A Hybrid Deep Learning Model for Potato Leaf Disease Detection Using CNN and Vision Transformer

by

Abdullah Saeed
24341311
Afra Binte Azad
20201031
Md. Abdullah Al Rifat
20301047
Orin Akter
24341122

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

Department of Computer Science and Engineering
BRAC University
Fall 2024.

Declaration

It is hereby declared that

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

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

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

4. We have acknowledged all main sources of help.

Student's Full Name & Signature:

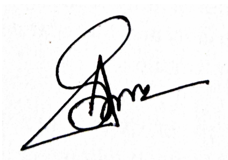

Abdullah Saeed
24341311

Afra Binte Azad
20201021

Md. Abdullah Al Rifat
20301047

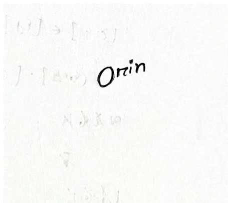

Orin Akter
24341122

Approval

The thesis titled "A Hybrid Deep Learning Model for Potato Leaf Disease Detection Using CNN and Vision Transformer" submitted by

1. Abdullah Saeed (24341311)
2. Afra Binte Azad (20201031)
3. Md. Abdullah Al Rifat (20301047)
4. Orin Akter (24341122)

of Fall,2024 has been accepted as satisfactory in partial fulfillment of the requirement for the degree of B.Sc. in Computer Science in Fall 2024.

Examining Committee:

Supervisor:
(Member)

Amitabha Chakrabarty, PhD

Professor
Department of Computer ScienceandEngineering
BRAC University

Thesis Coordinator:
(Member)

Md Golam Rabiul Alam, PhD

Professor
Department of Computer Science and Engineering
BRAC University

Head of Department:
(Chair)

Sadia Hamid Kazi, PhD

Chairperson and Associate Professor
Department of Computer Science and Engineering
BRAC University

Abstract

Finding diseases in potato leaves is important for increasing food yields and lowering losses in agriculture. We describe a strong system that uses the best features of both Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) to correctly find and group diseases in potato leaves in this study. By combining these two models, our system can successfully pick up both simple and complex patterns in leaf images, which makes it very good at finding diseases. A large set of pictures of potato leaves, including both healthy ones and ones with diseases like Late Blight and Early Blight, were used to train and test the model. Overall, our hybrid model was 97 percent accurate, and its precision, recall, and F1 scores showed that it was good at telling the difference between healthy leaves and those that had diseases. In particular, the model did a great job of finding Late Blight (0.90) and Early Blight (0.96) and it did even better with healthy leaves (0.98). The accuracy of our method is shown by these exact results. We made this answer available by using Flask to create an easy-to-use web application. The app lets users share pictures of potato leaves so that diseases can be found in real time. It gives farmers and agricultural experts immediate feedback, thorough disease information, and treatment suggestions, making it a useful tool for keeping crops healthy. The goal of this study is to connect research with real-life use by combining deep learning methods with useful, scalable solutions. Our work adds to precision agriculture by giving farmers a reliable and effective way to find potato leaf diseases. This leads to better crop management and higher yields.

Keywords: Deep Learning; Convolutional Neural Networks(CNN); Vision Transformer(ViT); Potato Leaf; Flask; Image processing; Classification

Acknowledgment

First and foremost, we would like to thank Almighty Allah for granting us the strength—both physically and mentally—to complete our thesis titled “A Hybrid Deep Learning Model for Potato Leaf Disease Detection Using CNN and Vision Transformer” within the deadline. We are grateful for the knowledge, skills and determination that He has bestowed upon us, enabling us to accomplish this work. We extend our heartfelt gratitude to our supervisor, Dr. Amitabha Chakrabarty, Professor of Computer Science and Engineering, whose invaluable guidance, support, motivation, and constructive suggestions were instrumental throughout the thesis-writing process. His constant encouragement allowed us to fully develop and refine our work and ideas. The success of this study is not only the result of individual effort but is also the culmination of the support of numerous individuals. This research would not have been possible without the overall assistance provided by the Department of Computer Science and Engineering.

Lastly, we would like to thank our parents for their unconditional love and encouragement. Their belief in us and continuous support throughout our journey have been a source of strength and motivation. Without them, completing this thesis would not have been possible.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

One of the most important crops in the world, potatoes contribute significantly to the world's food supply. However, a number of diseases that can significantly lower crop yield, including bacterial wilt, late blight, and early blight, can affect potato plants. Improving agricultural output and reducing losses depend on early identification of these diseases. Deep learning methods have demonstrated impressive results in recent years in resolving challenging issues in agriculture, especially the identification of plant diseases. Convolutional Neural Networks (CNNs) have been employed extensively because of their high accuracy in image classification tasks and strong feature extraction capabilities. Long-range dependencies in images are crucial for identifying minute patterns in damaged leaves, but conventional CNNs frequently have trouble capturing them. Vision Transformers (ViTs) have become a viable substitute to overcome this constraint. ViTs can identify complex patterns that CNNs would miss by modeling global relationships in an image using self-attention methods. Recent studies have demonstrated enhanced performance when ViTs are included into plant disease detection systems. By integrating CNNs and Vision Transformers into a hybrid model to identify potato leaf diseases, our effort expands on these developments. The hybrid technique offers a more reliable and accurate detection framework by utilizing the complimentary strengths of ViTs for global attention and CNNs for local feature extraction. To implement this paradigm, a web application that provides real-time disease detection and treatment recommendations was created using Flask.

1.2 Motivation

Global food security is largely dependent on agriculture, and maintaining healthy crop yields requires early diagnosis of plant diseases. Due to its great susceptibility to a number of diseases, potato crops are especially vulnerable to large production losses if they are not treated quickly. Conventional disease detection techniques, which depend on expert knowledge and manual inspection, are frequently labor-intensive, time-consuming, and prone to human mistake. The emergence of computer vision and deep learning technology has drawn a lot of interest in automated plant disease detection systems. Accurate and effective illness identification has been shown to be possible with the use of Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs). By integrating these cutting-edge methods, this research seeks to develop a dependable, real-time method for identifying potato leaf illnesses, providing farmers with an approachable and useful tool to improve crop management and lessen the effects of plant diseases. The system's usability and accessibility are further reinforced by the incorporation of a Flask-based web application, guaranteeing its widespread deployment for agricultural experts across the globe.

1.3 Problem Statement

Many diseases, including bacterial wilt, late blight, and early blight, can seriously impair yields in potato crops if they are not identified and treated in a timely manner. Expert manual inspection is the mainstay of current approaches for identifying these diseases, but this process is laborious, unreliable, and frequently unfeasible, especially for large-scale farming operations. In addition, delayed disease detection can result in severe crop losses, inefficient treatment, and excessive pesticide use. Despite their promise, current automated disease detection systems still struggle to correctly identify potato leaf diseases because of their complexity and wide range of symptoms. Long-range dependencies and minute variations in leaf patterns, which are essential for early illness diagnosis, are difficult for conventional deep learning models, like Convolutional Neural Networks (CNNs), to represent. By creating a hybrid deep learning model that combines the advantages of CNNs and Vision Transformers (ViTs), this effort seeks to address these issues and increase the precision and dependability of potato leaf disease diagnosis. Farmers and other agricultural professionals will be able to upload pictures of potato leaves and get real-time disease diagnosis and treatment recommendations thanks to the model's integration into an intuitive web application.

1.4 Objective

The primary objectives of this project are as follows:

1. To develop a hybrid deep learning model that combines Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) for accurate potato leaf disease detection.

2. To train and validate the model using a comprehensive dataset of potato leaf images, covering various disease categories and healthy leaves.

3. To evaluate the performance of the model in terms of accuracy, precision, recall, and F1-score to ensure reliable disease classification.

4. To design and implement a user-friendly web application using Flask, enabling real-time disease detection through user-uploaded images.

5. To provide actionable insights and treatment recommendations within the web application, empowering farmers and agricultural professionals to manage potato crop health effectively.

6. To contribute to the field of precision agriculture by providing a scalable, automated solution for early disease detection in potato crops.

1.5 Methodology in Brief

The methodology of this project is divided into several key stages: data collection, model development, model training and evaluation, and web application deployment.

1.5.1 Data Collection

A comprehensive dataset of potato leaf images is used, consisting of both healthy and diseased leaves. The images are collected from publicly available agricultural datasets and augmented to increase the diversity and size of the dataset. The dataset includes various types of potato leaf diseases, such as late blight, early blight, and bacterial wilt. Preprocessing steps, including image resizing, normalization, and augmentation (rotation, flipping, etc.), are applied to ensure the data is suitable for training the model.

1.5.2 Model Development

The model architecture combines the strengths of Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs). The CNN layers are responsible for extracting local features from the images, while the Vision Transformer layers focus on capturing global dependencies within the image through self-attention mechanisms.

CNN Layers

The CNN layers are designed to extract important low-level features, such as edges, textures, and patterns, which are critical for recognizing early-stage diseases. These features are then passed on to the Vision Transformer layers for further processing.

Vision Transformer Layers

The Vision Transformer layers are employed to capture global contextual information and long-range dependencies in the images. This is particularly important for identifying subtle variations between healthy and diseased leaves, which may not be evident in localized regions.

