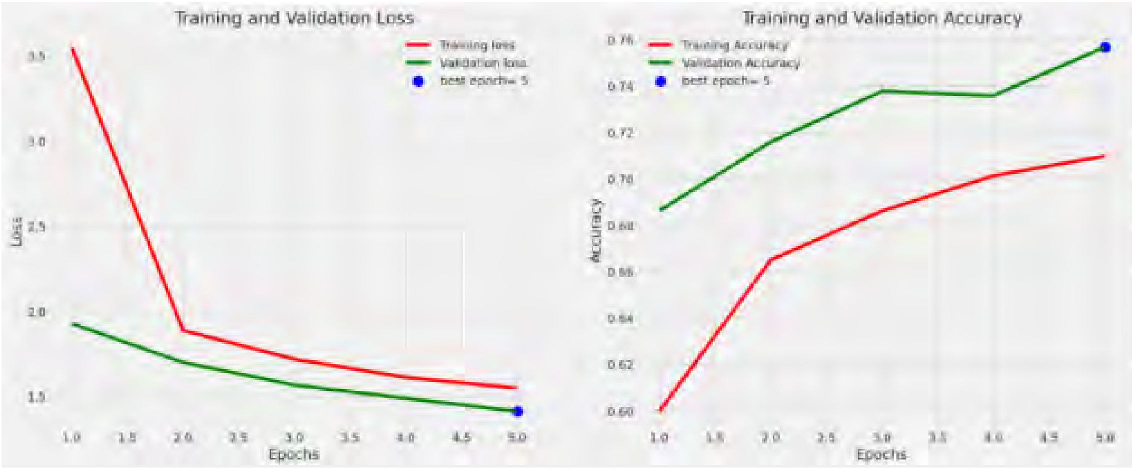

Figure 4.5: Training and Validation Loss & Accuracy graph for ResNer50

This graph illustrates the training and validation loss and accuracy of the ResNet50 model over five epochs. The training loss (red line) commences at approximately 3.5 and consistently declines to around 1.5 by the fifth period, signifying progressive learning. The validation loss (green line) initially falls below the training loss, indicating superior initial generalization, with both losses converging with time, suggesting consistent model learning. The training accuracy (red line) begins at 0.64 and steadily grows to 0.70, reflecting gradual improvement, whereas the validation accuracy (green line) initiates at 0.68 and escalates to 0.75 by the fifth epoch, signifying effective generalization and reliable performance on unique data.

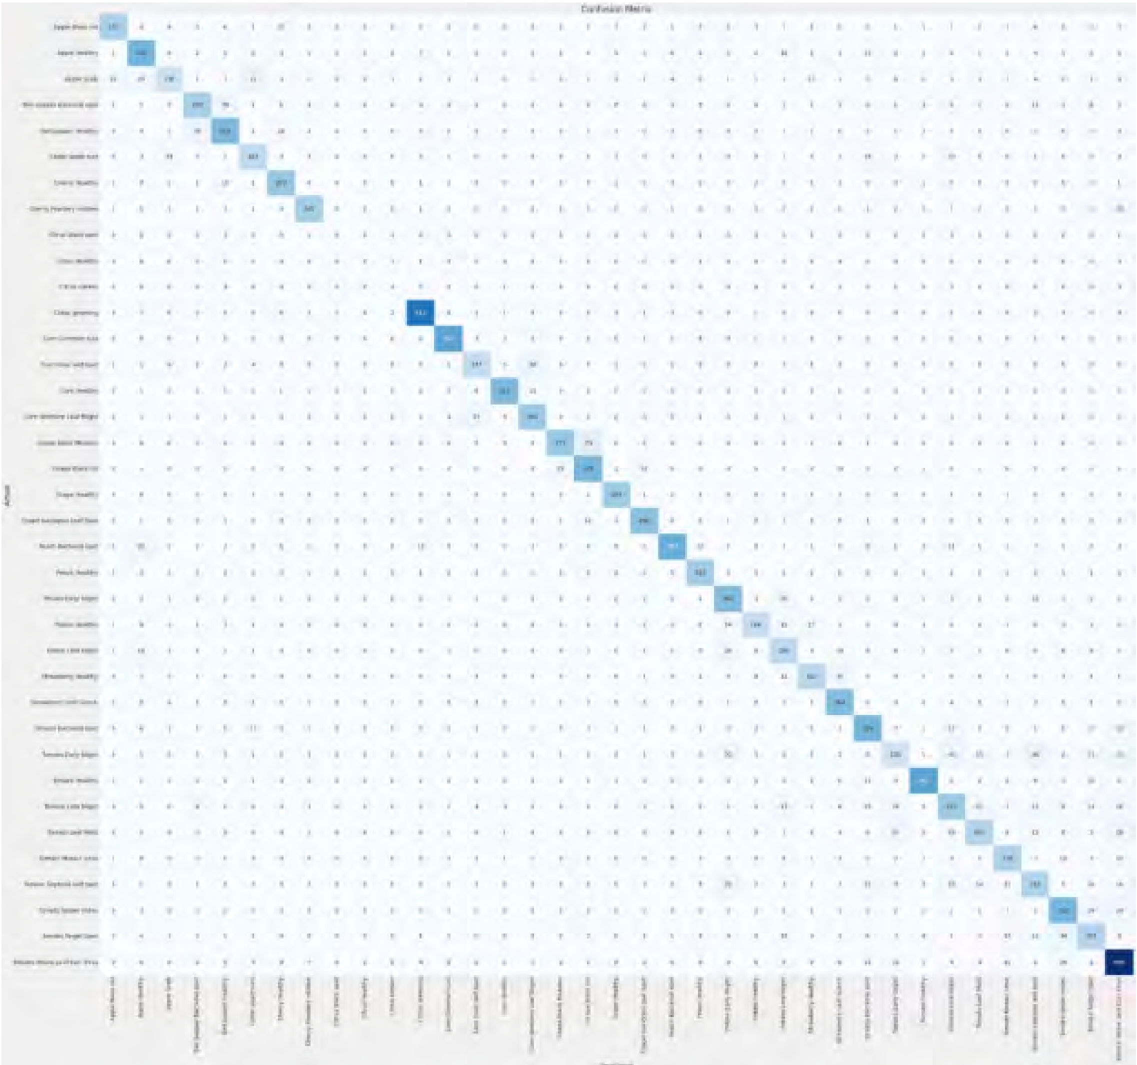

Figure 4.6: CONFUSION MATRIX FOR ResNet50

The confusion matrix illustrates the efficacy of the ResNet50 model in categorizing different plant diseases. Categories which include Citrus Greening (515 accurate predictions) and Tomato Healthy (382 accurate predictions) have great precision, reflecting robust model effectiveness in these domains. Nonetheless, certain misclassifications arise, exemplified by Apple Black Rot being mistaken for Apple Scab (19 instances) and Corn Common Rust being associated with Corn Northern Leaf Blight (20 instances), highlighting difficulties due to the visual similarities of the diseases. Strawberry Leaf Scorch is also incorrectly classified as Strawberry Healthy in 13 instances. The model is proficient in identifying diseases but has challenges in differentiating between identical diseases, especially within the same species.

4.4 MobileNetV2:

- **Training Loss:** 0.6527
- **Training Accuracy:** 92.66%
- **Validation Loss:** 0.5979

- **Validation Accuracy:** 94.18%
- **Analysis:** With a low validation loss and a high validation accuracy of 94.18%, MobileNetV2 demonstrates good performance. It also demonstrates significant generalization. Since it is not as complicated as deeper models like ResNet or VGG, MobileNetV2 is able to be efficient without compromising performance.
- **Precision:** 0.94
- **Recall:** 0.91
- **F1 Score:** 0.93
- **Accuracy:** 94.18%
- **Summary:**With a precision of 0.94, MobileNetV2 clearly has few false positives and does a decent job at disease classification. With a recall of 0.91, it's safe to say that the model finds the majority of diseased plants; however, it could miss a few. Achieving a balance between recall and precision, MobileNetV2 achieved an F1 score of 0.93. Ideal for real-time applications, this model has a comparatively minimal complexity, great performance, and an outstanding accuracy rate of 94.18

