EFFICIENT AND INTERPRETABLE DEEP LEARNING FOR MANGO LEAF DISEASE DETECTION USING KNOWLEDGE DISTILLATION AND XAI

BY

Nasir Uddin
ID: 242-25-040

This Report Presented in Partial Fulfillment of the Requirements for The Degree of Masters of Science in Computer Science and Engineering

Supervised By

Dr. Arif Mahmud
Associate Professor & Associate Head
Department of CSE
Daffodil International University

Co-Supervised By

Dr. Abdus Sattar
Associate Professor & Director, MSc. in CSE
Department of CSE
Daffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY

DHAKA, BANGLADESH

SEPTEMBER 2025

APPROVAL

This Project/Thesis titled **"Efficient and Interpretable Deep Learning for Mango Leaf Disease Detection Using Knowledge Distillation and XAI"**, submitted by **Nasir Uddin**, ID No: **242-25-040** 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 M.Sc. in Computer Science and Engineering and approved as to its style and contents. The presentation was held on 13-09-2025.

BOARD OF EXAMINERS

Chairman

Dr. Sheak Rashed Haider Noori
Professor and Head
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Internal Examiner

Dr. Md. Zahid Hasan
Associate Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Internal Examiner

Dr. Naznin Sultana
Associate Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

External Examiner

Mr. Nazibur Rahman
Head of IT Infrastructure
Networld Bangladesh PLC

DECLARATION

I hereby declare that this research has been done by me under the supervision of **Dr Arif Mahmud,** Associate Professor & Associate Head, Department of CSE, Daffodil International University. I also declare that neither this project nor any part of this project has been submitted elsewhere for award of any degree or diploma.

Supervised by:

Dr. Arif Mahmud
Associate Professor & Associate Head
Department of CSE
Daffodil International University

Co-Supervised by:

Dr. Abdus Sattar
Associate Professor & Director, MSc. in CSE
Department of CSE
Daffodil International University

Submitted by:

Nasir Uddin

Nasir Uddin
ID: 242-25-040
Department of CSE
Daffodil International University

ACKNOWLEDGEMENT

First, I express my heartfelt thanks and gratitude to Almighty Allah for His divine blessing, which makes it possible to complete the final year project/internship successfully.

I would like to express my deepest gratitude and profound indebtedness to **Dr. Arif Mahmud**, **Associate Professor and Associate Head**, Department of Computer Science and Engineering, Daffodil International University, Dhaka, Bangladesh, for his invaluable supervision and guidance throughout the course of this project. His profound knowledge, insightful expertise in the field of Deep Learning, and genuine enthusiasm for research have been a constant source of inspiration. I am especially grateful for his patience, scholarly advice, constructive criticism, and continual encouragement. His meticulous supervision, valuable suggestions, and kind support in reviewing and refining this work at every stage have been instrumental in the successful completion of this project.

I would like to express my heartfelt gratitude to **Dr. Sheak Rashed Haider Noori, Head of the Department of CSE**, for his kind assistance in completing our project, as well as to the other faculty members and staff of the CSE department at Daffodil International University.

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

Finally, I must acknowledge with due respect the constant support and patience of my parents.

ABSTRACT

Mango leaf diseases are a major cause of crop loss and can only be minimized through accurate and on-time disease identification. This paper proposes an effective and explainable deep learning model of mango leaf disease that makes use of logit-based Knowledge Distillation alongside temperature scaling and explainable artificial intelligence (XAI) strategies. The teacher network, DeiT-S, was pre-trained using a large collection of real-life and augmented images across eight disease categories with the validation accuracy of 99.14%. A small student model, MobileNetV3-small was then trained with a combined loss, where KL-divergence between the softened outputs of the teacher (soft targets, temperature T=4.0) and the ground-truth labels (hard targets) are combined with a cross-entropy loss. This method allowed the student to transfer the knowledge acquired by the teacher with high efficiency, so the training accuracy was 99.60%, and the validation was 99.46%. To allow better interpretability, Grad-CAM was used to visualize decision-making areas on leaf images. Lastly, an intuitive web-based application was created which could be used to upload leaf images and obtain predictions about the disease along with visual descriptions. The proposed framework is highly accurate and computationally efficient and transparent, which suggests its applicability in the practical implementation in agricultural disease management.

Keywords: Deep Learning, Knowledge Distillation (KD), Lightweight Neural Networks, Leaf Disease Detection, Temperature Scaling, Explainable AI (XAI).

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1
INTRODUCTION

1.1 Introduction

Mango (Mangifera indica): Mango is a widely produced tropical fruit in the world with a production of over 78 million metric tons in a year. Mangoes play an important role in global food security since they provide key vitamins and minerals as a very important source of food. Moreover, mango is a significant economic base in most countries, particularly in South Asian where it sustains the lives of millions of farmers [16,8]. However, of major concern in controlling leaf diseases adversely affecting mango crops worldwide is poor yield and loss of fruit quality. Diseases like Anthracnose, Powdery Mildew and Bacterial Canker are among those which are the most harmful and their financial losses affected globally are about USD 5 billion every year [16]. In places such as India, where mangoes play an important role in the agricultural industry, losses as a result of such diseases reach up to USD 100 million every year [1]. Therefore, disease control is yet another important element in maintaining the production of mangos and food security.

Conventional approaches of diagnosing mango leaf diseases are mainly through manual examinations by agronomy experts. The methods are however burdensome, time-consuming and subject to subjective errors especially in early-stage disease detection [6,7]. The precision of these manual observances is up to 65 percent [7], and considering the size of the mango plantations, limited numbers of trees are evaluated per day by the experts causing delayed detection of diseases. This inability is also added by the lack of skilled agriculturalists who reside in the countryside where exposure to skill is scarce [1]. Analysis of diseases may therefore not be complete and yield can be reduced by up to 15 percent during production [7]. Finding an automated, precise, and scalable automated disease detection platform has therefore become a matter of high priority.

The newest implementations of Deep Learning (DL) have potential applications in plant disease classification, as it presents automated and precise methods of classifying large quantities of data [6,13]. The most effective of them, and the one that has seen the widest application to date, is Convolutional Neural Networks (CNNs). CNN-based methods usually need a large amount of labeled data to be trained on but these can be hard to find in agricultural scenarios. Moreover, most of these models would require substantial computational power in order to be useful in real-time applications in computationally limited settings [20,2]. To address these difficulties, the concept of transfer learning has become an eminent alternative. Using pre-trained models, transfer learning decreases the amount of labeled data required and accelerates training, which makes it more viable in case of limited resources required to detect mango leaf disease [19] .

In spite of the effectiveness of CNNs and transfer learning, two critical issues still exist, namely, efficiency of these models and their interpretability. Even the modern and precise models are too large and complex to be implemented on low-resource hardware that are prevalent in the rural environment [20]. Also, deep learning models are very opaque and hence cannot be trusted much in agricultural decision-making which requires clarity and openness.

1.2 Motivation

As mango is an important commercial fruit particularly in places like South Asia. However, mango cultivation is seriously affected by different types of leaf diseases like Anthracnose, Powdery Mildew, and the Bacterial Canker. These diseases cause severe loss of production and characteristics of mango fruits which poses threats to farmers. The existing ways of detecting diseases are slow since thorough investigation is through the eyes, and detecting through human eyes is slow and vulnerable to errors. This evokes a great need for an automated disease detection system that can be able to give timely and accurate diagnosis, and help in reducing losses incurred on crops.

The proposed research solves these issues by combining the tools of Knowledge Distillation (KD) and Explainable AI (XAI). Knowledge Distillation provides the means of distilling knowledge in the teacher model to the student model where the knowledge can be transferred in a much smaller and simpler but effective student model, and the system can be deployed in low-resource settings. AI algorithms such as Grad-CAM and LIME can be used to increase transparency by making farmers able to visualize the particular parts of the leaf image that was influential in the disease classification. Such interpretability increases user confidence as well as provides farmers with actionable insights to gain insight into the symptoms and action they should take.

1.3 Objectives

The key goal of this study would be building a secure and scalable deep learning model in plant disease detection by resolving the issues of accuracy, scalability, interpretability, and data privacy of the existing models. The overall aims are:

- Design a lightweight deep learning model that utilizes the ViT framework and the mechanisms of self-attention in order to identify local and global details of mango leaf diseases.
- Integrate Knowledge Distillation (KD) with a teacher model based on DeiT (Data-efficient Image Transformer) and a student model with MobileNet_V3 that allows overcoming the limits of traditional deep learning models in terms of computational efficiency and model size to achieve high accuracy levels in the classification of mango leaf diseases.
- Introduce explainability AI (XAI) systems, e.g., Grad-CAM to visualize the model predictions, and make them more transparent and increase trust in the end users .
- Validate the model in various datasets in order to check the generalizability and flexibility to a variety of agroecological areas which are relevant to ensure that the problem of domain shifts and class imbalance has been addressed.
- Used the optimized model through a web interface to provide real-time disease detection and associated probability scores and visual explanations.

This paper will propose a solution that not only aims at augmenting the classification accuracy but also at making the solution more practical towards the taking of mango diseases in resource-limited environments. The end result is a reliable, real-time and interpretatively available product that can be used by the mango farm-level community and agricultural experts in handling disease outbreaks more easily and, hence, making the production of mangoes sustainable globally.

1.4 Methodology

The suggested methodology of mango leaf disease identification starts with preprocessing of the input image into resized, normalized, tensors to standardize all the leaves. Generalization and overfitting are addressed through a variety of data augmentation techniques, which include noise addition, rotation, flipping, change in brightness, and scaling. After the preprocessing, the dataset will be divided into training (80%), validation (10%), and testing (10%) to guarantee a fair and stable evaluation.

A framework inspired by a teacher-student architecture through knowledge distillation is then used to attain high accuracy and computational efficiency. We use the DeiT-Small (DeiT-S) model as the teacher and MobileNetV3-Small as the student since it is specifically optimized to be lightweight in terms of the inference process. The student model is trained using a combination of cross-entropy and KullbackLeibler divergence as the hybrid loss function to learn with reference to the ground-truth labels as well as the soft predictions provided by the teacher. This method translates knowledge of the larger model to the smaller one, cutting the size of the model and the time taken doing inference without compromising accuracy. Here, figure 1 represents the overview of the proposed methodology.

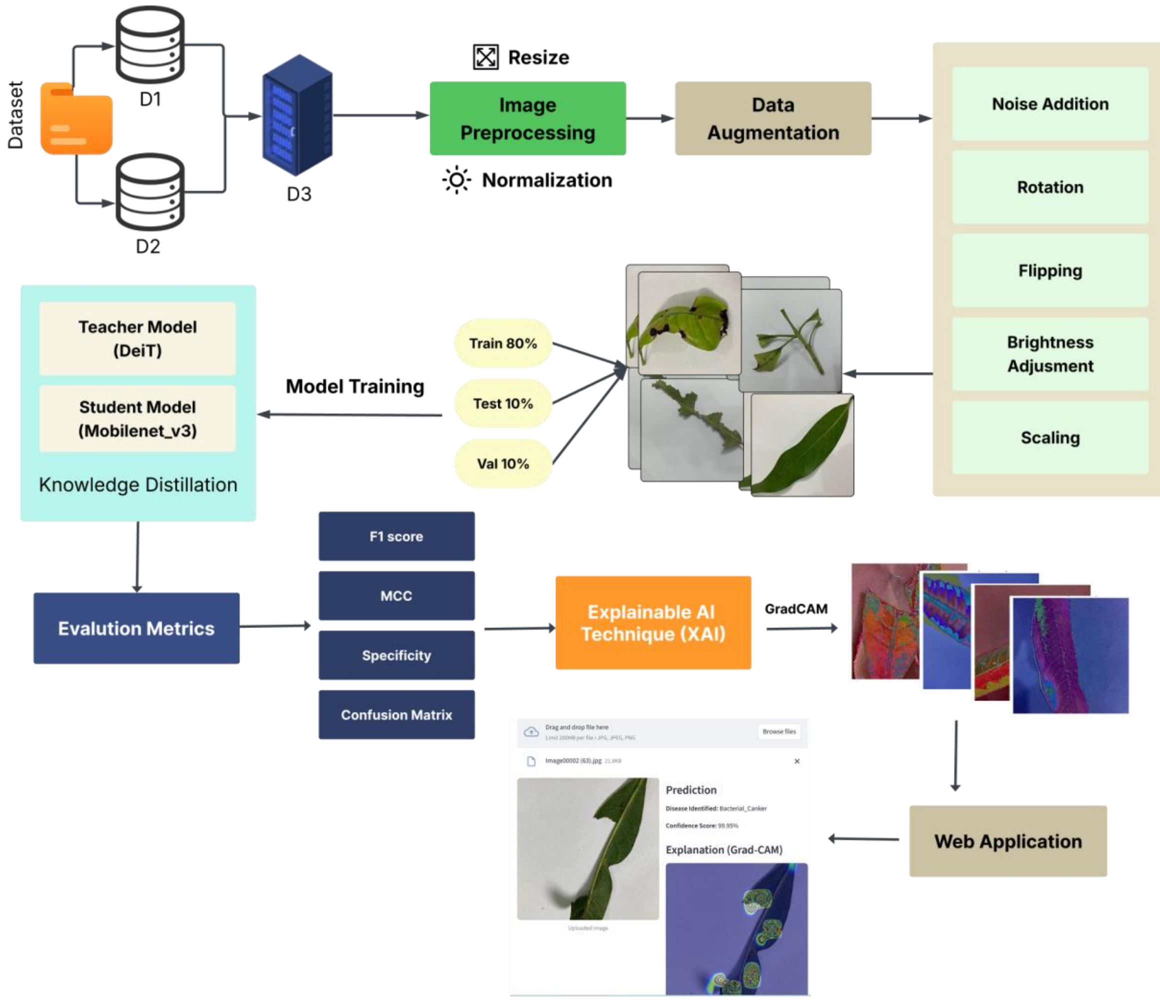

Figure 1. Overview of the proposed methodology.

To certify performance, the models are tested on various measures other than the basic accuracy, such as F1-score, specificity, Matthews Correlation Coefficient (MCC), and confusions matrices. The metrics offer more insight into their performance of classification in various disease classes. To further increase the interpretability, an explainable AI (XAI) technique is incorporated by means of Gradient-weighted Class Activation Mapping (Grad-CAM), which shows the decisive areas of sick leaves that contributed to model predictions.

Lastly, a user-friendly web application is deployed, with the optimized student model used in it. The app allows farmers and other stakeholders in mangoes to take pictures of mango leaves and get on-demand predictions as well as Grad-CAM heatmaps detailing how the prediction was made.

This whole-process pipeline has strong predictive precision, computational cost-effectiveness, and interpretability, which makes it feasible in realistic farming disease control processes.

1.5 Research Question

The main research question that informed this research is:

What is the effectiveness of combining Knowledge Distillation (KD) and Explainable AI (XAI) to enhance the efficiency, accuracy, and comprehensibility of deep learning models and mango leaf disease detection?

The following aspects are going to be explored by this question:

- Efficiency and Precision: How can knowledge distillation be used to maximise the computational efficiency of deep learning models to detect mango leaf disease that meet or improve their accuracy even in constrained agricultural settings of low resources?
- Interpretability: How would a model based on explainable AI (XAI) methods, e.g., Grad-CAM, increase the transparency and trustworthiness of disease detection models, such that farmers and agricultural experts understand and can implement the predictions made?

By solving these elements, the proposed research will offer a solution to an early and accurate mango leaf disease condition with sufficient scaling and trustworthiness to be used by farmers to make sound decisions and implement AI-based technologies to manage crop health.

1.6 Expected Outcome

The anticipated results of this proposal are the following ones:

- More Efficient Deep Learning Models: It is anticipated that the research will be able to develop a deep learning model on mango leaf disease detection that will not compromise the levels of accuracy, yet the model will be computationally efficient. Through the Knowledge Distillation (KD), the model would expect to act as an equally complicated and gigantic model, yet to have a substantially smaller size as well as constrained essentiality. This will make it possible to provide real time detection of diseases in resource limited devices like smart phones making it affordable to farmers in resource-constrained settings.

