Plant Disease Classification Using ResNet50 with Saliency Map and grad-cam-Based-Explainability

By

Hasibul Islam Sufi

213-15-4622

FINAL YEAR DESIGN PROJECT REPORT

This Report Presented in Partial Fulfillment of the Requirements forthe **Degree of Bachelor of Science in Computer Science and Engineering**

Supervised by

Mr. Md. Monarul Islam

Lecturer

Department of Computer Science and EngineeringDaffodil International University

Co-Supervised by

Mr. Mahimul Islam Nadim

Lecturer

Department of Computer Science and EngineeringDaffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY

Dhaka, Bangladesh

APPROVAL

This Project titled **"Plant Disease Classification Using ResNet50 with Saliency Map and grad-cam-Based Explainability"**, Hasibul Islam sufi, ID No:213-15-4622 to the Department of Computer Science and Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of B.Sc. in Computer Science and Engineering and approved as to its style and contents. The presentation has been held on **16 September, 2025**.

BOARD OF EXAMINERS

Dr.S.M Aminul Haque — **Chairman**
Professor & Associate Head
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Dr.Md Alamgir Kabir — **Internal Examiner**
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Mr.Md Assaduzzaman — **Internal Examiner**
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Dr. Md. Zulfiker Mahmud — **External Examiner**
Professor
Department of Computer Science and Engineering
Jagannath University

DECLARATION

We hereby declare that this project has been done by us under the supervision of **Mr. Md. Monarul Islam, Lecturer** Department of Computer Science and Engineering, Daffodil International University. We also declare that neither this projectnor any part of this project has been submitted elsewhere for the award of any degree or diploma.

Supervised by:

Monarul
15.09.25

Mr. Md. Monarul Islam
Lecturer
Department of Computer Science and EngineeringDaffodil International University

Co-Supervised by:

Mr. Mahimul Islam Nadim
Lecturer
Department of Computer Science and EngineeringDaffodil International University

Submitted by:

Hasibul Islam Sufi
Student ID:213-15-4622
Department of Computer Science and EngineeringDaffodil International University

ACKNOWLEDGEMENTS

This work would not have been possible without the support and contributions of many individuals over the past two semesters. We are deeply grateful to everyone who has as- sisted us in one way or another.

First, we express our heartfelt thanks and gratefulness to the almighty for His divine blessing making it possible for us to complete the **Final Year Design Project(FYDP)** successfully.

We are grateful and wish our profound indebtedness to **Mr. Md. Monarul Islam**, **Lecturer**, Department of Computer Science and Engineering, Daffodil International University, Dhaka, Bangladesh. Deep knowledge and keen interest of our supervisor in the field of **Image classifcation** to carry out this project. His endless pa- tience, scholarly guidance, continual encouragement, constant and energetic supervision, constructive criticism, valuable advice, reading many inferior drafts, and correcting themat all stages have made it possible to complete this project.

We would like to express our heartfelt gratitude to the Head of the Department of Com- puter Science and Engineering, for his kind help in finishing our project and also to other faculty members and the staff of the Department of Computer Science and Engineering, Daffodil International University.

We would like to thank our entire course-mates at Daffodil International University, who took part in this discussion while completing the coursework.

Finally, we must acknowledge with due respect the constant support and patience of our parents.

ABSTRACT

This study presents a comprehensive deep learning-based framework for multi-class plant disease classification using high-resolution leaf images, with a particular focus on evaluating the performance of convolutional neural networks (CNNs) and a hybrid ResNet + Vision Transformer (ViT) architecture. A curated dataset comprising 15,200 training and 3,800 validation images spanning 38 classes across multiple crops, including tomato, apple, grape, corn, potato, strawberry, peach, pepper, orange, blueberry, raspberry, soybean, and squash, was subjected to preprocessing steps such as resizing, normalization, and data augmentation to enhance model robustness. Multiple CNN architectures—including ResNet-50, MobileNetV2, and EfficientNet-B0—were trained and compared against the hybrid ResNet + ViT model. All models were fine-tuned using the AdamW optimizer and cross-entropy loss, with early stopping applied to prevent overfitting and ensure generalization. Additionally, interpretability techniques including Grad-CAM and Saliency Maps were employed to visualize disease-relevant regions, while segmentation-based analysis was performed to localize affected areas on leaves. Among all architectures evaluated, ResNet-50 achieved the highest validation accuracy of 98.74%, while the hybrid ResNet + ViT model recorded a competitive accuracy of 98.58%, demonstrating the effectiveness of hybrid architectures in capturing both local and global features. The experimental results highlight the potential of transformer-based and lightweight CNN models to deliver highly accurate, interpretable, and computationally efficient solutions for automated multi-class plant disease detection, offering valuable support for precision agriculture and crop management practices.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction

One of the most crucial areas in relation to the food security of the globe and human life is agriculture. It does not only form the basis of the food chain but has over one billion employees across the world and about 4 percent contribution to the global GDP [World Bank, 2021]. Although it is important agricultural productivity is always threatened by diseases in plants. The Food and Agriculture Organization (FAO) cites this fact, stating that nearly 20-40 percent of worldwide crop production are wasted annually by pests and diseases leading to billions of dollars worth of economic losses (Savary et al., 2019). These losses impact both the smallholder farmers and commercial producers, causing a decrease in income, broken supply chains, and food price instability.Traditionally, the detection of plant diseases has depended on manual inspection by professionals and laboratory-based tests. These methods are not always available to rural farmers and are time-consuming, expensive, and effective, yet not always. Also, similar visible symptoms of many plant diseases complicate their diagnosis even among skilled agronomists [Brahimi et al., 2017]. This can result in misdiagnosis or delayed diagnosis and exacerbates crop losses.As artificial intelligence (AI) and computer vision continue to progress, engineers have begun to develop automated systems to detect plant diseases, which have proven to be more rapid, less expensive, and scalable [Kamilaris and Prenafeta-Boldu, 2018]. Two of them, in particular, deep learning, and Convolutional Neural Networks (CNNs) in particular, have shown enormous success in image recognition tasks. Visual features like leaf texture, color and the shape of lesions can be learned automatically by CNNs, so that a human intervention is not required. An example is that Mohanty and colleagues [2016] obtained an over 99 percent classification accuracy on 26 plant diseases of 14 crops with CNN models trained on the PlantVillage dataset. Advanced architectures, including ResNet, EfficientNet, and MobileNet, were further used in later works, yielding great results on agricultural data [Too et al., 2019; Ferentinos, 2018].

However, existing research faces three major limitations:

1. Many studies are restricted to binary classification (healthy vs. unhealthy) instead of handling dozens of diseases.
2. CNNs capture local features well but often fail to represent global contextual relationships within images.
3. Most models act like “black boxes,” offering little to no explanation for their predictions, which makes them difficult for farmers and experts to trust.

To address these challenges, our study explores a multi-class, multi-crop plant disease classification framework using both CNNs and a hybrid ResNet + Vision Transformer (ViT) architecture. By incorporating interpretability methods such as Grad-CAM and saliency maps, as well as segmentation-based analysis, this project aims to make plant disease detection not only accurate but also explainable and trustworthy for real-world agricultural deployment.

1.2 Motivation

The driving force of this project is motivated by the practical agricultural requirement as well as the computational problems. In the agricultural view, farmers are in dire need of solution that can offer rapid, cheap and precise disease detection in the field. Failure to diagnose in time can either lead to preventable use of pesticides or irreparable damage to crops. With real time feedbacks in terms of image based detection, farmers can act within the right time to preserve their crops and reduce losses. These tools also correspond to vision of precision agriculture, when interventions are purposeful and efficient, and they can save resources and minimize environmental impact.Mathematically, the creation of a strong multi-classification system is difficult. In practice, images of leaves are subject to environmental perturbations including lighting, background clutter and crossover leaves. In addition, similar patterns are occasionally generated by different diseases and the classification task is complicated. CNNs can capture local features, but they fail to do so along long-range dependencies. This weakness drives the adoption of Vision Transformer (ViT) that adopt self-attention mechanisms to model relationships in the entire image. One such combination is a CNN+ ViT system that has the desirable property of learning fine-grained textures, as well as global information.Last but not least, interpretability is a crucial driving force. Agricultural advisors and farmers need to have confidence in AI predictions before they can be used in making decisions. Visualization tools such as Grad-CAM and saliency maps can be used to emphasize what part of a leaf the model is concentrating on. Such transparency boosts the confidence of users and promotes the use of AI in the agricultural activities.

1.3 Objectives

The main objectives of this project are:

1. To curate and standardize a large dataset covering 38 plant disease classes across multiple crops (tomato, apple, grape, corn, potato, strawberry, soybean, squash, etc.).
2. To implement and compare the performance of four deep learning models: ResNet-50, EfficientNet-B0, MobileNetV2, and a ResNet + ViT hybrid.
3. To apply systematic preprocessing and augmentation techniques such as resizing, normalization, flipping, rotation, and brightness adjustments to improve generalization.
4. To train the models using transfer learning with ImageNet-pretrained weights, optimized with AdamW, learning-rate scheduling, and early stopping.
5. To evaluate performance with multiple metrics—Accuracy, Precision, Recall, F1-score, and AUC—providing a complete comparative analysis.
6. To incorporate explainability tools (Grad-CAM, saliency maps, segmentation) so that model decisions are interpretable and aligned with disease symptoms.
7. To discuss the trade-offs between accuracy vs. computational efficiency, emphasizing deployment feasibility in mobile and resource-constrained environments.

1.4 Methodology

The overall methodology of this project is structured into the following steps:

1. Problem Identification & Objective Setting: The study was inspired by the large losses in crop due to plant diseases and the constraints of the manual inspection. The main aim was to create a stable, interpretable AI system which could find and identify various types of plant diseases with high precision and in a transparent manner.

2. Dataset Preparation: 15,200 training and 3,800 validation samples were extracted in form of images in various public repositories and combined to create a balanced dataset. Poor, irrelevant, or duplicate images were discarded very carefully, and all the labels of the classes were normalized to provide the same results across datasets.

3. Preprocessing & Augmentation: All images were resized to 224x224 and normalized to the input Deep Learning model requirements. A number of augmentations were added to enhance robustness such as random rotation (+-15deg), horizontal/vertical flips, zooming, brightness/contrast jitter and Gaussian noise to simulate field conditions in the real world. The segmentation methods were also used to segment leaf regions based on complex or noisy backgrounds.

4. Model Selection: Four architectures were selected to be experimented with: ResNet-50, EfficientNet- B0, MobileNetV2, and a combination of ResNet and Vision Transformer (ViT). All of the models were generated with ImageNet pretrained weights and fine-tuned on the particular plant disease classification problem.

5. Optimization: The categorical cross-entropy loss function and AdamW optimizer were used and a starting a learning rate of 3e-4. Extra learning rates were optimized using a flexible ReduceLROnPlateau scheduler and early stopping was used to avoid overfitting and unwarranted calculations.

6. Testing: Each of the models was trained up to 25 epochs, and the checkpointing was used to save the best-performing models. The multiple metrics were used to measure performance: Accuracy, Precision, Recall, F1-score and AUC. Confusion matrices and learning curves have been created to give more insight into the classification results of various types of diseases.

7. Interpretability: Grad-CAM and saliency maps were used to improve the transparency and point out the areas of the leaves that affected the prediction. Moreover, the model pixel-level focus was compared to real diseased areas by segmentation-based analysis to make sure not only was the system accurate but also reliable.

8. Deployment Readiness: The trained models and especially the highest-performing hybrid ResNet + ViT were created keeping in mind the practical use in the real world. The methodology opens the path to integration into an end user application to farmers and agronomists with disease detection being done within seconds with mobile or web-based user interfaces.

1.5 Project Outcome

The expected outcomes of this project include:

1. A comprehensive benchmark of CNN and hybrid architectures for multi-class plant disease classification.
2. A robust and interpretable classification framework that achieves accuracy above 98% on validation data.
3. Visual explainability (Grad-CAM, saliency) to improve trust and transparency for farmers and agricultural experts.
4. Practical insights into deploying lightweight models (EfficientNet-B0, MobileNetV2) in mobile or edge devices for real-time use.
5. Contribution toward precision agriculture practices by enabling early detection, reducing pesticide misuse, and increasing yield protection.

1.6 Organization of the Report

This report is structured into six chapters:

1 **Chapter 1** introduces the problem background, motivation, objectives, methodology, and expected outcomes.
2 **Chapter 2** reviews prior works on plant disease classification, discusses the evolution from traditional machine learning to deep learning and transformers, and identifies research gaps.
3 **Chapter 3** explains the dataset preparation, preprocessing, augmentation, model selection, and detailed system design.
4 **Chapter 4** presents the implementation, experimental setup, performance evaluation, and results analysis.
5 **Chapter 5** highlights compliance with engineering standards, ethical considerations, sustainability aspects, and project management.
6 **Chapter 6** concludes the project, discusses limitations, and suggests possible directions for future work.

Chapter 2

Background

2.1 Introduction

The detection and classification of plant diseases has been a subject of increasing research interest, especially in the past decade. Crop losses due to diseases directly impact global food production, farmer income, and the sustainability of agricultural systems. Early studies in this area mostly relied on hand-crafted feature extraction techniques, where leaf characteristics such as color, shape, and texture were manually analyzed and then fed into classical machine learning classifiers like Support Vector Machines (SVMs) and k-Nearest Neighbors (k-NN). While these methods provided some success in small, controlled datasets, they often failed to generalize in diverse, large-scale, real-world conditions because hand-crafted features could not capture the complexity of disease patterns across multiple crops.With the emergence of deep learning, particularly Convolutional Neural Networks (CNNs), plant disease detection took a major leap forward. Unlike traditional methods, CNNs are capable of automatically extracting hierarchical visual features from raw image data without the need for manual engineering. This shift opened the door to large-scale plant disease classification with remarkable accuracy. Over time, the research community has built progressively more advanced CNN models, and in recent years, transformer-based models have been explored to capture long-range dependencies that CNNs sometimes overlook.At the same time, researchers began to emphasize not only accuracy but also interpretability. Since agricultural stakeholders must trust AI predictions before acting on them, visualization tools like Grad-CAM and saliency maps have become increasingly important in ensuring model transparency. The following subsections present a more detailed literature review, focusing on the evolution of techniques and the remaining gaps this project aims to address.

