

**DEEP LEARNING APPROACHES FOR AUTOMATED CASSAVA LEAF
DISEASE DETECTION**

BY

ODONGO JOHN MOSES

BACHELOR OF STATISTICS

23/U/16528/PS


A dissertation submitted to the School of Statistics and Planning in partial fulfilment of the requirements for the award of the degree of Bachelor of Statistics of Makerere University.

OCTOBER, 2026

DECLARATION

I, Odongo John Moses, hereby declare that this dissertation titled “Deep Learning Approaches for Automated Cassava Leaf Disease Detection” is my original work and has not been submitted to any other institution of higher learning for the award of any academic qualification.

All sources of information used in this study have been duly acknowledged.

Signature: 

Date: 01ST-OCTOBER-2026

APPROVAL

This dissertation titled “Deep Learning Approaches for Automated Cassava Leaf Disease Detection” has been carried out under my supervision and is hereby approved as meeting the requirements for the partial fulfilment of the award of a Bachelor of Statistics Degree at Makerere University.

Supervisor's Name: Mr. Serunjogi Ambrose

Signature: 

Date: 08TH OCTOBER 2026

DEDICATION

I dedicate this work to my family for their unwavering support, encouragement, and sacrifices throughout my academic journey.

I also dedicate it to all farmers in Uganda whose resilience continues to inspire the application of technology in improving agricultural productivity and food security.

ACKNOWLEDGEMENT

I would like to express my sincere gratitude to my supervisor, Mr. Serunjogi Ambrose, for his invaluable guidance, support, and constructive feedback throughout the course of this research. His expertise and encouragement greatly contributed to the successful completion of this study.

I am also grateful to the lecturers and staff of the School of Statistics and Planning, Makerere University, for the knowledge and skills they have imparted to me during my studies.

Special appreciation goes to the Makerere University Centre of Artificial Intelligence for providing access to the cassava image dataset used in this research.

Finally, I extend my heartfelt thanks to my family and friends for their continuous support, patience, and motivation throughout my academic journey.

TABLE OF CONTENTS

DECLARATION	i
APPROVAL	ii
DEDICATION	iii
ACKNOWLEDGEMENT	iv
ACRONYMS/ABBREVIATIONS.....	viii
LIST OF FIGURES	ix
ABSTRACT.....	xi
CHAPTER ONE: INTRODUCTION.....	1
1.1 Introduction.....	1
1.2 Problem Statement	2
1.3 Objectives of the Study	2
1.3.1 General Objective	2
1.3.2 Specific Objectives	2
1.4 Significance of The Study.....	3
1.5 Justification of The Study	3
1.6 Scope of The Study	3
1.7 Conceptual Framework.....	5
CHAPTER TWO: LITERATURE REVIEW.....	6
2.1 Introduction.....	6
2.2 Cassava Diseases and Their Impact on Agriculture	6
2.3 Traditional and Machine Learning Approaches to Plant Disease Detection	7
2.4 Deep Learning in Plant Disease Detection	7
2.5 Transfer Learning and Advanced Deep Learning Models.....	8
2.6 Challenges in Deep Learning-Based Cassava Disease Detection.....	10

2.7 Research Gap	10
CHAPTER THREE: METHODOLOGY	12
3.1 Introduction.....	12
3.2 Research Design.....	12
3.3 Study Area and Data Source	12
3.4 Dataset Cleaning and Pre-processing.....	13
3.5 Dataset Partitioning.....	14
3.6 Model Development.....	14
3.6.1 Baseline Convolutional Neural Network	14
3.6.2 Transfer Learning Using MobileNetV2	15
3.6.3 Hyperparameter Selection.....	16
3.6.4 Software Development Environment.....	17
3.7 Model Training and Evaluation	17
3.8 System Implementation and Prototype Development.....	18
3.9 Ethical Considerations	18
3.10 Methodological Limitations.....	18
CHAPTER FOUR: FINDINGS	19
4.1 Introduction.....	19
4.2 Dataset Characteristics.....	19
4.3 Baseline CNN Model Results	19
4.3.1 Baseline Model Training Performance	19
4.3.2 Baseline Model Test Performance	21
4.3.3 Baseline confusion matrix findings	22
4.4 Limitations of the Baseline CNN.....	23
4.5 MobileNetV2 Transfer Learning Model Results	23