- XAI to Promote Enhanced Interpretability: The second anticipated result is the creation of an interpretable model using the techniques of Explainable AI (XAI) grad-CAM, and LIME. This will enable transparency in the decision making process of the model so that the farmers can understand what components of the leaf image have had effect in the detection of the particular disease. Such interpretability will engender trust in AI systems and ease its application in real-life within farmlands.
- Practical Implementation in the Real Agricultural Life: The desired ultimate result is that the model will be judged effective and subsequently implemented in actual farming set-ups. They give the system an ability to identify many types of mango leaf diseases, and the system should be very accurate, meaning that disease control will improve, and the practice of sustainable farming will be beneficial.

1.7 Report Layout

Chapter 1 has the introductory part, research objectives, and research questions of the study. Chapter 2 will have a literature review explained in detail. Chapter 3 explains the methodology, what is going to be the framework and processes used in the research. Chapter 4 presents and discusses experimental outcomes. Chapter 5 discusses the sustainability plan, the implications of the society, and the environmental implications of the work along with the ethical reasoning on the same. Lastly, Chapter 6 brings the study to an end, and points to possible future research avenues.

CHAPTER 2
BACKGROUND

2.1 Introduction

This chapter gives a detailed overview of the literature already available that uses deep learning models to detect diseases in plants, and specifically it addresses the use of deep learning models in mango leaf disease classification. The review reflects the history of development of methodologies, starting with standard CNN-like approaches and moving to new trends in lightweight structures, hybrid learning, and explainable AI. A substantial number of deep learning methods have been tested in the previous papers and found to have a promising prospect on image-based diagnosis of diseases in crops like mango, tomato etc. CNNs have been particularly striking since they are able to extract spatial structures of leaf images. Notwithstanding these achievements, important challenges still exist, such as inability to be computationally efficient, explainable, and cover a wide range of disease classes, which precludes implementation in practical agricultural applications. Moreover, attainment of high classification accuracy is the focus of many studies, without giving much thought to the practicality of execution in the mobile context and in real-time implementation or accessibility among farmers. This chapter evaluates these contributions in a critical manner with a view of identifying their strengths and weaknesses and the research gaps that serve as motivation of the proposed system. In this proposed method, CNNs are used to achieve high accuracy, efficient hybrid learning algorithms have been employed to define robustness, and explainable AI has been used to ensure transparency.

2.2 Literature Review

The use of deep learning, especially convolutional neural networks (CNNs) in mango leaf disease classification and detection has been gaining increased awareness in recent years. This trend is mainly fuelled by the need of high-resolution agriculture and early detection of diseases to minimize crop loss and improve crop output forecasts. Research specialized solutions currently available rely on a variety of methods, which might be divided into ensemble-based and hybrid systems, custom architectures, transfer learning methods, explainable AI (XAI) and data augmentation.

Ensemble and hybrid algorithms are broadly used to enhance the ability of the results of a series of models that would augment their classification accuracy. Such as, an ensemble-based CNN was suggested by Bezabh et al. [1] who combined the lightweight CNN and hybrid CNN-SVM structures, yielding an accuracy of 99.87% on their test data. Although this figure improved, their work indicated that real-time mobile deployment was not backed and that little study had been done towards the hybrid combinations (e.g., decision trees or SVMs). Similarly, custom deep learning architecture has also exhibited an impressive potential. Mahmud et al. [2] designed a custom framework, called DenseNet78, serving as a specialized variant of DenseNet, with accuracy of 99.44 per cent, which exceeded those of traditional pre-trained models. Similarly, Pratap and Kumar [6] merged YOLOv8 with the African Buffalo Optimization (ABO), producing an accuracy of 98.69%. These examples of personalized strategies show how strategic models can be used to deal with the issues of crop disease detection.

Transfer learning has also been an overarching theme with researchers modifying existing pre-trained models such as VGG16, ResNet, EfficientNet to the mango leaf disease detection problem. Efforts by fine-tuning pre-trained networks have been demonstrated to reach an aggregate of 99.8% accuracy [3]. Although it is so powerful, this method has two limitations, and the first one is its insensibility to differentiate visually similar disease classes, whereas another limiting factor is its need of large annotated datasets. In a bid to work out concerns over interpretability, other studies have resorted to using explainable AI. MangoLeafXNet [19] employs Grad-CAM, Saliency Maps, and LIME, to provide visual explainability on the classification decisions. This increases the trust and transparency in the deep learning system which will make them more accessible to end-users, e.g. farmers, who might not be tech-savvy.

Preprocessing and data augmentation is another area of very high interest. Techniques of augmentation have also been exploited to enhance the diversity of datasets and the robustness of models especially in situations of small training data. Areal et al. [16], in their work, contrast-enhanced after affine transformations, and reached the 96.80% accuracy mark. However, overfitting and lack of variance in the real world are still the major threats to generalization.

Generally, the involved researches show high accuracies that are in fact above 99 percent in classification matters [2], [5]. The increased application of hybrid systems, ad hoc architectures and explainable frameworks like MangoLeafXNet [19] indicates the further development of the field to more precise and explainable systems. There is a simultaneous emphasis on real-time implementation and mobile usability so it is becoming more and more important that any real-world agricultural implementation of such content is practical. As one such example, Chowdhury and Ahmed [20] were able to show that Convention Transformer models could be augmented in runtime to be capable of operating even on low-end devices, closing the gap between research-only performance and applicability in the field.

Table 1: Comparative analysis of existing methodologies in plant disease classification across different domains.

Ref	**Model**	**Dataset**	**Contribution**	**Challenges**
Bezabh et al. [2024]	Ensemble CNN (Hybrid Models, GoogLeNet, VGG16)	Dataset acquired from the Merawi Agriculture Research Centre and the surrounding areas of Bahir Dar city.	Achieved 99.87% accuracy; integrated K-Means clustering and CNNs for disease detection.	High compute cost, limited dataset diversity, reliance on fully connected classifiers.
Mahmud et al. [2024]	DenseNet78 (lightweight DenseNet variant)	MangoleafBD	Achieved 99.47% accuracy for healthy leaves, 99.44% for diseased; focuses on disease segmentation.	Image segmentation issues, large annotated datasets required, limited disease scope.
Ünal & Türkoğlu [2025]	Pre-trained CNNs (Darknet19, Xception, etc.) + ML Classifiers	MangoleafBD	Achieved 99.8% accuracy with SVM; used multiple pre-trained models for deep feature extraction.	Limited datasets, difficulty in distinguishing visually similar diseases, sensitivity to image quality.

Preeya & Raja [2024]	Gated Recurrent Network (GRU) with Grey Wolf Optimization	PlanVillage	Achieved 99.87% accuracy for mango leaf disease detection using optimized GRU.	Narrow focus, potential overfitting, real-time deployment challenges.
Ramadan et al. [2023]	Vision Transformers (ViT), Big Transfer, CNNs (VGG16, ResNet)	MangoleafBD	ViT and Big Transfer models achieved 100% accuracy; CNNs varied from 92–99.75%.	Overfitting risks, scalability issues, and lack of real-world validation.
Pratap & Kumar [2024]	CNNs (VGG16, MobileNet, YOLOv8, EfficientNet)	MangoleafBD	YOLOv8 with African Buffalo Optimization achieved 98.69% accuracy.	Narrow disease focus, reliance on high-quality datasets, no real-time deployment.
Thaseentaj & Ilango [2024]	Transfer-learning ResNet50 with Grad-CAM	South-Indian mango leaf dataset	Achieved 92.8% accuracy; used Grad-CAM for interpretability.	Limited dataset, need for better interpretability and real-time usability.
Demba Faye et al. [2024]	Pre-trained CNN classifier with data augmentation	MangoleafBD	Achieved 97.80% accuracy, f1_score > 0.9; combined data augmentation techniques for enhanced accuracy.	Small dataset, overfitting, and generalizability concerns.

Saleem et al. [2024]	AgirLeafNet (NASNetMobile, Few-Shot Learning, ExG)	Multi-crop leaf dataset	100% accuracy for potato, 92% for tomato, and 99.8% for mango leaves.	Dependency on high-quality data, low interpretability, and limited generalizability.
Budania et al. [2024]	CNNs (VGG16, VGG19) with image preprocessing	Region-specific dataset	Achieved high accuracy (exact value unspecified); used image preprocessing techniques.	Model specificity, limited datasets, no consideration of environmental factors.
SB Porna et al. [2024]	Standalone CNNs and Hybrid (ResNet50V2 + EfficientNetB1)	MangoleafBD	Hybrid models achieved 100% test accuracy across all metrics.	Potential overfitting, no cross-location validation, and limited real-world testing.
Shaik Thaseentaj et al. [2023]	Customized CNN for mango leaf disease detection	South-Indian mango leaf dataset	Achieved 93.34% accuracy, outperforming traditional CNN models.	Lack of field-captured variability and dataset generalization.
Amirtha Preeya et al. [2024]	Multi-objective CNN (MOCNN) optimized via PPGO	PlantVillage mango-leaf dataset	Achieved ~96% accuracy, outperformed various models in disease detection.	Hyperparameter tuning needed, difficulty with complex feature distinction, real-world testing missing.

Lavanya et al. [2025]	Self-supervised learning with BYOL, modified ResNet	Mango-leaf dataset	Achieved 98.11% accuracy, demonstrated reduced annotation effort with self-supervised learning.	Limited dataset diversity, no evaluation on field-captured images, generalization untested.
Menaka Radhakrishnan [2025]	CNN (AlexNet, VGG16, ResNet) + Vision Transformers	MangoleafBD	ResNet + AlexNet ensemble achieved 99.97% accuracy, with Grad-CAM used for explainability.	Limited real-world validation, unclear sample size and class balance.
Sarika Khandelwal et al. [2024]	CNNs (MobileNetV2, EfficientNetV2) with PCA and SVC	MangoleafBD	Achieved 99.83% accuracy with advanced machine learning techniques.	Need for larger datasets, more transfer learning exploration.
Ahmad Ramadan Auni et al. [2024]	Transfer-learning MobileNetV2 and DenseNet201	MangoLeafBD	Combined CNNs achieved 99.25% accuracy, outperforming previous methods.	Dataset lacks field variability, no real-world validation.
P. Mathur et al. [2024]	CNNs (InceptionV3, MobileNetV3, ResNet50)	Mango leaf dataset	Achieved 100% accuracy, misclassifications were focused on three disease types.	Misclassification challenges for specific classes, need for differentiation between visually similar diseases.

Md. Eshmam Rayed et al. [2025]	MangoLeafXNet (Deep CNN with XAI)	Mango Pest Classification, MangoleafBD, MLDID,	Achieved 99.8% accuracy, with robust validation and high interpretability.	Dataset details missing, unclear real-world generalizability.
Rafi Hasan et al. [2024]	EfficientViT b0, TinyViT 5m, and MobileViT s	MangoleafBD	Achieved 99.43% accuracy, hybrid models outperformed others.	Challenges with inter-class similarities, need for field-condition validation.

2.2 Gap Analysis

In spite of the excellent contributions of the existing studies in the use of deep learning in detecting mango leaf diseases, there are still gaps. First of all, the use of small, limited amounts of data limits the scope of the models and makes it hard to differentiate between visually similar symptoms in varying real-life circumstances [3], [5]. Secondly, the majority of the models are constructed and tested in clean-room conditions, but they have not been optimized to be utilized in the field; hence they tend to suffer the effects of delays, limited hardware capabilities, and variable environments [6], [16]. Lastly, even though models have been effective in diagnosing disease in a particular crop/region, there is a lack of scalability and flexibility to other crops, regions, and diseases. This compromised generalization reminds of the necessity of more solid alternatives resistant to different circumstances and requiring fewer preprocesses to be performed, as well as more understandable approaches [9]. It will be important to fill these gaps and come up with practical scalable and field ready systems in precision agriculture.

2.3 Justification of The Current Studies

As seen in the literature reviewed, there is an increasing possibility of using deep learning in the detection of mango leaf diseases but there are still a lot of gaps to fill when it comes to diversity of the data used, real-time capability, and model generalization. These constraints pose sufficient justification to continue to investigate the development of more muscular, efficient, and interpretable models that can be deployed in real agricultural environments. Moreover, it is evident that the models taking into consideration low-resource settings also need to be developed with that in mind because not all existing models satisfy this requirement. My research endeavour tries to fill these gaps by not only endeavouring to make models more accurate through knowledge distillation but also making them more interpretable with explainable AI methods, which are vital in creating trust and usability in real-life applications.

This review of the existing literature reveals the necessity to further innovation in the model architecture and deployment mechanism, and specifically on overcoming real-world farming problems of detecting diseases.

CHAPTER 3
RESEARCH METHODOLOGY

3.1 Overview

This chapter contains a descriptive procedure relating to an efficient and pragmatic approach to mango leaf disease detection through deep learning. The main target is to come up with a solution that will be accurate and efficient to operate at the same time in balancing the actual field and those who will use it, the farmers. The research design is organized into a number of consecutive steps where the first ones pertain to the data collection and preprocessing, and the following steps are model training and evaluation.

The technique described in the present study is based on Knowledge Distillation (KD) used to achieve an optimal performance-efficiency balance, which is impossible when conventional mobile-based systems are used. Within this context, an effective DeiT-Small (teacher) model is used to train a low-level model MobileNetV3-Small (student) such that the latter model learns the high-level representation capacity of the former model, but also retains its small size in order to be used in web applications. This solution is able to eliminate most of the problems that have been cited as hurdles by other publications, such as constrained device resources, low transparency, and real-time response requirements.

To make predictions intelligible and establish it to be trustworthy, the system will employ the Explainable AI (XAI) methodologies and tools which can visualize and explain the decision-making processes to end-users and hence gain their trust. The student model with the best performance achieved with the KD framework is deployed in a web application, which is constructed with Streamlit. The application is user-friendly, lite, and makes it convenient to access by the farmers across a wide variety of environments.

Lastly, the chapter provides a step-by-step process of how experimental work is completely carried out with details on data preparation and model training including KD and how it is deployed on a web application.

3.2 Proposed Methodology

The proposed architecture model concentrates on the Teacher-Student model that uses Knowledge Distillation (KD) to classify mango leaf diseases efficiently and accurately. Here, the DeiT-Small

(Vision Transformer) model is selected as the teacher that has been pre-trained on big repositories and fine-tuned on the classification of diseases. Although this model is highly accurate, it is computationally demanding, and therefore, this model cannot be used as an application in resource-limited settings. To deal with this, a reduced and more efficient mobile student model, MobileNetV3 Small, is presented. The student model is targeted on mobile devices, low-latency applications and is lightweight. Knowledge Distillation is used where the student picks up patterns on the teacher model by matching its logits but it also learns the real labels through a combination of a soft and a hard railway. This makes the student model accurate yet it does not require the relatively intensive calculation of the teacher model.

The overall procedure starts with image preprocessing, in which the input pictures are resized and normalized to become compatible to each of the models. This web application is then built to permit users to upload mango leaf images where the program will predict the disease, give a confidence score, and create an explanation using Grad-CAM. Visual explanations using Grad-CAM show the user the areas of the image that helped the model make the prediction, adding useful interpretability to the user.

The model is evaluated by different parameters such as accuracy, precision, recall, F1-score and specificity. It is projected that the two metrics would evaluate the teacher and student models. The model performance is also verified with further metrics such as ROC-AUC to coordinate its performance with multi-class classification problems.