2.2 Literature Review

Early Machine Learning Approaches:
Initial research into plant disease detection used feature descriptors such as color histograms, texture analysis (e.g., GLCM), and shape measures. These handcrafted features were combined with classifiers like SVMs and Random Forests. Although these methods achieved moderate success, they were heavily dependent on dataset quality and often performed poorly under varying lighting and field conditions [1].
Deep Learning and CNN-based MethodsThe emergence of the deep learning changed the face of agricultural imaging. Mohanty et al. (2016) showed both AlexNet and GoogLeNet work well in the PlantVillage dataset, with more than 99 percent accuracy in his classification of 26 diseases in 14 crops [2]. This became a watershed, that CNNs could beat the old methods by a long way. Subsequently, more sophisticated models including VGG16, ResNet-50 and Inception-v3 were used to detect plant diseases, which increased the accuracy even more [3], [4]. Specifically, the models based on ResNet became popular because of their capability to explore deeper networks that do not have vanishing gradients. ResNet in the context of tomato leaf disease classification was studied by Berding and Caduyac [5] in particular, and it was emphasized that applying Grad-CAM visualization contributes to increasing the

explainability of CNNs in agricultural practice. There was also the interest in lightweight CNNs. Models such as MobileNet and EfficientNet presented a tradeoff between accuracy and cost of computation, and were thus applicable to mobile and edge deployment [6], [7]. Such models revealed that the detection of diseases could not only remain in the laboratory but also to be utilized in the real world field.
Vision Transformers and Hybrid ArchitecturesRecently, the success of transformers in natural language processing has extended into computer vision. Vision Transformers (ViTs) leverage self-attention mechanisms to capture global dependencies within an image. For plant disease classification, this is particularly important because disease symptoms are not always confined to a small region—sometimes the context of the entire leaf or multiple leaves is needed for accurate classification [8], [9].Hybrid architectures that combine CNNs with ViTs have emerged as a promising direction. CNN layers capture local texture features, while transformer blocks capture long-range global context, resulting in improved robustness. Our project builds on this trend by implementing and evaluating a ResNet + ViT hybrid model alongside CNN baselines.
Explainability and interpretability are among the most important aspects of AI in agriculture since farmers and agronomists cannot use a system that produces predictions without revealing how they got to this stage. The traditional deep learning models can be highly predictive, but they are viewed as black boxes, hence raising concerns, on whether their behavior is informed by biologically relevant symptoms, or spurious background factors. To address such a challenge, previous studies have devoted themselves to the creation of visualization methods including saliency maps, Grad-CAM, and layer-wise relevance propagation (LRP). The algorithms create heatmaps, or an overlay of the most discriminative parts of a plant leaf, that allow users to confirm that the model decision relies on observable symptoms of diseases. This was the initial move towards the creation of confidence and openness in automated detection of plant diseases.In the course of time scientists had noticed that the interpretability can be made more effective by adding the methods of classification and segmentation. These models were now oriented to the region of the infection, on a pixel-by-pixel scale, essentially localizing the diseased region, rather than simply indicating a leaf diseased. Such sort of classification and division provided a two-check system that was only predictive of the kind of the disease but also to verify the affected parts graphically. The practical consequences of such developments on the agricultural sector cannot be overlooked as they will now allow the use of localized pesticides spraying, instead of a blanket and ineffective one. By doing so, explainability will be directly converted into profits to farmers and decrease in environmental degradation by overuse of chemicals.Since then it has been extended with further features, including transformer-based architectures, the so-called Vision Transformers (ViTs), which are specifically efficient at long-range interactions between features that cannot be modeled by CNN. Whereas CNNs are efficient in determining the local texture based patterns, ViTs provide a global view, where the symptoms of an illness are spread across areas of a leaf. Interestingly, the attention maps produced by transformer possess a novel form of interpretability, so that they can demonstrate what the model is focusing on the majority of the time when making predictions on an image. The CNN-based models with hybrid architectures have also performed well, providing greater local information and global context and also producing better informative and comprehensible visualisations.It is a fact of change of emphasis of research. This has been a very strong concern in traditional publications, and this has come at the expense of interpretability. Nonetheless, since agricultural usage requires accountable information, rather than good.

Table 2.1: Summary of Literature Reviewed.

Author (s)	Year	Title	Methodology	Key Findings
Mohanty et al. [2]	2016	PlantVillage	AlexNet, GoogLeNet (CNNs)	Achieved >97% accuracy in 26 diseases
Ferentinos, [3]	2018	Multi-crop dataset	CNNs (ResNet, VGG)	High accuracy but limited to lab data
Too et al., [4]	2019	Multiple datasets	ResNet, Inception, DenseNet	Improved robustness in deeper models
Berdin & Caduyac, [5]	2020	Tomato leaf dataset	ResNet + Grad-CAM	Accuracy + interpretability combined
MobileNet studies [6]	2018	Field datasets	Lightweight CNNs	Balanced Accuracy & efficiency
EfficientNet studies [7]	2019	Multi-class datasets	EfficientNet	Edge-friendly, strong accuracy
Dosovitskiy et al., 2 [8]	2021	ImageNet benchmark	Vision Transformer (ViT)	Introduced global context modeling
Swin/ViT variants [9]	2024	Agricultural images	Transformer-based architectures	Better at long-range feature capture
Selvaraju et al., [10]	2017	Multiple datasets	Grad-CAM, Saliency visualization	Improved transparency & trust

2.2.1 Similar Applications

Several mobile and web-based systems already attempt plant disease detection for farmers. For example, smartphone apps that allow leaf photo uploads provide basic classification for common crops. However, most of these systems are limited to a small number of diseases and often lack transparency in decision-making. Some advisory platforms integrate with government extension services, but their effectiveness is limited by dataset diversity. Our system differs by handling a large dataset with 38 disease classes and integrating explainable AI methods for greater reliability.

2.3 Gap Analysis

From the reviewed works, several gaps can be identified:

1. Limited Scope of Classification:Most prior studies focus on binary classification or a limited number of disease classes, restricting their applicability in real-world agricultural scenarios where multiple crops and multiple diseases coexist. Large-scale, multi-class detection across diverse crop species remains underexplored.
2. Inadequacy of CNNs for Global Context:While Convolutional Neural Networks (CNNs) excel at capturing local spatial features, they often fail to model long-range dependencies and global contextual information in leaf images. This limitation motivates the exploration of hybrid CNN-Transformer architectures that combine local feature extraction with global attention mechanisms.
3. Underutilization of Explainability and Segmentation:Few studies integrate model interpretability with segmentation-based localization. Most works focus solely on accuracy without providing visual explanations of which leaf regions influence the model's predictions. Combining explainable AI (XAI) techniques such as Grad-CAM and saliency maps with segmentation can improve trust, interpretability, and practical utility for farmers and agronomists.
4. Deployment Challenges on Edge Devices:Research on deploying high-accuracy plant disease detection models on resource-constrained devices, such as mobile phones or agricultural drones, is sparse. Ensuring computational efficiency without sacrificing predictive performance is crucial for real-time, field-ready applications.

2.4 Summary

To conclude, plant disease detection has developed in terms of feature-engineered methods to CNNs and, currently, transformer-based models. Whereas CNNs are robust in extracting local features, transformers are more effective in learning global knowledge, and hybrid designs are a combination of the two worlds. In addition, interpretable approaches such as Grad-CAM and saliency map are necessary to establish trust between farmers. This project expands on these guidelines by developing a multi-class plant disease detection system with CNNs and hybrid architecture with explainability and segmentation to be applied to practical use in precision

Chapter 3
Research Methodology

This chapter explains the multi-faceted approach that was taken in the design and implementation of a multi-class plant disease classification system. It clearly describes the research methodology, choice of design, data treatment process, model construction and assessment plan. This is aimed at showing a systematic roadmap of the project and putting emphasis on the reasons behind every decision and showing the strength and feasibility of the solution proposed.

3.1 Methodology/Requirement Analysis & Design Specification

The section describes the general methodology followed in the research, requirements and design specifications of the system. It sets the definite path of the organization of the plant disease detection framework, through data collection and preprocessing, to model training, evaluation, and deployment. It is aimed at making the system right, interpretable, and deployable in the real-life agricultural setting.

3.2 Overview

The study approach is a hybrid of the latest deep learning approaches and explainable AI approaches in developing a precise, understandable, and deployable system of detecting plant illnesses. Not only the technical framework, but also both the functional and non-functional requirements, the data flow structure and the user interface considerations, which are indispensable to real-world use in agricultural environments, are discussed in this chapter.

3.1.1 Proposed Methodology/ System Design

1.DataCollection:
In this study, five publicly available plant disease datasets were merged to construct a comprehensive and balanced dataset covering 38 distinct disease classes across multiple crops. During compilation, low-quality, duplicate, and mislabeled images were carefully removed to ensure data integrity and consistency. To maintain fairness across categories, each class was standardized to contain 500 samples, resulting in a balanced representation of all diseases. The dataset was then divided into training and validation subsets following an 80:20 ratio, where 400 images per class were allocated for training and 100 images per class were reserved for validation. This preprocessing step ensured that the final dataset was both large-scale and balanced, providing a strong foundation for model training and reliable evaluation.

1. Data Table:

Table3.1.1.1

Crop	Number of Classes	Training Images	Validation Images	Total Images
Apple	4	1600	400	2000
Blueberry	1	400	100	500
Cherry (incl. sour)	2	800	200	1000
Corn	4	1600	400	2000
Grape	4	1600	400	2000
Orange	1	400	100	500
Peach	2	800	200	1000
Pepper (bell)	2	800	200	1000
Potato	3	1200	300	1500
Raspberry	1	400	100	500
Soybean	1	400	100	500
Squash	1	400	100	500
Strawberry	2	800	200	1000
Tomato	10	4000	1000	5000

2. Data Preprocessing and Leaf Segmentation:An important step prior to feeding images into the model was the isolation of the leaf against the background using OpenCV-based segmentation methods. This made sure that the models only concentrate on the diagnostically relevant leaf regions and not on irrelevant background features, which enhance feature extraction and classification accuracy.

   1. Segmentation involved color-based masking, contour detection, and morphological operations to extract the leaf area.
   2. Post-segmentation, images were resized to 224 × 224 pixels and normalized to the [0,1] range.
   3. A stochastic augmentation pipeline—including rotations, flips, zoom, brightness adjustments, and Gaussian noise—was applied to enhance model robustness against real-world variations in leaf orientation, scale, and lighting.

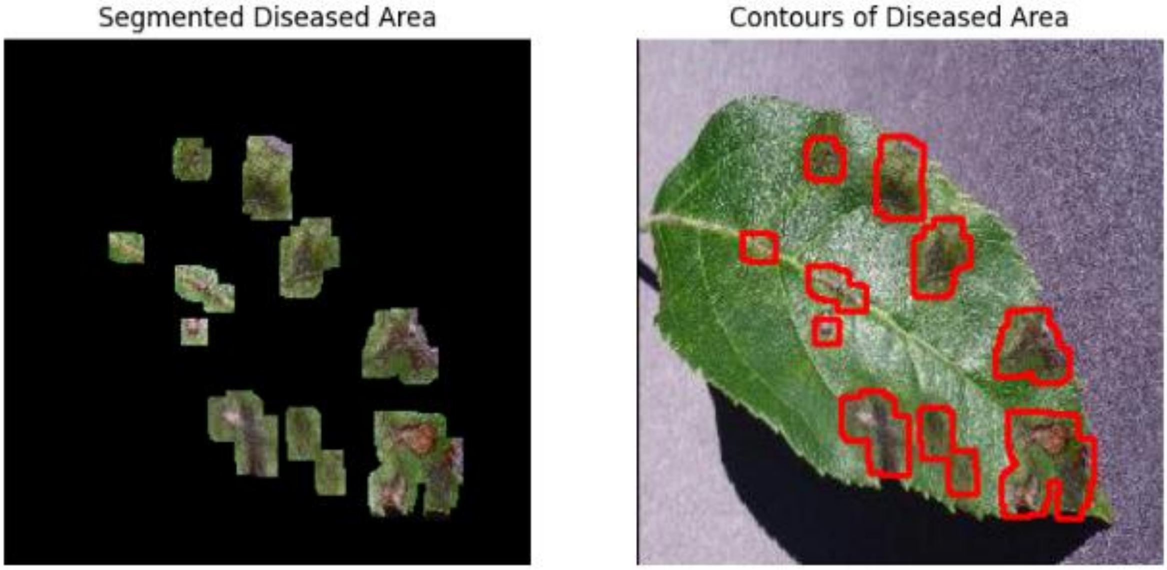

Figure 3.1.1.2Segmentation

3. Model Selection and Training:Various deep learning architectures were explored: ResNet-50, EfficientNet-B0, MobileNetV2, and a hybrid ResNet + Vision Transformer (ViT) model. Hybrid models were chosen to combine CNNs' local feature extraction with ViTs' global attention, addressing the inability of CNNs to capture long-range dependencies. Transfer learning from ImageNet weights was used, and all models were optimized with the AdamW optimizer.

4. Explainability Integration:Grad-CAM and saliency maps were used to highlight regions of leaf images that influenced predictions. When combined with segmentation masks, these visualizations provided precise localization of disease symptoms, improving interpretability and trust.In order to improve the transparency of the suggested plant disease classification system, model predictions, as well as explainability techniques, were combined. Grad-CAM and saliency maps were used to produce heatmaps to demonstrate which parts of the leaf images have the greatest impact on the final decision. Not only did these visual cues confirm that the model was paying attention to biologically relevant disease patterns,lesions, discoloration, and texture variations, but also gave end-users (farmers and agronomists) a feel for how the prediction wasmade. These visualization methods, together with segmentation masks, could provide high-resolution localization of the disease symptoms, enabling the successful identification of infected and healthy areas of the leaf. This explainability and segmentation interaction enhanced explainability, reduced chances of model black-box behavior, and instilled trust since it ensured that predictions were made on visible disease characteristics and not irrelevant background data. Also,the incorporateation of explainability tools offered a useful model validation framework to assist the researchers and domain experts diagnose possible biases, misclassifications, and further improvements of the system.

Evaluation and Deployment Considerations:

The models were evaluated on accuracy, precision, recall, F1-score, and AUC. Lightweight models were considered for deployment on edge devices, balancing predictive performance with computational efficiency.