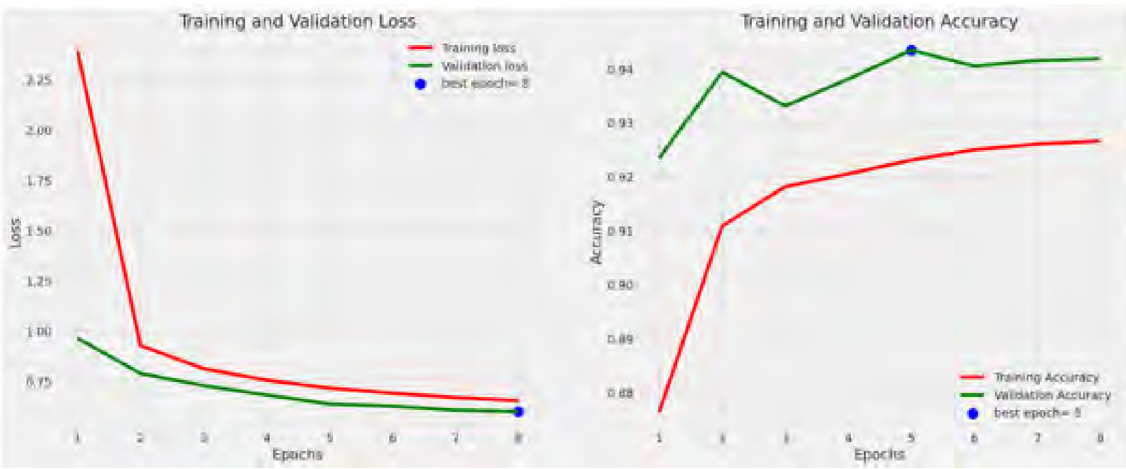

Figure 4.7: Training and Validation Loss & Accuracy graph for MobileNetV2

This graph illustrates the training and validation loss and accuracy of the MobileNetV2 model over 8 epochs. The training loss (red line) commences at approximately 2.25 and significantly declines to around 0.75 by epoch 8, suggesting that the model is acquiring knowledge from the training data effectively. The validation loss (green line) declines, beginning at 1.0 and aligning with the training loss around epoch 7, indicating that the model generalizes effectively to unseen data without considerable overfitting. The training accuracy (red line) commences at roughly 0.88 and progressively climbs to approximately 0.92 by the eighth epoch, but the validation accuracy (green line) begins at a higher value of about 0.93 and advances to around 0.94 by epoch 8. A minor decline occurs at epoch 5; however, the model demonstrates consistent performance on validation data, indicating better generalization.

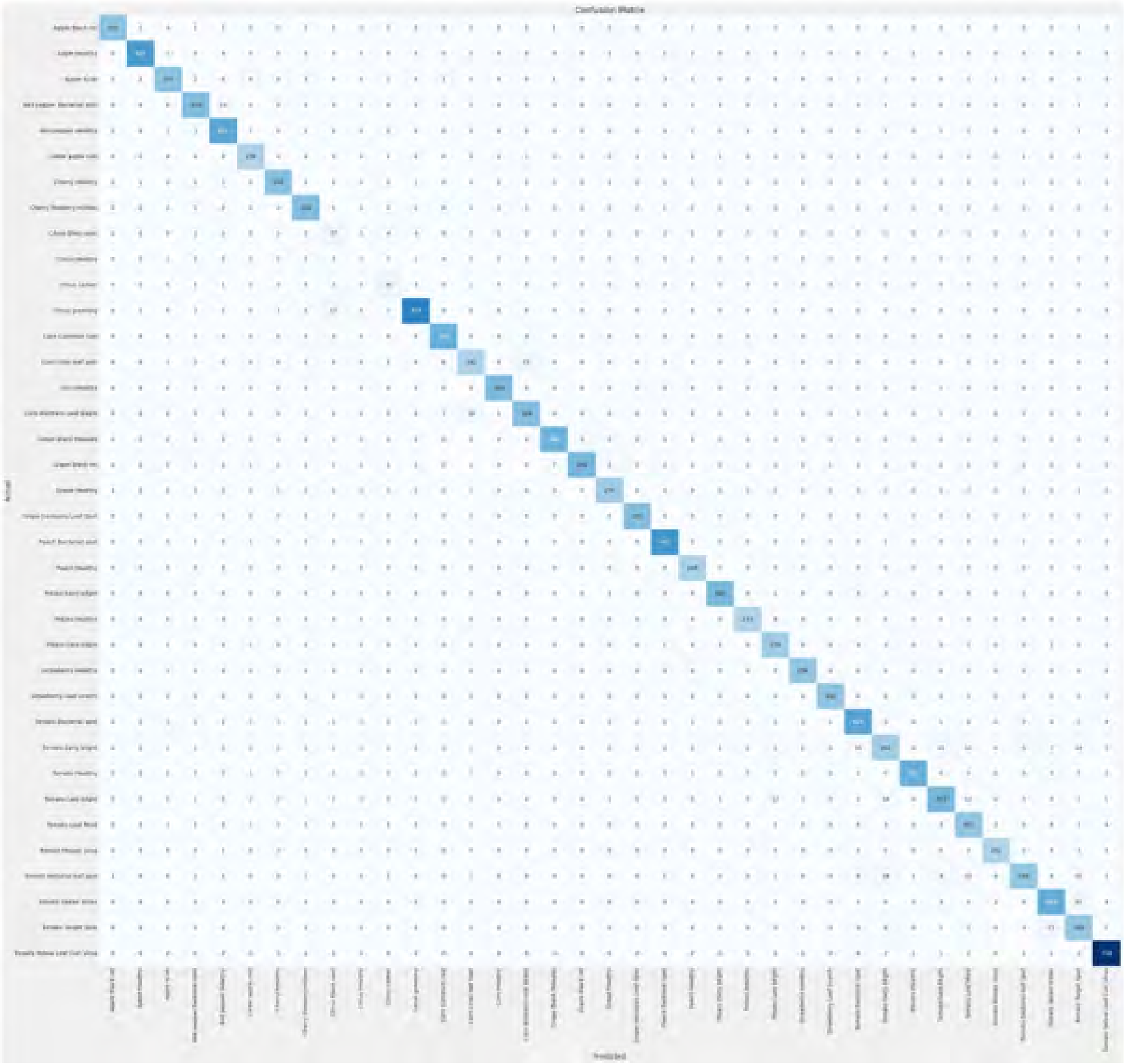

Figure 4.8: Confusion Matrix for MobileNetV2

The confusion matrix illustrates the accuracy of the MobileNetV2 model in categorizing plant diseases. Categories such as Tomato Healthy (381 accurate predictions) and Citrus Greening (522 accurate predictions) illustrate the model's substantial precision in these instances. However, certain categories exhibit ambiguity, with Apple Black Rot incorrectly classified as Apple Scab in 14 instances and Corn Northern Leaf Blight misidentified as Corn Common Rust in 27 instances. Strawberry Leaf Scorch is an established category with 306 accurate predictions; yet, it is often conflated with Strawberry Healthy. Although the model effectively differentiates among most diseases, specific categories, particularly within the same plant species, demonstrate misclassification, indicating a necessity for additional modification to further distinction.

4.5 InceptionV3:

- **Training Loss:** 0.7727
- **Training Accuracy:** 91.86%

- **Validation Loss:** 0.7080
- **Validation Accuracy:** 93.38%
- **Analysis:** When compared to VGG16, InceptionV3 performs admirably, particularly in validation accuracy (93.38%). This model might not have shown good performance as well as others, such as MobileNetV2 or the proposed hybrid model, due to slightly larger training and validation losses.
- **Precision:** 0.93
- **Recall:** 0.92
- **F1 Score:** 0.92
- **Accuracy:** 93.38%
- **Summary:**Just like VGG16, InceptionV3 has a high precision of 0.93 in identifying diseased plants. It does a decent job of identifying most sick plants, with a recall of 0.92. The F1 score of 0.92 shows that, like VGG16, InceptionV3 keeps an acceptable mix of recall and precision. Impressively, the model's accuracy (93.38%) is on par with, if not better than, VGG16's performance.

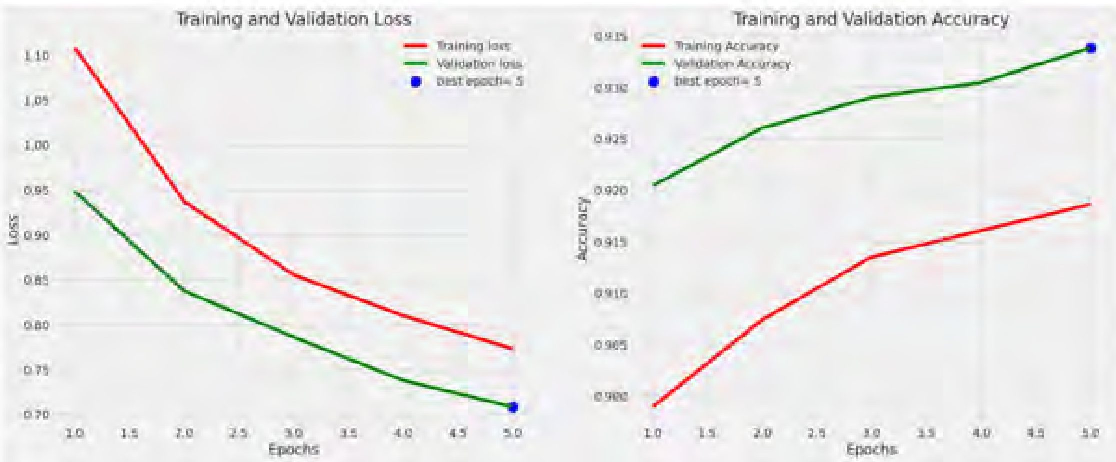

Figure 4.9: Training and Validation Loss Accuracy graph for InceptionV3

This graph illustrates the training and validation loss and accuracy of the InceptionV3 model over 5 epochs. The training loss (red line) initiates at approximately 1.1 and consistently declines to roughly 0.75 by the fifth epoch, indicating the model's gradual progress in fitting the training data. The validation loss (green line) begins at approximately 0.95 and consistently declines, achieving roughly 0.70 by the fifth epoch, suggesting effective model generalization without indications of overfitting. The training accuracy (red line) initiates at approximately 0.91 and progressively ascends to 0.92, whereas the validation accuracy (green line) commences at a higher value of about 0.93 and experiences a more noticeable improvement, attaining nearly 0.935 by the conclusion. The constantly elevated validation accuracy indicates that the model is generalizing proficiently to additional data. InceptionV3 has robust learning efficacy, demonstrated by the decline in both training and validation loss, alongside an enhancement in accuracy throughout the epochs.