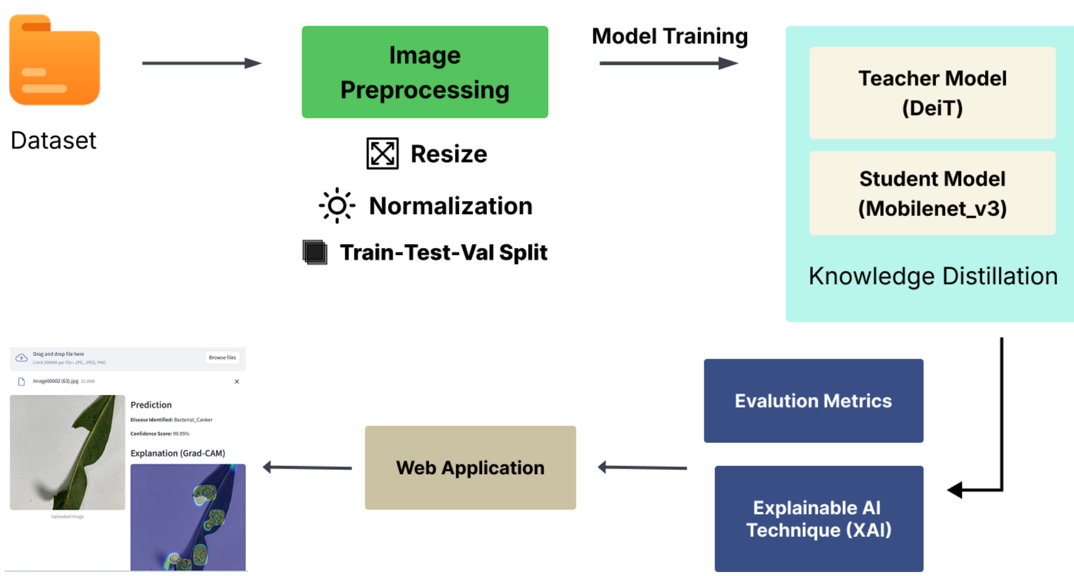

Figure 2. The Methodological procedure.

3.3 Data Description

To train and test classifiers, initially, one open access dataset and one real time collected dataset were used in this research. The publicly available dataset consists of seven mango leaf diseases, as well as the healthy leaves, whereas the real-time harvested dataset consists of five mango leaf diseases, as well as the healthy leaves. The visual symptoms presented by each of the diseases differ in important ways, with Black necrotic patches with Anthracnose, water-soaked lesions that later transform into canker-like lesions in Bacterial Canker and pimple-like projections as presented by Gall Midge. A consolidated dataset was also made between the two datasets.

To make the training process varied and to guarantee its strength, the sets were further split into two versions: a well-balanced dataset and the one that is imbalanced. The balanced dataset was generated by making sure that equal number of samples were there in each class, whereas the imbalanced dataset was generated in a manner that made individual classes have a far greater amount of samples than others in order to reflect real distribution. This division has permitted an out-and-out assessment of the performance of the model in both balanced and suitably unbalanced cases. Table 2 shows the distribution of experimental dataset classes.

Table 2: Class distribution of experimental datasets.

Class	**D1(Real Time Collected)**	**D2 (Public Data)**	**D3 (Combine of D1+D2)**
Anthracnose	618	800	1400
Bacterial_Canker	0	800	1400
Cutting_Weevil	530	800	1400
Die_Back	516	800	1400

Gall_Midge	458	800	1400
Healthy	437	800	1400
Powdery_Mildew	496	800	1400
Sooty_Mould	0	800	1400

3.3.1 Dataset (D1)

Mango Leaf Disease Dataset (D1) was obtained in mango orchards in Gazipur and Kishoreganj, Bangladesh. This dataset contains the images, taken with a Nothing 3A camera, originally having the resolution of 4080 pixels by 3072 pixels. The images were then resized to 512 x 512 to process the dataset as much as possible in order to carry out more procedures. There are six classes in the dataset that are: Cutting Weevil, Die Back, Gallmidge, Powdery mildew, Anth, and Healthy.

Its major distinction to other datasets is that it is not well balanced, with different numbers of images in the different classes. The images are provided in the format of .jpg which makes this data a good fit in the studies of the plant pathology and machine learning evaluation tasks of disease studies.

Figure 3. Representative samples of the collected Dataset (D1).

3.3.2 Dataset (D2)

The Mango Leaf Disease Dataset (D2) can be used in identifying diseased and normal mango leaves. The images in total amount to 6,400, and they are divided into 8 categories: 7 types of diseases and one of healthy leaves. The diseased classes are Sooty Mould, Powdery Mildew, Cutting Weevil, Gall Midge, Bacterial Canker, Anthracnose and Die Back. All the classes are balanced; they comprise 800 images.

The pictures were obtained in various mango plantations in Kushtia and Dhaka, Bangladesh. Also to be consistent in case and reduce memory usage, the image size was reduced to 240 pixels square by 240 pixels. All images were taken using an iPhone-SE so that there were consistent device specifications. A researcher can use the given file as the main asset towards the analysis of leaves in plant pathology and data science in general to contribute to the research on mango-related leaves diseases.

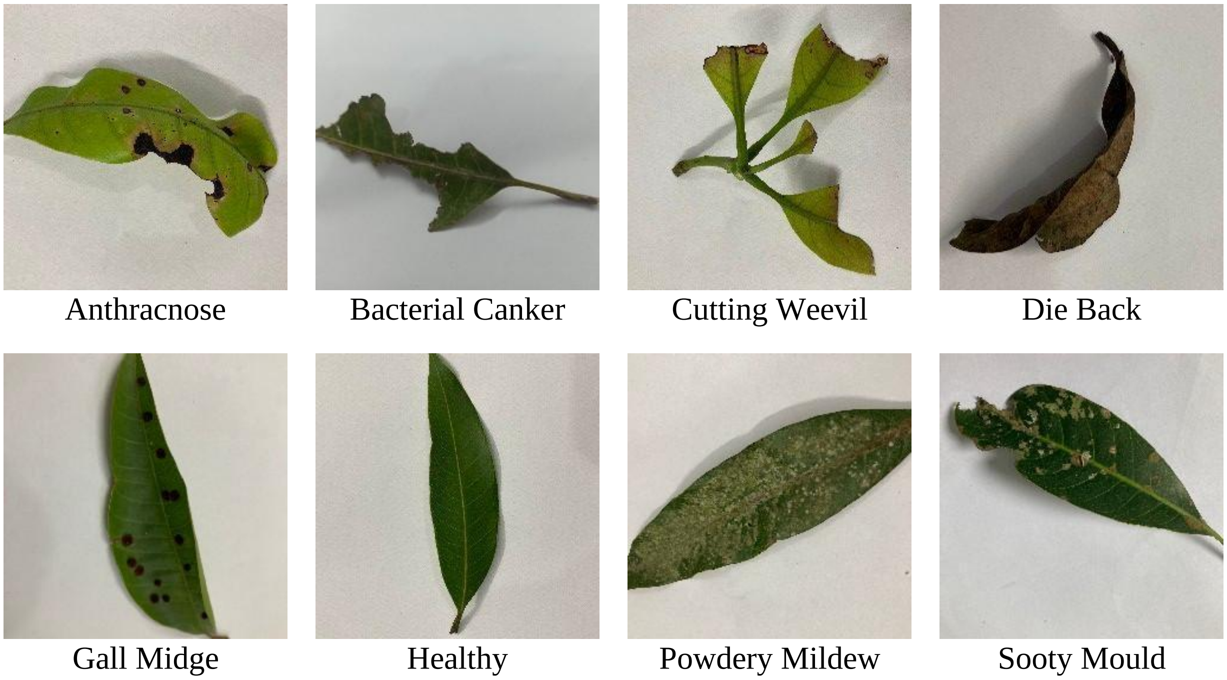

Figure 4. Representative samples of the public Dataset (D2).

3.3.3 Dataset (D3)

Finally, we created a merged dataset (D3) of both datasets D1 and D2 and established a single sourced dataset to be used in our study. Because both data had common class features, such a combination produced a more organized and complete dataset. The D3 dataset has 9,260 images which are combinations of the original datasets as follows; 3,055 images are drawn out of D1 and 6,400 out of D2. The dataset also involves more variety when it comes to the variations of diseases and environmental conditions because it has sampled a wider area of Bangladesh, i.e. Dhaka, Kushtia, Gazipur and Kishoreganj.

Eight categories are listed in D3 Weevil, Die Back, Gall Midge, Powdery Mildew, Anthracnose, Bacterial Canker, Healthy, and Sooty Mould. The break-up of classes is skewed where the total number of images is not evenly distributed. This heterogeneous data set allows neutral model training with the presence of large amounts of items in each category. In particular, it has 997 images of Cutting Weevil, 1316 images of Die Back, 1043 images of Gall Midge, 1127 images of Powdery Mildew, 1418 images of Anthracnose, 1237 images of Healthy Leaves, 1129 images of Bacterial Canker and 993 images of Sooty Mould.

Having balanced less distribution, a large overall size, and higher class diversity, the D3 dataset is a resourceful addition to deep learning-based plant disease detection.

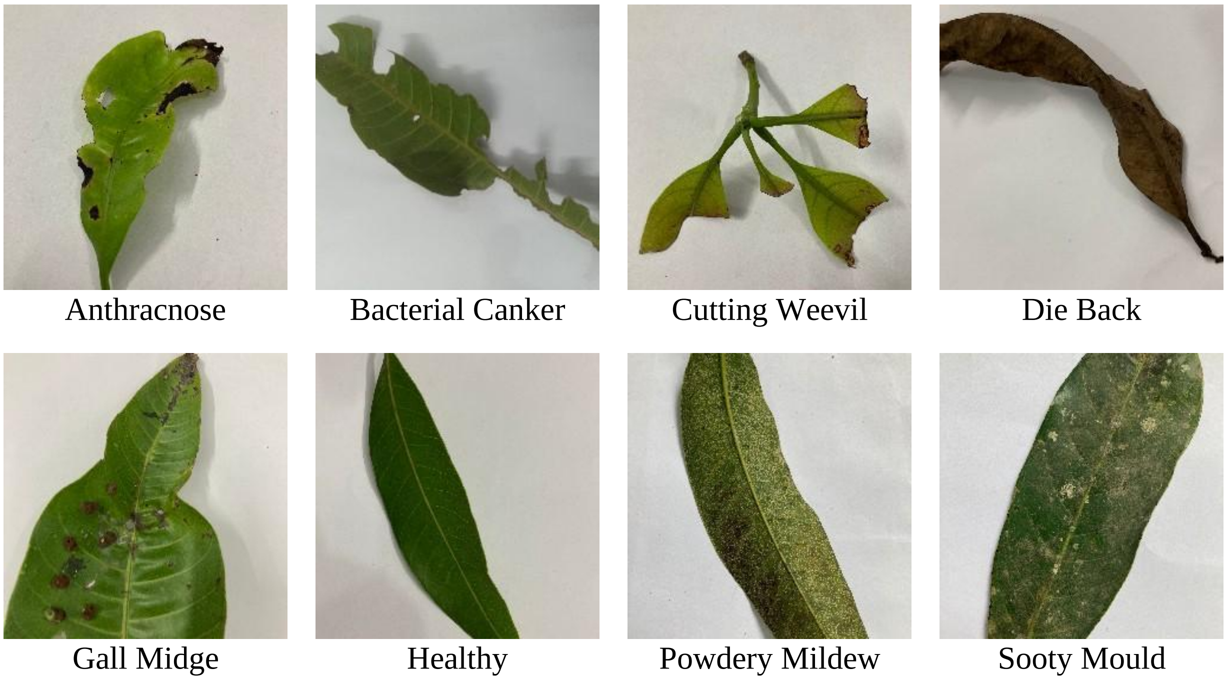

Figure 5: Representative samples of the merged Dataset (D3).

3.4 Data Preprocessing

3.4.1 Resizing

All images were resized to 224 × 224 pixels in order to guarantee uniform size of the input for the deep learning models. Bilinear interpolation, a technique that computes the weighted average of the intensities of the four closest neighbouring pixels from the original image, was used in the resizing process to estimate the intensity of a pixel in the resized image. Let I(x,y) stand for the original image's pixel's intensity at coordinates (x,y). Equation c(1) uses the weighted sum of the four adjacent pixels to calculate the intensity at the corresponding coordinates (x′,y′) in the resized image, represented as I′(x′,y′). The relative distances between the target pixel and the nearby pixels at coordinates (xi, yj) are used to calculate the interpolation weights, wij.

$$I' \ (x', y') = \sum wij \ . I \ (\mathrm{xi, yj}) \quad (1)$$

3.4.2 Normalization

The images used as input were standardized and made similar by applying Normalization. This is converted to scale the pixel values of the image to make the model learn faster. In particular, the pixel values, into the range of 0-255, are centered around the mean of 0 and the standard deviation of 1. One of the benefits of this is that it helps accelerate the rate of convergence in training and it does not bias the model to some ranges of pixel intensity.

In the present work, the mean and SD of the normalized values were the ones that are commonly used in pre-trained models, namely that of the ImageNet dataset. Normalization process is accomplished on the three channels of RGB (Red, Green, Blue) and with the following values:

Mean: [0.485,0.456,0.406]

Standard Deviation: [0.229,0.224,0.225]

This indicates that every pixel in the picture is modified as indicated:

$$I'(x,y) = \frac{I(x,y) - \mu}{\sigma} \tag{2}$$

Where:

- $I(x,y)$ is the original pixel intensity at position (x,y).
- μ is the mean intensity for the corresponding color channel,
- σ is the standard deviation for the corresponding color channel,
- $I'(x,y)$ is the normalized pixel intensity.

This normalization procedure makes the images amenable to the pretrained models and makes them more efficient due to the reduction of bias and variance across different datasets.

3.4.2 Data Augmentation

To maximize the potential of the model to generalize under different real-world scenarios in agricultural set-ups, a set of data augmentation procedures has been applied. Table 1 shows the values of those techniques applied in this study. The techniques chosen were also strategically picked to resemble natural variances that could exist in the images of the plant leaves to enable the model to become more robust and flexible. Images were also randomly rotated in a bid to make the model invariant to the position of the leaves, a fact that enabled it to be able to label leaves accurately, irrespective of how they are oriented. Also, vertical as well as horizontal flipping was added so that the model could recognize features that do not depend on orientation, allowing recognition of upside-down views as well as left-right symmetries.

Table 3: Data augmentation techniques and parameters.

Technique	**Parameter**	**Value**
Flip	Direction	Horizontal, Vertical
Rotation	Angle	Between -15° and +15°
Shear	Horizontal & Vertical Shear	±10°
Scaling	Scale factor (sx,sy)	0.5 – 1.5
Brightness	Scaling factor (β)	0.8 - 1.2

Improvements to brightness also came about and the model would be able to accommodate the various lighting situations which may arise in more practical environments like sunlight or light indoors. Images were upscaled on their resized versions, to ensure that the model outperformed on its recognition of both, smaller details, such as lesions and, the overall shape of a leaf. Finally, Gaussian noise to simulate the usual imperfections in an image that will tend to exist in non-ideal image collecting conditions was added. This simulation on noise served to increase the robustness of this model to flawed or erroneous data. The combination of these augmentation techniques brought about better performance of the model and robustness such that it could manage the variability and complexity that is usually encountered in plant leaves when they have been photographed in different settings. Figure 6 provides the illustration of the impact of these techniques, displaying an example of an unaugmented plant leaf image and comparing it to the multiple augmented versions, i.e., the rotated, flipped, brightness-adjusted, scaled, and noise-added versions thereof.