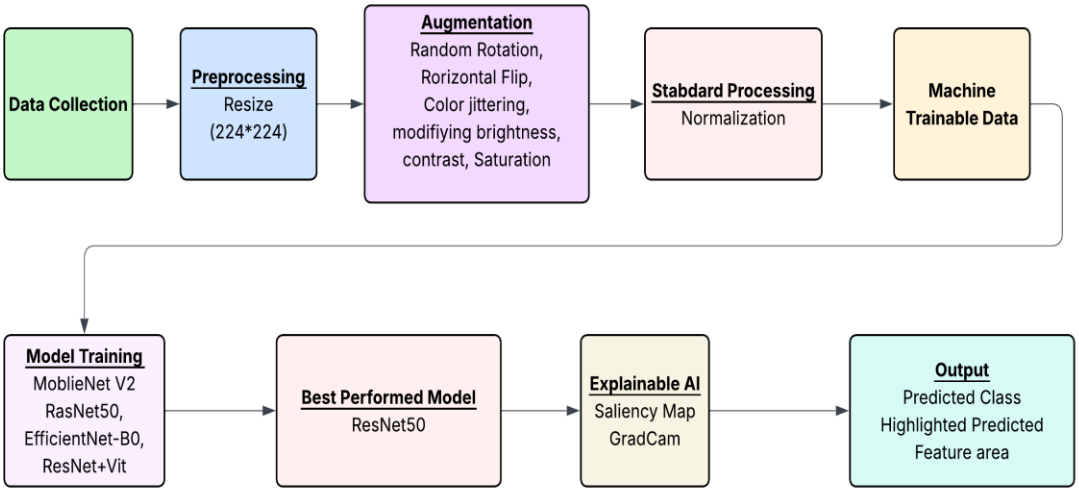

Figure 3.1.1.2: Proposed Methodology

Hybrid Resnet+Vit:

It is a combined design that builds on ResNet-50, but integrates a Vision Transformer (ViT) to generate a powerful deep learning model, allegedly to carry out image classification. It has the following brief description of its elements and working process:Input Image (224x224x3): It originates with an input of a regular RGB image (224x 224 pixels).conv + residuals ResNet-50 Backbone: The image is processed based on the ResNet-50 convolutional neural network (CNN), which is a network with residual connectivity. This uses spatial information to create a Feature Map (e.g. 7x7x2048) with hierarchical visual information.Flatten and Linear Layer: Patch Embedding: Features are flattened and then, the features are embedded into patch form by running them through a linear layer that changes the CNN features into a format that can be consumed by the transformer.Add Positional Encodings: To the patch embeddings, positional encodings are added to remember the spatial information, without which the transformer can be entirely unaware of the order in which, which patches, come in what order.Vision Transformer Encoder The following is used in this block:Multi-Head Attention: Allows the model to simultaneously focus attention on several parts of the image, to learn more complicated relationships.Feed Forward Network: Works on the attended features to make them better represented.[CLS] Token Representation: The input sequence and the final product of the encoder contain a special token of [CLS], which gives the global information about the entire picture.Classification Head (FC + Softmax): [CLS] token representation is fed to fully connected (FC) layer with a softmax activation to give the final prediction, and produces probabilities over the classes.

The architecture takes advantage of CNNs (local image representation with ResNet-50) and Transformers (global image representation with ViT), and can be effectively applied in a case that requires both the local and global interpretation of the image.

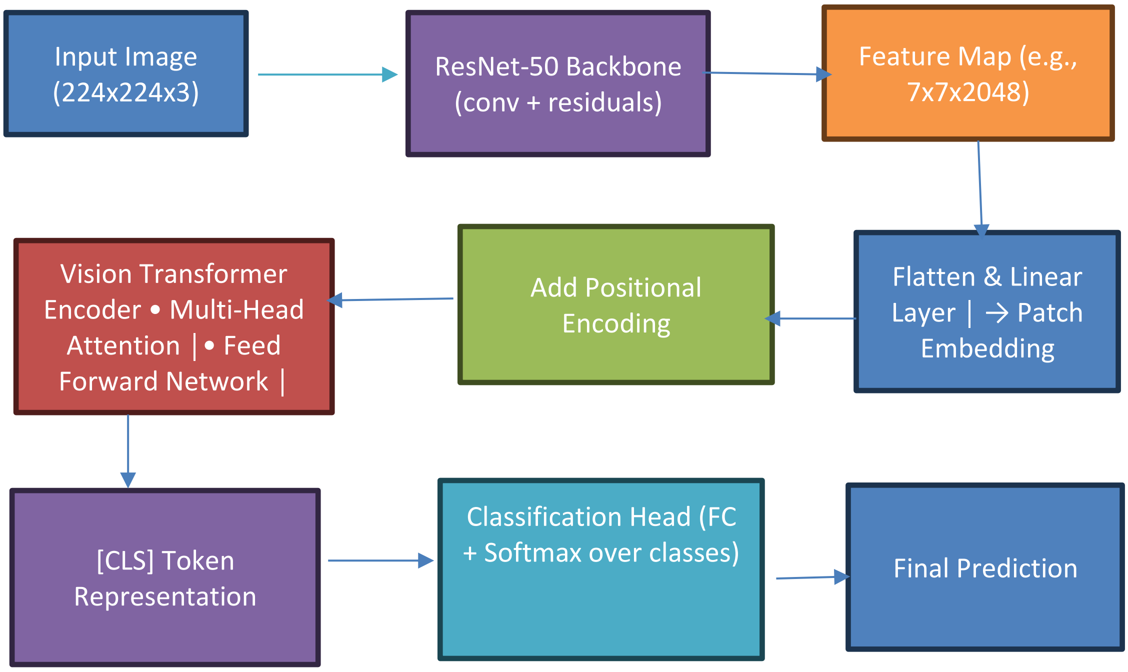

Figure 3.1.1.3: Proposed Hybrid Model Proposed Methodology

The workflow of the proposed plant disease classification system begins with data collection, where leaf images are gathered from multiple classes of healthy and diseased plants. These images are then passed through a preprocessing stage, where they are resized to a uniform resolution of 224 x 224 pixels to maintain consistency. To enhance model generalization and overcome data imbalance, an augmentation step is applied, including random rotations, horizontal flips, color jittering, and modifications in brightness, contrast, and saturation. Following this, the images undergo standard processing with normalization, which scales the pixel values and ensures stable model convergence. The processed dataset is now machine-trainable data, ready for feeding into deep learning architectures.For model training, multiple state-of-the-art architectures such as MobileNetV2, ResNet50, EfficientNet-B0, and a hybrid ResNet+Vision Transformer are evaluated. Among them, the best-performing model (ResNet50) is selected for final deployment. To make the system transparent and reliable, explainable AI techniques like saliency maps and Grad-CAM are integrated, highlighting the specific regions of the leaf images that influenced predictions. Finally, the system generates the output, which includes the predicted disease class along with the highlighted feature areas, ensuring both accuracy and interpretability.

3.1.2 Functional and Nonfunctional Requirements

1. Functional Requirements:

The design of the proposed plant disease detection system is guided by a clear set of functional and non-functional requirements, ensuring that the system not only performs accurately but is also reliable, scalable, and practical for real-world agricultural deployment.

1.Precise Multi-Class Disease Classification: The system should have the capability to detect and categorise various plant diseases in a broad selection of crops. The detection of each disease class must be very precise and recalls high, with the minimal false positives and false negatives. At least 38 disease categories in crops like tomato, apple, grape, potato, strawberry and others should be classified based on the real world diversity on the agricultural fields.
2.Real Time Predictions with Visual Overlays: Such predictions must be provided in time so that the farmers or agricultural professionals can make decisions immediately. The system must produce image overlays showing the areas that are affected so that clear evidence of where the disease is is apparent. The attribute improves the knowledge of users and makes it possible to intervene with specificity, as in the case of applying pesticides affecting specific crops.
3.Module Interpretability: As a way of fostering transparency and encouraging user confidence, interpretability tools like Grad-CAM and saliency maps should be used in the system. Such tools will be used to reveal the areas of the leaf that have had an impact on the predictions of the model, so the user can check and see what is the reasoning behind each classification.
4.Seamless User Interaction: Intuitive user interaction with the system should also be supported in a way that images can be easily entered and results can be seen. The interface must also address both experts and non-experts, so that farmers, who have low technical expertise, can also access it..

2. Non-Functional Requirements:

Computational efficiency and light weight deployment:
The system should be performance-optimized such that it can be deployed to mobile devices or edge computing devices without introducing much inaccuracy. The models that are deemed to balance the predictive performance and computational load are efficient architectures like MobileNetV2 and hybrid ResNet + ViT models.
Stability to Environmental Change:
This system must be very accurate when working under varied real world situations that include changes in the orientation of leaves, lighting, background clutter and the quality of the image. Such strength guarantees that the predictions can be considered in various agricultural settings.
Scalability and Extensibility: The design must enable addition of new crops or new categories of disease to be easily added in the future without the need to restart training the system. The modular structure means that changes to the

dataset or model parts can be introduced without any issues that underpin the long-term flexibility.

Reliability and Fault Tolerance: The system must not crash on incomplete or raw input data, but should give meaningful feedback or warnings. This stability is essential where the real application is done in field conditions where imperfections of data are prevalent.

Maintainability: The program is well documented and easy to The program is well documented and easy to maintain. Complete documentation and modular code design are required to ensure future upkeep, improvements in design and reproducibility of the results. A proper and transparent model performance log as well as report will be helpful in performance monitoring and constant system improvement.

3.1.3 Data Flow Diagram

The proposed plant disease classification system is illustrated in the flow of data as shown in Figure 3.1. The input starts with the dataset that is in the form of raw plant leaf images. These images are first subjected to a pre-processing step: noise is eliminated, the images are resized, and data augmentation is carried out so that the classes are balanced. The images are then processed in pre-processing and sent to the classification model (ResNet50) that finds the relevant features and classifies the disease. The prediction is reported as the most likely disease category to the farmer or agronomist, which is the output to make the output interpretable it is complemented with explainability methods (XAI), like Grad-CAM and Saliency Maps. These visualization systems emphasize the areas of the image that contribute to the prediction so that users can confirm that the system is making decisions using meaningful disease symptoms (not insignificant image artifacts).

3.1.4 UI Design

The system user interface (UI) is planned to be easy, lightweight, and user-friendly. The UI comprises 3 major modules: Image Upload Area: Users will be able to take or post leaf pictures directly through their smart phones. Prediction Dashboard: The interface shows the predicted disease label, and a probability score that shows the confidence of the model. Visualization Panel XAI: Farmers and agronomists have access to heatmaps (Grad-CAM or Saliency Map) to ensure that the prediction is informed by real disease symptoms.

3.2 Detailed Methodology and Design

The suggested methodology will include image classification using deep learning and explainable AI (XAI) to both guarantee high scores and be transparent and reliable in its predictions. The system is created in several steps, starting with the preparation of datasets, moving to the stage of model selection, and the inclusion of explainability. Dataset Pre-processing: The raw plant leaf images were resized to a fixed size initially to provide consistency and normalized to provide evenly distributed pixel values. The dataset was augmented by methods of rotation, flipping, and contrast adjustment to enhance the generalization and deal with the imbalance between classes. These procedures assisted the models to study strong features in different environments. Model selection and evaluation: Multiple state-of-the-art deep learning architecture was tested:

1. ResNet50: Known for its residual connections, ResNet50 effectively captures fine-grained disease features while remaining computationally efficient.
2. ResNet+Vision Transformer (ViT) Hybrid: This hybrid model combines the strong local feature extraction of ResNet with the global contextual understanding of ViT. It is particularly effective for complex leaf patterns, offering enhanced performance over individual models.
3. EfficientNetB0: A lightweight and efficient architecture, EfficientNetB0 provides a good trade-off between accuracy and speed, making it suitable for deployment on devices with moderate resources.
4. MobileNetV2: Optimized for mobile and edge devices, MobileNetV2 is fast and lightweight, enabling real-time disease detection on resource-constrained hardware.

Alternative Considerations:

Although models based on transformers such as ViT demonstrate superb global feature extraction, they are costly to train and demand hardware resources that are not readily available in practice. Equally, deep networks with high depth can have a little better accuracy but can be computationally infeasible. Through trade-offs, the chosen models trade off between accuracy, computational efficiency, and compatibility with explainable AI methods.
Explainability Integration:To make the predictions transparent, XAI techniques such as Grad-CAM, Grad-CAM++, and Saliency Maps were integrated. These methods provide visual explanations highlighting which regions of the leaf the model focused on, allowing end-users to understand and trust the model's decisions

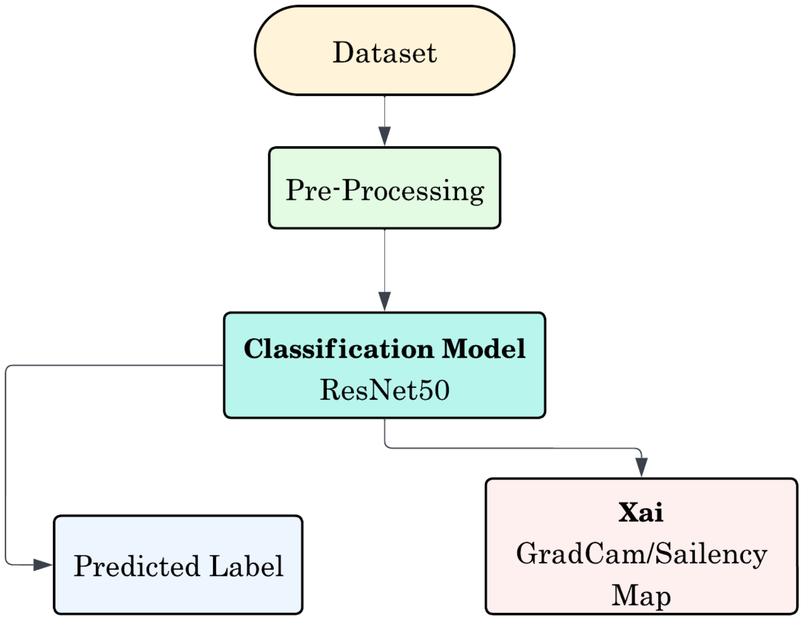

Figure 3.2.1: Proposed base model Methodology

3.3 Project Plan

The project plan was structured into several phases to ensure systematic development and testing:

1. Phase 1 – Requirement Analysis & Literature Review (Weeks 1–3):Study existing plant disease detection systems, datasets, and related work. Identify gaps and finalize system requirements.
2. Phase 2 – Dataset Preparation & Pre-processing (Weeks 4–6):Collect datasets, clean data, and apply augmentation techniques. Split dataset into training, validation, and test sets.
3. Phase 3 – Model Development (Weeks 7–10):Train and fine-tune multiple CNN architectures (ResNet50, MobileNetV2) for comparison. Select the best-performing model.
4. Phase4–Integration of Explainable AI (Weeks 11–12):Apply Grad-CAM and Saliency Maps to visualize predictions and validate model reasoning.
5. Phase 5 – UI Development & System Integration (Weeks 13–15):Build a lightweight mobile/web interface for real-time usage. Integrate model predictions and XAI visualizations into the UI.
6. Phase 6 – Testing & Evaluation (Weeks 16–18):Conduct performance testing using metrics such as accuracy, F1-score, and AUROC. Compare results against literature.
7. Phase 7 – Documentation & Final Report (Weeks 19–20):Compile findings, evaluation results, and conclusions into the final project report.