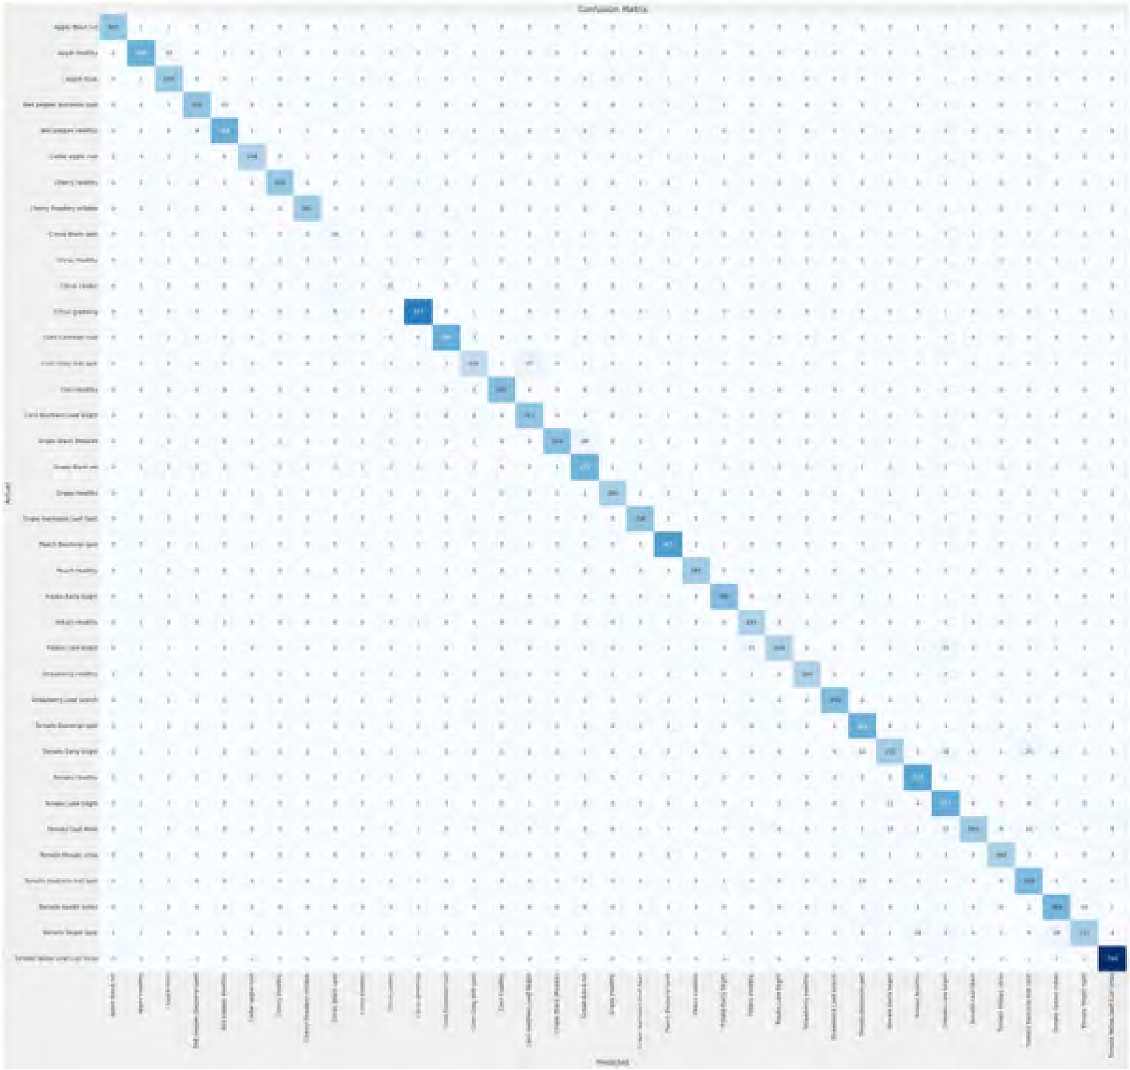

Figure 4.10: Confusion Matrix for InceptionV3

The confusion matrix demonstrates the efficacy of the InceptionV3 model across several plant disease classifications. Significant accuracy is observed in categories such as Citrus Greening (557 correct predictions) and Tomato Healthy (428 correct predictions), reflecting robust performance in these fields. Nonetheless, confusion arises, illustrated by Apple Black Rot being incorrectly classed as Apple Scab in 21 instances, and Corn Northern Leaf Blight being misidentified as Corn Common Rust in 20 instances. Strawberry Leaf Scorch has 343 accurate predictions but is occasionally misidentified as Strawberry Healthy. The model proficiently differentiates most diseases; nonetheless, there exists some overlap in visually identical categories, especially within the same species, suggesting potential for enhancement in discrimination.