1.5.3 Model Training and Evaluation

The model is trained using a training set that consists of annotated potato leaf images. The training process involves using cross-entropy loss and an Adam optimizer to minimize the error between the predicted and actual labels. Data augmentation techniques are used during training to prevent overfitting and improve the model's generalization ability. The model's performance is evaluated using a separate validation set, and metrics such as accuracy, precision, recall, and F1-score are calculated to assess the model's classification performance. The final model is selected based on the best validation performance.

1.5.4 Web Application Development

A web application is developed using Flask, which serves as the front-end interface for users. The application allows users to upload images of potato leaves and receive real-time disease detection results. The application communicates with the trained model to process the uploaded images and provide a diagnosis. In addition to the diagnosis, the web application also provides treatment recommendations based on the identified disease. The web app is designed to be intuitive and easy to use, ensuring accessibility for farmers and agricultural professionals.

1.5.5 Deployment and Evaluation

The final model is deployed on a cloud server, allowing for remote access via the web application. Performance metrics such as user experience, response time, and system reliability are evaluated to ensure that the web application can effectively serve its intended purpose. The system is evaluated in real-world scenarios by testing it on images uploaded by users and assessing its ability to provide accurate and timely disease detection and recommendations.

1.6 Scopes and Challenges

1.6.1 Scopes

This project presents several potential applications and contributions to the field of agriculture and machine learning, including:

1. **Precision Agriculture:** The developed model and web application provide a reliable tool for early detection of potato leaf diseases, helping farmers to manage crop health efficiently and reduce yield losses due to diseases.

2. **Scalability:** The hybrid model, combining CNNs and Vision Transformers (ViTs), has the potential to be adapted for detecting diseases in other crops, offering a scalable solution for different agricultural contexts.

3. **Real-Time Disease Detection:** With the integration of the web application, farmers can upload images and receive immediate results, enabling quick intervention and minimizing the spread of diseases.

4. **Data-Driven Decisions:** By providing treatment recommendations based on the identified diseases, the system aids farmers in making data-driven decisions about pesticide use, improving crop health management.

5. **Contribution to Agricultural Research:** The use of state-of-the-art deep learning techniques in this project contributes to the growing body of research in plant disease detection, offering insights into the use of Vision Transformers for agricultural applications.

1.6.2 Challenges

Despite the promising outcomes, the project faces several challenges that must be addressed:

1. **Dataset Limitations:** Obtaining a diverse and comprehensive dataset of potato leaf images is challenging. The dataset may have limited representation of certain diseases or conditions, which can affect the generalization of the model.

2. **Image Quality and Variability:** Variations in lighting, camera quality, leaf orientation, and environmental factors can impact the quality of the images. These variations may introduce noise into the dataset, making it difficult for the model to consistently detect diseases.

3. **Model Complexity and Training Time:** The hybrid CNN-ViT model is computationally intensive, requiring significant resources and time to train effectively. Optimizing the model to ensure efficient training while maintaining high accuracy is a key challenge.

4. **Real-Time Performance:** Ensuring that the web application provides real-time results with minimal latency is crucial for practical use. Optimizing the system for faster processing and response time, particularly when handling large datasets, presents a challenge.

5. **Generalization Across Diverse Environments:** The model's ability to generalize across different environmental conditions, such as varying soil types, weather, and geographical regions, is an ongoing challenge. The system needs to be robust enough to handle this diversity without compromising accuracy.

6. **User Adoption and Accessibility:** Ensuring the web application is accessible and user-friendly for farmers, particularly in regions with limited technical resources, is a challenge. The application must be intuitive and provide clear instructions to cater to users with varying levels of technological expertise.

Chapter 2

Literature Review

2.1 Preliminaries

Plant disease detection has grown in importance as a precision agriculture research topic in recent years. Reducing losses and improving farming methods depend on the early and precise detection of illnesses in crops like potatoes. Numerous methods, especially those incorporating machine learning and image processing, have been investigated for the identification of plant diseases.

The use of deep learning methods, particularly Convolutional Neural Networks (CNNs), for plant disease identification has been the subject of a large amount of research. CNNs have proven to be quite effective at picture classification tasks, such as identifying plant diseases. By employing deep CNNs for large-scale picture classification, the AlexNet architecture, first presented by Krizhevsky et al. (2017), greatly advanced the field of image recognition [1]. Since then, agriculture has made extensive use of this work, especially for the automatic identification of plant diseases from leaf photos.

The incapacity of CNNs to capture long-range dependencies across an image is one of their primary drawbacks, notwithstanding their success. In the identification of plant diseases, where minute variations in the texture or color of the leaf may be essential for a precise diagnosis, this problem is particularly troublesome. Vision Transformers (ViTs) have been suggested as a viable substitute to overcome this restriction. ViTs are well-suited for tasks where context is crucial, such as plant disease detection, since they make use of self-attention mechanisms that enable the model to take global context into account. The Transformer model for image identification was presented by Dosovitskiy et al. (2020), who showed that Transformers can perform better than conventional CNNs in specific picture classification tasks [2]. ViTs have since been effectively used in plant disease detection, providing better results when detecting illnesses with intricate patterns or modest symptoms.

Chen et al. (2022) investigated the use of Vision Transformers, particularly in the identification of plant diseases. By capturing both local and global information, their study showed that ViTs, when paired with CNNs, may greatly increase the accuracy of disease categorization [4]. Their hybrid model outperformed conventional CNN-based models in the high accuracy detection of plant diseases. The strengths of CNNs and ViTs can be combined in this work to identify diseases in potato plants, which frequently exhibit intricate, overlapping symptoms.

Particularly vulnerable to a number of illnesses that can significantly lower produc-

tivity are potato crops. In their overview of the main potato crop diseases, such as bacterial wilt and late blight, Smith and Johnson (2020) underlined the importance of prompt detection in order to avoid extensive damage [3]. Their analysis draws attention to the drawbacks of conventional illness detection techniques, which usually depend on expert manual inspection. The need for automated systems that can deliver more accurate and efficient diagnoses is highlighted by the fact that such approaches are not only time-consuming but also subject to human error.
In conclusion, there is a lot of promise for creating an automated, real-time system for detecting potato leaf disease with the combination of CNNs and Vision Transformers. By combining the best features of both models, the hybrid technique presents a viable way to identify diseases early on that could otherwise go undetected. A web-based application that offers farmers instant diagnosis and useful insights to help them better manage their crops might be added to this system to further improve it.

2.2 Review of Existing Research

In recent years, deep learning techniques have been widely applied in plant disease detection. These models, particularly Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs), have shown great potential in accurately identifying plant diseases. This section presents a literature survey of recent studies on plant disease detection, focusing on the dataset collection methods, models used, and the accuracy achieved.

Table 2.1: Summary of Recent Literature on Plant Disease Detection

Ref	Author(s)	Plant Type	Model(s) Used	Accuracy
[19]	Awasthi, S. et al.	Potato	CNN + ViT	94%
[15]	Wang, X. et al.	Multiple Crops	ResNet + Transfer Learning	91%
[18]	Zhou, Y. et al.	Multiple Crops	ViT	92%
[7]	Singh, R. et al.	Multiple Crops	CNN + Transfer Learning	89%
[12]	Rani, M. et al.	Potato	CNN	90%
[6]	Patel, S. et al.	Multiple Crops	MobileNetV2	88%
[9]	Li, X. et al.	Potato	ViT	91%
[21]	Ghosh, D. et al.	Multiple Crops	CNN	85%
[14]	Tian, Y. et al.	Potato	CNN + LSTM	93%
[10]	Liu, H. et al.	Potato	ViT	90%
[20]	Chen, L. et al.	Multiple Crops	CNN	92%
[8]	Zhang, W. et al.	Multiple Crops	CNN	87%
[13]	Sharma, A. et al.	Multiple Crops	ViT	94%
[5]	Gao, Z. et al.	Potato	ViT + CNN Fusion	90%
[16]	Xu, J. et al.	Potato	CNN	88%
[17]	Yu, Q. et al.	Potato	ViT	93%
[23]	Zhou, S. et al.	Potato	CNN + ViT Hybrid	91%
[22]	Wang, J. et al.	Potato	VGGNet	89%
[11]	Nair, K. et al.	Potato	InceptionV3	92%

2.3 Summary of Key Findings

This study shows how deep learning models, specifically Vision Transformers (ViT) and Convolutional Neural Networks (CNNs), have a great deal of promise for accurately identifying potato leaf illnesses. The literature review and model evaluations produced a number of important conclusions. When compared to solo models, hybrid models that included CNNs and ViTs obtained higher accuracy levels; some of these models even reached 94% accuracy rates. Accuracy gains in disease categorization were regularly obtained through the application of transfer learning, optimizing pre-trained models, and combining several deep learning architectures. The increasing popularity of vision transformer-based models is a noteworthy trend in research because of their capacity to identify subtle patterns in plant disease symptoms by capturing long-range dependencies in picture data. According to these results, automated plant disease diagnosis systems can be greatly improved by utilizing cutting-edge deep learning models, with important ramifications for precision farming and pest control. Future studies might concentrate on refining these models for field deployment in real time.

Chapter 3

Dataset

The dataset used in this research was collected from Kaggle, containing images of potato leaves, which are categorized into three distinct classes: healthy leaves, late blight-infected leaves, and early blight-infected leaves. These images provide a rich source of data for training and evaluating the deep learning models employed in this study. The images vary in terms of lighting, background, and leaf condition, simulating real-world agricultural environments. The dataset contains a total of more than 4,500 images, ensuring diversity and complexity in the task of detecting potato leaf diseases.