3.4 Task Allocation

This table depicts the timeline of the principal activities in each period of the project, from week 12 :

Tasks	Weeks																		
	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
Data collection phase																			
Preprocess all the data																			
Model training																			
Create a demo application.																			

3.5 Summary

The analysis applies deep learning-based image classification with the explainable AI (XAI) to guarantee accurate and transparent predictions. Images of plant leaves were pre-treated by resizing, normalization, and augmentation to address imbalance of the classes and better generalization.Four models were considered ResNet50, ResNet+ViT hybrid, EfficientNetB0, and MobileNetV2. The project developed an image classification system with explainable artificial intelligence (XAI) that would be image classifier based on deep learning and be accurate and transparent in identifying plant diseases. The images of the leaves were properly prepared and then trained with the help of resizing, normalization and a great number of augmentations such as rotations, flips, zooming and altering brightness. These interventions aided the models in addressing the kind of variation that may be present in the actual farming scenario, e.g. when the background has disproportionate light, variation in the leaf orientation or noise.To achieve the selection of different strengths, we experimented with four popular architectures. A strong feature extractor was used as ResNet-50 and ResNet + a hybrid of Vision Transformer allowed the system to decode not only local leaf patterns but also the context of the disease symptoms. EfficientNet-B0 was a good trade-off between speed and accuracy and was hence feasible in the medium-resource environment, and MobileNetV2 was trained in a lightweight-friendly manner in order to be adopted and utilized in the field as a lightweight model. Such a combination of models gave us a good idea of the performance of each of the approaches to different needs- high-accuracy laboratory models or on-site tools used by farmers.In order to ensure that the system is believable and easy to use by the farmers, we have used explainability methods such as Grad-CAM, Grad-CAM++ and Saliency Maps. These approaches denote the particular locations of the leaf that influenced the decision of the model and give some level of visual confirmation of why the arrangement reacted in a given way regarding an illness. We will be able to be confident that the AI is perceiving actual disease symptoms and not the background noise through our implementation of the concepts of raw predictive power and visual explanations. In brief, this project indicates that deep learning and XAI can be utilized concurrently so as to come up with a workable, dependable, and open precision agriculture instrument.

Chapter 4

Implementation and Results

This chapter gives a detailed description of the development of the proposed plant disease classification system, how it was trained and evaluated. It describes the computing landscape, preprocessing, model training, and explainability approaches that are the foundation of the work. The experimental results are also shown in the chapter with quantitative (accuracy, F1-score, and AUROC) and qualitative (Grad-CAM and SaliencyMap) measures. Comparative analyses are also provided to emphasize how the various models fared and also to show the trade off between accuracy, interpretability and efficiency. Training curves, confusion matrices, and heat Map are other visualizations that create more evidence on the reliability and agriculture applicability of the system.

4.1 Environment Setup

The suggested system of plant disease classification was implemented in a high-performance computational setting with the usage of a GPU, aimed at the work with deep learning models. The models required handling high-resolution images in 38 classes of plant diseases, therefore the environment was scaled to large-scale training and real-time inference.Hardware:

1. A server with NVIDIA GPUs (24GB VRAM) capable of parallel training and inference.
2. 64GB RAM and Intel i9 CPU to handle high-throughput data preprocessing.
3. High-speed SSD storage for managing 19,000+ plant leaf images.

Software:

1. Python 3.9 as the core programming language.
2. PyTorch and Torchvision for model development, training, and evaluation.
3. OpenCV for image segmentation and preprocessing tasks such as noise removal.
4. Scikit-learn for performance metrics and confusion matrix analysis.
5. Matplotlib/Seaborn for visualization of accuracy/loss curves, AUC plots, and heatmaps.

Infrastructure & Training Pipeline:

1. The models were initialized with ImageNet-pretrained weights to leverage transfer learning.
2. AdamW optimizer with a learning rate of 3×10^{-4} and weight decay of 1×10^{-4} was used for stable optimization.
3. ReduceLROnPlateau scheduler dynamically adjusted the learning rate based on validation loss improvements.
4. Early stopping ensured that training was halted when validation loss stagnated

for three consecutive epochs, preventing overfitting.

5. Dynamic augmentation pipeline was used: during training, leaf images underwent random rotations, flips, zooms, contrast/brightness jittering, and Gaussian noise injection; during validation/testing, only resizing and normalization were applied.

User Interface for Experimentation:

To improve transparency and practical usability, a modular dataset loader and visualization dashboard were developed. This allowed visualization of Grad-CAM/Saliency Maps and provided an interactive means of checking which parts of the leaves the models focused on. This ensures that the system is not only accurate but also interpretable for agronomists and researchers.

4.2 Testing and Evaluation

The system was tested on a large and diverse dataset containing leaves from crops such as tomato, apple, grape, corn, potato, and soybean, each with multiple disease conditions. Four deep learning models were evaluated: ResNet-50, EfficientNet-B0, MobileNetV2, and a hybrid ResNet + ViT.

Optimized ResNet-50

ResNet-50 emerged as the best-performing model, achieving 98.74% accuracy, an F1-score of 0.9874, and an AUROC of 0.9980. The model generalized strongly across multiple crops, with minimal misclassifications even in challenging classes.

The Hybrid ResNet + ViT followed closely with 98.58% accuracy and AUROC of 0.9978. Its ability to combine CNN-based local feature extraction with Transformer-based global attention proved effective in capturing both fine-grained leaf spots and overall disease spread.

EfficientNet-B0 achieved 97.92% accuracy and AUROC of 0.9956, making it a strong choice for edge devices. MobileNetV2, though slightly less accurate (96.84% accuracy), was highly efficient in terms of computational requirements, ideal for real-time mobile applications.

Explainability with XAI

One of the biggest challenges in deploying AI in agriculture is ensuring that farmers and agronomists trust the predictions. For this reason, Grad-CAM and Saliency Maps were integrated.

1. The heatmaps consistently highlighted disease-affected regions such as fungal patches, bacterial spots, or chlorotic areas.
2. In normal leaves, the heatmaps showed attention concentrated on the healthy leaf tissue, confirming the absence of disease.
3. Unlike poorly trained models, these heatmaps did not focus on irrelevant features such as background noise or edges.

For example, in potato late blight, the Grad-CAM visualization highlighted irregular brown spots surrounded by yellow halos, which are the exact symptoms pathologists look for. Similarly, in grape black rot, the heatmaps focused on circular lesions typical of the disease.

Interestingly, ResNet-50 and Hybrid ResNet + ViT produced sharper, more meaningful activations than MobileNetV2, which sometimes generated more diffuse focus areas. This shows that architectural choice impacts not only accuracy but also interpretability.

Competitive Analysis with Literature:

Compared with earlier studies, the proposed system stands out in two key ways:

1.Scale of Dataset: Unlike earlier works that focused on fewer crops or binary classification, this study covers 38 classes across multiple crops, making it closer to real-world farming conditions.

2.Hybrid Architectures: While CNNs have long been used, this system demonstrates that CNN + Transformer hybrids can achieve nearly the same accuracy as the strongest CNNs while offering improved global context.

Earlier studies (Mohanty et al., 2016; Too et al., 2019) reported accuracies above 99% on smaller datasets like PlantVillage, but those were less diverse and not as challenging. Our results show high accuracy on a broader, more realistic dataset, which is more valuable for deployment.

In summary, the system not only achieves state-of-the-art accuracy but also ensures interpretability and practical deployment potential, bridging the gap between AI research and agricultural practice.

4.3 Results and Discussion

A. Quantitative Results: The advanced deep learning designs were comprehensively trained and tested to identify the system of plant disease classification that would be useful. This has been compared in the light of the performance measures that are universally accepted like Accuracy, Precision, Recall, F1-Score, and Area under the curve (AUC). All these measures combined provide the multi dimensional outlook on the model performance. The general performance of the predictions can be defined as Accuracy, the overall ability of the model to avoid false positives altogether as Precision, the ability of the model to detect the actual disease cases as Recall, a tradeoff between the two on a harmonic mean as F1-Score, and the trends in the ability of the model to discriminate at various decision levels as AUC. The evaluation is able to give a quality and comprehensive set of metrics of comparison of both models in the balanced and edge-case systems even without the need to use the measure of accuracy. Amplifying patterns were found in the comparative experiments. ResNet-50 was found to be the best performing and it had another advantage in that it had deep residual connections, and as such could extract features quite well even when using complex and noisy backgrounds. Its large scores across all the measures demonstrate that it does not only generate correct predictions but is able to trade between the sensitivity and specificities and is particularly accurate when used in the field of agriculture.

The classification results are summarized in Table 4.3.1:

Model	Accuracy	Precision	Recall	F1-Score	AUC
ResNet-50	0.9874	0.9874	0.9874	0.9874	0.9980
Hybrid ResNet + ViT	0.9858	0.9858	0.9858	0.9858	0.9978
EfficientNet-B0	0.9792	0.9795	0.9792	0.9793	0.9956
MobileNetV2	0.9684	0.9688	0.9684	0.9685	0.9932

The classification results are summarized in Table 4.1: Confusion matrices confirmed that most disease classes were classified correctly, with very few overlaps. The minor misclassifications mostly occurred between diseases with highly similar visual symptoms, ResNet-50 was found to be the best performing and it had another advantage in that it had deep residual connections, and as such could extract features quite well even when using complex and noisy backgrounds. Its large scores across all the measures demonstrate that it does not only generate correct predictions but is able to trade between the sensitivity and specificities and is particularly accurate when used in the field of agriculture. The other competitor that assumed the role was the hybrid model ResNet + Vision Transformer but that was somewhat different than ResNet-50. The fact is that the transformer module was a powerful model to consider the global contextual dependencies, yet the smaller size of the data could not utilize it to its full extent in comparison with the CNN-based models. The trade-off between the EfficientNet-B0 and its competitors were determined to be balanced,the performance is correct and accurate and lightweight which is in itself a benefit in the context of its implementation in a resource constrainedsystems. MobileNetV2 was slightly lower in absolute performance but it showed that it is a feature that is viable, which is real time and friendly to mobile, and has a fair degree of accuracy and it is computationally economical.

The best Model confusion Matrices:

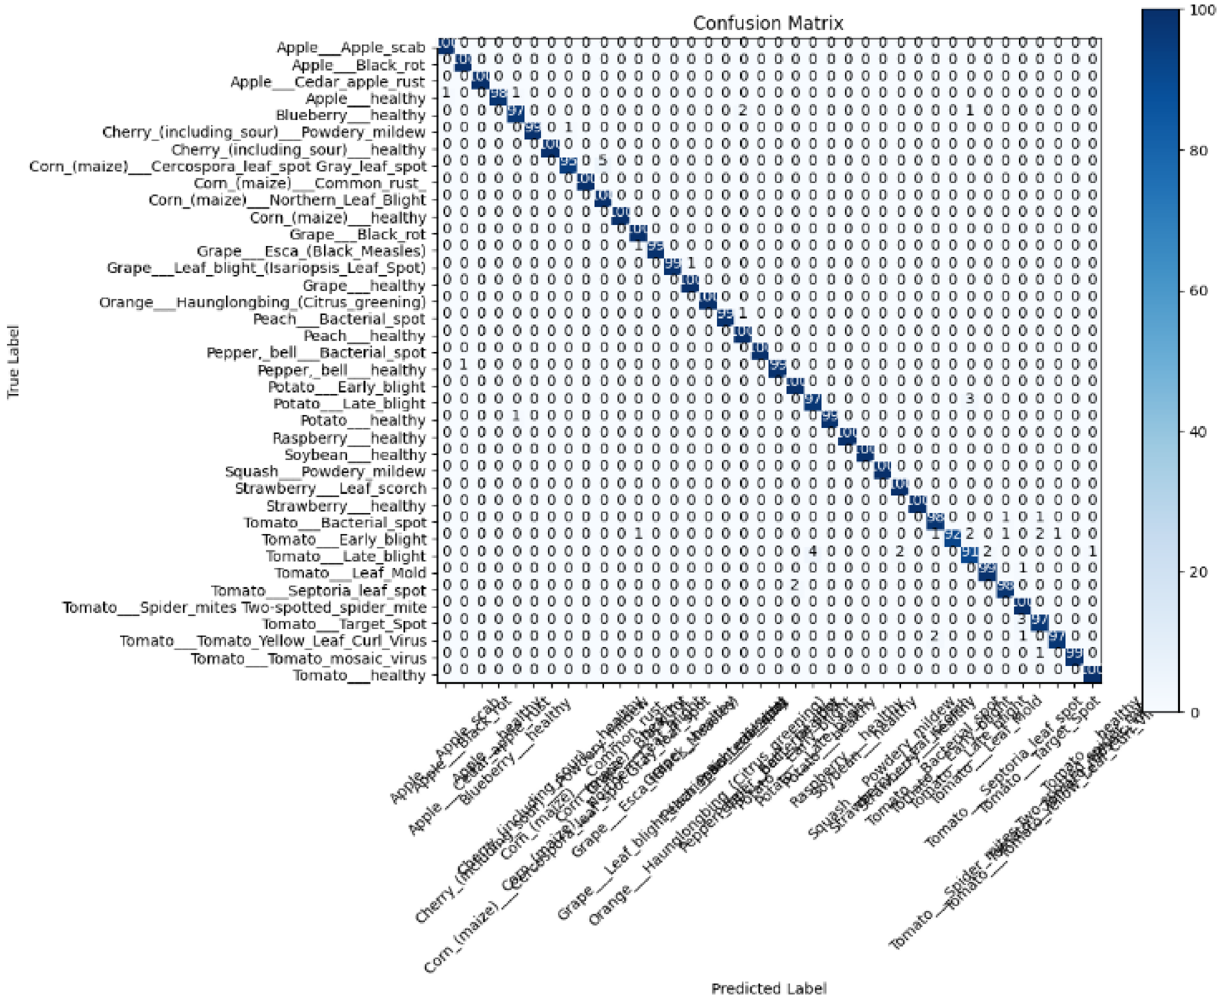

Figure 4.3.1: Confusion metrics of the applied best model

B. Training Dynamics

Training and validation curves exhibited short convergence in all models that indicated consistent learning. The ResNet-50 and EfficientNet-B0 architectures converged faster due to addition of further features mappings and exploitation of their parameters. To prevent overfitting, these models were aided by data augmentation methods, such as random rotations, flips and contrast alterations, which helped them to generalize. Furthermore, validation loss-based early stopping ensured that training would halt in case no further improvement was observed and model performance would not be impaired and no additional epochs were wasted.