4.6 Propesed Hybrid Model (DenseNet201+ResNet152):

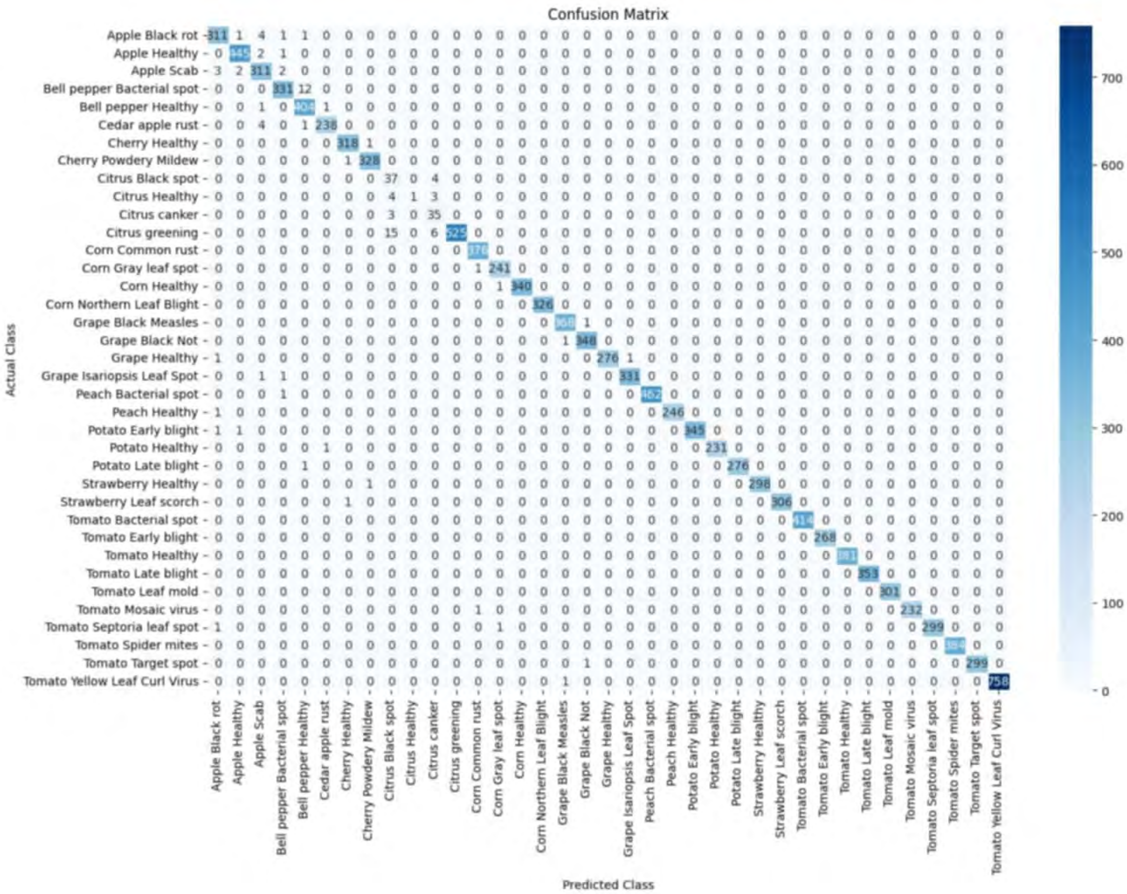

Figure 4.11: Confusion Matrix of DenseNet201+ResNet152 before Attribute Dropping

The confusion matrix above shows the correct classification on the diagonal and misclassifications in the off-diagonal cells for the hybrid model before dropping the attributes. Here, we can observe most classes obtain high accuracy with the marked presence of samples correctly classified along the diagonal. It can be seen that the model does best in the following categories: "Tomato Yellow Leaf Curl Virus" with 758 correct predictions, "Tomato Spider Mites" with 384 correct predictions, and "Apple Healthy" with 445 correct predictions, as evidenced by the color of the darker squares along the diagonal. However, there are some areas of concern, especially regarding the three categories under the Citrus class that do very poor detection. For example, Citrus Black Spot only had 37 predictions; Citrus Healthy had only 1 prediction; and Citrus Canker had only 35 predictions; all the other images were misclassified. Only the Citrus Greening showed some promising result with 525 predictions. The confusion between the three categories of Citrus may be due to similar textural visual patterns, such as leaf discoloration or leaf texture, not different enough for the model to appropriately categorize.

```
164/164 [==============================] - 145s 838ms/step
Classification Report:
----------------------
                               precision    recall  f1-score   support

              Apple Black rot       0.99      0.98      0.99       305
                Apple Healthy       0.99      0.99      0.99       420
                   Apple Scab       0.97      0.98      0.98       346
     Bell pepper Bacterial spot       1.00      0.98      0.99       357
          Bell pepper Healthy       0.98      0.99      0.99       377
             Cedar apple rust       0.96      0.99      0.97       256
               Cherry Healthy       0.99      0.99      0.99       331
        Cherry Powdery mildew       1.00      1.00      1.00       285
            Citrus Black spot       0.89      0.89      0.89        38
               Citrus Healthy       0.60      0.50      0.55         6
                Citrus canker       0.89      0.78      0.83        41
              Citrus greening       0.99      0.99      0.99       548
             Corn Common rust       0.99      0.98      0.99       381
          Corn Gray leaf spot       0.95      0.98      0.96       244
                 Corn Healthy       1.00      1.00      1.00       359
    Corn Northern Leaf Blight       0.98      0.99      0.98       359
          Grape Black Measles       0.99      1.00      1.00       340
              Grape Black rot       1.00      0.99      0.99       354
                Grape Healthy       0.99      0.99      0.99       268
   Grape Isariopsis Leaf Spot       1.00      1.00      1.00       334
         Peach Bacterial spot       1.00      0.99      0.99       472
                Peach Healthy       0.98      0.98      0.98       242
          Potato Early blight       0.99      0.98      0.98       370
               Potato Healthy       0.98      1.00      0.99       248
           Potato Late blight       0.97      0.97      0.97       379
           Strawberry Healthy       1.00      1.00      1.00       261
      Strawberry Leaf scorch       0.99      1.00      1.00       344
        Tomato Bacterial spot       0.99      0.97      0.98       414
          Tomato Early blight       0.94      0.99      0.96       359
               Tomato Healthy       0.99      0.97      0.98       410
           Tomato Late blight       0.98      0.96      0.97       431
             Tomato Leaf Mold       0.95      0.98      0.96       367
          Tomato Mosaic virus       1.00      0.96      0.98       276
     Tomato Septoria leaf spot       0.97      0.96      0.97       436
          Tomato Spider mites       0.99      0.97      0.98       368
           Tomato Target Spot       0.93      0.99      0.96       361
 Tomato Yellow Leaf Curl Virus       1.00      0.99      0.99       777

                     accuracy                           0.98     12464
                    macro avg       0.97      0.96      0.97     12464
                 weighted avg       0.98      0.98      0.98     12464
```

Figure 4.12: Classification Report OF DENSENET201+RESNET152 before Attribute Dropping

This classification report is introduced here consists of precision, recall, f1 score, and the actual score of the predicted images. Also, this report demonstrates the insignificance of the three Citrus classes (Citrus Black Spot, Citrus Healthy, and Citrus Canker) in this study. The issue with this three classes was that they had very little amount of images within them, resulting in bad prediction rate for the models. As we can observe, the values of the three classes are significantly lower than the other classes with more images. Also, the amount of images predicted from these attributes are low as well. Which results in making visualization of confusion matrix more complex and time-consuming.

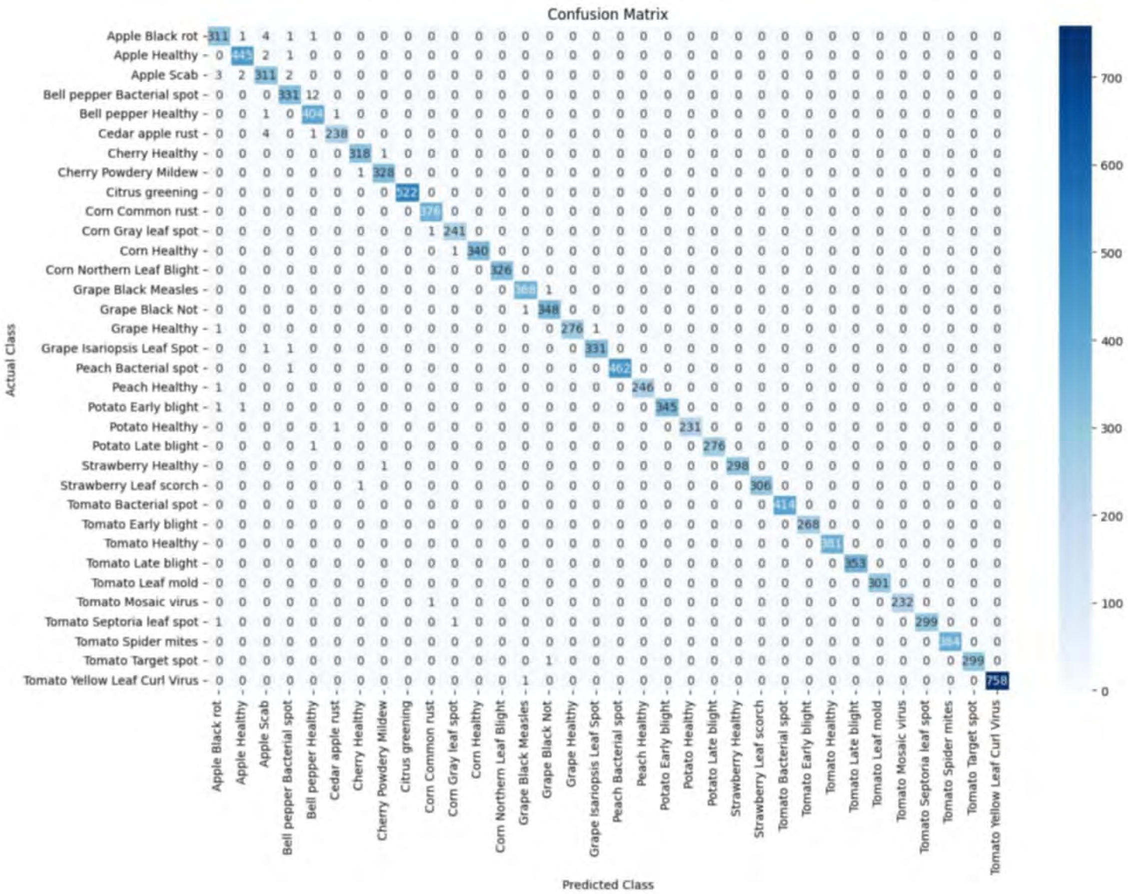

Figure 4.13: Confusion Matrix of DenseNet201ResNet152 after Attribute Dropping

This confusion matrix confirms the performance of the hybrid model staying unchanged even after attribute dropping, which indicates that the attributes that were dropped had insignificance in the model's performance. The absence of the Citrus Black Spot, Citrus Healthy, and Citrus Canker is important, as these classes have already shown very poor detection in the previous matrix by showing very low or zero correct predictions: 37, 1, and 35, respectively. It is likely that this removal indicates that these attributes were not meaningfully contributing to the model's performance and thus were dropped. This highlights the limited significance these classes had in influencing the overall model, confirming that their presence was not beneficial for accurate predictions.

4.7 Further Evalution of DnseNet201+ResNet152

- **Training Loss:** 0.0028
- **Training Accuracy:** 99.91%
- **Validation Loss:** 0.0839
- **Validation Accuracy:** 98.35%
- **Analysis:** When compared to other models, this suggested hybrid model achieves far better results in terms of loss and accuracy. With a training

accuracy of 99.91%, it's safe to say that the model has mastered nearly every pattern in the training set. The model's 98.35% validation accuracy and incredibly low validation loss of 0.0839 demonstrate its exceptional generalizability to unseen data.

- **Precision:** 0.97
- **Recall:** 0.96
- **F1 Score:** 0.97
- **Accuracy:** 98.35%
- **Summary:**There are fewer false positives since the proposed hybrid model obtains the best accuracy (0.97), which means it nearly never wrongly labels healthy plants as sick. With a recall of 0.96, it outperforms other models in identifying sick plants. This model outperforms all others in terms of precision and recall, with an F1 score of 0.97, demonstrating the best balance between the two. With an impressive 98.35% accuracy rate, the hybrid model clearly outperforms all others on this dataset when it comes to generalizing new data.