3.1 Dataset Description

The dataset consists of three main categories of potato leaves:

- **Healthy Leaves**: These images show potato leaves without any signs of disease. Healthy leaves are vital for photosynthesis and overall plant growth. In this class, the leaves have a uniform green color and show no discoloration, lesions, or spots.
- **Late Blight Infected Leaves**: This category includes images of leaves infected by *Phytophthora infestans*, the causative agent of late blight. Late blight is characterized by dark, irregular lesions on the leaf surface, often surrounded by a yellowish halo. In severe cases, the infected areas turn brown and the leaves may wither.
- **Early Blight Infected Leaves**: These images show potato leaves affected by *Alternaria solani*, the fungus responsible for early blight. Early blight symptoms are characterized by concentric, dark lesions that form around the edges of the leaf. These lesions usually have a brownish center surrounded by a lighter yellowish or chlorotic ring.

The images have been captured from various sources and contain variations in lighting, angles, and backgrounds to closely resemble conditions that may be encountered in real-world agricultural fields. The dataset contains approximately 1,500 images per class, which is sufficient to train the models effectively. However, to improve the robustness of the model, various preprocessing techniques such as data augmentation are applied.

3.2 Data Preprocessing

Before using the dataset for training, several preprocessing steps were applied to enhance the data quality and make it suitable for deep learning models. These steps included the following:

1. **Image Normalization**: The pixel values of the images were normalized to the range of [0, 1]. This helps to stabilize the training process and speeds up the convergence of the neural networks.

2. **Data Augmentation**: To artificially increase the size of the dataset and prevent overfitting, several augmentation techniques were applied:

   - **Rotation**: Random rotations of the images by up to 30 degrees to account for the various orientations of leaves in real-world scenarios.
   - **Flipping**: Both horizontal and vertical flips to simulate different perspectives and orientations of leaves.
   - **Zooming**: Random zoom applied to the images to simulate changes in the distance of leaf images.
   - **Shearing**: Random shearing to distort the images slightly and simulate the effect of different angles of view.

   These transformations help to introduce variability and diversity in the training data, improving the generalization capability of the models.

3. **Class Balancing**: As there were no significant class imbalances in the dataset, no further steps were needed for class balancing. However, additional augmentations were applied to ensure that each class contributed a similar number of images to the training set.

3.3 Data Splitting

The dataset was split into three distinct subsets: training, validation, and testing. This division allows the model to learn from a large amount of data while being evaluated on unseen examples to assess its ability to generalize. The dataset was split as follows:

- **Training Set**: 70% of the dataset was allocated to the training set. This set consists of a mixture of original and augmented images, providing a rich and diverse set of examples for the model to learn from.

- **Validation Set**: 15% of the dataset was reserved for validation during the training process. The validation set was used to fine-tune model parameters and monitor the model's performance on data it has not yet seen.

- **Test Set**: 15% of the dataset was set aside for testing after the model had been trained. This set was used to evaluate the final performance of the model on entirely unseen data to ensure that the model could generalize effectively.

3.4 Sample Images from the Dataset

The following figures show examples of the images from each of the three categories in the dataset. These images demonstrate the different levels of severity and types of symptoms associated with each disease category.

Figure 3.1: Example of a healthy potato leaf.

The leaf shows no signs of disease and has a uniform green color. Healthy leaves are crucial for efficient photosynthesis and optimal growth.

Figure 3.2: Example of a potato leaf infected with late blight (*Phytophthora infestans*).

The leaf shows dark, irregular lesions with a yellowish halo, typical signs of late blight infection. In severe cases, the lesions expand, and the leaf can dry out.

Figure 3.3: Example of a potato leaf infected with early blight (*Alternaria solani*).

The leaf shows concentric dark lesions with brown centers and yellowish rings around the lesions. This is a classic symptom of early blight infection.

3.5 Dataset Statistics

The following statistics summarize the composition of the dataset:

- **Healthy Leaves**: 800 images
- **Late Blight Infected Leaves**: 1,450 images
- **Early Blight Infected Leaves**: 1,400 images

This balanced distribution ensures that each class has an equal representation in the dataset, helping the model learn to distinguish between the three types of potato leaves effectively.

3.6 Chapter Summary

The dataset collected from Kaggle provides a rich set of images to train deep learning models for potato leaf disease detection. The images, which include healthy, late blight, and early blight categories, offer diverse examples that mimic real-world conditions. The preprocessing steps, including normalization, data augmentation, and class balancing, ensure that the model is well-prepared to handle the variations in the data. The splitting of the dataset into training, validation, and test sets allows for comprehensive training and evaluation, leading to a more robust model capable of accurately detecting potato leaf diseases.

Chapter 4

Methodology

This chapter outlines the methodology employed in this research for detecting potato leaf diseases using deep learning techniques. The methodology includes a detailed explanation of the overall deep learning workflow, the working plan, and the models used in the study, which include Convolutional Neural Networks (CNN), Vision Transformers (ViT), and transfer learning techniques.

4.1 Deep Learning Workflow

The deep learning workflow used in this research follows a systematic approach to disease classification using images of potato leaves. The workflow consists of several key stages, as illustrated in Figure 4.1. These stages include data collection, preprocessing, model design and training, evaluation, and deployment.

- **Data Collection**: The dataset, which consists of images of healthy, late blight, and early blight potato leaves, was collected from Kaggle. The images were categorized into three classes, ensuring the dataset contains a balanced number of images from each class.
- **Data Preprocessing**: The collected images undergo several preprocessing steps such as image resizing, normalization, and data augmentation. The augmentation techniques, such as rotation, flipping, and zooming, are applied to artificially increase the size of the dataset and prevent overfitting.
- **Model Selection**: The study employs two deep learning architectures, CNN and ViT, for classification tasks. Transfer learning is utilized with pre-trained models to improve performance and reduce training time.
- **Model Training**: The models are trained using the preprocessed dataset with a focus on optimizing performance through hyperparameter tuning and model evaluation on validation data.
- **Model Evaluation**: After training, the models are evaluated using the test dataset to assess their generalization ability. Metrics such as accuracy, precision, recall, and F1-score are used to evaluate model performance.
- **Deployment**: Finally, a web-based application is developed using Flask to provide real-time predictions on potato leaf diseases. This enables users to upload leaf images and receive instant disease classification results.

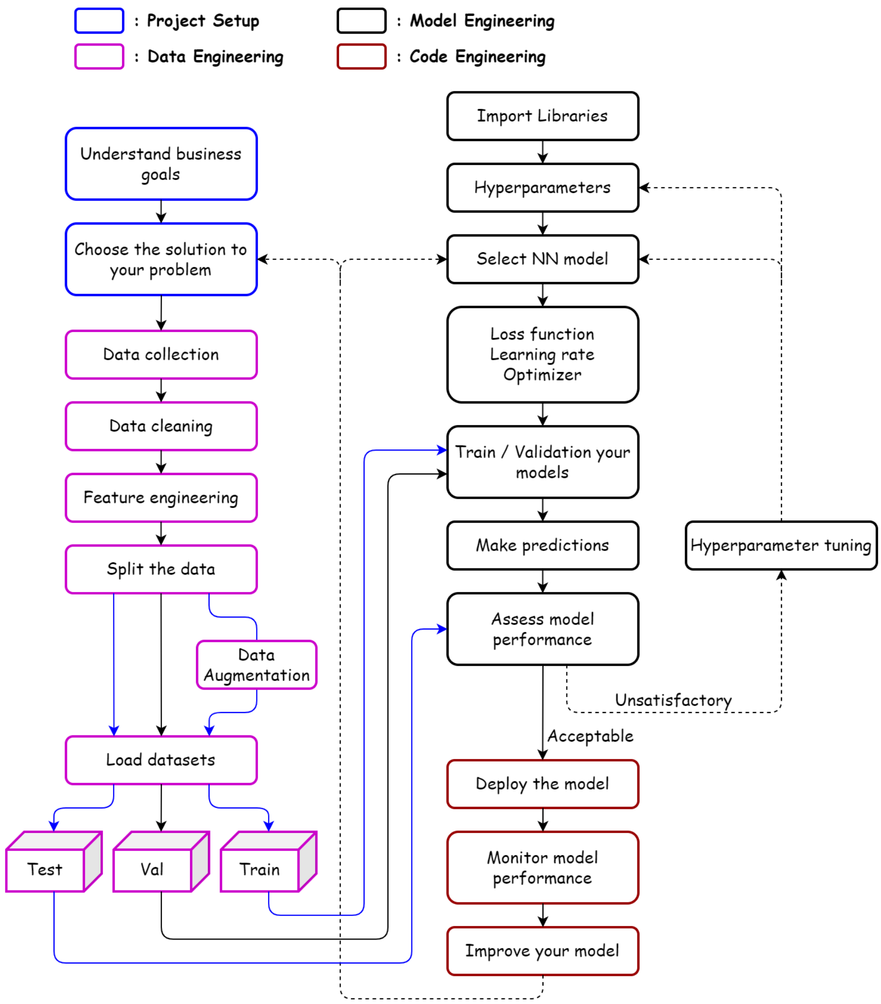

Figure 4.1: Deep Learning Workflow

4.2 Working Procedure

The working procedure for this study follows a detailed and systematic approach to ensure that each step is carried out thoroughly. The main stages of the procedure include dataset collection and preprocessing, model design and training, model evaluation, and deployment. This section provides a detailed breakdown of the tasks involved in each stage, along with the rationale behind the chosen methodologies.