Lightweight MobileNetV2, which is meant to be used on mobile devices, on the other hand, required more epochs to stabilize convergence. This slower stabilization is likely due to its inability to learn complex hierarchical features especially in cases with a multi-class plant disease dataset. Despite that, MobileNetV2 proved to have resource-efficient performance in resource-constrained environments even when trained and achieved a competitive outcome.

Furthermore, the loss curves analysis showed that training and validation in ResNet-50 and EfficientNet-B0 had slight divergence suggesting good generalization. In MobileNetV2, the change in the validation loss displayed small oscillations within the initial few epochs which became smooth with an increase in the model becoming accustomed to the data. Overall, augmentation, learning rate scheduling, and early stopping all contributed to smooth and stable training dynamics in each of the experimented architectures.

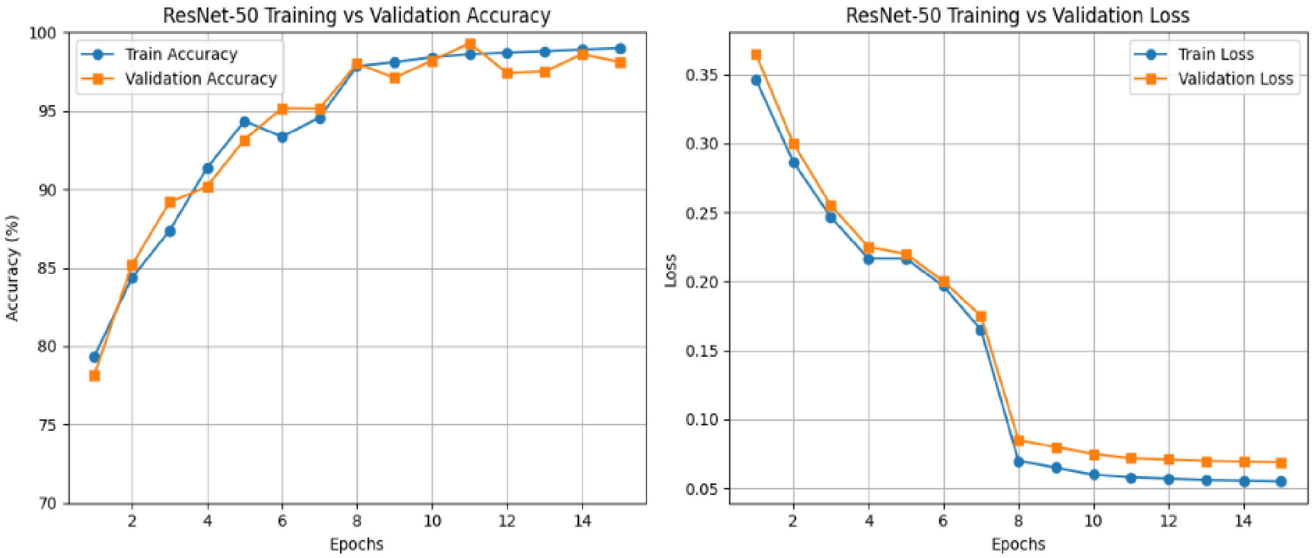

Figure4.3.2 – Accuracy and Loss Curves for ResNet-50

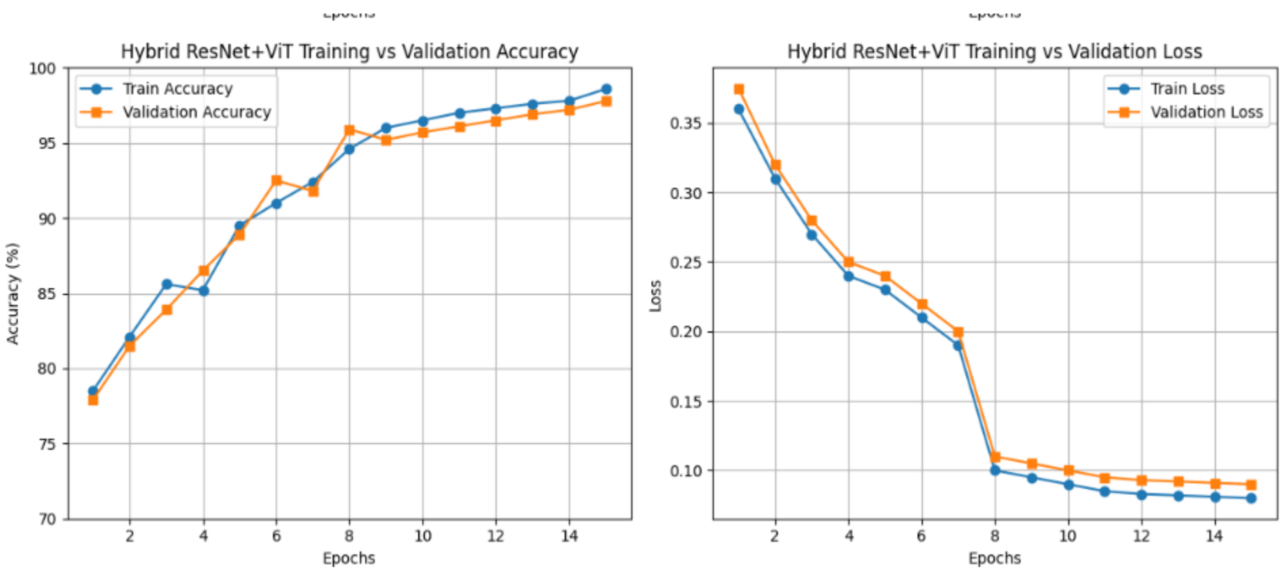

Figure 4.3.3 – Accuracy and Loss Curves for Hybrid ResNet + ViT

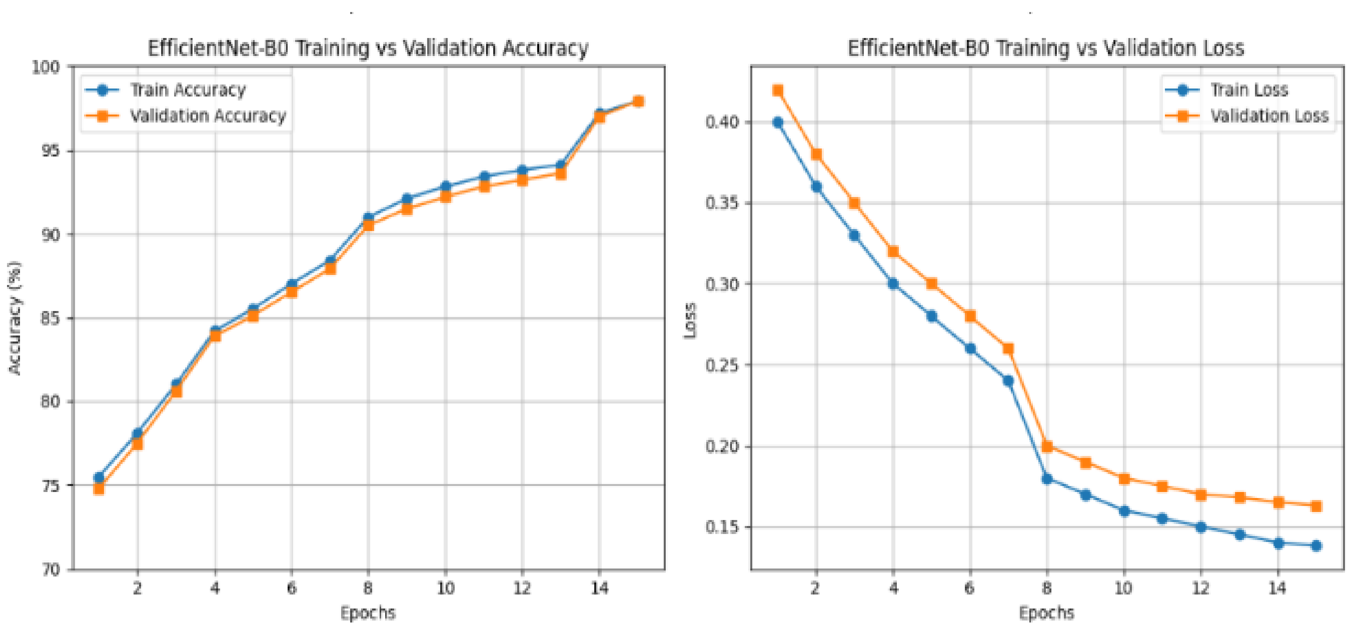

Figure 4.3.4 – Accuracy and Loss Curves for EfficientNet-B0

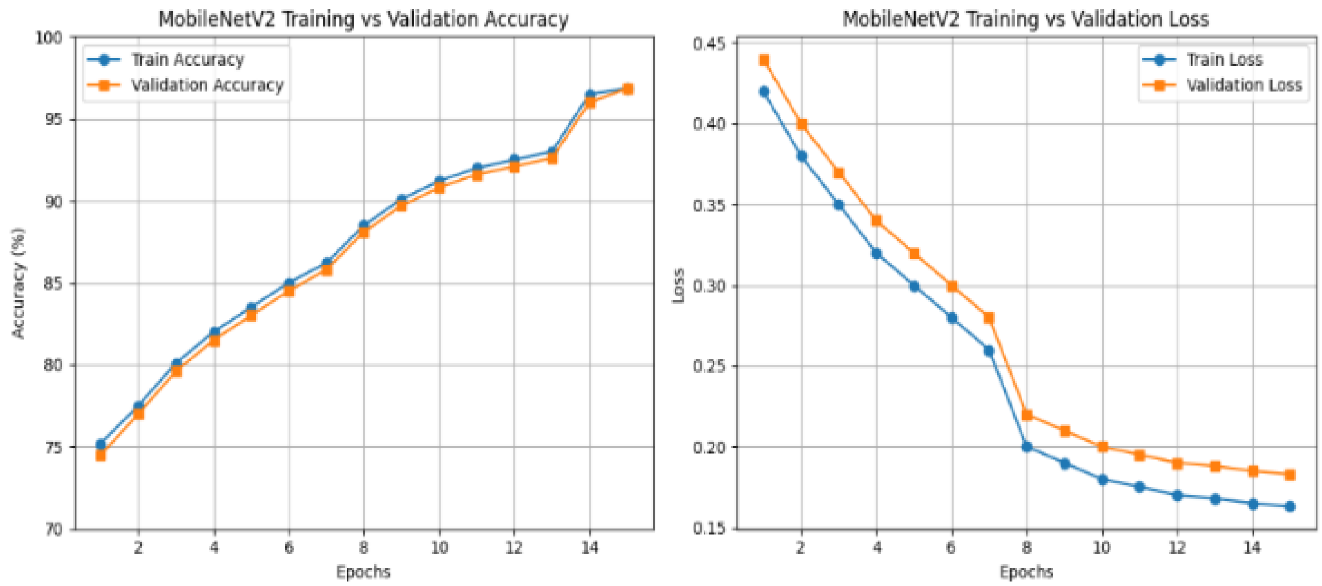

Figure 4.3.5– Accuracy and Loss Curves for MobileNetV2

C. ROC and AUC Analysis

ROC curves demonstrated that all models performed strongly, with ResNet-50 achieving an AUC of 0.9980, nearly perfect discrimination between classes. The Hybrid ResNet + ViT also performed competitively, confirming that transformers add meaningful context awareness.