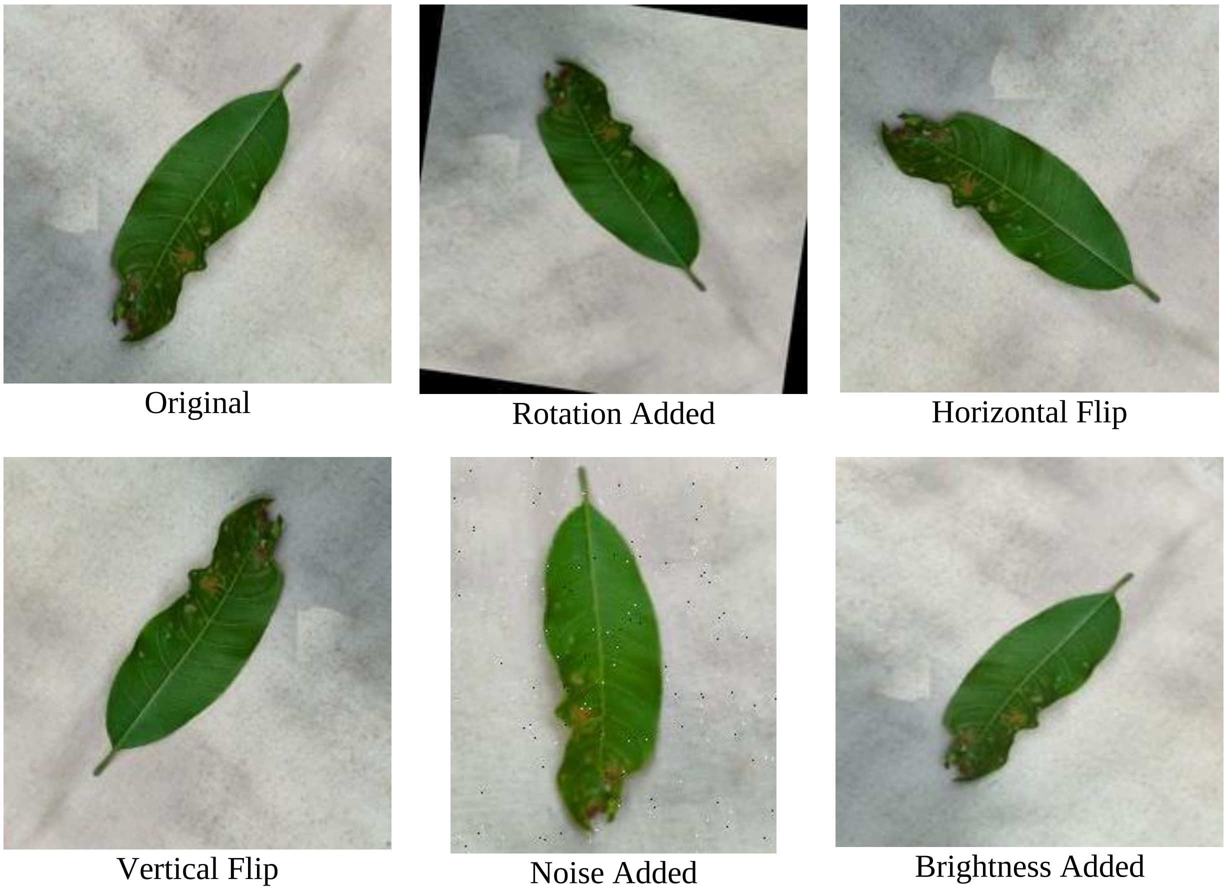

Figure 6: Visualization of each step of augmentation technique.

3.5 Model Description

3.5.1 Teacher Model: DeiT-S

The Vision Transformer (DeiT-Small, or DeiT-S) is a new solution to the problem of image classification with innovative solutions to incorporate the self-attention mechanism of transformers, as opposed to convolutional networks. Although convolutions allow CNNs to extract useful features locally, they are unable to gather information globally across an image. To resolve this issue, transformers have become an attractive alternative as they can represent a global dependency across the whole image. This is especially advocated in fine-grained tasks like mango leaf disease detection where minute disease cues in a large tree architecture need to be captured.

In our study, we choose a replica of Vision Transformers, DeiT-S, as the teacher model in a Knowledge Distillation framework. DeiT-S performs well because its strength lies in capturing long-distance spatial connections between image patches, and these types of images have complex and frequently subtle patterns in images of mango diseases.

The DeiT-S takes an RGB image that has undergone resizing to 224 x 224 pixels. This picture is then broken into disjointed patches of 16 x 16 pix with a total of 196 patches (14 x 14 grid). Every patch is flattened and put through a teachable linear projecting layer, known as the patch embedding layer (patch_embed.proj). This projection passes the raw pixel data into a high-dimensional feature space that can be processed by a transformer.

The patch embedding is then followed by a positional embedding that encodes the spatial relationships between the patches to learn to preserve the information about the original arrangement of coordinates in spite of the patch-wise tokenization operation. Figure 7 shows the architecture of the entire focal model of the DeiT teacher.

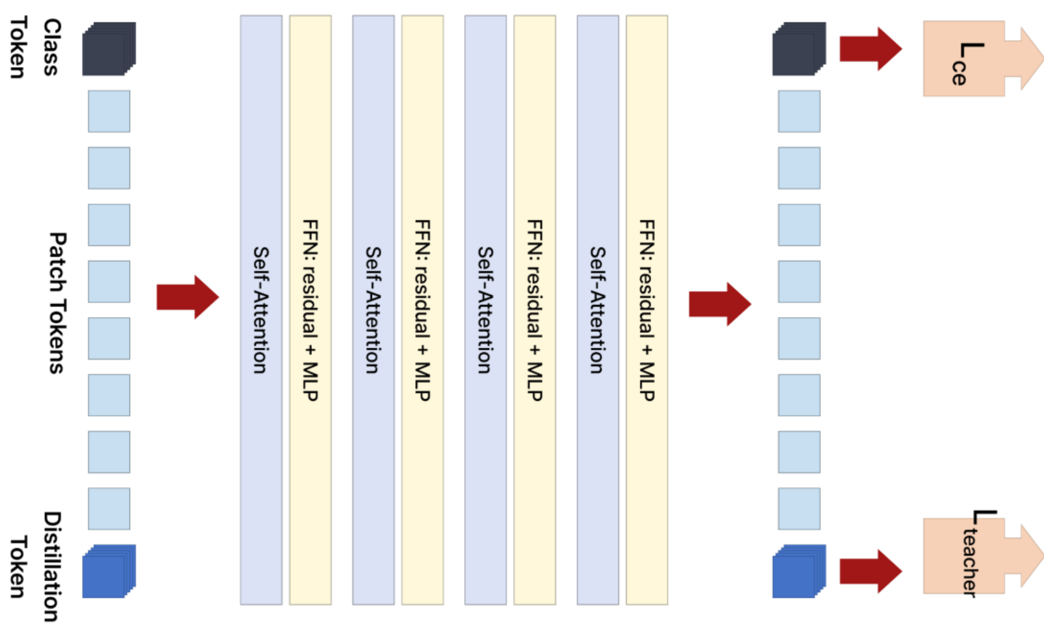

Figure 7: Architectural overview of DeiT-S teacher model.

The DeiT-S model uses 12 transformer encoder blocks (blocks.0 to blocks.11), stacked blocks of key components as shown below:

- Layer Normalization: Layer normalization (norm1 and norm2) introduces us with the custom of initiating each block by layer normalization which enhances training stability and gradient flow.
- Multi-Head Self-Attention (MHSA): The attention module (attn) employs 8 attention heads to capture the different feature interaction among all patches parallel. Since the query, key and value vectors are obtained through a linear transformation (qkv),a scaled dot-product attention formation procedure is employed. The projections of the attention then go through dropout layers which diminish overfitting.
- Residual Connections: Drop connections (drop_path1, drop_path2) encloses attention and feedforward subversions to enable us to train deeper models without becoming degraded.
- Feedforward Network (MLP): Each block has a multi-layer perceptron (mlp) with two fully connected layers (fc1 and fc2) which is separated by GELU activation function (act). The extra dimension in the MLP is increased to 3072 which has enough capacity to capture complicated representation.
- Dropout Layers: Dropout is being used following the attention and MLP layers to overcome overfitting.

This architecture gives the model the capacity to gradually sharpen patch embeddings with references to contextual information across the whole image, thus, bettering its discriminative power.

The result of processing by all the blocks of transformers is subjected to the final layer normalization (norm). Pieces of information are aggregated through a classification token ([CLS]), which is added at the input stage to the patch sequence and is the state of which provides a summary representation of the entire image at the output stage.

The classification token is processed with a fully connected linear layer (head) and the outputs are logits with length equal to the number of target disease classes (eight in this study). During inference, the logits are subsequently converted into probabilities by using a softmax. Nevertheless, the structure of the DeiT teacher model was described in Figure 8.

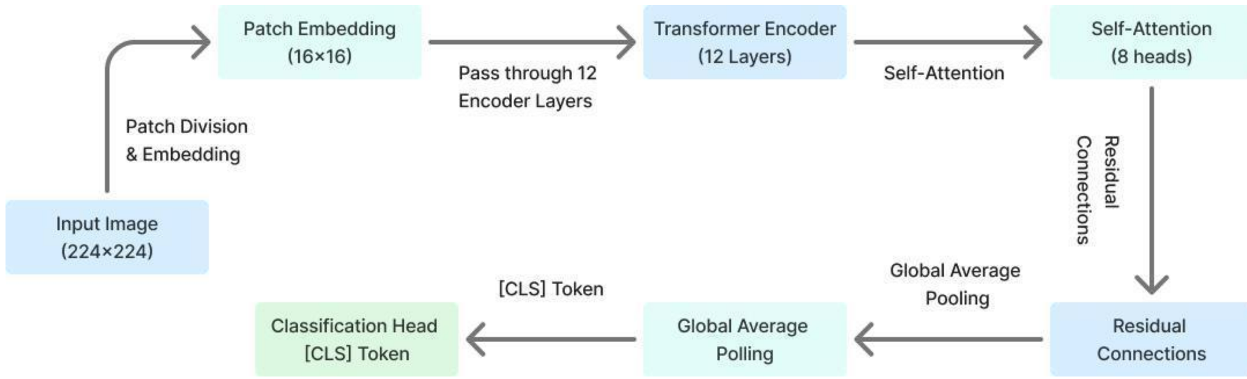

Figure 8: Overview of DeiT-S Architecture used as teacher model.

3.5.2 Student Model: MobilenetV3 - S

In the case of the student model, we use a lightweight CNN (MobileNetV3) that is used extensively in mobile and edge computing because it has better performance in regard to speed as well as memory consumption. In particular, we adopt the MobileNetV3-Small version, which represents the compromise between the accuracy and computing speed. MobileNetV3 architecture was created with a priority in the optimization of performance on mobile devices with the help of depthwise separable convolutions, squeeze-and-excitation blocks, and hard-swish activations.

The network takes input mango leaf images of size 224 x 224 and adopts the staged feature extraction paradigm approach with a first stage stem convolution, a few inverted residual bottleneck blocks, and the final short classification head.

The architecture is initiated by a stem convolution layer (conv_stem) after which it follows with batch normalization (bn1) and a hard-swish activation to rapidly grow the low-level spatial features. The body of feature extractions is divided into five sequential block groups (blocks[0] to blocks[5]), comprising one or more inverted residual modules framework. These modules use dimensionality expanding pointwise convolutions (conv_pw), depthwise convolutions (conv_dw) as a spatial filter, and dimensionality reducing projections (conv_pwl).

We also include Squeeze-and-Excitation (SE) modules within bottleneck blocks (se.conv_reduce, se.conv_expand, se.gate) to dynamically re-weight channel-wise feature attention thereby promoting a more accurate identification of minor changes in disease texture on mango leaves. After each convolution, batch normalization layers are added just to stabilize training, and the activation functions (ReLU or hard-swish) in all the layers are chosen as in the original MobileNetV3-Small regarding computational efficiency.

Spatial downsampling occurs in select points through stride-2 depthwise convolutions where we simultaneously downsample in spatial scales whilst also increasing the channel-depth. To counter information loss during downsampling, we have employed anti-aliasing layers (aa) and DropPath regularization (drop_path) is applied between blocks to improve generalization. Figure 9 indicates the entire architectural diagram of the student model of MobileNetV3 small.

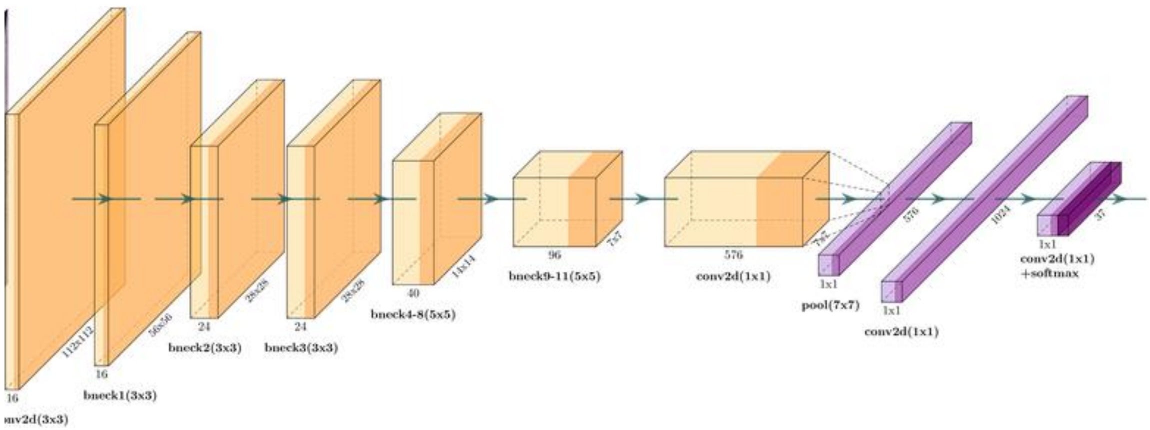

Figure 9: Overview of MobileNetV3 small Architecture used as student model.

The output flow of the final bottleneck stage is connected to a global average pooling layer that extracts all of the spatial information into a single feature vector, which then undergoes a lightweight conv_head, batch normalization layer (norm_head) and activation layer (act2). The last and the fully connected classifies using a softmax activation to obtain class probabilities with respect to the target classes of mango leaf disease.

This design contains high classification performance effectively despite the low floating-point operations to the minimal number of parameters, as well as achieves appropriate application to real-time, in-field applications. In the knowledge distillation framework, the student model MobileNetV3-Small not only gets to learn using the ground-truth labels but also gets to learn using the softened version of the full-capacity teacher model thereby allowing the student to not only remain competitive in accuracy but also be of small size due to the small architecture.

3.5.3 Knowledge Distillation Loss

To move the knowledge between the high-capacity teacher model and the lightweight student network we are going to use knowledge distillation (KD) framework. Unlike ordinary training, where the learner only observes one-hot labels the KD method provides it with an extra supervision signal that is based on softened predictions made by the teacher. Such softened outputs are given by passing a temperature parameter TTT on the logits of the teacher and student model, then returning probability distributions that are less jagged but illuminating on relating inter-class similarities.

Figure 10 shows an example of the teacher network with soft labels produced via temperature-scaled softmax, and the student network with the resultant soft and hard predictions. The soft predictions are contrasted with the soft labels set by the teacher to calculate the soft loss in terms of the Kullback-Leibler (KL) divergence, and the hard predictions with the ground-truth labels to calculate the hard loss by the cross-entropy. The last training goal is the weighted combination of the two factors of direct supervision and transferred knowledge of the teacher and leads the student to learn both the direct supervision and through the knowledge of the teacher transferred over.

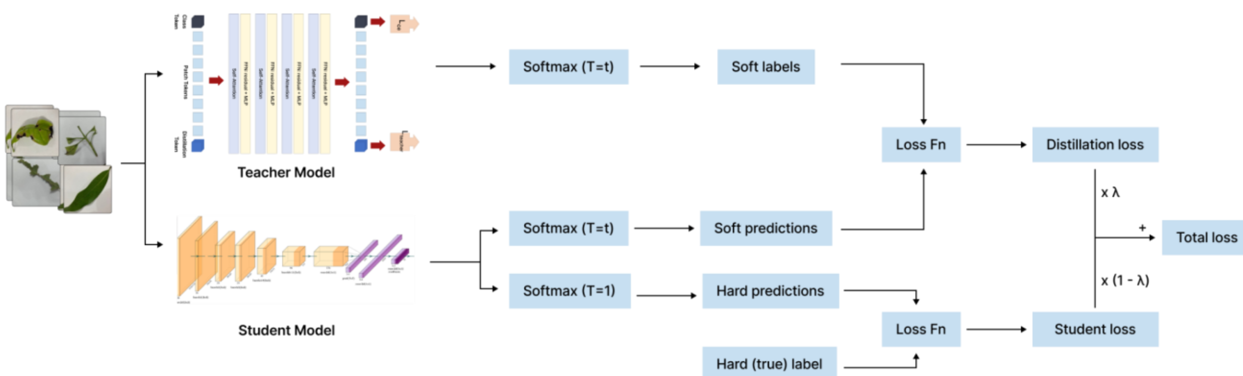

Figure 10: Knowledge Distillation model architecture.