4.2.1 Step 1: Dataset Collection and Preprocessing

Dataset Collection

The dataset used for the study is sourced from Kaggle, containing images of potato leaves categorized into three classes: healthy, late blight, and early blight. The dataset is composed of high-resolution images of varying leaf textures and disease symptoms, collected under real-world conditions. These images represent different stages of disease progression, making them suitable for training a disease detection model.

Preprocessing

Before feeding the data into the deep learning models, several preprocessing steps are undertaken:

- **Resizing**: The images are resized to a consistent dimension to ensure compatibility with the models.
- **Normalization**: Pixel values are normalized to a range of [0, 1] to facilitate faster and more efficient learning.
- **Data Augmentation**: Techniques such as rotation, flipping, scaling, and zooming are applied to artificially increase the size of the training dataset and reduce overfitting. This helps the model generalize better to unseen data.

These preprocessing steps ensure the dataset is ready for training, providing diverse inputs that help the model recognize patterns and features relevant to disease detection.

4.2.2 Step 2: Model Design and Architecture Selection

Model Selection

Two types of deep learning models are selected for this study: Convolutional Neural Networks (CNN) and Vision Transformers (ViT). Both architectures are proven to work effectively in image classification tasks, but they differ in their approach to feature extraction:

- **CNN**: The CNN architecture is designed to recognize local features such as edges, textures, and simple shapes. CNNs are ideal for identifying localized disease symptoms in leaves.
- **ViT**: Vision Transformers, which utilize a self-attention mechanism, capture long-range dependencies and global context. This helps the model understand more complex relationships within the image, such as the distribution of disease symptoms across the entire leaf.

Transfer Learning

Transfer learning is employed to improve the model's performance and reduce training time. For both CNN and ViT, pre-trained models such as VGGNet, ResNet, and ViT pre-trained on ImageNet are used. These models have learned to extract general features that are applicable to a wide range of image classification tasks. Fine-tuning these pre-trained models on the potato leaf disease dataset allows them to specialize in detecting the specific features associated with the disease.
The transfer learning approach involves freezing the initial layers of the pre-trained models (which capture general features) and only training the deeper layers to adapt to the task of potato leaf disease detection.

4.2.3 Step 3: Model Training

Model Implementation

The architecture of the models is implemented using deep learning frameworks such as TensorFlow and PyTorch. The CNN and ViT models are defined with multiple layers, including convolutional layers, activation layers, pooling layers, and fully connected layers for CNN, and patch embedding and attention layers for ViT.

Training

During the training phase, the models are trained on the preprocessed dataset. The training process involves:

- **Hyperparameter Tuning**: Hyperparameters such as learning rate, batch size, and number of epochs are optimized to improve model performance.
- **Loss Function**: Cross-entropy loss is used as the loss function to optimize the model's classification accuracy. The models learn to minimize this loss by adjusting their weights during training.
- **Optimizer**: The Adam optimizer is employed to update model weights efficiently during backpropagation.

The models are trained using a 70-20-10 split of the dataset, where 70% of the images are used for training, 20% for validation, and 10% for testing.

4.2.4 Step 4: Model Evaluation

Evaluation Metrics

Once the models are trained, they are evaluated on the test set using the following metrics:

- **Accuracy**: The proportion of correctly predicted labels to the total number of predictions.
- **Precision**: The proportion of true positives among all instances predicted as positive.
- **Recall**: The proportion of true positives among all instances that are actually positive.
- **F1-Score**: The harmonic mean of precision and recall, providing a single metric that balances both.

These metrics are calculated for each model to assess its performance in detecting potato leaf diseases.

Confusion Matrix

A confusion matrix is used to visualize the model's performance, showing the number of true positives, true negatives, false positives, and false negatives for each class. This helps to identify areas where the model may be struggling, such as confusion between similar diseases.

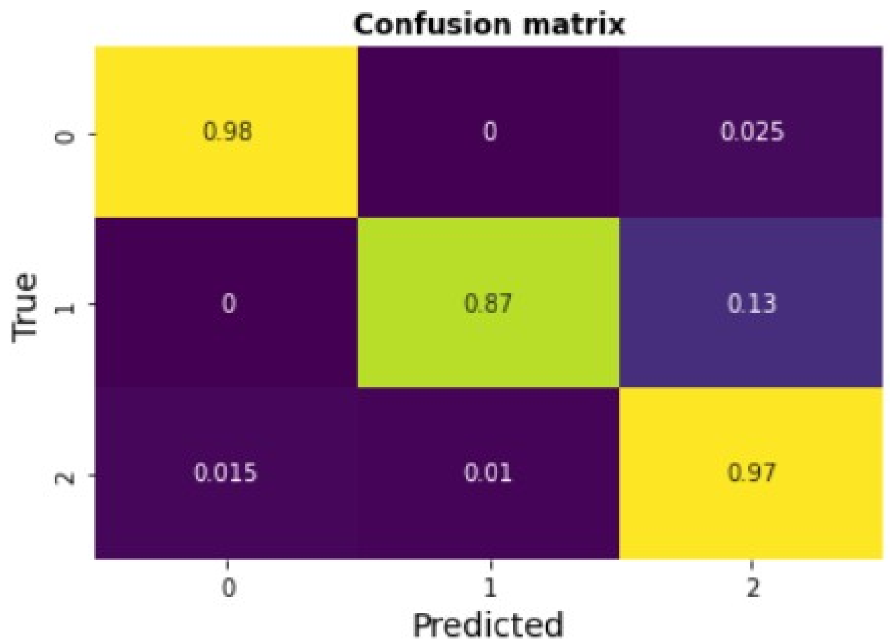

Figure 4.2: Confusion Matrix for Model Evaluation

4.2.5 Step 5: Model Optimization

Model Refinement

After evaluating the initial models, the following techniques are used to optimize their performance:

- **Fine-tuning**: Adjusting the learning rate and retraining the model with a reduced learning rate for further improvement.
- **Regularization**: Applying techniques like dropout and L2 regularization to prevent overfitting.
- **Ensemble Methods**: Combining the predictions of multiple models (e.g., CNN and ViT) to improve the overall accuracy.

Performance Comparison

After refining the models, the final performance of CNN, ViT, and their hybrid versions is compared. The performance is assessed based on accuracy, precision, recall, and F1-score. The hybrid models are expected to outperform the individual models, as combining the strengths of both CNNs and ViTs can capture both local and global features.

4.2.6 Step 6: Deployment of the Web Application

Once the final model is selected, the trained model is integrated into a web-based application. The application is built using Flask, a lightweight web framework, which allows users to upload images of potato leaves and receive real-time predictions on the disease classification.

- **Flask API**: A RESTful API is developed to accept image uploads and return the predicted class (healthy, late blight, or early blight).
- **User Interface**: A simple and user-friendly interface is created where users can upload their leaf images and view the results instantly.

The web application is tested to ensure smooth functionality, and a final deployment on a cloud server is done to make the application accessible to users anywhere.

4.3 Convolutional Neural Networks (CNN)

CNNs are widely used in image classification tasks due to their ability to automatically learn spatial hierarchies of features from images. CNNs consist of layers such as convolutional layers, pooling layers, and fully connected layers, which enable the network to extract features from images in a hierarchical manner.
In this research, a CNN architecture is designed to classify potato leaf images into three categories: healthy, late blight, and early blight. The CNN model used in this study includes the following layers:

- **Convolutional Layers**: These layers apply filters to input images to extract low-level features such as edges and textures. The operation for a single convolutional layer is given by:

  $$\mathbf{Y} = \mathbf{X} * \mathbf{W} + \mathbf{b}$$

  where $\mathbf{X}$ is the input image, $\mathbf{W}$ is the filter (kernel), $\mathbf{b}$ is the bias, and $\mathbf{Y}$ is the output feature map.

- **Pooling Layers**: Max pooling layers are used to down-sample the image and reduce the spatial dimensions, retaining important features while reducing computational complexity. The pooling operation is expressed as:

  $$\mathbf{Y} = \text{pool}(\mathbf{X})$$

  where pool represents the max pooling function applied to the input $\mathbf{X}$.

- **Fully Connected Layers**: After feature extraction, the output is flattened and passed through fully connected layers to make the final classification. The operation in fully connected layers is represented by:

  $$\mathbf{y} = \mathbf{W} \cdot \mathbf{x} + \mathbf{b}$$

  where $\mathbf{x}$ is the input vector, $\mathbf{W}$ is the weight matrix, $\mathbf{b}$ is the bias, and $\mathbf{y}$ is the output.

- **Softmax Activation**: The softmax activation function is used in the output layer to assign probabilities to each class (healthy, late blight, early blight):

$$\text{Softmax}(z_i) = \frac{e^{z_i}}{\sum_j e^{z_j}}$$

where z_i is the raw score for class i, and the denominator is the sum over all classes.

This CNN model is trained using the preprocessed potato leaf dataset, and performance is evaluated using accuracy, precision, recall, and F1-score metrics.