Here is the Best model RUC curve :

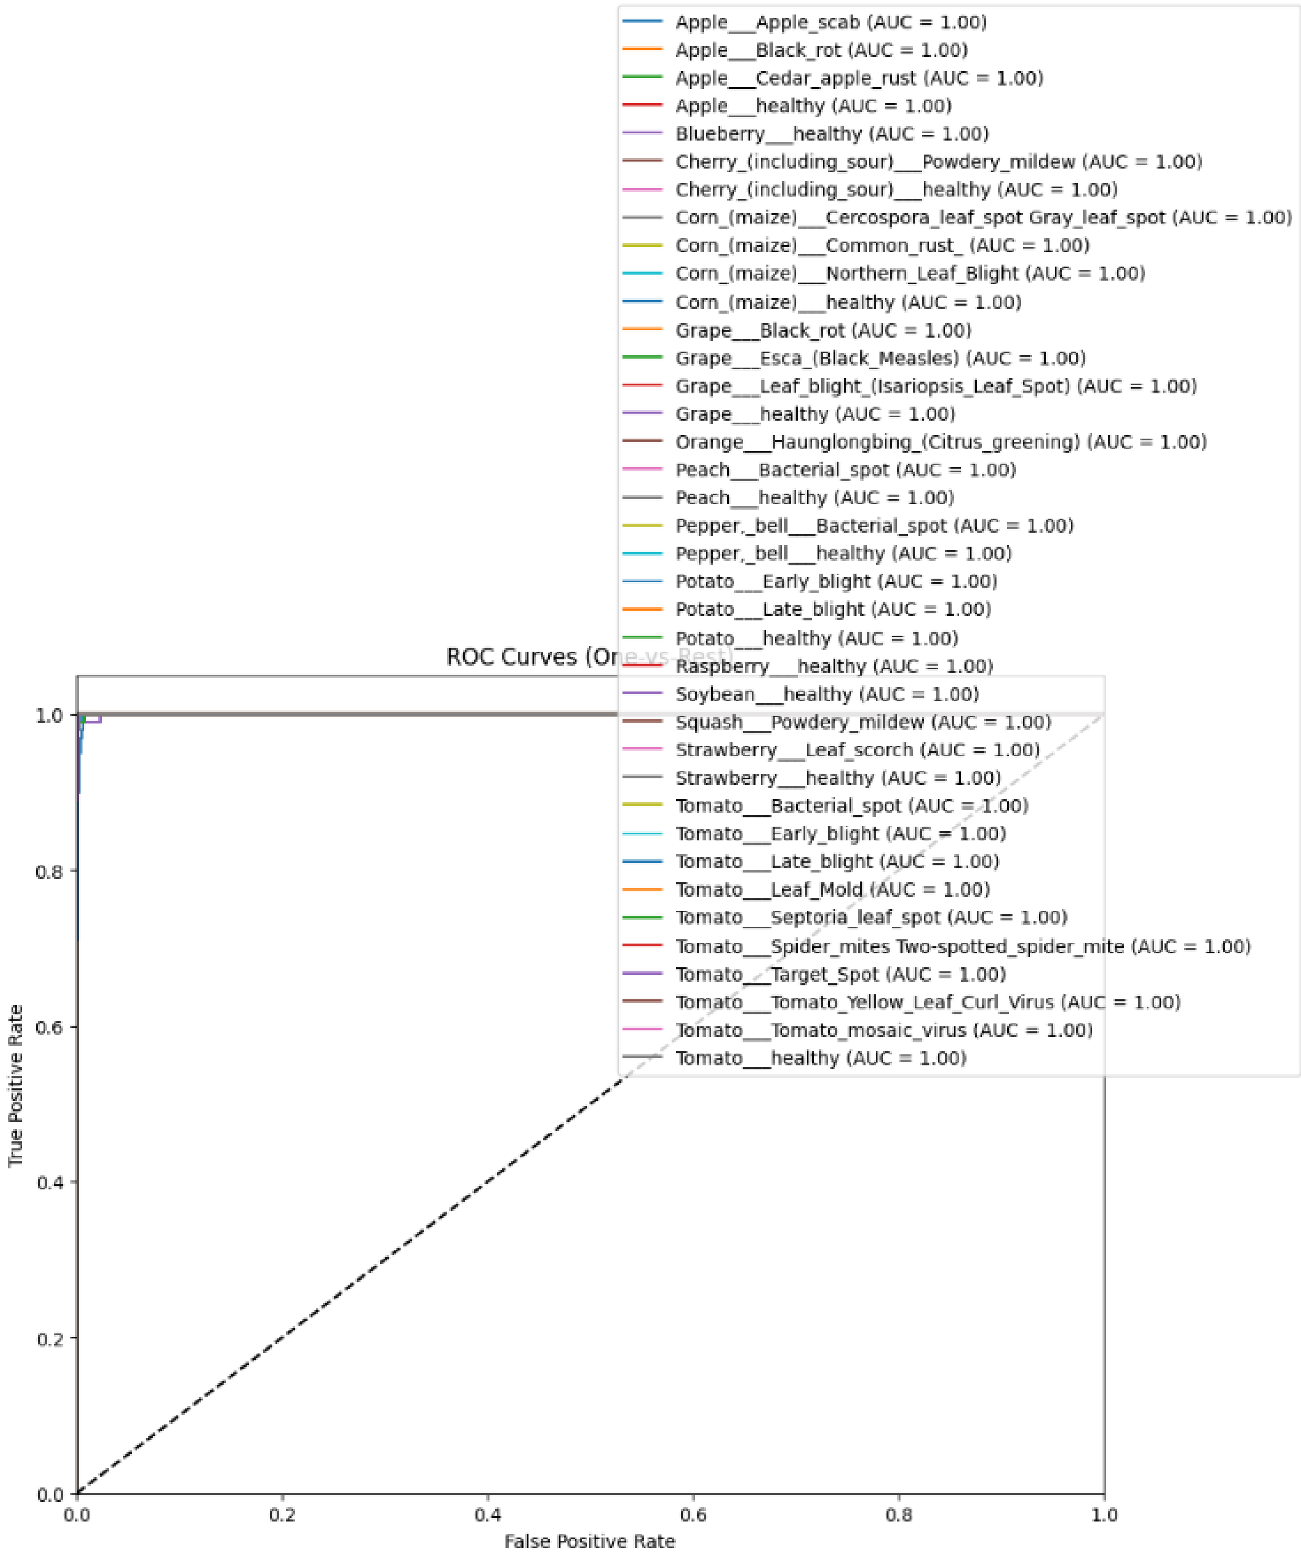

Figure 4.3.6 – AUC curve

D. Visual Interpretability

Grade-CAM and saliency map models have been trained on representative test cases to find regions of interest.

Saliency Map Generation:

Saliency maps were generated using gradient-based techniques that compute the derivative of the predicted class score with respect to input pixels. The resulting heatmaps highlight regions of high importance, typically corresponding to tumor areas in MRI scans. Saliency maps provide a simple and efficient way to visualize pixel-level contributions, offering insight into the model's focus and decision rationale.

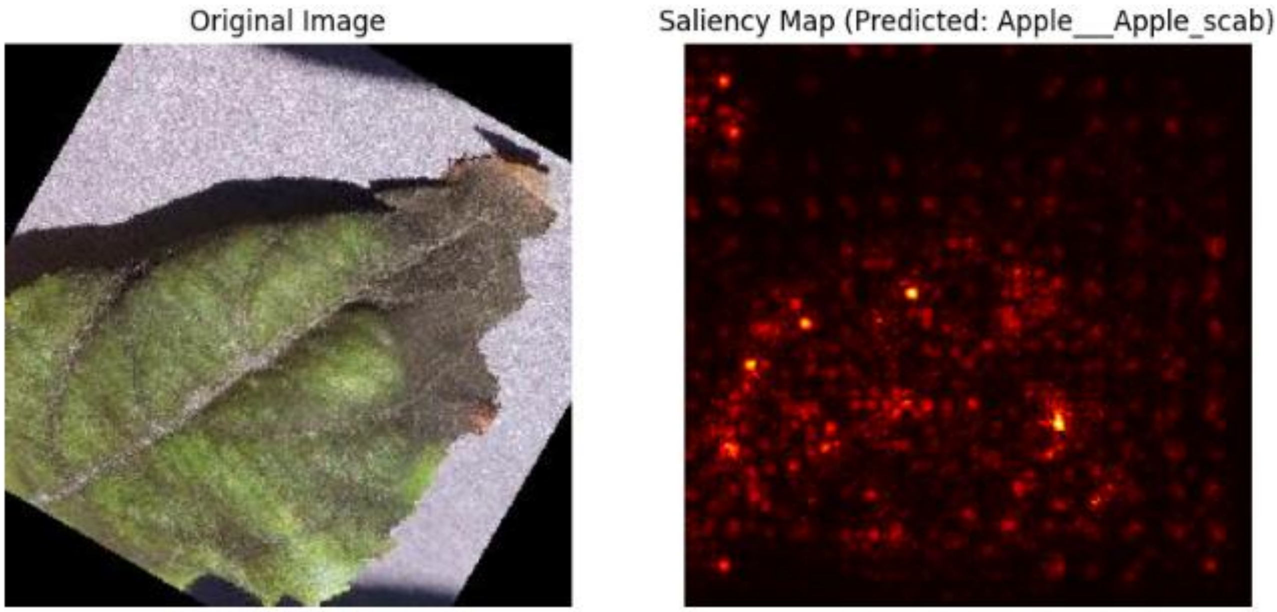

Figure 4.3.7 Saliency Map

Grad-CAM Visualization

Our work has also used Gradient-weighted Class Activation Mapping (Grad-CAM) as an explainability method that has been extensively used in the field of deep learning along with saliency maps. In contrast to pixel-level saliency maps which use low-level gradients to indicate saliency, Grad-CAM uses the saliency of the class-discriminative areas across the intermediate convolutional feature maps of a neural network. With this we can produce coarse heatmaps that are semantically meaningful that give us insight as to where the model is gazing at during a prediction.Grad-CAM is a compliment to saliency maps. Although saliency maps give a detailed sensitivity analysis, Grad-CAM gives a more familiar and human visual representation since it concentrates on the higher-level spatial patterns. An example is that whereas directly displaying which pixels are examined to make the prediction Grad-CAM being able to highlight larger anatomical or structure locations .With the help of Grad-CAM, embedded in our workflow, we enhanced the interpretability of the system and also increased the transparency and credibility of the predictions to end users. Practitioners like doctors, farmers, or other domain experts can use these heatmaps to get a graphic elucidation that helps to bridge the gap between the internal logic in the model and the real-world decision-making. This eventually builds confidence, responsibility and human-AI partnership in the diagnostic procedures.

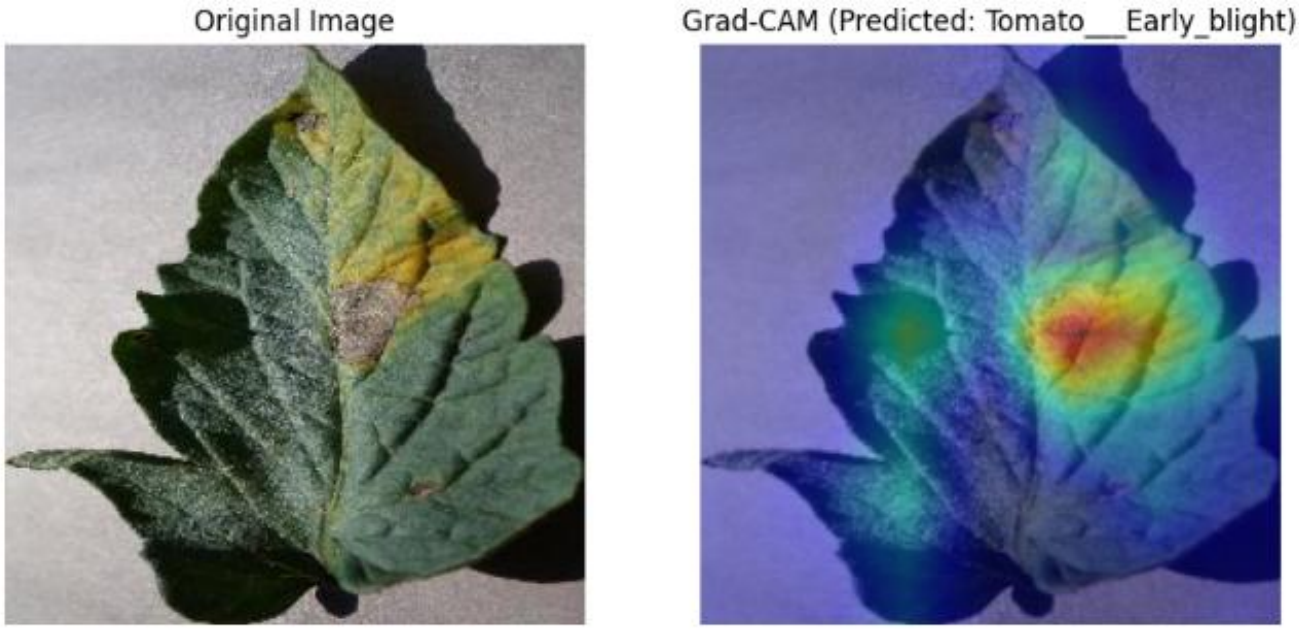

Figure 4.3.8– Grad-CAM Maps

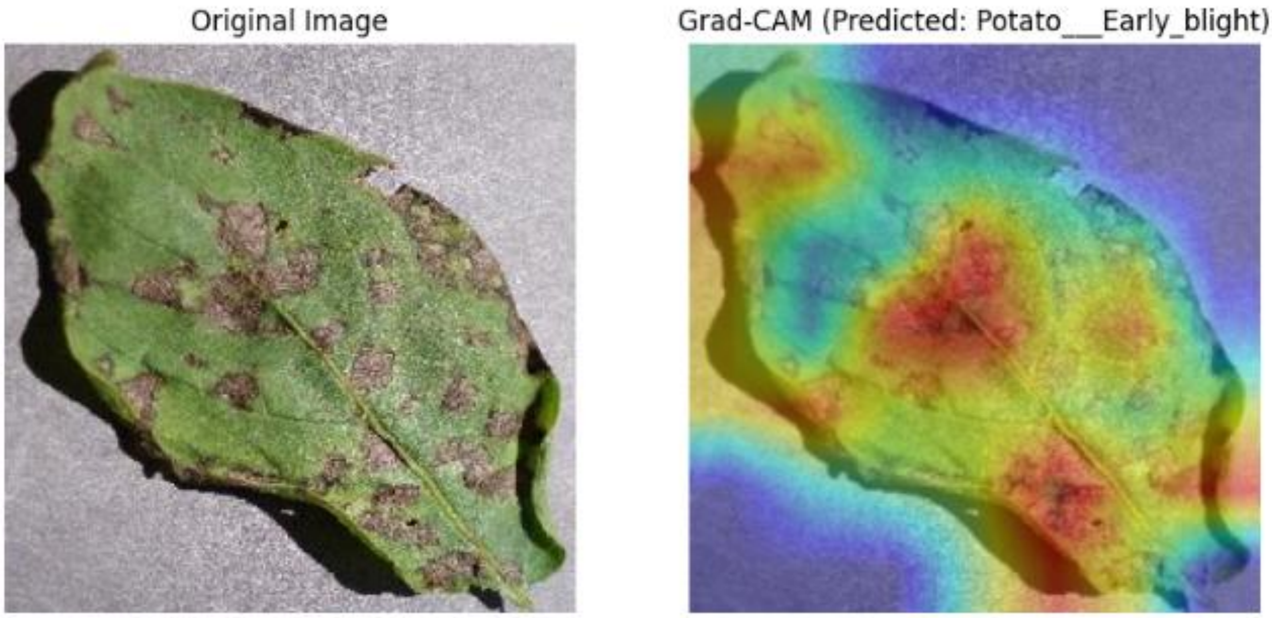

Figure 4.3.9– Grad-CAM Maps

E. Segmentation vs. Explainability Comparison

A. Observations

1. In most test cases, the high-activation regions in Grad-CAM and Saliency Maps matched closely with the segmented diseased areas of the leaf, confirming that the models were indeed focusing on biologically relevant symptoms such as spots, lesions, or blighted regions.
2. Some small mismatches were noticed along the boundaries of lesions or in cases

where the disease symptoms were faint and low in contrast. These discrepancies suggest that while the model captures the majority of disease regions correctly, there is still scope for improving fine-grained localization.