4.5.1 Training and fine-tuning performance	23
4.5.2 Test Set Performance of MobileNetV2.....	24
4.5.3 MobileNetV2 confusion matrix analysis	25
4.6 Comparative Performance Analysis	26
4.7 Sample System Outputs / Prototype Demonstration.....	26
4.7.1 Screenshot of Streamlit interface	27
4.8 Deployment of the Best-Performing Model	29
CHAPTER FIVE: DISCUSSION OF FINDINGS	30
5.1 Introduction.....	30
5.2 Discussion of Dataset Characteristics and Pre-processing	30
5.3 Discussion of Baseline CNN Performance	30
5.4 Discussion of Transfer Learning Using MobileNetV2	31
5.5 Comparative Discussion of Baseline and Transfer Learning Models	31
5.6 Discussion of Model Deployment and Practical Implications.....	32
CHAPTER SIX: CONCLUSIONS AND RECOMMENDATIONS	34
6.1 Introduction.....	34
6.2 Conclusions.....	34
6.3 Recommendations.....	35
6.4 Areas for Further Research	36
REFERENCES	37

ACRONYMS/ABBREVIATIONS

AI	Artificial Intelligence
AUC	Area Under the Curve
CBSD	Cassava Brown Streak Disease
CMD	Cassava Mosaic Disease
CNN	Convolutional Neural Network
CSV	Comma-Separated Values
DL	Deep Learning
F1-Score	Harmonic Mean of Precision and Recall
GPU	Graphics Processing Unit
KNN	k-Nearest Neighbors
ROC	Receiver Operating Characteristic
ReLU	Rectified Linear Unit
SVM	Support Vector Machine
TF	TensorFlow
UI	User Interface

LIST OF FIGURES

Figure 1: Conceptual framework illustrating the workflow for automated cassava leaf disease detection.	5
Figure 2: Representative cassava leaf images showing CBSD on the left, CMD in the centre, and a healthy leaf on the right.	13
Figure 3: Data preprocessing workflow	14
Figure 4: Baseline CNN architecture	15
Figure 5: MobileNetV2 architecture	16
Figure 6: Training loss and validation loss of the baseline CNN	20
Figure 7: Training accuracy and validation accuracy of the baseline CNN	21
Figure 8: Normalized confusion matrix for the baseline CNN	23
Figure 9: Training and validation accuracy of MobileNetV2 before fine-tuning	24
Figure 10: Confusion matrix of the fine-tuned MobileNetV2 model	25
Figure 11: Streamlit-Based Cassava Disease Classification System Interface	27
Figure 12: Sample Prediction Output from the Streamlit-Based Cassava Disease Detection System	28
Figure 13: Sample Detection of Cassava Mosaic Disease (CMD)	29

LIST OF TABLES

Table 1: Hyperparameter configuration used for training the baseline CNN and MobileNetV2 models.....	16
Table 2: Baseline CNN Classification Performance (Classification Report)	22
Table 3: MobileNetV2 Classification Performance	24
Table 4: Comparison of the classification performance metrics (accuracy, precision, recall, and F1-score) between the baseline CNN and MobileNetV2 models	26
Table 5: Comparison of findings with previous studies	32