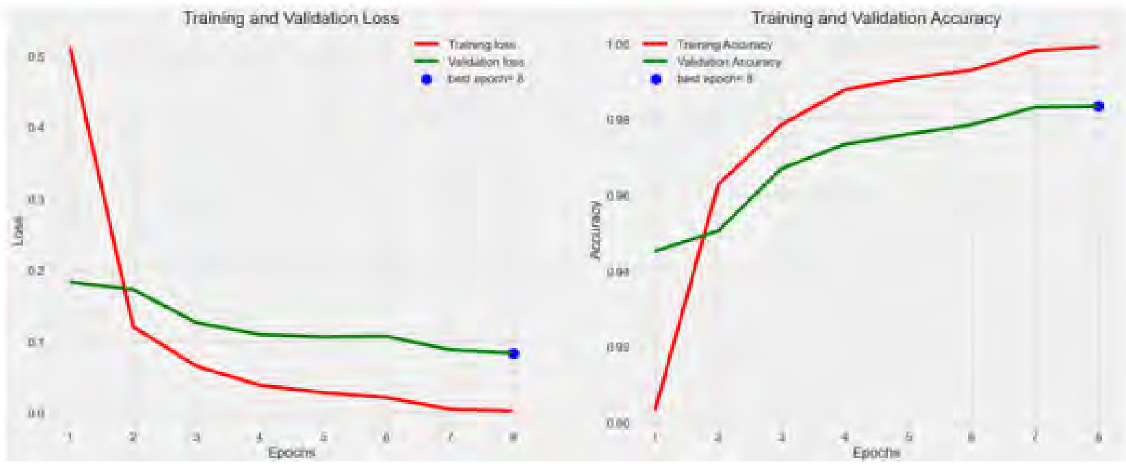

Figure 4.14: Training and Validation Loss & Accuracy graph for DenseNet201+ResNet152

This graph illustrates the training and validation loss and accuracy for the primary model (DenseNet201+ResNet152) across 8 epochs. The training loss (red line) begins at approximately 0.5 but declines promptly approaching zero by the 8th epoch, signifying that the model is successfully integrating the patterns in the training data. The validation loss (green line) consistently declines over the initial epochs, afterwards stabilizing after the third epoch, achieving approximately 0.09 by the eighth epoch, marginally exceeding the training loss, indicating effective generalization to unique data. The training accuracy (red line) commences at approximately 92% and escalates swiftly, approaching 100% by the 8th epoch, signifying proficient learning. The validation accuracy (green line) commences at 94%, steadily enhancing to 98.2% by the 8th epoch, although it maintains post the 5th epoch, indicating that the model has likely reached optimal performance. The optimal validation accuracy, indicated by the blue dot, occurs at the 8th epoch, signifying the model's optimum performance characterized by persistent convergence, minimal losses, and improved accuracy.

4.7.1 SUMMARY OF ALL RESULTS

The evaluation metrics such as accuracy, precision, recall, and F1-score provide a comprehensive assessment of model performance across all CNN architectures. VGG16, VGG19, MobileNetV2, ResNet50 and InceptionV3 demonstrate solid accuracy, particularly in differentiating various leaf disease categories. However, some overfitting is observed, as shown by the gap between training and validation performance in certain models. Precision and recall metrics give information about the tradeoffs that models make between false positives and false negatives; the F1-score gives a comprehensive performance assessment. Models with high precision yield the highest correctness in their identification of diseases, but those with higher recall make sure not to miss the true disease cases. The confusion matrices over these models highlight key areas of misclassifications and pinpoint the classes where the models failed. In particular, the hybrid model that combined DenseNet201 with ResNet152 showed a better confusion matrix, as was indicated from the classification report before dropping the attribute. The report showed poor detection for Citrus classes: Citrus Black Spot, Citrus Healthy, and Citrus Canker. After dropping attributes, these classes were removed; this is evidence that they contributed little to the overall performance of the proposed model. This adjustment led to more profound visualization of our desired confusion matrix. Additionally, by achieving better results in accuracy, precision, recall, and F1-score, the hybrid model has outperformed all the other models. In particular, the hybrid model in this paper has aptly handled imbalanced data for fewer misclassifications, especially after the adjustment of attributes; it is therefore emerging as a strong candidate for real-world application in plant disease detection. Indeed, its generalization over diverse plant disease categories makes it the most viable even with difficult-to-detect categories.

Chapter 5

Implementation

5.1 Implementation

5.1.1 The Reasoning Behind Real-Time Implementation:

Thus, the prime strength of this paper will be the real-time implementation of the hybrid model in practical agriculture. Real-time detection of plant diseases will highly influence crop health management. Applying the model in a natural environment will help farmers and agronomists make rapid decisions by avoiding time-consuming manual analyses. This functionality is further enhanced through the website interface we developed, where a user can upload or drop any image from the dataset and get a very accurate disease prediction instantly. Not only does this approach serve to use the model for its practical utility, but it also affords an accessible, user-friendly avenue whereby non-technical users can benefit from state-of-the-art deep learning techniques. By integrating into a web platform, we would like to contribute to the popularization of AI-based solutions in agriculture, giving these powerful tools to support plant health monitoring as much as possible.

5.1.2 Implementation Procedure

The plant disease detection model, encompassing 34 distinct classes of plant diseases, was successfully deployed using TensorFlow Serving within a Docker containerized environment. The deployed model used in this example was TensorFlow Serving on a containerized environment in Docker, while training on detecting 34 classes of plant diseases. First, the model trains on an image dataset containing both healthy and diseased leaves. After that, the model is exported and saved to a format compatible with TensorFlow Serving. This included the creation of the "saved_model.pb": file along with the accompanying 'variables' directory that TensorFlow Serving needs.

Model serving was enabled via the Docker environment by leveraging the containerized solution provided by TensorFlow Serving. This ensures that the model can be effectively served through a RESTful API capable of reaching real-time predictions. To this aim, a Docker container was exploited for wrapping the TensorFlow serving environment, with the purpose of ensuring the portability and scalability of the deployment phase on different infrastructures.

Once the serving infrastructure was up, a web application using Flask was created. Its role would be to perform the work of middleware between the user interface and model API. The Flask application processed the incoming request from the front-end and forwarded it to the TensorFlow Serving API. The REST API of the model was used in order to send image data for prediction and get the corresponding results, including the class predicted and confidence score of the disease.

The interface is designed in such a way that the user can upload pictures of leaves through a friendly interface, using HTML, CSS, and JavaScript. Afterward, the images were passed on to the backend, Flask, for processing, which in turn sent the images to the API exposed by TensorFlow Serving for the model's inference. Finally, on the front end, it displays the model's prediction of the detected plant disease and the confidence level of such prediction. All that put together yielded one interactive, smooth ecosystem for plant disease detection, while Docker and TensorFlow Serving provided scalability, robustness, and efficient handling of model predictions in real time.

This deployment process highlights an efficient integration of machine learning model serving using modern web technologies and containerized environments, making it suitable for real-time applications in agricultural diagnostics.