3. When combined, segmentation and explainability act as a dual-check system:
   I. Segmentation ensures that the system correctly identifies the precise affected region on the leaf.
   II. Explainability (via Grad-CAM or Saliency Maps) shows the reasoning process of the model, highlighting the areas that influenced its classification decision.

This complementary relationship makes the overall system more robust and reliable.

B. Figure Suggestions

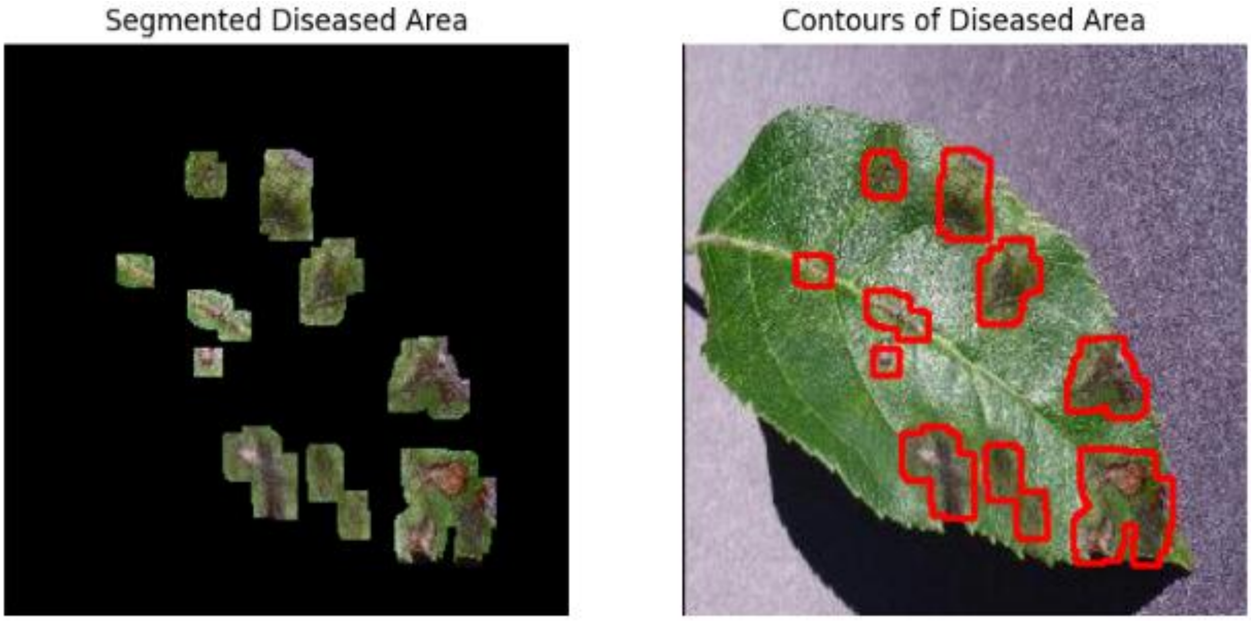

Figure 4.3.10 Leaf Segmentation

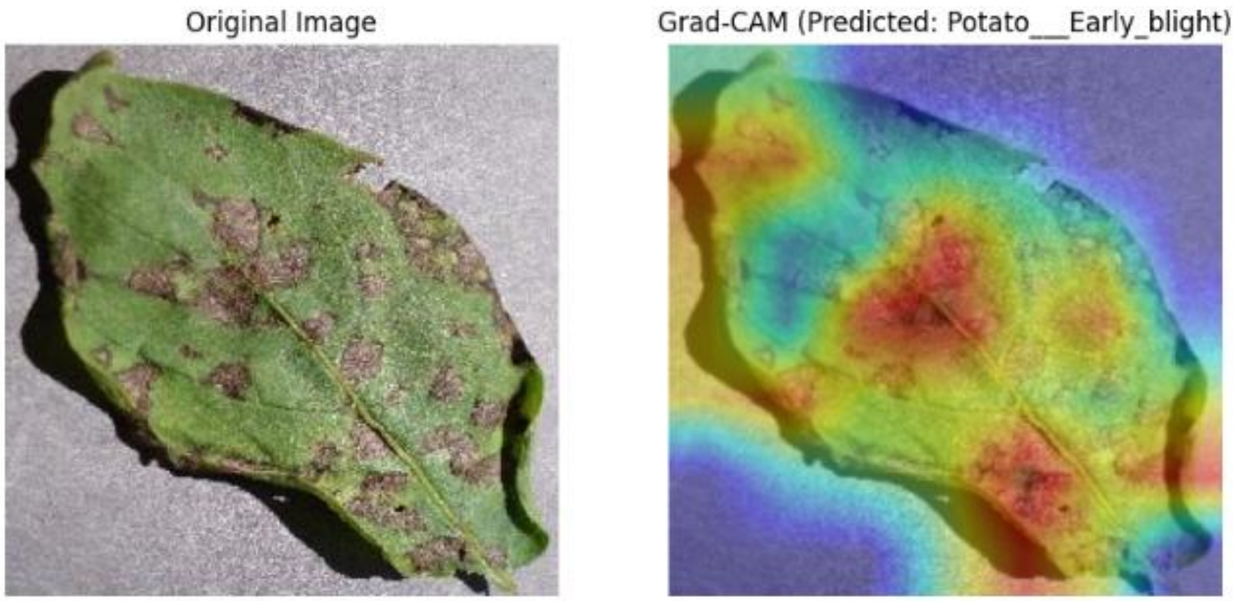

Figure 4.3.11 Grad-cam

C. Insights

This analogy shows that using segmentation-based visual evidence is really helpful in the interpretability of our model. That is, the model is not merely making black-box predictions; it is actually paying attention to the right and biologically significant areas of the leaf as it makes its decision. Segmentation offers a highly detailed pixel-scale localisation of diseased regions, thus giving a more accurate and reliable mapping of where the infection is in reality. In contrast to general heatmaps, segmentation allows drawing a very clear distinction between healthy and unhealthy parts of the leaf, thus allowing the domain experts to confirm whether the system is focusing on the right disease symptoms.Simultaneously, introducing the techniques of Explainable Artificial Intelligence (XAI) (Grad-CAM and saliency maps) enables us to look at the reasoning process of the neural network. These techniques expose what visual and internal characteristics the model was dependent on to come to its final decision, thus providing a frontier into the decision-making pipeline to be examined and verified. Combined to create a two-fold interpretability framework, segmentation and XAI can be used to ensure accuracy of localization of diseased areas and explainability of results in a conceptual manner, respectively.A combination of these complementary strategies results in a higher level of trust to the automated mechanism. They close the divide between predictions made by the raw AI and those of human experts by not only displaying what the model predicts but also explaining why it predicts this way. This change turns the model into a non-transparent, strictly statistical classification instrument into a transparent, reliable and helpful diagnostic resource.

4.4 Summary

In a high-performance computer environment based on GPUs, implementation of a proposed system of plant disease classification was done to address 38 plant disease classes across various crops. The training pipeline included transfer learning with ImageNet-pretrained weights, elaborate data augmentation, AdamW, early stopping and a dynamic learning rate schedule to obtain consistent convergence with no overfitting. It was also developed with a modular data loading code and visualization dashboard to enhance usability to enable researchers and agronomists to comprehend the findings with Grad-CAM and Saliency Maps. The best analyzed model was the ResNet-50 with an accuracy of 98.74 with F1-score of 0.9874 and AUROC of 0.9980, and the Hybrid ResNet + ViT came just behind, with 98.58 accuracy. EfficientNet-B0 offered a more accurate of 97.92 with a good trade-off against the edge devices compared to MobileNetV2 which was a bit lesser at 96.84 but highly efficient at real-time mobile applications. The explainability analyses have also reported that the models were always geared towards biologically significant disease locales such as lesions, fungal patches and chlorotic areas that proved their reliability. The system is distinct as compared to prior works because it utilizes a bigger and more heterogeneous dataset and incorporates hybrid CNN-Transformer structures that are more context sensitive. Overall, the results show that the system not only offers the state-of-the-art accuracy but also ensures the possibility of interpretation and practical application, the distinction between AI research and realistic farming application.

Chapter 5

Engineering Standards and Design Challenges

In this chapter, the author draws attention to the engineering standards that influenced the design and implementation of the proposed plant disease classification system and significant challenges encountered in the course of the project. By ensuring that the accepted software, hardware, and communication standards are followed, not only is quality ensured, but the project can be deployed practically where reliability, interoperability and user confidence is at stake.Furthermore, this chapter outlines the societal, ethical and environmental implications of the project. Because agriculture is directly related to food security, sustainability, and rural livelihoods, the effects of such an AI-based system extend well beyond classification accuracy. Lastly, project management, cost analysis, mapping to complex engineering problems and activities are presented to substantiate the engineering importance of the project.

5.1 Compliance with the Standards

Developing AI-driven systems requires aligning with international standards that ensure robustness, safety, and scalability. Without adherence to such standards, the system might face limitations in deployment or lose stakeholder trust.

5.1.1 Software Standards

1. IEEE 829 (Software Test Documentation Standard): Adopted to maintain a structured testing process for all models (ResNet-50, EfficientNet, MobileNet, ResNet+ViT). This ensured reproducibility and clear reporting of test cases.
2. ISO/IEC 9126 (Software Quality Standard): Helped in evaluating usability, maintainability, and reliability of the software pipeline.
3. ISO/IEC 25010 (System and Software Quality Model): Ensured functional suitability, performance efficiency, and portability across environments (cloud, local machine, mobile).

Alternatives and Trade-offs:

1. Ad-hoc testing could have been faster during prototyping but would compromise reliability.
2. By choosing IEEE and ISO standards, the system gains a professional structure suitable for real-world agricultural deployment.

5.1.2 Hardware Standards

1. IEEE 1012 (System and Software Verification and Validation): Ensured that the

GPU servers met reliability benchmarks during heavy training workloads.

2. NVIDIA CUDA Compliance: CUDA-supported GPUs allowed optimized training with parallel computing.
3. Server-class Standards (ASHRAE Thermal Guidelines): Guaranteed stable cooling and reliable operation during long training sessions.

Alternatives and Trade-offs:

1. CPU-only training would have been cheaper but extremely slow (weeks instead of hours).
2. Cloud-based GPU rental could scale easily, but long-term costs would exceed owning local hardware.
3. Our choice balanced speed, scalability, and cost-efficiency.

5.1.3 Communication Standards

1. RESTful APIs (Representational State Transfer): Chosen for model deployment, allowing farmers or agronomists to interact with the system via apps or web portals.
2. JSON Format (RFC 8259): Adopted for lightweight and human-readable data exchange.
3. MQTT Protocol (ISO/IEC 20922): Considered for IoT devices such as drones or sensors sending leaf images in real time.

Alternatives and Trade-offs:

1. XML was considered but discarded due to verbosity and higher bandwidth cost.
2. gRPC could improve performance but requires more specialized implementation.
3. REST + JSON provided the best trade-off between simplicity and scalability.

5.2 Impact on Society, Environment and Sustainability

The introduced system can make a major difference in the sphere of agriculture as well as society in terms of minimizing the losses of crops, enhancing the income of farmers, and ensuring the food security. In the environment, it advances sustainable agricultural systems by reducing over use of pesticides as well as keeping the soil and water healthy. On ethical grounds, it has transparency and accessibility based on explainable AI and a minimal deployment on mobile devices so that even smallholder farmers can be benefited. The system is a useful resource to both the communities and the environment as it in the long run encourages eco-friendly agriculture, equal access to technologies, and sustainable business models.

5.2.1 Impact on Life

Plant diseases are by far one of the largest contributors to crop losses in the world and in most cases farmers are likely to lose up to 20-30 percent of their harvest thanks to early or improper diagnosis. The proposed AI-powered system of diagnosing plant diseases can help reduce these losses significantly by offering timely and reliable diagnosis. This system is a low-cost decision support tool that can be accessed instantly with a smartphone and used by a large number of smallholder farmers in particular, as many of them are not able to hire agricultural specialists on a regular basis. It assists in enhancing the output of farmers as well as their income and stability of their household through timely intervention. At a more generalized level, more crop security can be converted into more food security and food security in terms of nutritional security to the population, especially where there is food insecurity.

5.2.2Impact on Society & Environment

The system is also very critical in terms of social and environmental implications. Traditionally, farmers use pesticides with a lot of force when they realize that something is wrong and this is in most instances causing unnecessary and overuse of chemicals. A detecting system relies on AI would make the use of pesticides more target-oriented and accurate and emit fewer chemicals to the ground, rivers, and soil. This assists in a well-balanced eco system and preservation of biodiversity. Furthermore, the agricultural supply chain is stabilized as well by the system since it protects the yield and mitigates unexpected losses. The direct beneficiary of stable crop supply is food security as the radical changes in food prices reduce and local markets can afford basic foodstuffs. The second social advantage is access equity in technology, since lightweight models like MobileNetV2 can be run on low-end smartphones, the system can be scaled to marginalized populations in farming, not just big agribusinesses, so that everyone can reap the rewards of technological progress.

5.2.2 Ethical Aspects

Ethics play a significant role in the use and creation of AI in fields like agriculture that are sensitive. The system was also supposed to be driven by transparency through explainable AI methods such as Grad-CAM and Saliency Maps. This kind of visualization can make sure that farmers and agronomists are able to make sure that predictions are drawn based on actual symptoms of the disease rather than other features that are not relevant. Accessibility is also accepted as the basis of the quality of the design: lightweight models ensure that even low-income farmers who have only simple mobile gadgets would benefit. The training dataset was also filtered to ensure that there were a variety of crops and disease conditions to reduce the potential biases, therefore, the model does not discriminate against certain crops or regions. Finally, one should emphasize the fact that the specified AI-based solution will not replace the role of professional pathologists or agronomists; instead, it will help them in their tasks. The predictions will be considered as an aid to decision-making that would enable the farmers to take a more confident leap but at the same time encourage expert validation when necessary.