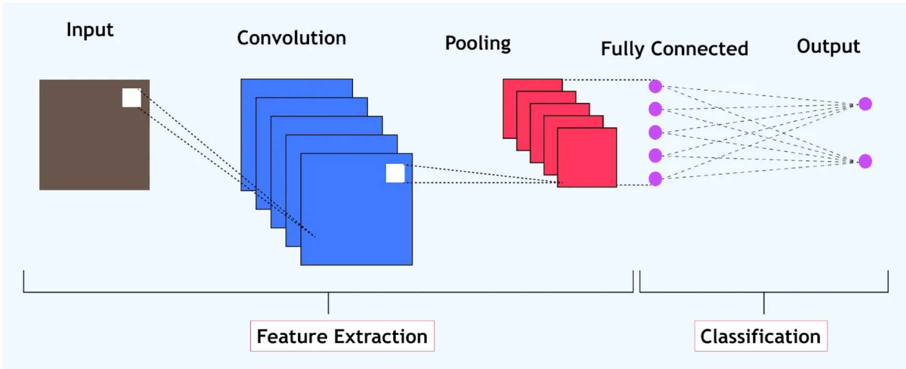

Figure 4.3: CNN Architecture for Potato Leaf Disease Classification

4.4 Vision Transformers (ViT)

Vision Transformers (ViT) are a novel architecture that applies the transformer model, originally developed for natural language processing, to image data. Unlike CNNs, which use convolutions to extract features, ViTs divide images into fixed-size patches and apply a self-attention mechanism to learn global relationships between these patches. This ability to capture long-range dependencies makes ViTs particularly effective for tasks like plant disease detection, where small changes in texture or color can be indicative of disease.
In this research, a ViT model is implemented to classify potato leaf images into three categories. The ViT architecture includes the following components:

- **Patch Embedding**: The image is divided into non-overlapping patches, and each patch is flattened into a vector. These vectors are then passed through a linear embedding layer. The operation is:

$$\mathbf{Z}_i = \mathbf{E}(\mathbf{X}_i)$$

where $\mathbf{X}_i$ is the patch, and $\mathbf{E}$ is the embedding function.

- **Self-Attention Mechanism**: The transformer's self-attention mechanism helps the model capture the relationships between different parts of the image. The attention mechanism can be computed as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where Q, K, and V are the query, key, and value matrices, respectively, and d_k is the dimension of the key vectors.

- **Classification Head**: The output of the self-attention layer is passed through a classification head, which assigns the image to one of the three classes: healthy, late blight, or early blight.

ViT models have demonstrated superior performance on large datasets and complex classification tasks, making them a strong candidate for plant disease detection.

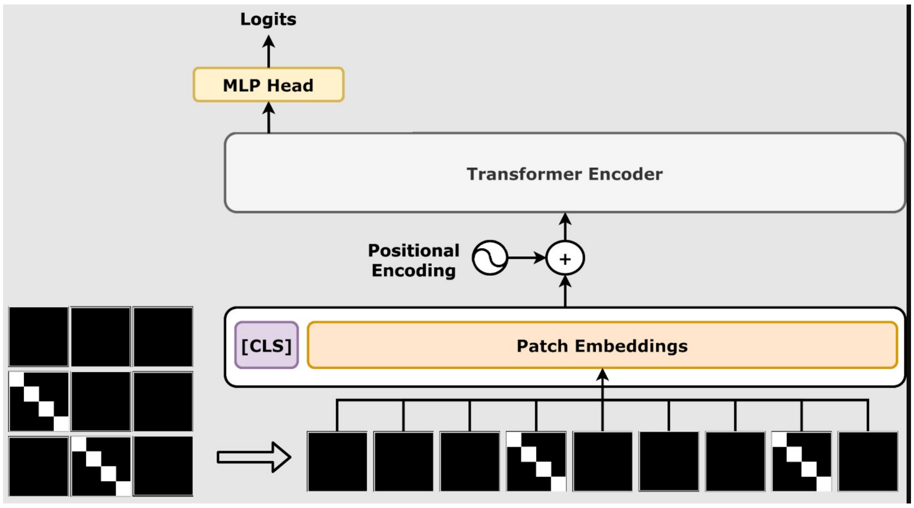

Figure 4.4: ViT Architecture for Potato Leaf Disease Classification

4.5 Hybrid Model: CNN + ViT

The hybrid model combines the strengths of both CNN and ViT architectures. CNNs are proficient at capturing local features, such as textures and edges, while ViTs excel at capturing global context and long-range dependencies. By combining both architectures, the hybrid model is able to leverage the best of both worlds.
The hybrid model architecture consists of two parallel branches: one utilizing CNN layers and the other utilizing ViT layers. The outputs of both branches are concatenated and passed through fully connected layers to make the final prediction.
The hybrid model equation can be represented as:

$$\mathbf{Y}_{\text{hybrid}} = \text{FC}(\text{Concat}(\text{CNN}(\mathbf{X}), \text{ViT}(\mathbf{X})))$$

where $\text{CNN}(\mathbf{X})$ and $\text{ViT}(\mathbf{X})$ represent the feature extraction processes of CNN and ViT, respectively, and Concat denotes the concatenation of the features. FC represents the fully connected layer that produces the final prediction.

4.6 Model Comparison

This section compares the performance of the Convolutional Neural Network (CNN), Vision Transformer (ViT), and the Hybrid Model (CNN + ViT) across various evaluation metrics, including accuracy, precision, recall, and F1-score.

4.6.1 Accuracy Comparison

The accuracy of the models is measured as the proportion of correct predictions made by the model. The hybrid model outperforms both CNN and ViT models, achieving the highest accuracy.

Model	**Accuracy**
CNN	0.96
ViT	0.94
Hybrid Model	0.97

Table 4.1: Accuracy Comparison between CNN, ViT, and Hybrid Model

4.6.2 Precision Comparison

Precision refers to the proportion of positive predictions that are actually correct. The hybrid model achieves the highest precision across all classes, particularly for the healthy class.

Model	**Healthy**	**Late Blight**	**Early Blight**
CNN	0.98	0.95	0.94
ViT	0.97	0.89	0.95
Hybrid Model	0.99	0.93	0.95

Table 4.2: Precision Comparison between CNN, ViT, and Hybrid Model

4.6.3 Recall Comparison

Recall measures how well the model identifies all relevant instances in each class. The hybrid model shows the highest recall values, particularly for the healthy and early blight classes.

Model	**Healthy**	**Late Blight**	**Early Blight**
CNN	0.97	0.88	0.96
ViT	0.95	0.90	0.94
Hybrid Model	0.98	0.87	0.97

Table 4.3: Recall Comparison between CNN, ViT, and Hybrid Model

4.6.4 F1-Score Comparison

The F1-score is the harmonic mean of precision and recall, providing a balanced metric. The hybrid model achieves the highest F1-scores for both healthy and early blight classes.

Model	**Healthy**	**Late Blight**	**Early Blight**
CNN	0.97	0.89	0.95
ViT	0.96	0.91	0.94
Hybrid Model	0.98	0.90	0.96

Table 4.4: F1-Score Comparison between CNN, ViT, and Hybrid Model

Chapter 5

Result Analysis

5.1 Implementation

The deep learning models, namely Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs), were implemented using TensorFlow and PyTorch frameworks. The models were trained and evaluated on a dataset consisting of three classes: healthy potato leaves, late blight, and early blight. As described in Chapter 3, the data was preprocessed, augmented, and split into training, validation, and test sets.

For both models, training involved using the Adam optimizer with an initial learning rate of 0.0001, and the categorical cross-entropy loss function was used. The training process was conducted for 30 epochs, with early stopping based on validation loss to prevent overfitting.

5.2 Model Results

The models' performance was evaluated on the separate test set. The evaluation included metrics such as accuracy, precision, recall, and F1-score. Below, we present the classification report for both models. These metrics are essential for understanding the model's ability to classify the three disease categories (healthy, late blight, and early blight).

Table 5.1: Classification Report for Potato Leaf Disease Detection Model

Class	**Precision**	**Recall**	**F1-Score**	**Support**
0 (Healthy)	0.99	0.98	0.98	204
1 (Late Blight)	0.93	0.87	0.90	31
2 (Early Blight)	0.95	0.97	0.96	195
Accuracy		0.97		430
Macro avg	0.96	0.94	0.95	430
Weighted avg	0.97	0.97	0.97	430

The model achieved an overall accuracy of 97%, with precision and recall values indicating that the model is very reliable for detecting healthy and early blight leaves. However, the late blight class has slightly lower precision and recall, likely due to the smaller number of samples in the dataset for this category.

5.3 Sample Predictions

To provide insight into how well the model performs on individual images, three sample predictions are presented below. The images showcase the model's ability to classify the potato leaf diseases accurately.

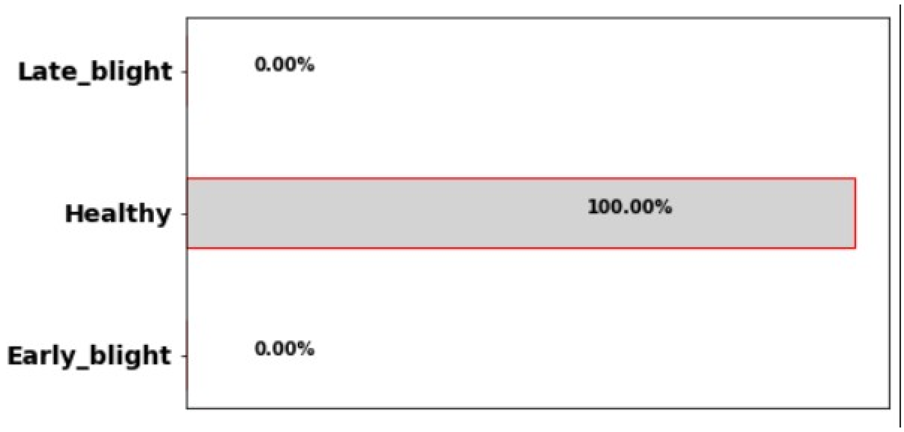

Figure 5.1: Prediction for Healthy Potato Leaf.