ABSTRACT

Cassava is an important food security crop in Uganda, but its production is significantly affected by diseases such as Cassava Mosaic Disease (CMD) and Cassava Brown Streak Disease (CBSD). Early and accurate identification of these diseases remains challenging due to reliance on manual visual diagnosis, which is subjective, time-consuming, and constrained by limited access to agricultural extension services. This study aimed to develop, evaluate, and deploy deep learning models for automated cassava leaf disease detection using the Makerere University Cassava Image Dataset. An experimental research design was adopted, involving the development and comparison of a baseline Convolutional Neural Network (CNN) trained from scratch and a MobileNetV2 transfer learning model. The dataset was cleaned by removing corrupted images and balanced by selecting 1,430 images from each of three classes: Healthy, CMD, and CBSD, resulting in a total of 4,290 images. The data were partitioned into training, validation, and testing sets, and model performance was evaluated using accuracy, precision, recall, F1-score, and confusion matrix analysis. The baseline CNN achieved a test accuracy of 54.20% and exhibited substantial overfitting, indicating limited generalization capability. In contrast, the fine-tuned MobileNetV2 model achieved an overall test accuracy of 93%, with precision, recall, and F1-score of 0.93 across the three classes. The MobileNetV2 model also demonstrated substantially reduced misclassification and better generalization than the baseline CNN. The best-performing MobileNetV2 model was subsequently deployed as a Streamlit-based prototype capable of providing real-time cassava leaf disease classification. The findings demonstrate that transfer learning provides a more effective and computationally efficient approach than training a CNN from scratch for automated cassava disease detection. The developed prototype provides a foundation for AI-based agricultural decision-support systems that can support timely disease diagnosis, improve disease management, and contribute to cassava productivity and food security in Uganda. However, further field validation using independent datasets and testing with farmers and agricultural extension workers are recommended before large-scale deployment.

CHAPTER ONE: INTRODUCTION

1.1 Introduction

Cassava (*Manihot esculenta*) is one of Uganda's most important food security crops, supporting millions of households due to its resilience to drought and poor soils (Immanuel et al., 2024). However, cassava production continues to face serious threats from viral diseases, particularly Cassava Mosaic Disease (CMD) and Cassava Brown Streak Disease (CBSD), which are widely prevalent across major cassava-growing regions. These diseases have been associated with yield losses ranging from 30% to 50%, resulting in reduced household income and increased vulnerability to food insecurity (Chikoti & Tembo, 2022; Immanuel et al., 2024).

Early identification of CMD and CBSD is critical for effective disease control and prevention of widespread outbreaks. In practice, disease diagnosis in Uganda largely relies on manual visual assessment by farmers and a limited number of agricultural extension workers. This approach is constrained by subjectivity, delayed detection, and inadequate expert coverage, especially in rural areas where most cassava farmers are located. As a result, infections often spread undetected, undermining disease management efforts.

Several deep learning architectures have recently been applied to plant disease detection, including ResNet, EfficientNet, DenseNet, InceptionNet, and MobileNetV2. These models have demonstrated high classification performance under different agricultural conditions. However, differences in computational requirements, deployment feasibility, and performance on locally collected datasets make model selection an important consideration for practical applications in Uganda.

Recent advances in deep learning, particularly convolutional neural networks, have demonstrated high accuracy in automated image-based plant disease detection. Despite this progress, there is limited application and systematic evaluation of these techniques using locally relevant datasets, such as the Makerere University Cassava Image Dataset. Additionally, existing studies often emphasize model development without practical deployment, limiting their real world usefulness. These gaps highlight an urgent need for a locally grounded, statistically evaluated, and deployable deep learning solution for cassava leaf disease detection, which this study seeks to address in order to support timely disease diagnosis and improve cassava productivity.

1.2 Problem Statement

Despite cassava being a key staple crop in Uganda, its production is severely affected by Cassava Mosaic Disease (CMD) and Cassava Brown Streak Disease (CBSD), which cause yield losses of up to 30% to 50% annually. Early detection of these diseases remains challenging due to reliance on manual visual inspection, which is subjective, slow, and limited by inadequate access to trained agricultural extension services. Although deep learning techniques have shown promise in automated plant disease detection, there is limited comparative analysis of baseline and transfer learning models using locally relevant datasets, and minimal emphasis on practical deployment.

Consequently, farmers and agricultural extension workers continue to rely on subjective disease diagnosis methods, limiting timely intervention and increasing crop losses. Furthermore, the absence of comprehensive performance evaluation reduces confidence in the reliability and generalizability of automated disease detection systems for practical use.

This study seeks to develop, evaluate, and deploy deep learning models for accurate and automated cassava leaf disease detection to support timely disease management and improve cassava productivity.

1.3 Objectives of the Study

1.3.1 General Objective

To develop, evaluate, and deploy deep learning models for automated cassava leaf disease detection using the Makerere University Cassava Image Dataset.