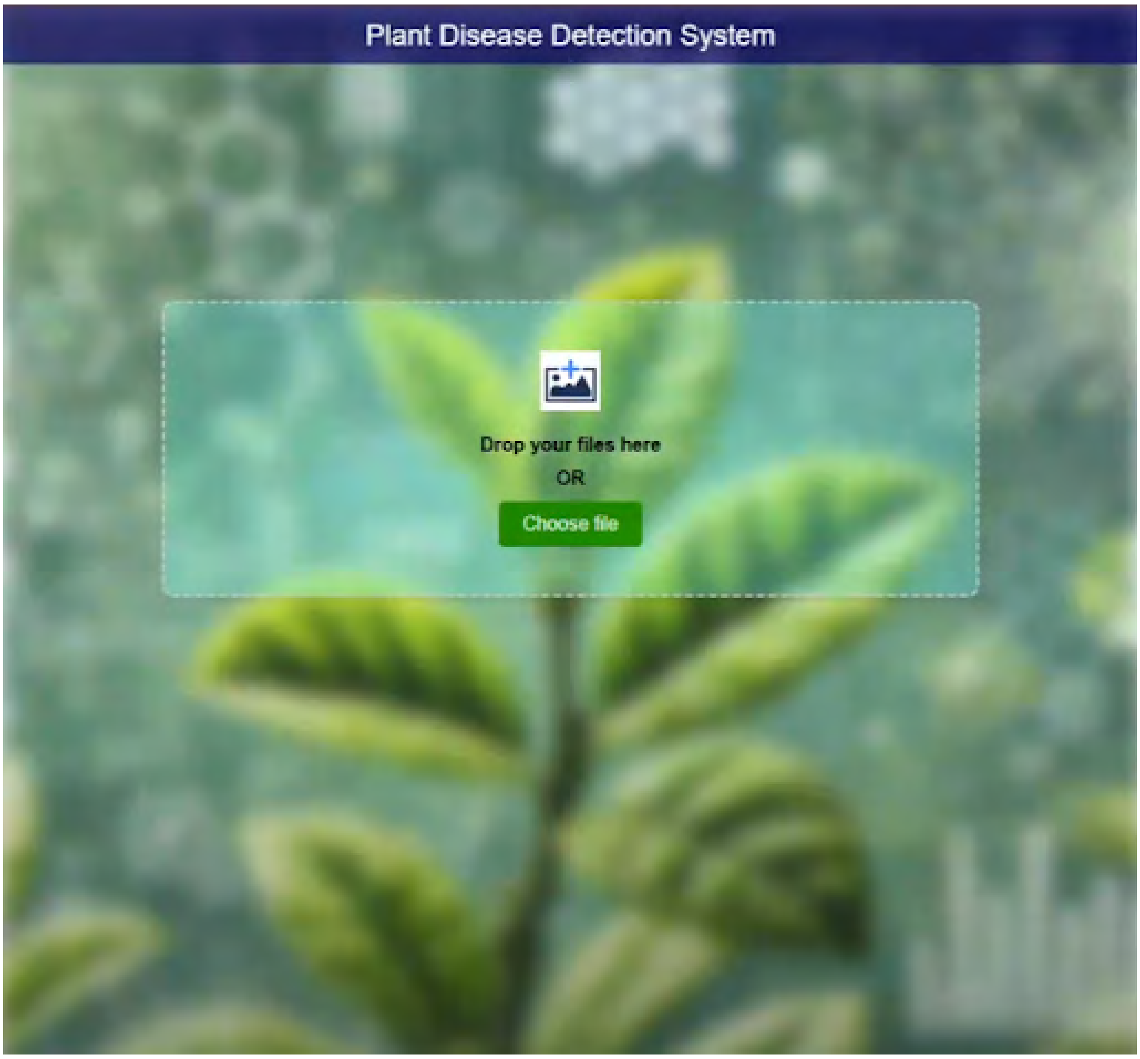

Figure 5.1: Visualization of Drag and Drop Interface

This is the initial interface of the system that we have developed for plant disease detection. One can drag and drop their single input file (image) in the interface or select from their directory. The system can be accessed very easily and it is also very user friendly. It is simplified for all sorts of users. Also, it is almost 98% accurate in data detection and analysis.

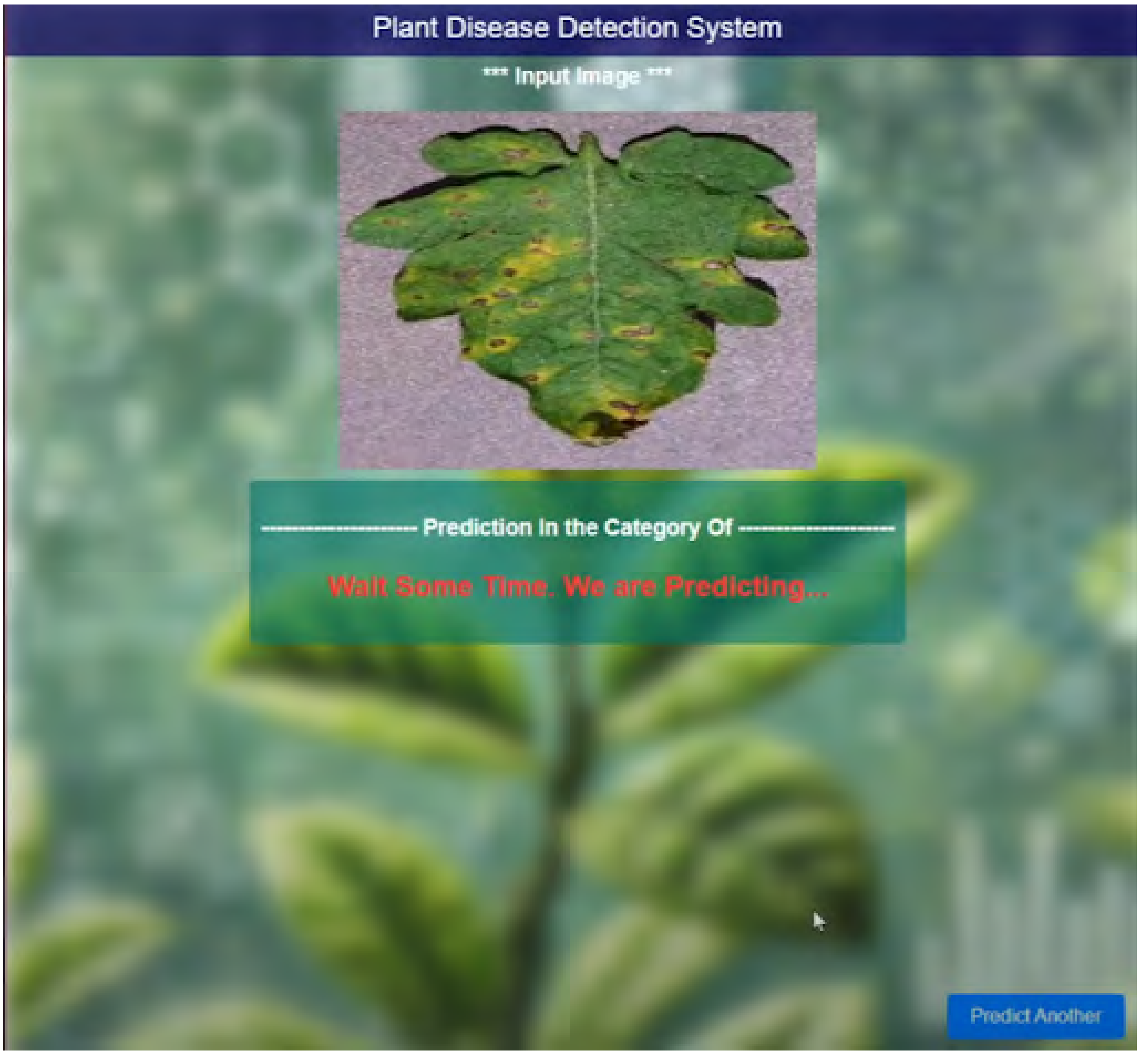

Figure 5.2: An Image of Tomato Septoria Leaf Spot being Processed

At first, we selected a tomato leaf from a class of the dataset. The system is almost instantaneous to detect which leaf class it belongs to. It takes a few seconds to process the image before showing the result.

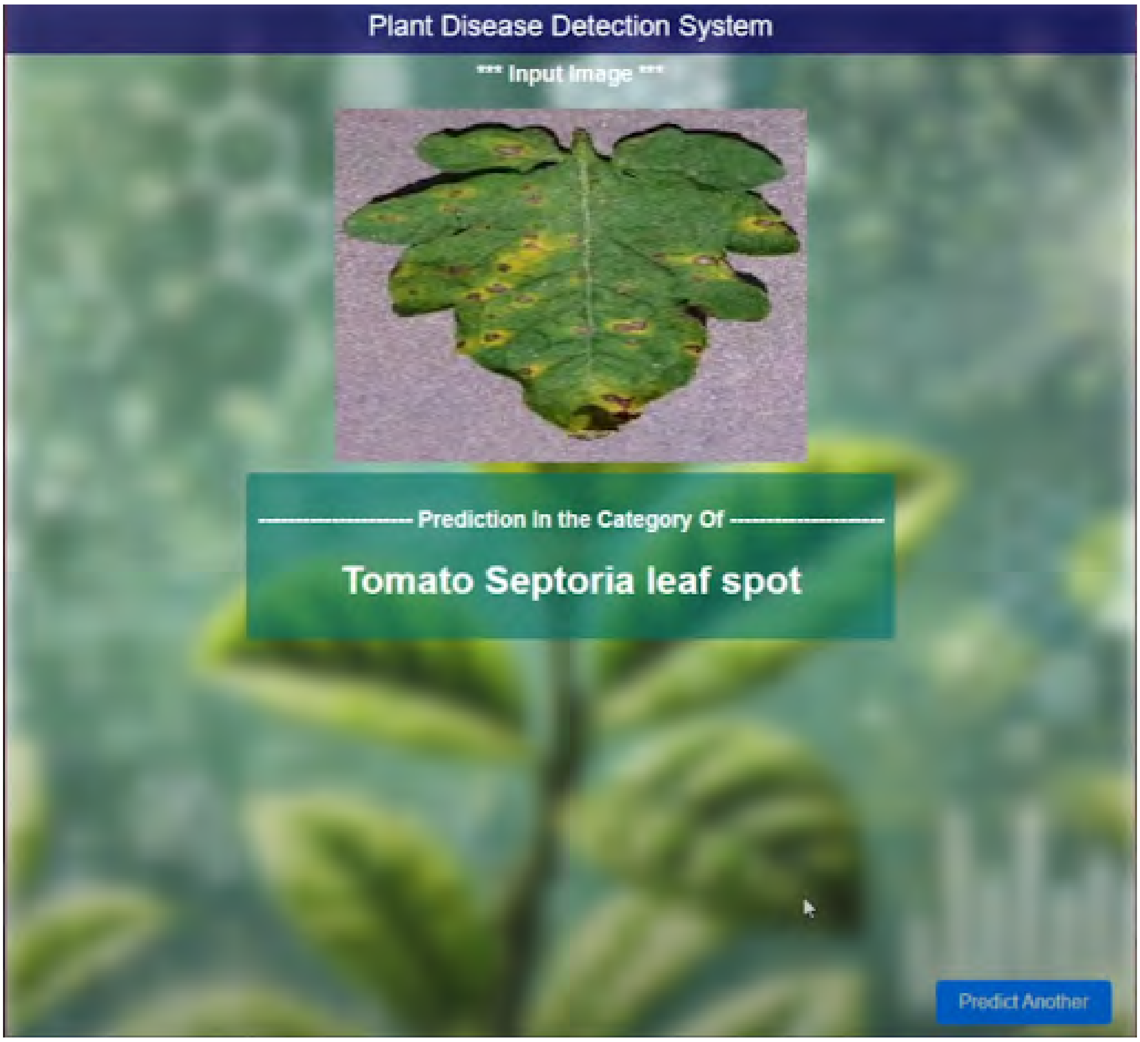

Figure 5.3: The System Correctly Predicting Tomato Septoria Leaf Spot

After a few seconds of time, it detects the leaf image as a tomato leaf and alongside it mentions the disease as well. We tried the same picture rotating left, right, 180*, mirroring and even cropping. The result was the same and the time it took to detect the leaf and the disease were also the same throughout the process.