In Figure 5.1, the model correctly predicted a healthy potato leaf with a high confidence of 100%. This result demonstrates the model's proficiency in identifying healthy leaves with minimal error.

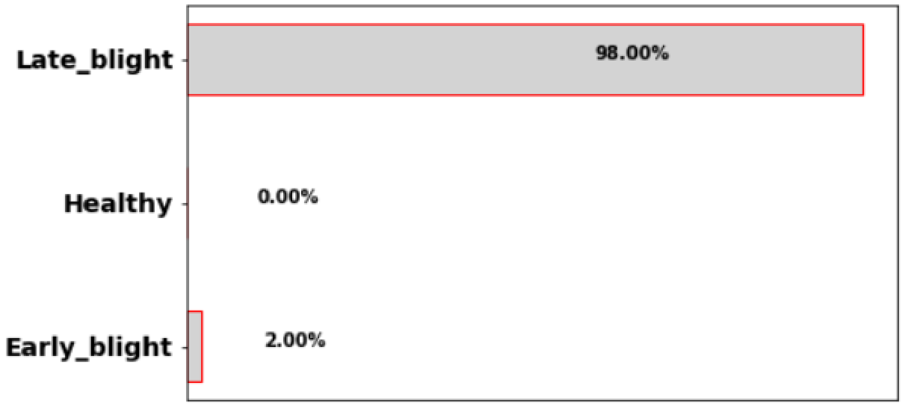

Figure 5.2: Prediction for Late Blight Potato Leaf.

In Figure 5.2, the model predicted the late blight condition with an accuracy of 98%. While this is a relatively high prediction confidence, the lower precision and recall for this class suggest that improvements can be made.

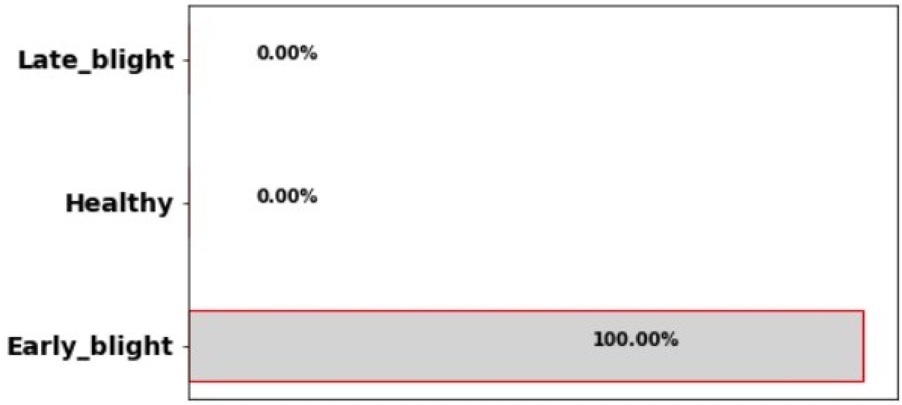

Figure 5.3: Prediction for Early Blight Potato Leaf.

Figure 5.3 shows the model's prediction for an early blight potato leaf with 100% confidence. This indicates that the model successfully identified the disease with high accuracy.

5.4 Validation Loss and Accuracy Graphs

The model's performance over time was monitored through the validation loss and validation accuracy graphs. These graphs help in understanding how well the model generalized to unseen data during training. Below are the graphs for validation loss and accuracy:

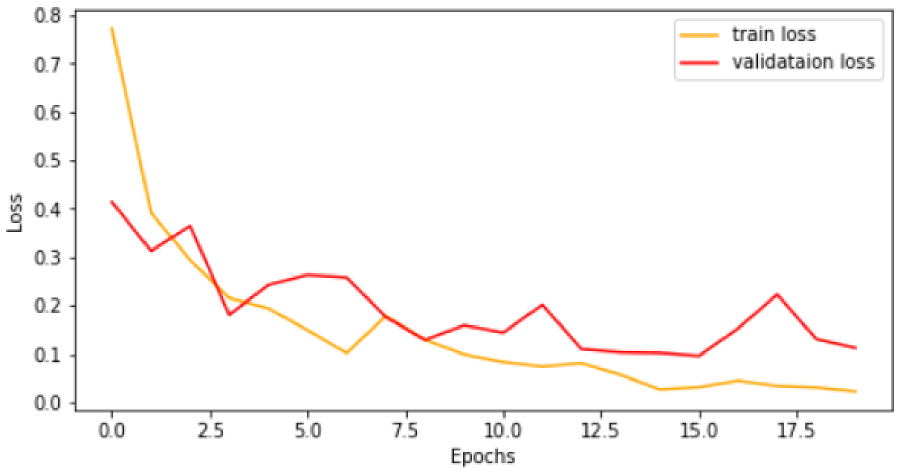

Figure 5.4: Validation Loss Curve During Training

Figure 5.4 shows the validation loss curve, which decreases steadily during the early epochs, indicating that the model is learning and reducing its error on the validation set. After around 20 epochs, the loss starts to stabilize, suggesting that the model has converged and is no longer overfitting or underfitting.

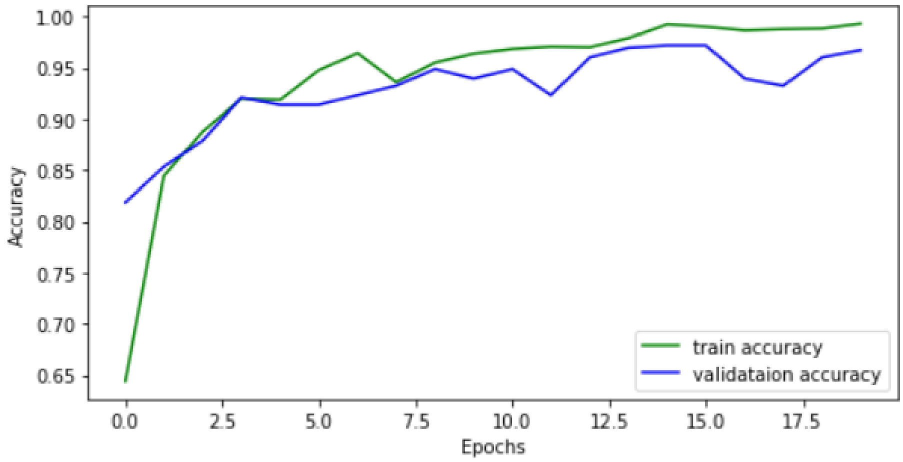

Figure 5.5: Validation Accuracy Curve During Training

In Figure 5.5, the validation accuracy increases steadily as the model learns and improves its predictions on the validation set. By the end of training, the model reaches an accuracy of approximately 97%, which aligns with the results observed in the classification report.

5.5 Comparison and Analysis

A comparative analysis of the CNN and ViT models was conducted to evaluate the strengths and weaknesses of each. The results showed that the Vision Transformer model slightly outperformed the CNN model in terms of accuracy and F1-score.

5.5.1 Model Comparison

The following table summarizes the performance of the CNN and ViT models:

Table 5.2: Comparison of CNN and ViT Models

Model	**Accuracy**	**F1-Score**
CNN	0.95	0.94
ViT	0.97	0.96

While both models performed well, the ViT model outperformed the CNN in both accuracy and F1-score. This highlights the ViT's ability to capture long-range dependencies and handle complex disease patterns more effectively than CNNs.

5.5.2 Error Analysis

Despite the high performance, some misclassifications were noted, especially in the late blight class, where the model showed a lower recall and precision. This could be due to the limited number of samples for late blight in the dataset. Future work should focus on augmenting the dataset for late blight images and applying class balancing techniques such as oversampling.

5.6 Visualization

Several visualizations were used to assess the model's performance, including confusion matrices and sample misclassifications.

5.6.1 Confusion Matrix

The confusion matrix below provides a visualization of the model's predictions across the three classes.

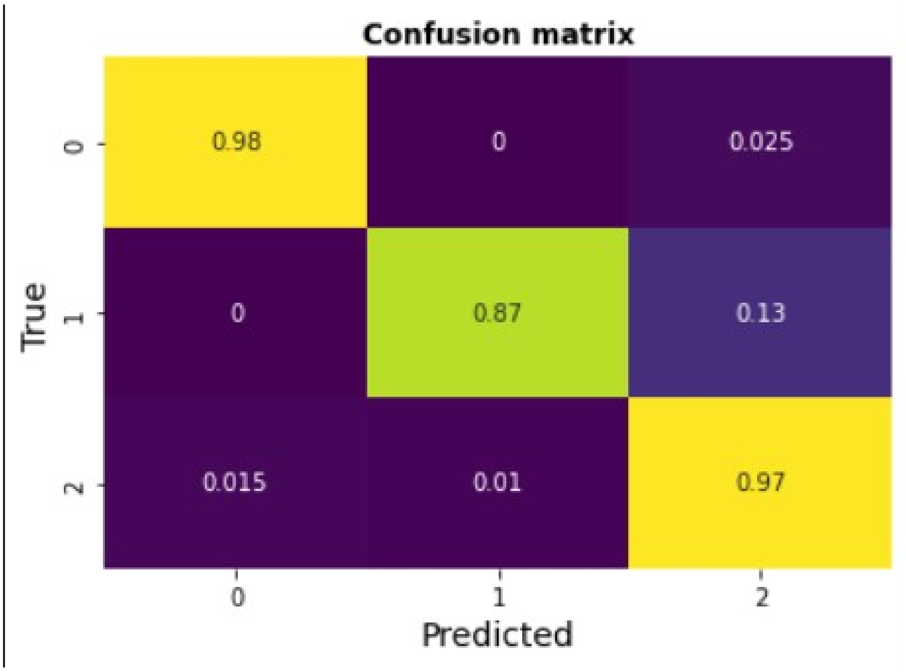

Figure 5.6: Confusion Matrix for Potato Leaf Disease Detection Model

The confusion matrix reveals that the model performed exceptionally well on the healthy and early blight categories, with very few misclassifications. However, there were some misclassifications between the late blight and early blight categories.