Nonetheless, The total distillation loss is a combination of two tasks: the cross-entropy loss between the predictions, made by the student, and the ground-truth labels, as well as the KullbackLeibler divergence between the softened teacher and student distributions. In formal terms, this can be stated as:

$$L_{KD} = \alpha\, T^2\, D_{KL}(p^T \,||\, p^S) + (1 - \alpha)\, L_{CE}(y, {y}^S) \tag{3}$$

The total loss pT and pS is the scaled output probability of the teacher and student addition, LCE in representing the cross-entropy loss with the actual labels, and the hyper-parameter alpha is the coefficient of cross-entropy to favor soft supervision.

Such a formulation aims at keeping the student under the guidance of two complementary classification systems: ground-truth labels insist on the correct classification, whereas the soft labels that the teacher provides inform about the relationships between the classes. To exemplify such a process, a schematic diagram may be invented with a teacher model outputting soft labels, followed by a student, which gen-erates its outputs, and both are inserted in the KD loss function. A visualization of this nature can help one to get a better understanding of how the two streams of supervision interplay to get the right student model.

CHAPTER 4
EXPERIMENTAL RESULTS AND DISCUSSION

The chapter includes the description of the experiment setup, performance evaluation of the models, and the explanation of the results with regard to explainability and deployment of the mobile application followed by the discussion on the findings therein.

4.1 Environment Setup

A standardized and optimized computational environment was needed to ensure that the mango leaf disease detection system is consistent and reproducible at every phase. The method applied a Knowledge Distillation (KD) model in which a pre-trained DeiT-S (Data-efficient Image Transformer-Small) model was taken as the teacher and a lightweight MobileNetV3-Small architecture was trained as a student model. In this design, the student model was able to inherit the representational power of the teacher (a transformer) and also maintain efficiency in the real world implementation.

All the experiments were performed in Google Colab (free tier), which allowed the use of a common GPU space to train and validate the model. The deep learning system was created on the basis of Python 3.10, using TensorFlow and Keras, as well as auxiliary packages: OpenCV, NumPy, Matplotlib, and scikit-learn. To be deployed, the TensorFlow Lite was incorporated into the Flutter SDK, which allowed mobile applications to make offline predictions during inference. The KD training configuration predisposed uniformity in assessment through the implementation of a common set of hyperparameters on the teacher and student model. To optimize between preservations of the disease features and computational efficiency table 3 summarizes the training process, the input images were resized to 224 x 224. The batch size was set at 16 in order to take advantage of the GPU memory without making any compromise on convergence. The training of the student model was performed on 20 epochs with the aim of learning distilled representations of the teacher model without overfitting. Adam optimizer was used because of its adaptive learning rate and speed of convergence as compared to standard stochastic gradient descent. To obtain stable training, a fixed learning rate value of 0.0001 was used throughout the training.

Table 4: Common parameter table for all experimented models.

Parameter	Range	Selected Value
Epochs	{10,15,20}	10
Optimizer	{Adam, AdamW,SGD}	Adam
Learning Rate	{0.01, 0.001, 0.0001}	0.0001
Validation method	{10-fold Stratified Cross-validation, 5-fold Cross-validation}	5-fold Cross-validation
Image Size	224 x 224	224 x 224
Batch Size	{10,16,32}	16

In all experiments, the datasets were repeatedly divided into three groups, namely: training, validation, and testing. Both the teacher (DeiT-S) and the student (MobileNetV3-S) models in the Knowledge Distillation framework used this approach of partitioning to make the assessment fair and unbiased. The training set was used to learn the model and recognize patterns whereas the validation set was used to monitor the accuracy and loss curves during training as a method to optimize hyperparameters. The last test set was entirely hidden throughout training and validation and was only used to assess performance in real-world conditions. This division allowed determining the final performance values, such as test accuracy, confusion matrices, and ROC curves, which can serve as a complex evaluation of the generalization abilities of the models.

4.2 Evolution Methods

To measure the performance of the experimental models in the multi-class mango leaf disease classification task, we considered three major metrics: Specificity, Precision-Recall AUC (PR AUC) and F1-Score. Given the natural class imbalance of the datasets, all metrics were calculated in the micro-averaging way, combining results across all classes to give an overall estimate of predictive performance.

Specificity measures how well the model can identify negative instances (i.e., samples that are not of a specific class). Under the micro-averaged formulation, the sum of all true negatives (TN) and false positives (FP) across classes is calculated before being divided by all the correct classes

(TNc) and false positives (FPc), in the form given in Equation 4, where TNc and FPc are the true negatives and false positives of a particular class.

$$Specificity = \frac{\sum TN}{\sum (TN+ FP)} \tag{4}$$

Accuracy is a percentage of correctly predicted instances (both positive and negative) to the total sample count. Although it gives a rough idea of overall model performance, accuracy is inaccurate in a situation where one of the classes is largely overrepresented in the dataset.

$$Accuracy = \frac{TP+TN}{TP+TN+ FP+FN} \tag{5}$$

Recall can be defined as a ratio of all positively identified records to the actual ones. In other applications like plant disease detection, there may be serious repercussions on failure to detect a diseased plant, and this makes this application particularly important.

$$Recall = \frac{TP}{TP+ FN} \tag{6}$$

Precision is computed as the ratio of the number of correctly predicted positives to the total number of all the predicted positives. In cases where false positives are very costly, precision is very important. In this paper, a precise model is one that, when the model predicts a disease, there is high reliability in its prediction.

$$Precision = \frac{TP}{TP+ FP} \tag{7}$$

Lastly is the F1-Score which is the harmonic mean of the precision and recall and it gives an unbiased appraisal of both measures. F1-Score is particularly useful in an imbalanced classification task, because it considers both false positives and false negatives and is therefore a more complete way of evaluating the effectiveness of the model.

$$F1\,Score = 2 \times \frac{precision \times recall}{precision+ recall} \tag{8}$$

4.3 Experimental Results & Analysis

This section has a detailed analysis of the experimental results achieved through a Knowledge Distillation (KD) framework, in which the DeiT-S model was used as the teacher and MobileNetV3-Small as the student. The comparison was carried out on several significant measures such as validation-test accuracy, classification reports, confusion matrices, and ROC-AUC curves to provide a multi-faceted analysis. Although the DeiT-S teacher model provided high accuracy, the MobileNetV3-Small student model, under the supervision of KD, showed competitive results with much lower computing requirements, which better fit the deployment to resource-constrained environments. The distilled MobileNetV3-Small maintained a trade-off between performance and efficiency and proved to be a practical solution due to its light weight and mobile friendliness. Moreover, to increase the interpretability of the model, Explainable AI (XAI) methods, specifically Grad-CAM, were applied to interpret the decision process by visual explanations and features.

4.3.1 Result of MobileNetV3 (Standalone)

As can be seen in the training outcomes of the MobileNetV3-Small model in our experiments, the accuracy of the training and validation process varies considerably in terms of epochs. The model began with a relatively low accuracy of 41.25 percent during the initial epoch with the train loss of 1.6115. As training proceeded, the accuracy of the model continued to improve, and the validation accuracy reached 94.00 % after the 10 th epoch, with the train loss at 0.8689, showing improved model convergence and learning.

Nevertheless, even with these enhancements, MobileNetV3-Small was not competitive with regard to accuracy against deeper models. This dilemma explains the efficiency-versus-representational-power tradeoff in lightweight models. Although MobileNetV3-Small is very efficient and fast in terms of speed and computation ability, it has a problem with identifying subtle patterns and small details in complex datasets. This can be seen quite clearly in the validation loss values which were relatively high compared to the lower training loss, which indicated that the model was not generalizing well on new/unseen data.

The errors in the confusion matrices (see Figure 11 and Figure 12) indicate that MobileNetV3-Small is ineffective in processing complex features when trained independently. These results confirm again the known problem in lightweight models, they do not capture the complex patterns in the data completely.

To address this performance bottleneck we deployed a Knowledge Distillation (KD) strategy using the rich, structured knowledge of a large-capacity teacher network. By distilling the lightweight model along this process, we expected to help it learn more discriminative and robust representations, hence produce better generalization performance and overcome the trade-off between efficiency and accuracy.

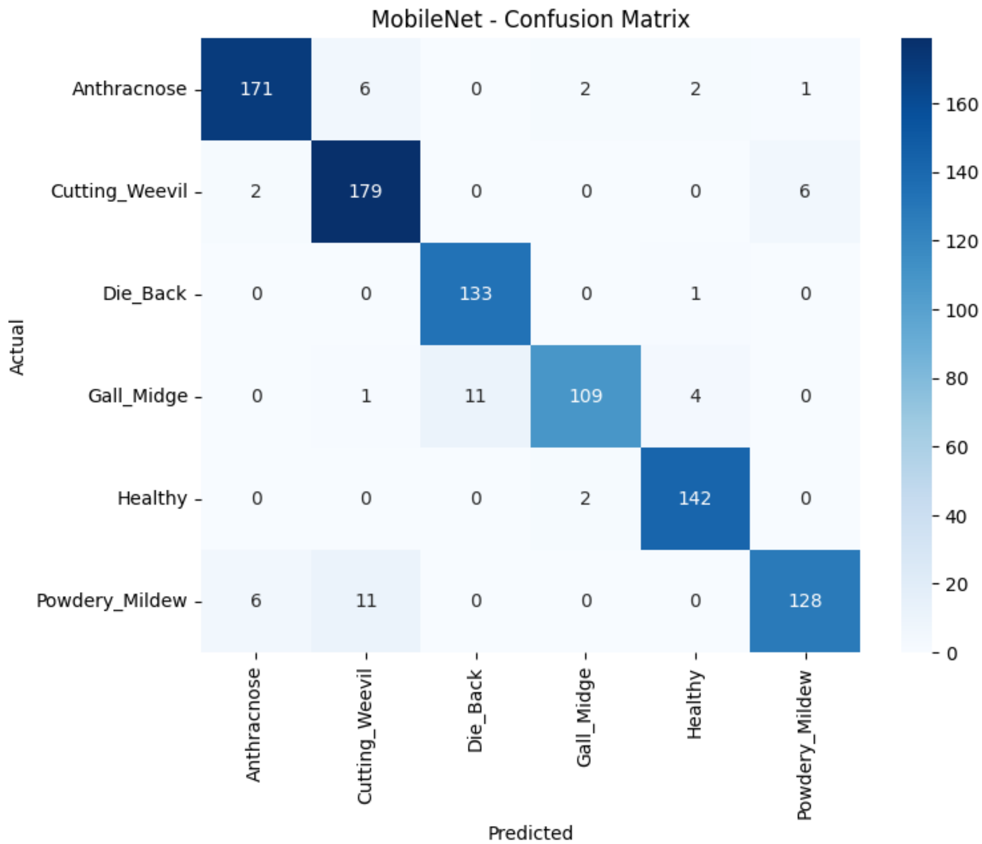

Figure 11: Confusion Matrix of MobileNetV3 on D1.

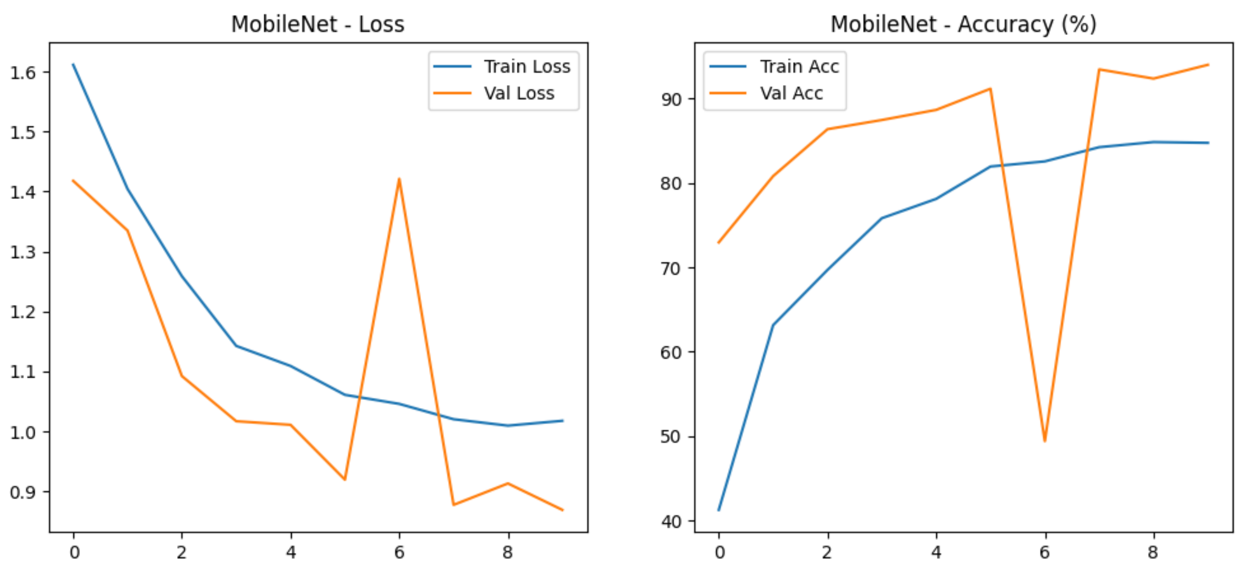

Figure 12: Accuracy & Loss Graph of MobileNetV3 on D1.

4.3.2 DeiT-S Teacher Model Results

The DeiT-S model, acting as the instructor in the KD setting, learned very well and reliably during the training procedure. As depicted in the epoch-wise results, the model quickly converged and the training accuracy and validation accuracy were 95.59 and 96.11 in the first epoch itself. Later epochs still improved on this showing, with validation accuracy reaching 99.57% at epoch 7, indicating that the teacher has a fantastic capacity to capture minute-level visual patterns in the data.

The loss plots during training and validations (Figure 13) show an evident negative slope, and there is no major deviation as it tends to converge. This fact supports the integrity and stability of the optimization process. Although slight variations were found to occur later (e.g., a temporary decrease in validation accuracy on epochs 8-9), the general trend matches that of a well-generalized model as opposed to one that overfits. This is further evidenced by the fact that the validation accuracies are very high above 98 percent in most of the epochs.

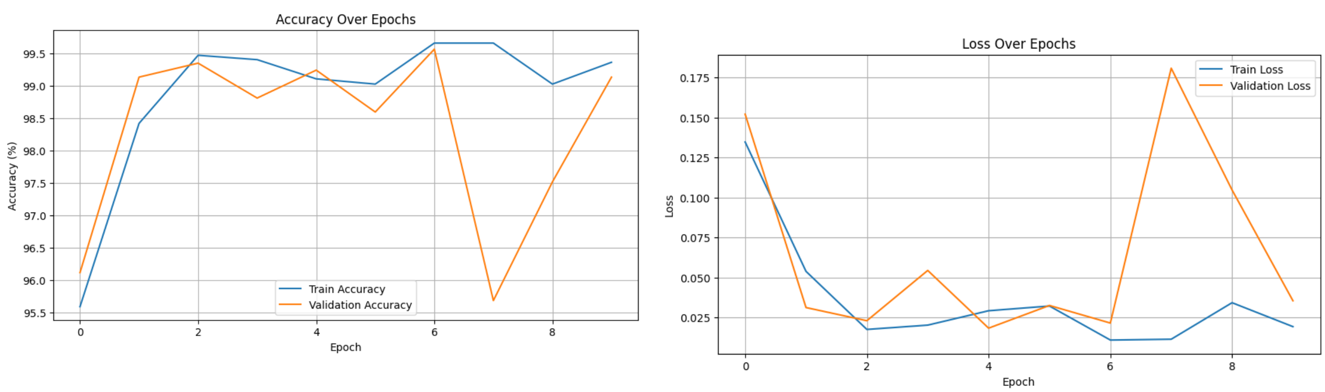

Figure 13: Accuracy & Loss Graph of Teacher Model on D3.

The confusion matrixes within folds (Figure 14) support the high discriminative power of the model, with the diagonal elements prevailing which indicates high levels of the correct prediction of all eight classes. Misclassifications were correctly limited, and in cases where they occurred, were largely confined to visually similar disease categories, and reflect the ability of the teacher to disambiguate subtle differences. This is also confirmed in the corresponding classification reports (Table 5) that depict strong precision, recall, and F1-scores across all classes further evidencing the balanced nature of the teacher.