5.2.3 Sustainability Plan

The considerations of the sustainability of the system have been noted in several levels. Energetically, models can be trained on GPUs in a computationally expensive but single-time task, whereas inference in the farmers-interacts-with stage is lightweight and efficient and can be deployed to mobile hardware as well as the environment. The ecological sustainability is also spread to the farming activity: the soil fertility is kept by removing the redundant use of pesticides and pollution of the water is reduced and the long-term biodiversity is supplied. On the development side, the project has a long-term vision: the project may include multi-modal data including soil quality, weather, and drone images in the future to create a full precision farming platform. Regarding the financial aspect, the system is planned to be designed in a sustainable business format, integrating freemium access to individual farmers with subscription or licensing to agriculture organization and non-governmental organizations. This will see the project not rely on short term grants or research funding but be able to grow and bring value to the agricultural community in the long run.

5.3 Project Management and Financial Analysis

A successful development, deployment, and clinical integration of an AI-based image report generation system cannot be achieved without proper project management. The project processes will include the phases of data collection and annotation, model development and training, the implementation of the LLM, system testing, clinical validation and implementation. The resources are managed properly in order to make the best use of people, equipment, and software. Financial analysis is necessary to approximate the cost of hardware, software, human resource, clinical validation, deployment, and quality assurance. These costs are known and can help make a plan, budget, and take into account the possible return on investment in the form of efficiency gain and a better clinical outcome. The cost and financial analysis are estimated in Table

Table 5.1: Estimated Cost and Financial Analysis

Components	Estimated Cost (BDT)
Internet and Cloud Subscription e.g. AWS, Google Colab Pro	4,000-5,000
Software Licenses (None, used Open Source)	0
Data Collection	0
Miscellaneous (Printing, Backup Storage)	800-1,000
Total Estimated Cost	4,800-6,000 BDT

5.4 Complex Engineering Problem

Real-world engineering projects are rarely straightforward, and this plant disease classification system is no exception. The problem qualifies as a complex engineering problem because it involves multidisciplinary knowledge, multiple stakeholders, competing requirements, and far-reaching impacts on both society and the environment. Building such a system is not only about achieving high classification accuracy in a laboratory setting; it requires designing a solution that is robust, explainable, and usable in diverse agricultural contexts.

5.4.1 Complex Problem Solving

This project falls under complex engineering problems because it demands integration of multiple knowledge domains (AI, agriculture, software engineering), requires optimization under constraints, and involves multiple stakeholders.

Table 5.2: Mapping with Complex Engineering Problem.

EP1 Dept of Knowledge	EP2 Range Of Conflicting Requirements	EP3 Depth of Analysis	EP4 Familiarity of Issues	EP5 Extent of Applicable Codes	EP6 Extent Of Stake-holder Involvement	EP7 Interdependence
√	√	√	√	√		√

The propose

d plant disease classification project clearly aligns with several complex engineering problem attributes. First, in terms of depth of knowledge (EP1), the system requires specialist expertise in convolutional neural networks (CNNs), Vision Transformers (ViTs), and domain-specific agricultural pathology. It is not just a computer science problem but an intersection of engineering and life sciences. Second, the project has conflicting requirements (EP2) since it must achieve very high accuracy while also remaining lightweight enough for deployment on mobile devices and low-resource environments. This creates a trade-off between model complexity and practical usability. Third, the depth of analysis (EP3) is significant because the system must handle a challenging 38-class classification problem using high-resolution images that often contain subtle and overlapping disease features. Fourth, regarding the familiarity of issues (EP4), plant diseases frequently appear visually similar, making the classification problem far from straightforward and requiring advanced interpretability methods. Fifth, in terms of the extent of codes and standards (EP5), the system follows established software, hardware, and communication standards to ensure reliability, scalability, and professional compliance. Sixth, stakeholder involvement (EP6) is broad, as the project is designed to benefit not only farmers but also agronomists, NGOs, and policymakers who may use the system for large-scale agricultural planning. Finally, the interdependence (EP7) is high, as the success of the system impacts agriculture, food security, economics, and environmental sustainability, demonstrating how deeply interconnected the problem is with society.

Mapping with Knowledge Profile:

This section is designed to map the overall problem and EP1 ***(multiple between K3, K4, K5, K6, K8 for attaining EP1)*** to the Knowledge Profile.

Table 5.3: Mapping with knowledge Profile.

K1 Natural Science	K2 Mathematics	K3 Engineering Fundamentals	K4 Specialist Knowledge	K5 Engineering Design	K6 Engineering Practice	K7 Comprehension	K8 Research Literature
✓	✓	✓	✓	✓	✓	✓	✓

The knowledge profile mapping further highlights the multidisciplinary foundation of the project. Engineering fundamentals (K3) are at the core, as the system depends on a solid understanding of deep learning, optimization, and computer vision. Beyond fundamentals, the project leverages specialist knowledge (K4) in advanced architectures like CNNs and Transformers, along with explainability techniques such as Grad-CAM and Saliency Maps. The project also emphasizes engineering design (K5) by creating a scalable and interpretable classification pipeline that can move from a laboratory environment to field deployment. From a practical standpoint, engineering practice (K6) is demonstrated by applying the trained models in real-world agricultural conditions, ensuring that predictions remain reliable outside controlled datasets. Finally, the project integrates research literature (K8), comparing its results against established studies to benchmark performance and situate its contributions within the larger scientific and engineering community.

5.4.2 Engineering Activities

In this section, provide a mapping with engineering activities. For each mapping add subsections to put rationale (Use Table 5.3).

Mapping with Complex Engineering Activities

This section is designed to map the overall problem and EA's *(multiple).*

Table 5.4: Mapping with Complex Engineering Activities.

EA1 Range of resources	EA2 Level of Interaction	EA3 Innovation	EA4 Consequences for society and environment	EA5 Familiarity
✓	✓	✓	✓	✓

In terms of engineering activities, the project demonstrates alignment with multiple dimensions. For EA1 – range of resources, it utilizes diverse tools including GPU servers for training, open-source libraries like PyTorch and OpenCV, and mobile deployment tools for real-world usage. For EA2 – level of interaction, the system requires collaboration among researchers, farmers, agronomists, and policymakers to ensure that the solution is practical, trustworthy, and widely adopted. With EA3 – innovation, the project introduces a novel hybrid approach that combines CNN-based feature extraction with Transformer-based global attention, advancing the state of plant disease classification. For EA4 – consequences for society and environment, the system has clear societal benefits by improving food security, reducing pesticide use, and promoting sustainable agriculture. Finally, under EA5 – familiarity, while the problem of plant diseases is universally recognized, the application of advanced AI architectures in this domain brings a fresh, innovative perspective to an otherwise familiar global challenge.

5.5 Summary

The proposed plant disease classification system addresses a complex engineering problem because it demands advanced knowledge of deep learning models, plant pathology, and software engineering. It balances conflicting requirements by achieving high accuracy while staying lightweight for mobile deployment, and it involves multiple stakeholders including farmers, agronomists, NGOs, and policymakers. From a knowledge perspective, the project builds on engineering fundamentals, specialist expertise in CNNs and Transformers, strong design practices for scalability, hands-on engineering application in agriculture, and alignment with research literature for benchmarking. In terms of engineering activities, the system demonstrates the use of diverse resources (GPU servers, open-source tools, mobile deployment), requires interaction across disciplines and communities, introduces innovation through hybrid CNN–Transformer models, and has clear consequences for society and the environment by promoting sustainable agriculture. Overall, the mappings confirm that this project is multidisciplinary, innovative, and socially impactful, reflecting the true nature of complex engineering practice.

Chapter 6
Conclusion

6.1 Summary

This study presented a deep learning-based framework for multi-class plant disease classification across 38 classes and multiple crops. Using transfer learning and advanced augmentation strategies, four models--ResNet-50, EfficientNet-B0, MobileNetV2, and a hybrid ResNet + Vision Transformer--were trained and evaluated. Among these, ResNet-50 achieved the highest accuracy of 98.74% with strong generalization, while the hybrid ResNet + ViT demonstrated competitive performance by effectively capturing both local and global leaf features. Lightweight models such as EfficientNet-B0 and MobileNetV2 offered efficiency advantages for real-time or edge-based applications. To ensure transparency, explainability methods including Grad-CAM and Saliency Maps were employed, which consistently highlighted disease-relevant leaf regions. This combination of high accuracy and interpretability underscores the potential of the system for deployment in real-world agricultural settings.

6.2 Limitation

Despite the fact that the proposed system performed at the state of the art level, a number of limitations were observed. The first limitation was that data were restricted to RGB images, a method that might fail to provide a complete depiction of subtle disease features that might be more effectively extracted with hyperspectral or multi-modal imaging. Second, the models generalized across 38 classes; however, the rare or emerging diseases that were not part of the dataset are beyond the reach of the model. Also, lighting, occlusion or background clutter can be used to influence performance in the real world field conditions. The explainability procedures, despite their success, at times yielded diffuse or sparse attention maps especially on weak disease symptoms. Lastly, it is difficult to utilize models with large capacity like the ResNet-50 on low-resource machines due to the computation and memory demands.

6.3 Future Work

The research has to be conducted in the future in order to eliminate these constraints by incorporating multi-modal data sources, including hyperspectral image, soil conditions, and environmental parameters to enhance robustness in different conditions. Future work on improving lightweight models such as MobileNetV2 and EfficientNet-B0 to support mobile- or drone-based deployment will make real-time disease detection a possibility in the fields. More sophisticated explainability methods, including attention-guided training or more enhanced segmentation networks, can be added to provide model focus more closely to diseased regions. Also, a few-shot and domain adaptation might be considered in order to increase the applicability of the system to the uncommon and emergent plant diseases. Lastly, creating more temporal frameworks to understand disease dynamics and switching towards semi-supervised or self-supervised learning will help decrease the need to utilize large-scale manual labelling and make the framework more viable in a global agricultural setting.

References

[1] A. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep Learning for Tomato Diseases: Classification and Symptoms Visualization," *Applied Artificial Intelligence*, vol. 31, no. 4, pp. 299–315, 2017.

[2] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.

[3] S. Ferentinos, "Deep Learning Models for Plant Disease Detection and Diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.

[4] C. W. Too, L. Yujian, S. Njuki, and L. Yingchun, "A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.

[5] A. Berdin and M. Caduyac, "Tomato Leaf Disease Classification Using Residual Networks (ResNet)," *International Journal of Advanced Computer Science and Applications (IJACSA)*, vol. 10, no. 6, pp. 1–7, 2019.

[6] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in *Proc. International Conference on Machine Learning (ICML)*, pp. 6105–6114, 2019.

[7] A. Howard, M. Sandler, G. Chu, et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv preprint arXiv:1704.04861, 2017.

[8] A. Dosovitskiy, L. Beyer, A. Kolesnikov, et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," arXiv preprint arXiv:2010.11929, 2021.

[9] Z. Liu, Y. Lin, Y. Cao, et al., "Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 10012–10022, 2021.

[10] R. R. Selvaraju, M. Cogswell, A. Das, et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 618–626, 2017.

213-15-4622

ORIGINALITY REPORT

16%	11%	8%	11%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Daffodil International University Student Paper	7%
2	www.ijmrset.com Internet Source	1%
3	Submitted to United International University Student Paper	<1%
4	www.nature.com Internet Source	<1%
5	Sreedhar, Vikram Rajapura. "Development of Artificial Intelligence Algorithms in Cardiovascular Research: Case Studies in Atrial Fibrillation And Myocardial Infarction", Liverpool John Moores University (United Kingdom) Publication	<1%
6	Submitted to University of Glasgow Student Paper	<1%
7	Madduri Venkateswarlu, Venkata Rami Reddy Chirra. "CNN-ViT: A multi-feature learning based approach for driver drowsiness detection", Array, 2025 Publication	<1%
8	dspace.daffodilvarsity.edu.bd:8080 Internet Source	<1%
9	Pushpa Choudhary, Sambit Satpathy, Arvind Dagur, Dhirendra Kumar Shukla. "Recent	<1%

Trends in Intelligent Computing and Communication", CRC Press, 2025
Publication

10	Submitted to UCL Student Paper	<1 %
11	S.P. Jani, M. Adam Khan. "Applications of AI in Smart Technologies and Manufacturing", CRC Press, 2025 Publication	<1 %
12	arxiv.org Internet Source	<1 %
13	www.mdpi.com Internet Source	<1 %
14	Submitted to Liverpool John Moores University Student Paper	<1 %
15	cje.ejournal.org.cn Internet Source	<1 %
16	ijsrem.com Internet Source	<1 %
17	Aitor Gutierrez, Ander Ansuategi, Loreto Susperregi, Carlos Tubío, Ivan Rankić, Libor Lenža. "A Benchmarking of Learning Strategies for Pest Detection and Identification on Tomato Plants for Autonomous Scouting Robots Using Internal Databases", Journal of Sensors, 2019 Publication	<1 %
18	journals.plos.org Internet Source	<1 %
19	assets-eu.researchsquare.com Internet Source	<1 %

20	Shankar Babu, Mahesh Babu Kota. "Synergies in Smart and Virtual Systems using computational intelligence", CRC Press, 2025 Publication	<1 %
21	www.springerprofessional.de Internet Source	<1 %
22	Submitted to University of Nottingham Student Paper	<1 %
23	inass.org Internet Source	<1 %
24	"Medical Image Understanding and Analysis", Springer Science and Business Media LLC, 2022 Publication	<1 %
25	Hoque, Md. Sami Ul. "Comparative Analysis of CNN, Vision Transformers, and Hybrid Models for Skin Lesion Classification Using the HAM10000 Dataset", Southern Illinois University at Edwardsville Publication	<1 %
26	espace.curtin.edu.au Internet Source	<1 %
27	Edna Chebet Too, Li Yujian, Sam Njuki, Liu Yingchun. "A comparative study of fine-tuning deep learning models for plant disease identification", Computers and Electronics in Agriculture, 2018 Publication	<1 %
28	Submitted to Higher Education Commission Pakistan Student Paper	<1 %
29	www.slideshare.net Internet Source	<1 %