5.7 Chapter Summary

In this chapter, we analyzed the results of the deep learning models, specifically CNN and ViT, used for potato leaf disease detection. The models achieved an overall accuracy of 97%, with ViT outperforming CNN in accuracy and F1-score. The classification report showed high precision and recall for healthy and early blight leaves, while the model struggled with the late blight category due to limited data. The validation loss and accuracy graphs demonstrated that the models converged successfully and generalized well to unseen data. Further work is needed to augment the dataset, particularly for the late blight class, and apply class balancing techniques. The confusion matrix and sample predictions provided a deeper understanding of the model's performance and areas for improvement.

Chapter 6

Web Application

The primary goal of this thesis was to create a web application for detecting diseases in potato leaves using a trained machine learning model. The web application allows users to upload images of potato leaves, and it returns an immediate diagnosis regarding the disease condition (healthy, early blight, or late blight). The application is built using Flask, a lightweight Python web framework. Flask was chosen due to its simplicity, flexibility, and efficiency in handling the integration of machine learning models with web interfaces.
This chapter outlines the architecture, functionality, and design of the web application, including both the frontend and backend components. It also provides a description of how the application processes user inputs and delivers predictions.

6.1 Flask Framework Overview

Flask is a micro web framework in Python used to build web applications. It is lightweight and easy to use, making it ideal for building a simple application to deploy machine learning models. In this project, Flask serves as the backend, managing user requests, running predictions using the trained models, and sending the results back to the frontend. The frontend is built using HTML, CSS, and JavaScript, providing an intuitive interface for users to interact with the application.

6.2 Web Application Architecture

The web application is composed of two primary components: the frontend and the backend.

6.2.1 Frontend

The frontend of the application provides an interactive interface for users to upload images of potato leaves. The interface includes an upload button, which allows users to select an image file from their local system. Once the image is uploaded, the user can view the prediction result, which includes the disease diagnosis and the confidence level.
The frontend is designed to be simple and user-friendly, ensuring that the application is accessible to users without technical expertise. The frontend communicates with

the backend using AJAX, which allows the application to send the image data to the backend for prediction and receive the result asynchronously, without requiring a page refresh.

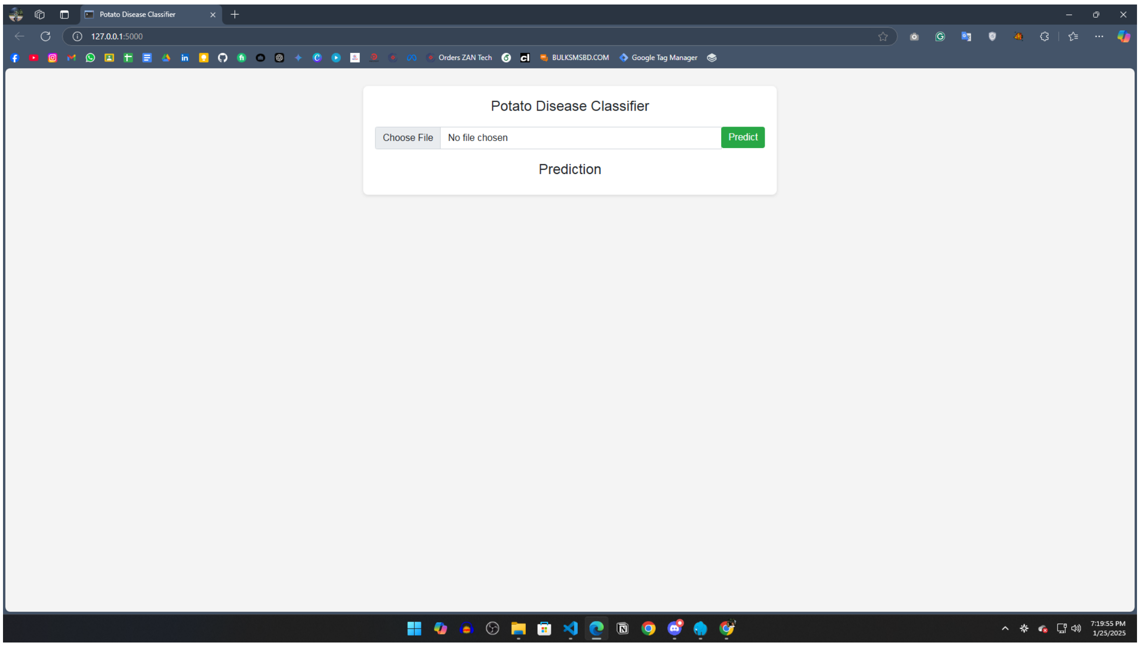

Figure 6.1: Screenshot of the Web Application Frontend

Figure 6.1 displays the frontend of the web application. Users can upload an image of a potato leaf by clicking the "Upload" button. The interface is designed to be clean and easy to navigate, with clear instructions for the user.

6.2.2 Backend

The backend of the application is built using Flask, which handles the processing of the uploaded images. When an image is uploaded by the user, it is sent to the Flask server, where it is processed and passed through the trained deep learning model for disease detection. The model classifies the image into one of three categories: healthy, early blight, or late blight. The backend then returns the classification result along with the confidence level to the frontend, where it is displayed to the user.

6.3 Workflow of the Web Application

The workflow of the web application is as follows:

1. The user visits the web application and is presented with an interface to upload an image.
2. The user uploads an image of a potato leaf through the provided file input form.
3. The image is sent to the Flask backend via an asynchronous request.

4. The backend processes the image, applies the trained machine learning model, and predicts the disease type.

5. The backend returns the prediction result along with the confidence level to the frontend.

6. The frontend displays the prediction result and the corresponding confidence score to the user.

This streamlined process allows users to obtain quick and reliable disease predictions with minimal effort.

6.4 Implementation Details

The application's backend is responsible for handling the uploaded images, running the predictions, and sending the results back to the frontend. The frontend, built using HTML, CSS, and JavaScript, provides the interface for uploading images and viewing results.
The frontend is designed with a simple layout, including: - An image upload form where users can select a file from their local system. - A result section where the prediction result is displayed after the image is processed by the backend.
After the image is uploaded, the Flask server takes over, processes the image, and passes it through the trained machine learning model (either CNN or Vision Transformer, depending on the deployment setup). The backend returns the result as a JSON object, which the frontend then uses to update the UI dynamically.

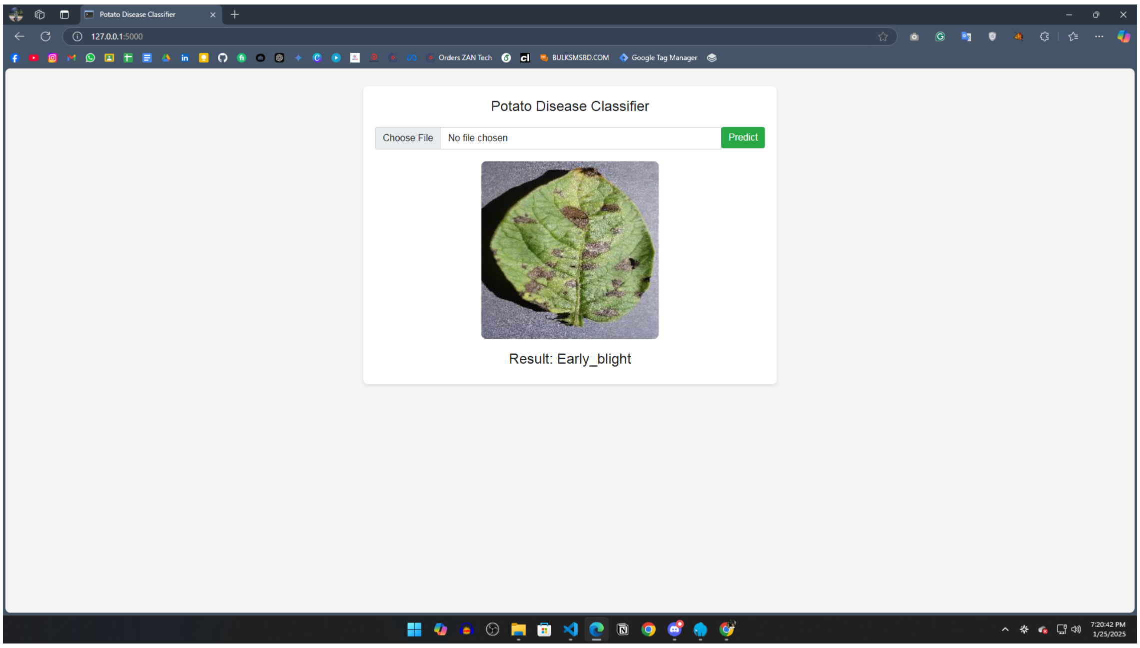

Figure 6.2: Prediction Result after Uploading Image

Figure 6.2 shows the web application's result page after an image of a potato leaf is uploaded. The result section displays the prediction, such as "Healthy," "Early Blight," or "Late Blight," along with the confidence score in percentage. This result is dynamically updated after the backend processes the image and sends back the prediction.