Figure 5.4: An Image of Cedar Apple Rust being Processed

Just like we detected a tomato leaf and the class it belongs to, we have decided to show another prediction for the system. To replace or predict a new image we have to click on the predict another button which is located at the bottom right corner of the system. After that we have selected another image from our data to predict its class and the category it belongs to.

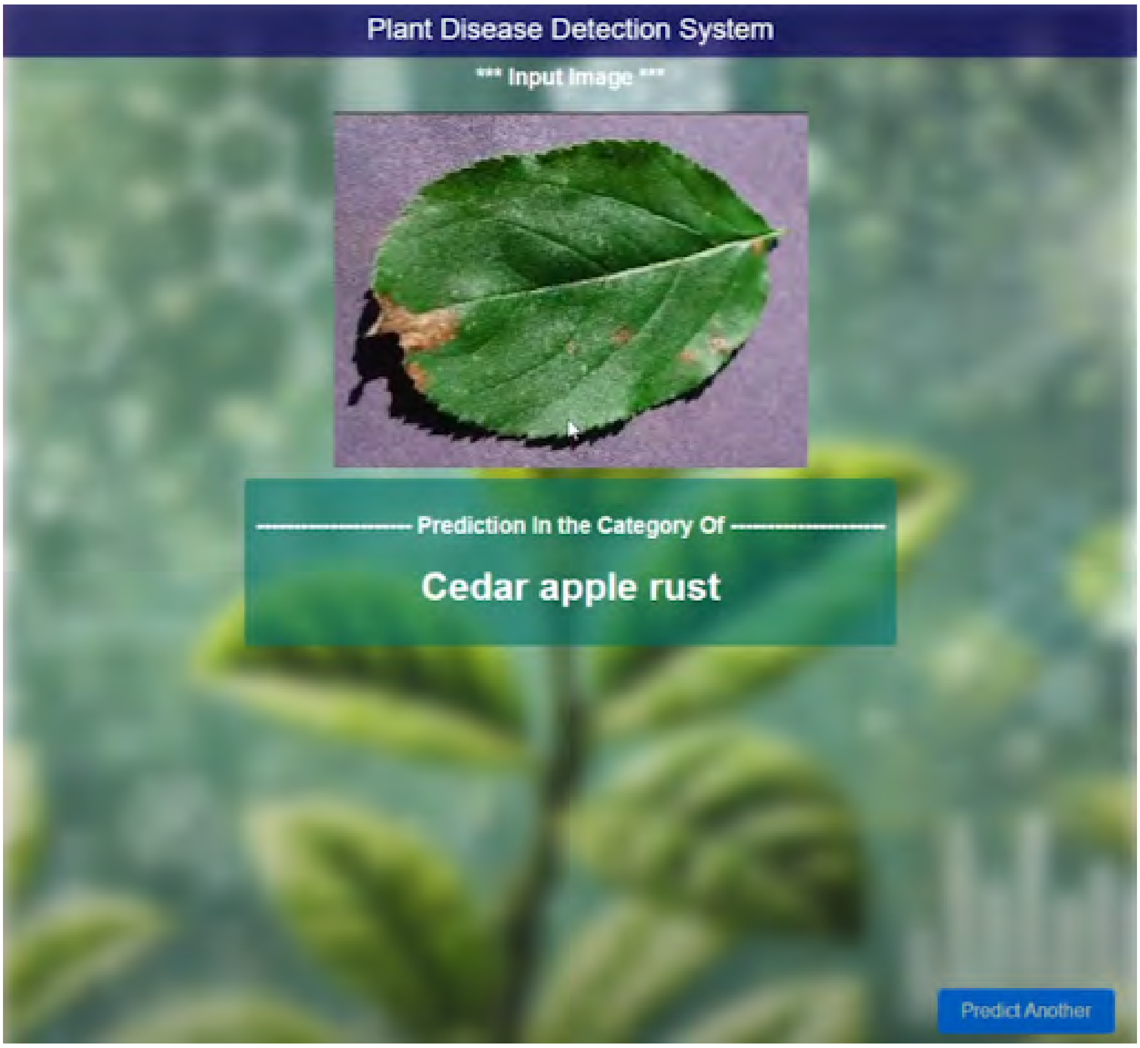

Figure 5.5: An Image of Cedar Apple Rust being Processed

Just like before, our system has declared the image and the leaf the plant belongs to. The image is an apple leaf image and the disease it has predicted was Cedar apple rust. We manipulated the image multiple times through different procedures and it still retained the exact class it belongs to.

Chapter 6

Impact on Environment and Sustainability

6.1 Impact on Environment

Lower Chemical Application: This means that plant diseases are recognised with a high degree of accuracy and in good time, enabling appropriate (timely) responses to be made, which could reduce the need for fungicides and insecticides. This decreases the impact on non-target organisms, such as predators and soil dwellers, along with reducing chemical waste into the ground or water bodies. Enhanced Crop Health: This technique provides an effective cure when disease attacks, hence ensuring the crops remain healthy, which leads to higher yields and minimal wastage. These are likely to be principal plant traits facilitating protection against disease transmission in healthy plants and the maintenance of biodiversity in agricultural ecosystems. Resources used: Better detection can lead to resources such as water, fertilizer, and land being more effectively utilized. By controlling disease outbreaks, farmers are able to reduce inputs used and pressure, which is currently driving agriculture towards more sustainable industry practices. Biodiversity preservation: The practice preserves the regional flora and fauna by preserving healthier crops that reduce chemical dependency. This becomes particularly important in diverse agricultural ecosystems.

6.2 Ethical Aspects

Data Privacy: Keeping farmers' data (crop health and farming practices) confidential when using machine learning solutions that require the collection of data Always ensure you have obtained an informed consent before any data collection. Heightened Access and Equity: All farmers (regardless of their financial status) need to have access to and use advanced technologies. This imbalance might result in unequal agricultural production and profit as smallholder or resource-constrained farmers are using their traditional methods, which do not include these technologies that will be used by the more sophisticated agribusinesses.

6.3 Technology Dependency

Too much reliance on technology for disease detection can lead to the loss of traditional farming skills. It is important for the conservation and valuation of traditional knowledge that there remains a balance between local farmers' practice and technical enhancement.

6.4 Sustainability Plan

Community Engagement: Involve local farmers and communities in the design of the model. Conduct training sessions and workshops to teach people how to support technology way better. This promotes local ownership and ensures that the technology meets community demands. Sustainability practices: ensure the integration of sustainable agricultural activities such as integrated pest management (IPM), crop rotation, organic farming, etc. with a disease detection model. This full-cycle approach can boost farm productivity and health on the whole. Research and Development – continue funding studies that improve the model and further explore when it may be useful for different crops or conditions Collaboration with agricultural specialists and organizations to evaluate the impact of this technology on soil health, biodiversity, and ecological services in the longer term. Overseeing and Evaluation: Set up a solid overseeing framework that permits checking the exhibition of the model in real-world application. To ensure that the technology serves sustainability, collect data on social justice/equity, economic benefits and environmental impacts. It will: Be involved in the formulation and lobbying for laws that support an adoption of agricultural technology — meanwhile maintaining sustainability, prudence to environmental resources, and fairness; Promote policy provisions that encourage sustainable environmental farming practices, making it accessible to smallholder farmers in easily usable tech solutions.

Chapter 7

Future Plan and Conclusion

7.1 Future Plan

7.1.1 Real-Time Detection System

Develop a real-time disease forecasting system by embedding the model in IoT devices such as drones (which can be part of smart farming equipment) or mobile applications. This will help the farmers to recognize diseases immediately in their fields and take faster action.

7.1.2 Multi-Class and Multi-Label Classification

Ability to detect a large number of diseases on one leaf or various portions of the plant (multi-label classification). Enhance the model to not just predict diseases but also identify their magnitude (early, mid, or late stages of infection) which can be helpful for farmers in terms of making management decisions.