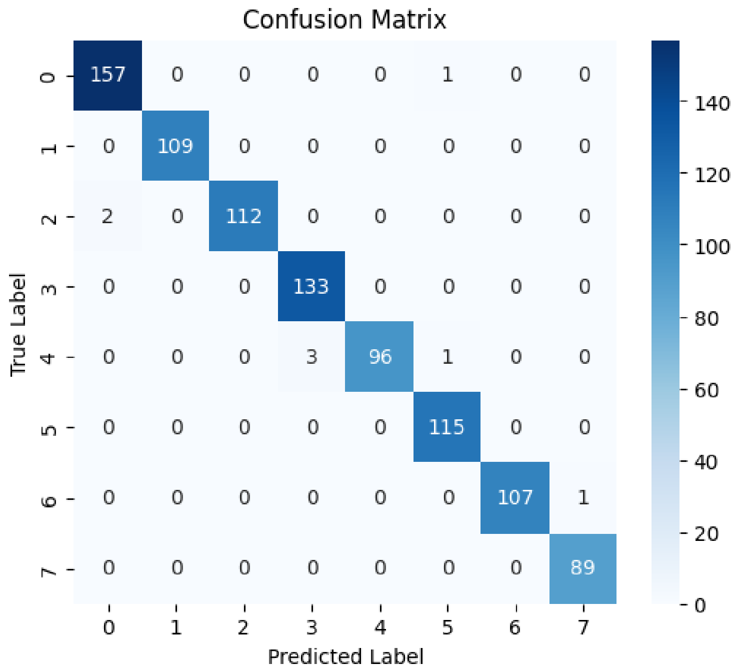

Figure 14: Confusion Matrix of Teacher Model on D3.

Table 5: Classification report of Teacher Model on D3 .

	Precision	**Recall**	**F1-score**	**Support**	**Specificity**
Anthracnose (0)	0.99	0.99	0.99	158	0.99
Bacterial_Canker (1)	1.00	1.00	1.00	109	1.00
Cutting_Weevil (2)	1.00	0.98	0.99	114	1.00
Die_Back (3)	0.98	1.00	0.99	133	0.99
Gall_Midge (4)	1.00	0.96	0.98	100	1.00
Healthy (5)	0.98	1.00	0.99	115	0.99
Powdery_Mildew (6)	1.00	0.99	1.00	108	1.00

Sooty_Mould (7)	0.99	1.00	0.99	89	0.99

Overall, the DeiT-S model proved to be a very stable and precise educator, being able to learn strong feature representations that not only lead the standalone performance but also serve as a powerful teacher network in imparting knowledge to the lightweight student MobileNetV3-Small model through knowledge distillation.

Moreover, after testing the teacher model (DeiT-S) on two new datasets- D1 (collected dataset) and D2 (public dataset), the performance was still very high. The teacher model performed well on both datasets with high training and validation accuracies as well as strong test performance indicating that the teacher model is capable of generalizing to different data sources. The accuracy/loss curve, confusion matrix and classification report visualizations of these datasets confirm that our model is stable and flexible. This consistency makes DeiT-S a credible teacher model, which makes it suitable to guide the lightweight MobileNetV3-S through knowledge distillation.

Figure 15: Accuracy & Loss Graph of Teacher Model on D1.

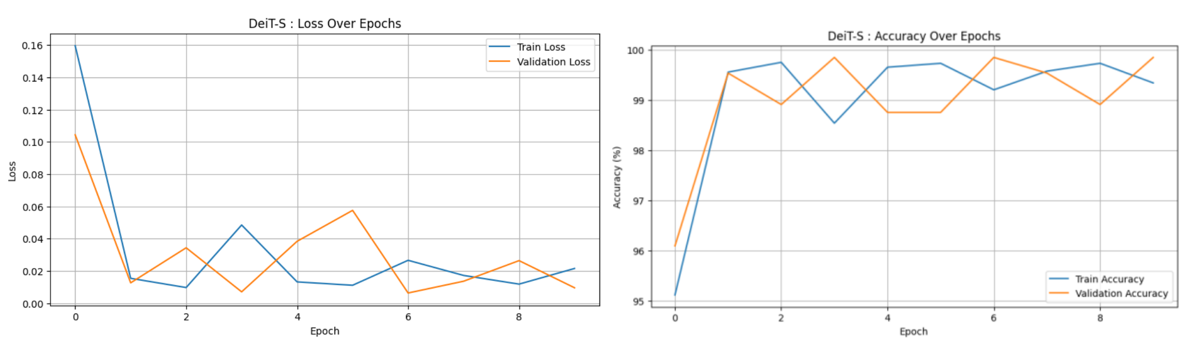

Figure 16: Accuracy & Loss Graph of Teacher Model on D2.

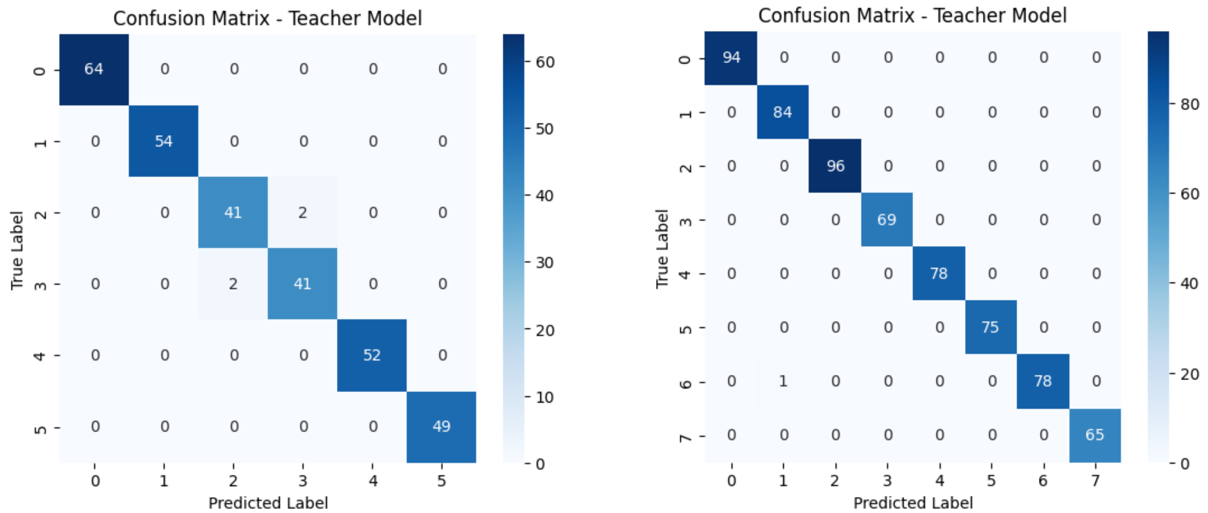

Figure 17: Confusion Matrix of Teacher Model on D1 & D2.

Table 6: Classification report of Teacher Model on D1 & D2 .

		Precision	**Recall**	**F1-score**	**Support**	**Specificity**
D1	Anthracnose (0)	1.00	1.00	1.00	64	1.00
	Cutting_Weevil (2)	1.00	1.00	1.00	54	1.00
	Die_Back (3)	0.95	0.95	0.95	43	0.99
	Gall_Midge (4)	0.95	0.95	0.95	43	0.99
	Healthy (5)	1.00	1.00	1.00	52	1.00
	Powdery_Mildew (6)	1.00	1.00	1.00	49	1.00
	Anthracnose (0)	1.00	1.00	1.00	94	1.00
	Bacterial_Canker	0.99	1.00	0.99	84	0.99

	(1)					
D2	Cutting_Weevil (2)	1.00	1.00	1.00	96	1.00
	Die_Back (3)	1.00	1.00	1.00	69	1.00
	Gall_Midge (4)	1.00	1.00	1.00	78	1.00
	Healthy (5)	1.00	1.00	1.00	75	1.00
	Powdery_Mildew (6)	1.00	0.99	0.99	79	1.00
	Sooty_Mould (7)	1.00	1.00	1.00	65	1.00

4.3.3 Result of Knowledge Distillation - MobileNetV3

To ascertain a strong and objective analysis of the Knowledge Distillation (KD) framework- wherein the DeiT-S teacher model was initially used to train the MobileNetV3-Small student model- a 5-fold cross validation strategy was thereafter applied on the distilled student model. Unlike one-train-test splits which may introduce a bias of a certain partition, by dividing the data into five identical subsets, K-Fold validation performs four training and one validation step with each fold in turn. This process involved using each sample to train and validate in all the folds, thus increasing the reliability of this evaluation as well as reducing variation of reported performance. The choice of k=5 allowed balancing the computational efficiency and the soundness of performance estimates, whereas the averaged k-fold means made it possible to learn more about the learning stability of the student model in the KD regime.

The post-training confusion matrices per fold given in Figure 17 also demonstrate the usefulness of the KD approach. In each of the five folds, the matrices exhibited extremely dominant diagonal elements, which indicate high class-specific accuracy and an overall decreasing number of misclassifications as compared to the standalone MobileNetV3-Small baseline. The few errors that were made were mostly within visually similar categories, where inter-class ambiguities are quite difficult even on high-capacity models. Notably, the consistency and invariance of the classification performance across folds supports the ability of the KD-informed MobileNetV3-Small to capture unbiased and balanced feature representations between the eight classes. The results are also supported by the respective classification reports, indicating that KD drastically increases the discriminative value of lightweight student models, without compromising their efficiency.

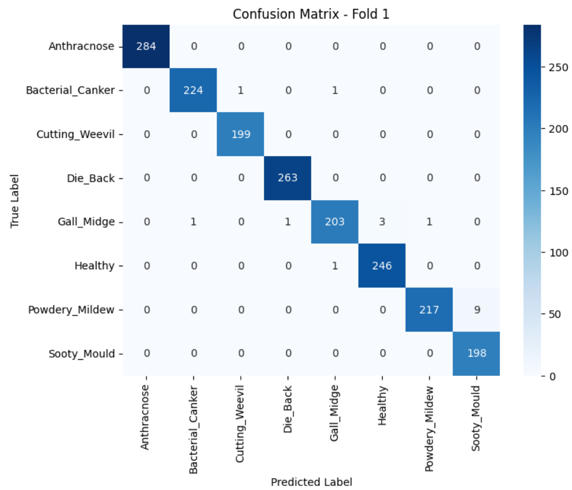

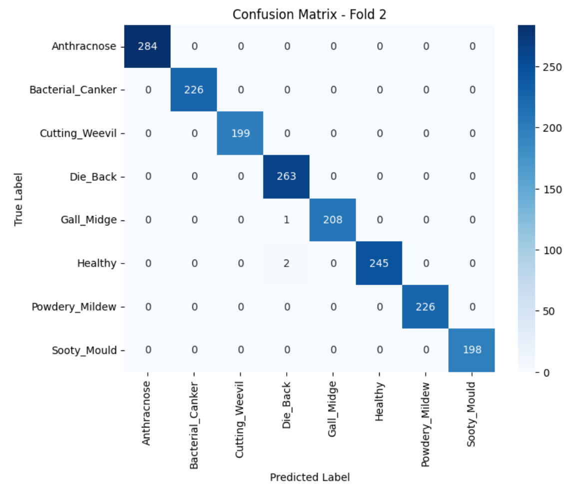

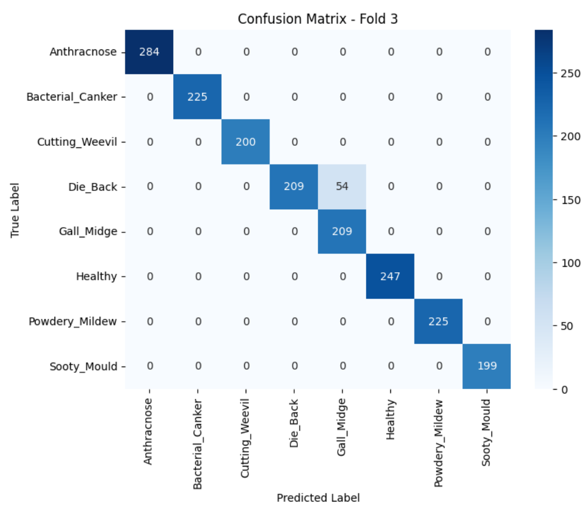

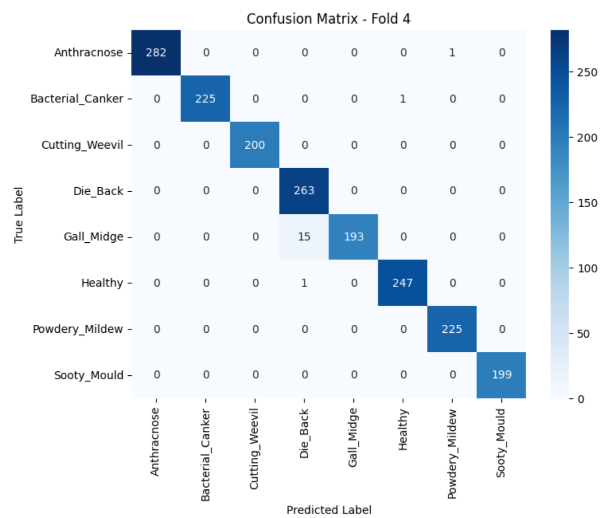

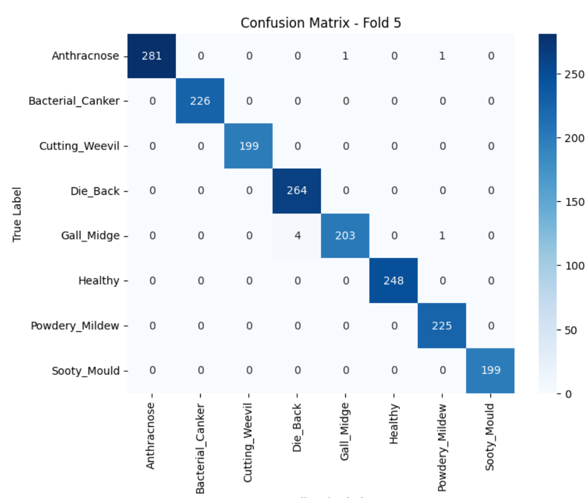

Figure 18: Confusion Matrices of KD Student Model Across Five-Fold Cross Validation.

Table 7: Fold-Wise Classification Reports of KD Student Model Across Five-Fold Cross-Validation.