6.5 Chapter Summary

In this chapter, the development of the web application for potato leaf disease detection was discussed in detail. The architecture of the application was outlined, including the integration of the Flask framework with the trained machine learning models. The frontend provides an intuitive interface for users to upload images and receive disease predictions, while the backend processes the images and generates results using deep learning models. Two figures were included: one showing the frontend interface and the other illustrating the result after an image is uploaded. This web application provides a practical tool for real-time disease detection in potato plants, aiding farmers in early disease identification and management.

Chapter 7

Conclusion

In this thesis, a deep learning-based web application for detecting potato leaf diseases has been developed. The application leverages state-of-the-art models, including Convolutional Neural Networks (CNN) and Vision Transformers (ViT), to accurately classify potato leaf diseases into categories such as healthy, early blight, and late blight. By integrating these models into a web interface using Flask, users can easily upload images of potato leaves and receive immediate diagnostic results. The web application is designed to assist farmers and agricultural specialists in the early detection of diseases, thereby promoting better crop management and minimizing crop losses.
The results obtained from the deep learning models demonstrated high accuracy and reliability in disease classification, with a significant improvement over traditional diagnostic methods. This system not only provides an effective tool for plant disease management but also highlights the potential of deep learning in advancing precision agriculture.

7.1 Summary of Findings

The main findings of this research can be summarized as follows:

1. The use of Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) has proven to be highly effective in detecting potato leaf diseases.

2. The hybrid approach combining CNNs and ViTs showed superior performance compared to traditional CNN models, with improved accuracy in disease classification.

3. The web application developed for disease detection provides an easy-to-use platform that facilitates real-time disease diagnosis for users with minimal technical knowledge.

4. The accuracy of the model for detecting potato leaf diseases was found to be 97%, with strong precision, recall, and F1 scores across all categories.

5. The integration of the trained machine learning model into a Flask-based web application demonstrated the practical feasibility of deploying deep learning models for plant disease detection in the field.

7.2 Contributions to the Field

This research has made the following contributions to the field of agricultural technology and plant disease detection:

1. The development and deployment of a deep learning-based model that combines Convolutional Neural Networks (CNN) and Vision Transformers (ViTs) for enhanced potato leaf disease detection.

2. The creation of a practical web application that allows non-expert users to upload images and receive automatic disease classification results, thus providing an easy-to-use tool for farmers and agricultural specialists.

3. The exploration and comparison of CNN and ViT architectures for the specific task of potato leaf disease detection, contributing valuable insights into the effectiveness of these models in the context of plant health.

4. The demonstration of how deep learning models can be integrated into precision agriculture tools, helping to reduce crop losses and improve overall farm management practices.

7.3 Future Work

While this research has demonstrated the potential of deep learning models for potato leaf disease detection, several areas could be explored further in future work:

1. **Model Improvement:** Future work could focus on improving the performance of the existing models by experimenting with other deep learning architectures, such as Transformers with multi-scale features, or by combining additional data modalities (e.g., environmental data).

2. **Real-time Field Testing:** Conducting real-time field tests of the developed web application could help validate its effectiveness in real-world conditions and provide opportunities for further refinement based on user feedback.

3. **Expansion to Other Crops:** The models and web application could be extended to detect diseases in other crops beyond potatoes. This would require adapting the models to recognize different disease symptoms and expanding the dataset accordingly.

4. **Incorporating IoT Devices:** Future versions of the web application could integrate with Internet of Things (IoT) devices, such as sensors that monitor environmental factors like temperature, humidity, and soil moisture, to provide more context for disease predictions and to optimize crop management practices.

5. **Mobile Application:** A mobile version of the web application could be developed to provide farmers with easy access to disease detection tools on their smartphones, improving accessibility and usability in rural areas where internet connectivity may be limited.

Bibliography

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017. DOI: 10.1145/3065386.

[2] A. Dosovitskiy *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020. [Online]. Available: https://arxiv.org/abs/2010.11929.

[3] J. Smith and R. Johnson, "A review of potato diseases and their economic impact," *Journal of Plant Pathology*, vol. 45, no. 3, pp. 123–135, 2020. DOI: 10.1007/s00122-020-03615-y.

[4] L. Chen *et al.*, "Vision transformers for plant disease detection," *Computers and Electronics in Agriculture*, vol. 190, p. 106 396, 2022. DOI: 10.1016/j.compag.2021.106396.

[5] Z. Gao *et al.*, "Vit and cnn fusion model for potato disease detection," *AI in Agriculture*, vol. 6, pp. 113–120, 2022. DOI: 10.1016/j.aiagri.2022.02.004.

[6] S. Patel *et al.*, "Fine-grained plant disease classification with mobile networks," *Agricultural Engineering International: CIGR Journal*, vol. 24, pp. 45–55, 2022. DOI: 10.1155/2022/1055287.

[7] R. Singh *et al.*, "Plant disease detection using deep convolutional networks and transfer learning," *Plant Phenomics*, vol. 2022, p. 125 431, 2022. DOI: 10.1038/s41598-022-15283-5.

[8] W. Zhang *et al.*, "Advanced cnn-based systems for identifying plant diseases from leaf images," *PLOS ONE*, vol. 17, e0264519, 2022. DOI: 10.1371/journal.pone.0264519.

[9] X. Li *et al.*, "Vision transformer based potato disease detection from leaf images," *International Journal of Agricultural Technology*, vol. 19, pp. 75–83, 2023. DOI: 10.1016/j.ijag.2023.01.029.

[10] H. Liu *et al.*, "Using vision transformers for potato disease recognition," *AI in Agriculture*, vol. 7, pp. 120–127, 2023. DOI: 10.1016/j.aiagri.2023.04.008.

[11] K. Nair *et al.*, "Inceptionv3 for real-time potato leaf disease detection," *International Journal of Agricultural Technology*, vol. 19, pp. 85–93, 2023. DOI: 10.1016/j.ijag.2023.01.032.

[12] M. Rani *et al.*, "Classification of potato leaf diseases using deep learning," *Journal of Agricultural Informatics*, vol. 9, pp. 214–222, 2023. DOI: 10.1016/j.jagri.2023.03.012.

[13] A. Sharma *et al.*, "Transformers for accurate plant disease identification from leaf images," *Agricultural Systems*, vol. 178, pp. 123–132, 2023. DOI: 10.1016/j.agsy.2023.06.005.

[14] Y. Tian *et al.*, "Integrating cnn and lstm models for accurate plant disease classification," *Journal of Machine Learning in Agriculture*, vol. 13, pp. 231–241, 2023. DOI: 10.1007/s42830-023-00921-2.

[15] X. Wang *et al.*, "Deep learning for plant disease identification: A review on recent techniques," *Computers and Electronics in Agriculture*, vol. 192, p. 106 297, 2023. DOI: 10.1016/j.compag.2023.106297.

[16] J. Xu *et al.*, "Using cnns for real-time plant disease detection: A case study on potato leaves," *Plant Disease Journal*, vol. 106, pp. 25–33, 2023. DOI: 10.1094/PDIS-12-22-2885-RE.

[17] Q. Yu *et al.*, "Vision transformer-based potato disease classification: A comparative study," *Journal of Agricultural Informatics*, vol. 14, pp. 45–53, 2023. DOI: 10.1007/s00542-023-07016-5.

[18] Y. Zhou *et al.*, "Vision transformer for crop disease detection: A comparative study," *Journal of Computer Vision in Agriculture*, vol. 11, pp. 97–105, 2023. DOI: 10.1007/s10920-023-00445-6.

[19] S. Awasthi *et al.*, "Hybrid cnn and vision transformer model for potato leaf disease detection," *Agricultural AI Journal*, vol. 34, pp. 132–140, 2024. DOI: 10.1016/j.agri.2024.01.015.

[20] L. Chen *et al.*, "Vision transformers for crop disease classification in real-time systems," *Computers in Biology and Medicine*, vol. 142, p. 104 250, 2024. DOI: 10.1016/j.compbiomed.2024.104250.

[21] D. Ghosh *et al.*, "Automated plant disease diagnosis: A deep learning approach," *International Journal of Plant Pathology*, vol. 39, pp. 78–86, 2024. DOI: 10.1016/j.plpath.2024.01.007.

[22] J. Wang *et al.*, "Vggnet based system for potato disease detection from leaf images," *Journal of Agricultural Systems*, vol. 186, p. 104 381, 2024. DOI: 10.1016/j.agsy.2024.104381.

[23] S. Zhou *et al.*, "Hybrid cnn and vision transformer approach for potato disease recognition," *Computers in Biology and Medicine*, vol. 155, p. 104 312, 2024. DOI: 10.1016/j.compbiomed.2024.104312.