7.1.3 Working hand in hand with Agricultural Scientists

Cooperate with plant pathologists and agronomists to validate the model's predictability in field conditions, incorporating enthusiasts' feedback toward making corrections on reference points. If you want to deploy your model at scale for real-time monitoring of plant health, consider collaborating with agricultural groups, organizations, governments, ns or farming enterprises.

7.1.4 IoT Technologies on the farm: Cloud-based agricultural platforms

Generating a platform based in the clouds to let farmers upload pictures to identify diseases gives it an interesting use of detection and prevention amongst poor areas. Integrate your model with Geographic Information System (GIS) data and monitor the disease spread over different sites, to adopt appropriate regional strategies for effective management of plant diseases.

7.1.5 Disease Prevention and Forecasting

Integrate this model with environmental data or climate variables to predict the likelihood of disease outbreaks. As a result, the farmers would be able to prepare for dangerous times. Investigate the model role in automated systems (e.g., smart spraying systems) that use disease detection along with pest and disease management strategies.

7.2 Conclusion

This research outlined a novel optimization of deep learning-based plant disease detection from leaf images while utilizing VGG16, VGG19, ResNet50, DenseNet201, and MobileNetV2 along with InceptionV3 as the models. One of its major contributions is the implementation aspects with a hybrid model using ResNet152 and DenseNet201 models, which were found to improve both accuracy as well as robustness in detecting COVID-19 [14]. Combining them enabled effective feature extraction and classification results, obtaining an accuracy greater than 98The results demonstrate the ability of deep learning models to transform plant disease detection systems, able to scale up with a high degree of accuracy and efficiency in real agricultural scenarios. This model has many applications, as it allows for the earliest possible detection of plant diseases and therefore an effective target to prevent their spread and reduce losses in crops. The combination of the hybrid model with another pre-trained model provided a more in-depth feature representation and resulted in achieving higher detection accuracy across different plant diseases. In short, this study can be described as a meaningful contribution to previous studies in the field of plant disease identification by optimizing deep learning model architecture using numerous pre-trained models. The successful deployment in conjunction with these models suggests that artificial intelligence can provide reliable, efficient, and scalable solutions to key agricultural dilemmas, thereby helping promote healthier agriculture and offering better food security worldwide.

References

[1] S. D. Khirade and A. Patil, “Plant disease detection using image processing,” in *2015 International Conference on Computing Communication Control and Automation*, 2015, pp. 768–771. DOI: 10.1109/ICCUBEA.2015.153.

[2] S. Ramesh, R. Hebbar, M. Niveditha, *et al.*, “Plant disease detection using machine learning,” in *2018 International Conference on Design Innovations for 3Cs Compute Communicate Control (ICDI3C)*, 2018, pp. 41–45. DOI: 10.1109/ICDI3C.2018.00017.

[3] P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang, “Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks,” *IEEE Access*, vol. 7, pp. 59 069–59 080, 2019. DOI: 10.1109/ACCESS.2019.2914929.

[4] S. Militante, B. Gerardo, and N. Dionisio, “Plant leaf detection and disease recognition using deep learning,” Oct. 2019, pp. 579–582. DOI: 10.1109/ECICE47484.2019.8942686.

[5] G. Zhou, W. Zhang, A. Chen, M. He, and X. Ma, “Rapid detection of rice disease based on fcm-km and faster r-cnn fusion,” *IEEE Access*, vol. 7, pp. 143 190–143 206, 2019. DOI: 10.1109/ACCESS.2019.2943454.

[6] M. Saleem, J. Potgieter, and K. Arif, “Plant disease classification: A comparative evaluation of convolutional neural networks and deep learning optimizers,” *Plants (Basel, Switzerland)*, vol. 9, no. 10, p. 1319, Oct. 2020. DOI: 10.3390/plants9101319.

[7] E. Ahmed, H. Walid, and A. Madallah, “Optimization of deep learning model for plant disease detection using particle swarm optimizer,” *Computers, Materials and Continua*, vol. 71, no. 2, pp. 4019–4031, 2021, ISSN: 1546-2218. DOI: https://doi.org/10.32604/cmc.2022.022161. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1546221821001417.

[8] E. Hirani, V. Magotra, J. Jain, and P. Bide, “Plant disease detection using deep learning,” in *2021 6th International Conference for Convergence in Technology (I2CT)*, 2021, pp. 1–4. DOI: 10.1109/I2CT51068.2021.9417910.

[9] L. Li, S. Zhang, and B. Wang, “Plant disease detection and classification by deep learning—a review,” *IEEE Access*, vol. 9, pp. 56 683–56 698, 2021. DOI: 10.1109/ACCESS.2021.3069646.

[10] P. Kaur, S. Harnal, R. Tiwari, *et al.*, “Recognition of leaf disease using hybrid convolutional neural network by applying feature reduction,” *Sensors*, vol. 22, no. 2, 2022, ISSN: 1424-8220. DOI: 10.3390/s22020575. [Online]. Available: https://www.mdpi.com/1424-8220/22/2/575.

[11] J. Pandian, V. D. Kumar, O. Geman, M. Hnatiuc, M. Arif, and K. Kanchanadevi, "Plant disease detection using deep convolutional neural network," *Applied Sciences*, vol. 12, no. 14, 2022. DOI: 10.3390/app12146982.

[12] D. Sachin, G. Tarun, and G. Dushyant, "Performance analysis of deep learning architectures for plant leaves disease detection," *Measurement: Sensors*, vol. 24, p. 100 581, 2022, ISSN: 2665-9174. DOI: https://doi.org/10.1016/j.measen.2022.100581. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S266591742200215X.

[13] F. Adnan, M. Awan, A. Mahmoud, H. Nobanee, A. Yasin, and A. Zain, "Efficientnetb3-adaptive augmented deep learning (aadl) for multi-class plant disease classification," *IEEE Access*, vol. 11, pp. 85 426–85 440, 2023. DOI: 10.1109/ACCESS.2023.3303131.

[14] H. Hamoud, I. Ahmed, and Osama, "Identification of olive leaf disease through optimized deep learning approach," *Alexandria Engineering Journal*, vol. 72, pp. 213–224, 2023, ISSN: 1110-0168. DOI: https://doi.org/10.1016/j.aej.2023.03.081.

[15] Hicham, B. Aissa, M. Noureddine, and Abdelbaki, "A novel approach for image-based olive leaf diseases classification using a deep hybrid model," *Ecological Informatics*, vol. 77, p. 102 276, 2023, ISSN: 1574-9541. DOI: https://doi.org/10.1016/j.ecoinf.2023.102276. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1574954123003059.

[16] Y. Mahmoud, N. Sakr, and M. Elmogy, "Plant disease detection and classification using machine learning and deep learning techniques: Current trends and challenges," Jan. 2023, pp. 197–217. DOI: 10.1007/978-981-99-4764-5_13.

[17] M. Yousef, J. Orlando, M. Liz, K. Chamandeep, S. Mohammed, and B. Kiran, "Leaf disease identification and classification using optimized deep learning," *Measurement: Sensors*, vol. 25, p. 100 643, 2023, ISSN: 2665-9174. DOI: https://doi.org/10.1016/j.measen.2022.100643. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S266591742200277X.

[18] E. Aldakheel, M. Zakariah, and A. Alabdalall, "Detection and identification of plant leaf diseases using yolov4," vol. 15, Apr. 2024. DOI: 10.3389/fpls.2024.1355941.

[19] G. Chinna, K. Hari, V. K.V., M. Alakananda, and A. Preethi, "Deep learning for rice leaf disease detection: A systematic literature review on emerging trends, methodologies and techniques," *Information Processing in Agriculture*, 2024, ISSN: 2214-3173. DOI: https://doi.org/10.1016/j.inpa.2024.04.006. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S221431732400026X.

[20] D. Joseph Susan, P. M. Pawar, and K. Chakradeo, "Real-time plant disease dataset development and detection of plant disease using deep learning," *IEEE Access*, vol. 12, pp. 16 310–16 333, 2024. DOI: 10.1109/ACCESS.2024.3358333.

[21] A. Mahmood and A. Khalaf, "Olive leaf disease detection via wavelet transform and feature fusion of pre-trained deep learning models," *Computers, Materials and Continua*, vol. 78, no. 3, pp. 3431–3448, 2024, ISSN: 1546-2218. DOI: https://doi.org/10.32604/cmc.2024.047604. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1546221824003242.

[22] P. Prasad and S. Agniraj, "Cross-domain adaptation techniques for robust plant disease detection: A dann-coral hybrid approach," *International Journal of Experimental Research and Review*, vol. 42, pp. 68–84, Aug. 2024. DOI: 10.52756/ijerr.2024.v42.007.

[23] C. Rohit and T. Bhalodia, "Detection and classification on plant disease using deep learning techniques," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 10, no. 3, pp. 365–375, May 2024. DOI: 10.32628/CSEIT2410333.

[24] V. Shwetha, B. Arnav, and L. Vijaya, "Leafspotnet: A deep learning framework for detecting leaf spot disease in jasmine plants," *Artificial Intelligence in Agriculture*, vol. 12, pp. 1–18, 2024, ISSN: 2589-7217. DOI: https://doi.org/10.1016/j.aiia.2024.02.002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2589721724000035.