		Precision	Recall	F1-score	Support	Specificity
F1	Anthracnose (0)	1.00	1.00	1.00	284	1.00
	Bacterial_Canker (1)	1.00	0.99	0.99	226	0.99
	Cutting_Weevil (2)	0.99	1.00	1.00	119	0.99
	Die_Back (3)	1.00	1.00	1.00	263	0.99
	Gall_Midge (4)	0.99	0.97	0.98	209	0.99
	Healthy (5)	0.99	1.00	0.99	247	0.99
	Powdery_Mildew (6)	1.00	0.96	0.98	226	0.99
	Sooty_Mould (7)	0.96	1.00	0.98	198	0.99
F2	Anthracnose (0)	1.00	1.00	1.00	284	1.00
	Bacterial_Canker (1)	1.00	1.00	1.00	226	1.00
	Cutting_Weevil (2)	1.00	1.00	1.00	199	1.00
	Die_Back (3)	0.99	1.00	0.99	263	0.99
	Gall_Midge (4)	1.00	1.00	1.00	209	1.00
	Healthy (5)	1.00	0.99	1.00	247	1.00
	Powdery_Mildew (6)	1.00	1.00	1.00	226	1.00
	Sooty_Mould (7)	1.00	1.00	1.00	198	1.00

F3	Anthracnose (0)	1.00	1.00	1.00	284	1.00
	Bacterial_Canker (1)	1.00	1.00	1.00	225	1.00
	Cutting_Weevil (2)	1.00	1.00	1.00	200	1.00
	Die_Back (3)	1.00	0.79	0.89	263	1.00
	Gall_Midge (4)	0.79	1.00	0.89	209	0.96
	Healthy (5)	1.00	1.00	1.00	247	1.00
	Powdery_Mildew (6)	1.00	1.00	1.00	225	1.00
	Sooty_Mould (7)	1.00	1.00	1.00	199	1.00
F4	Anthracnose (0)	1.00	1.00	1.00	283	1.00
	Bacterial_Canker (1)	1.00	1.00	1.00	226	1.00
	Cutting_Weevil (2)	1.00	1.00	1.00	200	1.00
	Die_Back (3)	0.94	1.00	0.97	263	0.98
	Gall_Midge (4)	1.00	0.93	0.96	208	1.00
	Healthy (5)	1.00	1.00	1.00	248	0.99
	Powdery_Mildew (6)	1.00	1.00	1.00	225	0.99
	Sooty_Mould (7)	1.00	1.00	1.00	199	1.00
	Anthracnose (0)	1.00	1.00	1.00	284	1.00
	Bacterial_Canker (1)	1.00	1.00	1.00	225	1.00
	Cutting_Weevil (2)	1.00	1.00	1.00	200	1.00
	Die_Back (3)	1.00	0.79	0.89	263	0.98

F5	Gall_Midge (4)	0.79	1.00	0.89	209	0.96
	Healthy (5)	1.00	1.00	1.00	247	1.00
	Powdery_Mildew (6)	1.00	1.00	1.00	225	1.00
	Sooty_Mould (7)	1.00	1.00	1.00	199	1.00

A table summarizing the classification performance of each of the five folds has been provided in Table 7. The results show that KD-based MobileNetV3-S had a high precision, recall, and F1-scores on all folds with little variance. This level of consistency across data splits provides compelling evidence of the accuracy as well as reliability of the student model when directed by the teacher (DeiT-S). Notably, the consistency of the F1-score demonstrates the balanced ability of the student model to recognize both classes of positive and negative without the inclination to any of the classes. The high correlation between the cross-fold and weighted averages further evidences the soundness of the distillation framework. Such fold-wise classification reports indicate that the KD method allowed the student model to learn generalized representations of the dataset and prevented overfitting to specific folds, therefore, increasing its credibility as an effective yet reliable lightweight classifier.

4.3.4 Visualizations after applying Explainable AI (XAI)

Grad-CAM was used to select correctly and misclassified mango leaf images to enhance the interpretability and transparency of the predictions of the Student model within the knowledge distillation framework. This visualization technique is crucial in knowing how the model reached its decisions, so that the users can visually verify that predictions were made on the basis of real symptoms of the disease.

The evidence in Grad-CAM heatmaps was strong that the student model consistently targeted the area of interest on mango leaf images. In correctly classified samples, the highlighted areas were consistent with disease-affected areas like spots, discolorations, fungal patterns, and these areas were closely similar to the human eye. These high-activation regions matched known diseased regions in the dataset, which showed that the model was not biased by irrelevant background noise, or the edge of the leaves.

In misclassification cases, Grad-CAM maps showed that the model still focused on disease-like patterns. Nevertheless, there was a high visual similarity between some classes of diseases that may have led to prediction confusion. These findings support that the attention mechanisms learned by student model is functional, and they contribute to trusting the model decision-making process and its performance in terms of reliability when deployed practically.

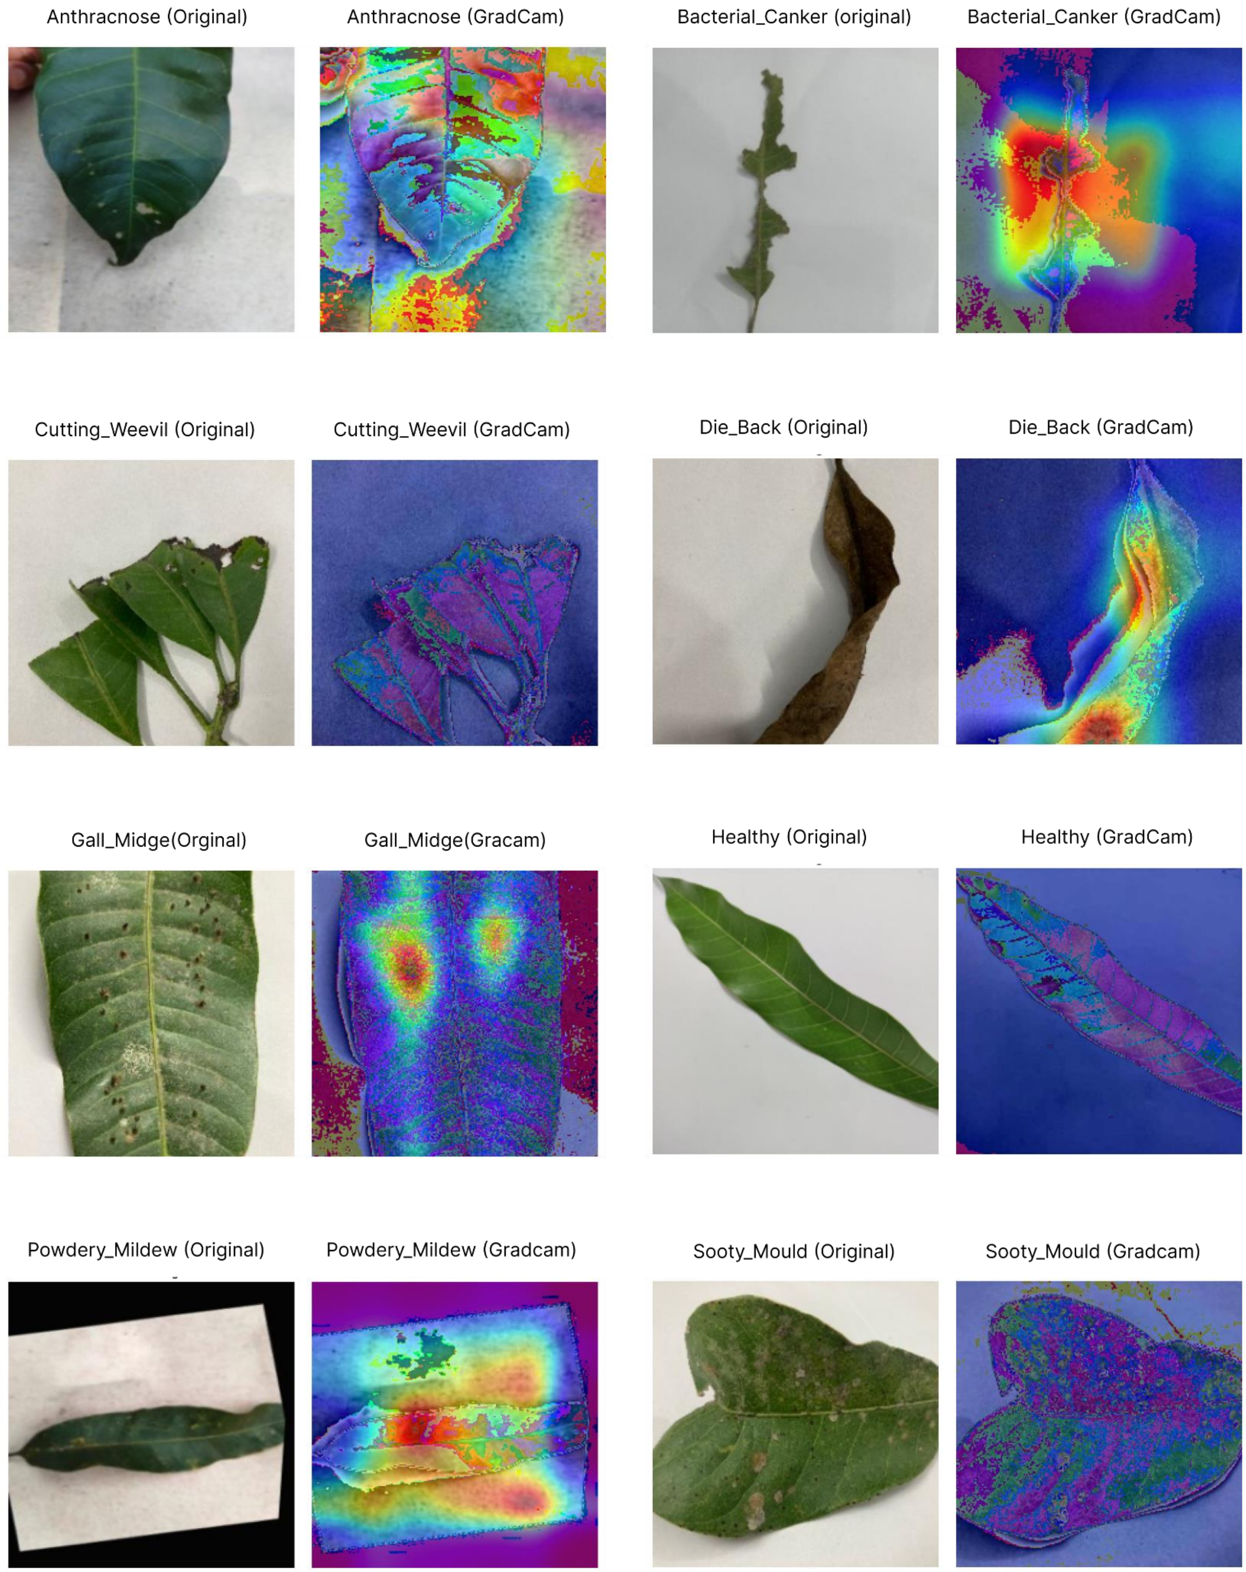

Figure 19: Grad-CAM highlights the disease-affected regions that were attended by the KD Student model during classifications on D3.

4.3.5 Results of Web Application Deployment

The web application makes use of a student model (MobileNetV3-Small) trained via Knowledge Distillation, so that the classification of diseases on mango leaves can be realized efficiently. It will be useful in aiding farmers, agricultural experts, and researchers with a quick and explainable manner of disease detection. They can upload pictures of mango leaves in JPG or PNG format which are then processed by the model that can give detailed predictions.

After an image has been uploaded, the application will show the predicted disease type with the confidence score. As shown in the screenshot above, the model is very confident that the leaf is infected with Anthracnose (99.99%), whereas the probabilities in all the other categories are negligible, which means that the model is very sure in its judgment.

An important aspect of the app is the Grad-CAM based heatmap which is presented together with the uploaded picture. The heatmap identifies the leaves areas that contributed to the prediction most allowing users to visually confirm the areas of disease appearance. This explainability aspect contributes to the trust in the automated diagnosis and helps to understand the model decision-making process in a better manner.

The interface is simplified and enables the user to drag and drop files, browse images, control such parameters as the size of the image, the opacity of the heatmap, and the selection of the target layer. This ability to be manipulated makes it simple to test several samples and compare the results and do so quickly.

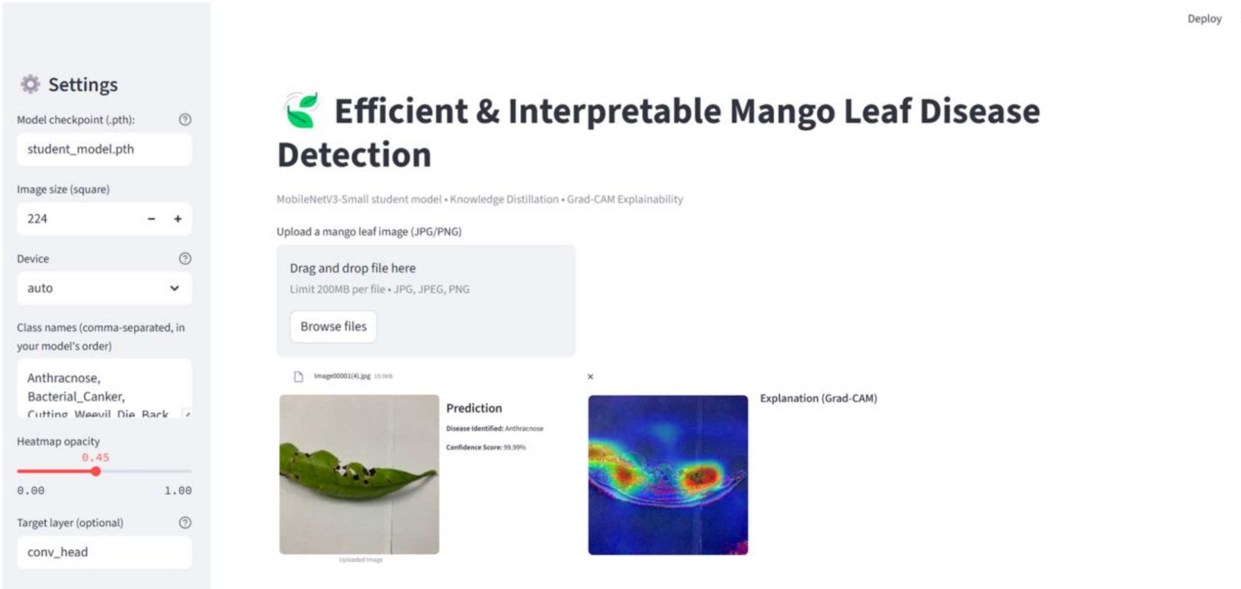

Figure 20: Mango leaf disease classification web application.

In the future, it is possible to monitor in real-time, predictive analytics, and combine with IoT-connected sensors to monitor environmental conditions continuously including temperature, humidity, and soil moisture. Such streams of data may be used in more elaborate predictive models to predict disease outbreaks. Also, by connecting the application to agricultural databases, it would be possible to perform region-specific disease risk mapping, and by using Firebase Cloud Messaging or SMS messaging, farmers could be warned of the possible threat. The addition of edge computing and GIS features would also be beneficial to increase responsiveness and scalability, such that the system can be used on a large scale to monitor agriculture and plan policies accordingly.

Being modular and scalable, this application closes the gap between sophisticated disease classification models and their application in practice, providing both accuracy and interpretability in a convenient form.

4.4 Summary

This chapter described the whole implementation and evaluation procedure of the suggested mango leaf illness recognition system in terms of a Knowledge Distillation (KD) framework. The teacher model (DeiT-S) had a good recognition ability and now provides a good baseline performance. Its processing expense however, pointed out the constraints of directly implementing such heavy structures in constrained systems. To eliminate this difficulty, the student model (MobileNetV3-Small) was trained in the KD paradigm. The student demonstrated competitive classification performance and an order of magnitude decrease in computation overhead, thus, the system is well suited to real-time deployment in a mobile setting.

The application of Explainable AI like Grad-CAM also helped to make the system more transparent in terms of how it arrives at the decision that it makes, which is essential in building trust in agricultural applications. Also, the experimentation on two separate datasets D1 (collected) and D2 (public) indicated that the teacher model had retained a high performance mark, which confirmed the generalization ability of the approach.

Lastly, the demonstration of the distilled student model within a web application using Streamlit proved that farmers and agricultural workers could be provided with user-friendly and easy-to-use disease detection that is offline and available in real time, bringing accessible and reliable technology to the field. The results reported in this chapter clearly show the feasibility and stability of the proposed solution.

CHAPTER 5
ENGINEERING STANDARDS & IMPACT ON SOCIETY, ENVIRONMENT AND SUSTAINABILITY

This chapter explains the engineering requirements observed, the ethical facts, the sustainability impact and the various problems it faced in the development and implementation of the systems.

5.1 Compliance with the Standards

5.1.1 Software Standards

The project software environment uses Google Colaboratory, a Google-based cloud computing environment which supports Jupyter notebooks. This platform allows the integration with any other cloud service easily and makes available strong computational capacity to train a deep learning model. The major programming language featured during the development is Python 3.x as it is highly-customizable and fits a variety of deep learning libraries.

Several notable libraries are taking advantage of the project, in terms of model development. TensorFlow has been used in the creation of deep learning models, particularly convolution neural network (CNN) models. Also, Keras, which is compatible with TensorFlow, is applied to the construction and learning process of models with the simplistic interface. TorchVision is used in model development tasks (mainly in the DeiT-S (DeiT-Small) teacher model and MobileNetV3-S student model in the Knowledge Distillation (KD) framework), especially image transformations. On the data handling side, NumPy and Pandas will help in effective data manipulation and data storage whereas OpenCV will be incorporated to provide enhanced image processing actions. Matplotlib and Seaborn are used to visualize the data and the model performance and they are displayed clearly and informatively. The datasets will be stored in Google Drive so that they are small, readily available in the Colab environment to use them in the project and manage large datasets easily.

5.1.2 Hardware Standards

The project hardware needs can be serviced through Google Colaboratory (Colab) as it offers virtual machines on free-tier subscriptions with different hardware specifications. Under this free pricing plan, the user is able to access up to 2 virtual cores of CPUs. More complex tasks can rely on access to NVIDIA GPUs available on Colab, including an example Tesla K80, T4, P100, or V100, units, which have the required computing power to train deep learning models.

In terms of memory, Colab has up to 12 GB RAM which is enough to carry out most deep learning tests. The disk space used in the Colab runtime environment is also undesirable and can be up to 15 GB. Such configuration allows the sufficient resources on the development and training of models, and tasks involving Knowledge Distillation (KD) can be performed without subscribing to Colab Pro.

5.2 Impact on Life

The introduction of a real-time mango leaf disease detector has greatly improved the lives of interested people in the agricultural sector, especially, the small holders. This AI-enabled solution presents itself as an auxiliary consultant to the underdeveloped regions where access to agricultural experts is limited and is able to identify diseases on mobile devices in only a couple of seconds. Due to early detection, farmers are able to take adequate measures, as problems can be solved with the help of specific treatments that prevent the spreading of diseases and possible decrease in crops. The tool also helps reduce the guess and check approach of the conventional method of using pesticide, relieving the pocket strain and health hazards involved by the wasteful use of chemical material. Exposure to chemical pesticides in the long-term may cause fatal illnesses such as respiratory diseases and skin diseases. This model can help achieve improved occupational safety by farmers because it allows them to make more specific and safe interventions.

Moreover, the tool enhances both the productivity and the stability of the income of the farmers by minimising the rate of crop failing. It helps to manage diseases more effectively, so farmers use available resources more sensibly and do not waste their money on inappropriate treatments and could enhance the overall quality of crops. The socio-economic ramification of such technology is significant in low income areas or rural areas, where every crop yields a lot. It not only guarantees the livelihood security of individuals but also builds resilience of the community because of the long term agricultural productivity that is required by the agricultural communities.

5.3 Impact on Society & Environment

On the societal level, the use of this AI-based disease detection system is able to bring third-world agricultural systems into the modern world. By stabilizing the health of agricultural products, it stabilizes the food delivery chain and the prediction of food prices, especially those products that are major necessities such as mangoes. Prevention of disease epidemics and their prevention enhances national food security, as the country will not have to depend on imports and this will contribute to the stability of the food market.

The system is also important in sustaining sustainable agriculture. Common methods in controlling pests and the disease processes usually involve the use of broad-spectrum of pesticides that may be detrimental to other nearby ecology. Escape due to untreated pesticides pollutes water, kills soil fertility, and disrupts biodiversity and natural pest predators and pollinators. AI-based solutions, on the other hand, promote the practice of precision agriculture, with disease diagnosis and treatment advice produced at a high level of accuracy and with minimal chemical usage and maximum efficiency. Not only does this lower the environmental pollution but also helps conserve native flora and fauna and embrace organic farming.

Also, by integrating real life data under different kinds of agricultural conditions, one can adjust the system to be applied to various ecological zones in a way that supports the technique of climate-resilient agriculture. This strategy would mean that agricultural activities will be sustainable and flexible to evolving environmental pressures to long-term agriculture productivity and conservation of the environment.

5.4 Ethical Aspects

The moral basis of the study is to protect data in terms of privacy and to have equality of access, and transparency of the models. With the agricultural sector becoming more and more digitized, questions of data privacy and misuse occur. In order to overcome the challenges, the study employs the principles of federated learning, which allows the use of decentralised statistical data in a decentralised fashion without centralising any sensitive data about crops. This methodology guarantees that the data of individual farmers would remain local, but still could be used to improve global models in a privacy- and collective intelligence-preserving manner.

The concept of inclusivity is also a major ethical guiding principle The system removes obstacles to uptake in underserved or resource-limited communities by using a lightweight mobile-compatible model MobileNetV3-small. The application can be useful even to farmers with low-end android devices or unstable internet connection, which is why all groups will be included in all the benefits of the technology.

Lastly, explainability is an important ethical factor. Conventional deep learning models have been regarded as the black boxes considering the fact that it is hard to understand the decision made by the model. This absence of transparency has the potential to erode trust, particularly in areas such as agriculture whereby AI decision-making can be of utmost importance. To counter this, the system takes advantage of Grad-CAM (Gradient-weighted Class Activation Mapping) which is a useful tool that offers visual and contextual explanation of model predictions. By explaining why a certain illness was detected, Grad-CAM generates trust on the AI-guided suggestions, and also enables users to make their own decisions.

5.5 Sustainability Plan

The holistic effect of this system depends on external vigilance as well as the involvement of the people to whom it benefits. The suggested sustainability plan is several-fold relating to environment and operational aspects. In an environmental perspective, the accuracy of the model is high, this reduces wastage of chemicals that may affect the quality of the soil, degrade water quality and other ecosystems around. By allowing farmers to base their decision-making on empirical evidence, instead of basing decisions on habitual grounds, the system fosters integrated pest management instead of practices associated with the promotion of more environmentally-friendly farming techniques.

Operationally the student model design is constructed to be versatile to a variety of crop type, geographical conditions. This is because as new disease images are acquired and model updates are pushed, the system is able to undergo an evolution without lengthy redevelopments. Such flexibility is further afforded by cloud based and federated learning which enables the model to improve without necessarily bringing sensitive data into a central place at the expense of incurring the cost of infrastructure and logistical overheads.

To attract community usage and adoption, the agricultural extension agencies, NGOs and local governments will be collaborated with. These alliances will make it easier to train, localize and raise awareness in farming communities. Such an approach to the system design allows it to continuously be updated and retrained on the model to provide it with increased precision and reliability over time and keep it relevant and practical to the farmers.

CHAPTER 6
CONCLUSION AND FUTURE WORK

6.1 Summary of the Study

The aim of this study was to find an efficient but high-performing deep learning algorithm applicable in the detection of mango leaf disease. In order to accomplish this, a Knowledge Distillation (KD) technique was used with a strong vision transformer (DeiT-S) as the teacher and a lightweight CNN (MobileNetV3-Small) as the student. Although the teacher model offered excellent representation of features and resulted in a high level of accuracy, its computational needs were not appropriate in real-time use in resource-limited settings. KD enabled the student architect to learn the useful knowledge in the teacher model and even attained the competitive accuracy as involving significantly fewer parameters and less inference time.

With a consistent experimental design, the distilled MobileNetV3-Small had excellent generalization and robustness with a variety of disease types, making it practical in an agricultural use case. To additionally increase interpretability and make users confident in AI, additional AI explainability methods (Grad-CAM) were used. These procedures went on to show that the predictions of this model were driven by disease-specific trends, and not incidental visual (noise).

Lastly, the optimized student model was integrated in a streamlit serviced web-based application, with friendly user interface and offline real-time prediction. This made sure that the rural farmer could also benefit in farming through connectivity generated by the system even in areas where connectivity was not readily available. This research, through its focus on uniting efficacy, knowledge transfer, interpretability, and real-world applicability, presents itself as a farmer focused, intelligent disease management tool to coalesce around precision agriculture.

6.2 Limitation

Though the suggested KD model performed admirably well in terms of accuracy and efficiency, a number of limitations of the current system are to be acknowledged. To begin with the model that is used, its performance is greatly influenced by the quality and variability of the training data. The combined (D3) dataset, although balancing on the eight classes, had been collected in specification of the regions of Bangladesh. This geographic biasness can limit the capacity of the model to generalize in other environmental circumstances, or outside the regions of the farm. Moreover, some minor misclassification corrections were made due to subtle similarities of some disease classes e.g. Powdery Mildew and Sooty Mould. These difficulties imply that, despite the good performance of KD to map wide disease patterns, additional strengths may be achieved by tuning it to bigger, more various and geographically diverse volumes of populations.

The absence of the mobile application implementation is one more limitation. The design does not provide the possibility of mobile use, but the output, i.e., lack of actual app, limits the system usability and access to end users, especially farmers who should be able to use the system offline and on the move.

Taken together and considered overall, these limitations point to possible future directions, namely, the need to create a mobile app and increase the size of the data to make this system more comprehensive, scalable, and adaptable to the needs of a changing agriculture.

6.3 Future Work

Based on the limitations and findings of the current investigation, some future possibilities are suggested that can be used to increase the functionality and impact of the system:

- Generalization: To make the Tiny-Net model more resistant, the next step will be to gather a larger and more extensive dataset based on several geographical zones and mango types. Such a solution will make the model more reliable to operate in various environmental conditions and under changes in leaf morphology that introduces some bias and decreases its real-life applicability.
- Disease severity estimation: In addition to predicting the type of disease, future models would also be predicting the stage of the infection related to severity. It would be possible to build in the severity estimation so that farmers could prioritize interventions based on the criticality of the disease and thereby be more effective at crop management.
- Further explainable AI: As Gad-CAM, LIME, and SHAP can be used, an open question in the future is whether an even more detailed explainable AI procedure, like Guided Grad-CAM++ or attention-based interpretability, can be applied. Such approaches can provide more detailed and local information on model predictions either at the pixel level or regional level.

- Mobile app development: As a future improvement, it has to be implemented to enhance availability and usability of a mobile version of the system. A mobile application would enable the farmer to run on field, real time disease detection even when offline, and thus the system would be more feasible and useful to farmers directly.

REFERENCES

[1] Bezabh, Y. A., Ayalew, A. M., Abuhayi, B. M., Demlie, T. N., Awoke, E. A., & Mengistu, T. E. *Classification of mango disease using ensemble convolutional neural network. Smart Agric Technol. 2024; 8: 100476*.

[2] Mahmud, B. U., Al Mamun, A., Hossen, M. J., Hong, G. Y., & Jahan, B. (2024). Light-weight deep learning model for accelerating the classification of mango-leaf disease. *Emerg. Sci. J*, *8*(1), 28-42.

[3] Ünal, Y., & Türkoğlu, M. (2025). Mango leaf disease detection using deep feature extraction and machine learning methods: A comparative survey. *El-Cezeri*, *12*(1), 35-43.

[4] Scientific, L. L. (2024). IDENTIFICATION AND CLASSIFICATION OF MANGO LEAVES DISEASE ON GREY WOLF OPTIMIZATION OF MULTIVARIATE GATED RECURRENT NETWORK. *Journal of Theoretical and Applied Information Technology*, *102*(23).

[5] Ramadan, S. T. Y., Sakib, T., Rahat, M. A., Mosharrof, S., Rakin, F. I., & Jahangir, R. (2023, December). Enhancing mango leaf disease classification: vit, bit, and cnn-based models evaluated on cyclegan-augmented data. In *2023 26th international conference on computer and information technology (ICCIT)* (pp. 1-6). IEEE.

[6] Pratap, V. K., & Kumar, N. S. (2024). Deep Learning based Mango Leaf Disease Detection for Classifying and Evaluating Mango Leaf Diseases. *Fusion: Practice & Applications*, *15*(2).

[7] Thaseentaj, S., & Ilango, S. S. (2024). Mango Leaf Disease Classification with Transfer Learning, Feature Localization, and Visual Explanations. *International Journal of Intelligent Engineering & Systems*, *17*(4).

[8] Faye, D., Diop, I., Mbaye, N., & Dione, D. (2023). A combination of data augmentation techniques for mango leaf diseases classification. *Glob. J. Comput. Sci. Technol*, 1-10.

[9] Saleem, S., Sharif, M. I., Sharif, M. I., Sajid, M. Z., & Marinello, F. (2024). Comparison of Deep Learning Models for Multi-Crop Leaf Disease Detection with Enhanced Vegetative Feature Isolation and Definition of a New Hybrid Architecture. *Agronomy, 14*(10), 2230.

[10] Budania, N. S. B. R. K., & Dhore, P. B. A Skillful Study of Mango Leaf Organization & Disease Finding By Using Deep Learning.

[11] Porna, S. B., Kabir, M. F., Rana, M. I. C., Sajol, M. S. I., Roy, T., Khan, M. A. U., ... & Bhavani, G. D. (2024, October). Hybrid Convolutional Neural Networks for Enhanced Detection of Mango Leaf Diseases. In *2024 IEEE 6th International Conference on Cybernetics, Cognition and Machine Learning Applications (ICCCMLA)* (pp. 547-552). IEEE.

[12] Thaseentaj, S., & Ilango, S. S. (2023). Deep Convolutional Neural Networks for South Indian Mango Leaf Disease Detection and Classification. *Computers, Materials & Continua*, *77*(3).

[13] Preeya V, A., Raja, S. P., Dhanalakshmi, B. K., Gururaj, H. L., Ravi, V., & Ravi, P. (2024). Perceptual Pigeon Galvanized Optimization of Multi-objective CNN on the Identification and Classification of Mango Leaves Disease. *The Open Agriculture Journal*, *18*(1).

[14] A. P.-M. and Machines and undefined 2019, “The pragmatic turn in explainable artificial intelligence (XAI),” SpringerA PáezMinds and Machines, 2019•Springer, vol. 29, no. 3, pp. 441–459, Sep. 2019, doi: 10.1007/s11023-019-09502-w.

[15] Radhakrishnan, M., Monish, N., Dev, P. S., Kesavan, N., & Thomas, N. S. (2025). Implementation of explainable Ai in deep learning methods for multiclass classification of plant diseases in mango lLeaves. *ELCVIA. Electronic letters on computer vision and image analysis*, *24*(1), 104-117.

[16] Khandelwal, S., Raut, A., Vyawahare, H., Theng, D., & Dhande, S. (2024). Optimizing Performance in Mango Plant Leaf Disease Classification through Advanced Machine Learning Techniques. *Engineering, Technology & Applied Science Research*, *14*(6), 18476-18480.

[17] Auni, A. R., & Sugiharti, E. (2024). Optimization of Mango Plant Leaf Disease Classification Using Concatenation Method of MobileNetV2 and DenseNet201 CNN Architectures. *Scientific Journal of Informatics*, *11*(4), 1023-1034.

[18] Mathur, P., Sheth, F., Goyal, D., & Gupta, A. (2024). Deep insight: Mathematical modeling and statistical analysis for mango leaf disease classification using advanced deep learning models. *Journal of Interdisciplinary Mathematics*, *27*(2), 317-342.

[19] Rayed, M. E., Jim, J. R., Islam, M. J., Mridha, M. F., Kabir, M. M., & Hossen, M. J. (2025). MangoLeafXNet: An Explainable Deep Learning Model for Accurate Mango Leaf Disease Classification. *IEEE Access*.

[20] Chowdhury, R. H., & Ahmed, S. (2024, December). MangoLeafViT: Leveraging Lightweight Vision Transformer with Runtime Augmentation for Efficient Mango Leaf Disease Classification. In *2024 27th International Conference on Computer and Information Technology (ICCIT)* (pp. 699-704). IEEE.

242-25-040

ORIGINALITY REPORT

15%	10%	9%	5%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Daffodil International University Student Paper	2%
2	www.mdpi.com Internet Source	2%
3	dspace.daffodilvarsity.edu.bd:8080 Internet Source	1%
4	arxiv.org Internet Source	<1%
5	Pushpa Choudhary, Sambit Satpathy, Arvind Dagur, Dhirendra Kumar Shukla. "Recent Trends in Intelligent Computing and Communication", CRC Press, 2025 Publication	<1%
6	Abdullah Al Noman, Amira Hossain, Anamul Sakib, Jesika Debnath et al. "ViX-MangoEFormer: An Enhanced Vision Transformer–EfficientFormer and Stacking Ensemble Approach for Mango Leaf Disease Recognition with Explainable Artificial Intelligence", Computers, 2025 Publication	<1%
7	www.ijsat.org Internet Source	<1%
8	Submitted to Meerut Institute of Engineering & Technology Student Paper	<1%
9	ddd.uab.cat Internet Source	<1%