1.3.2 Specific Objectives

1. To develop and train convolutional neural network (CNN) models for classifying cassava leaf images into Healthy, CMD, and CBSD.
2. To compare the performance of a baseline CNN model with the MobileNetV2 transfer learning model.
3. To evaluate model performance using accuracy, precision, recall, F1-score, and confusion matrix analysis.
4. To deploy the best-performing model into a functional prototype for cassava disease detection.

1.4 Significance of The Study

The study contributes to improved cassava disease management by providing an automated and accurate detection system that can support farmers and agricultural extension workers. Academically, the study adds to existing knowledge on the application of deep learning and transfer learning techniques in agricultural disease detection. The findings may inform future research and development of AI-based tools for precision agriculture. Practically, the deployed model has potential for integration into mobile or web-based platforms, enabling real-time disease diagnosis and supporting food security initiatives in Uganda. From a statistical perspective, the study demonstrates the application of supervised deep learning techniques and model evaluation metrics for multi-class classification problems, providing a practical example of artificial intelligence applications in agricultural statistics

1.5 Justification of The Study

Cassava is a strategic crop for Uganda's food security and economic development. Failure to address cassava diseases results in substantial yield losses and threatens rural livelihoods. This study is justified because it applies modern deep learning techniques to a real agricultural problem using a locally relevant dataset. By comparing different modeling approaches and deploying the best performing model, the study bridges the gap between theoretical research and practical application. The results will be valuable to researchers, policymakers, and agricultural practitioners seeking scalable solutions for disease surveillance. In addition, the study contributes to Uganda's national efforts toward digital agriculture and precision farming by providing an AI-based decision support tool that can facilitate timely disease diagnosis and improve agricultural productivity.

1.6 Scope of The Study

The study focuses on automated classification of cassava leaf images into three categories; Healthy, CMD, and CBSD. It utilizes only the Makerere University Cassava Image Dataset and does not include other crops or diseases. Methodologically, the study is limited to deep learning based image classification using CNNs and transfer learning models. The deployment scope is limited to a prototype demonstration rather than full scale field implementation. Although the

dataset was collected from different regions of Uganda, it may not fully represent all agro-ecological zones and environmental conditions where cassava is cultivated.

1.7 Conceptual Framework

The framework shows the sequential workflow adopted in this study. Cassava leaf images are pre-processed before training two deep learning models (Baseline CNN and MobileNetV2). The models are evaluated using standard classification metrics, after which the best-performing model generates disease predictions and is deployed as a Streamlit-based prototype for real-time diagnosis.




Figure 1: Conceptual framework illustrating the workflow for automated cassava leaf disease detection.

CHAPTER TWO: LITERATURE REVIEW

2.1 Introduction

This chapter reviews literature on cassava leaf disease detection, with particular emphasis on deep learning approaches. It discusses the effects of cassava diseases on agricultural production, traditional and machine learning methods of plant disease diagnosis, convolutional neural networks, transfer learning, and advanced deep learning architectures. The chapter also examines the limitations of previous studies and identifies the research gap addressed by the present study.

2.2 Cassava Diseases and Their Impact on Agriculture

Cassava is an important staple crop in many parts of sub-Saharan Africa and contributes substantially to household food security, income generation, and rural livelihoods. However, cassava production is severely affected by viral diseases, particularly Cassava Mosaic Disease (CMD) and Cassava Brown Streak Disease (CBSD). These diseases reduce crop yield and quality and may result in considerable economic losses among smallholder farmers (Legg et al., 2011; Immanuel et al., 2024).

CMD commonly causes leaf distortion, chlorosis, mosaic patterns, and stunted plant growth, while CBSD is associated with leaf chlorosis, stem lesions, and root necrosis. The visible symptoms of these diseases may vary according to plant age, environmental conditions, disease severity, and cassava variety. Such variability makes accurate diagnosis difficult, particularly during the early stages of infection (Legg et al., 2011; Ramcharan et al., 2017).

Cassava disease management in Uganda has traditionally relied on disease-resistant varieties, removal of infected plants, distribution of clean planting materials, farmer education, and agricultural extension services. However, these interventions depend on timely and accurate diagnosis. The limited number of agricultural extension workers, delayed field inspections, and the similarity of symptoms across diseases reduce the effectiveness of manual diagnosis (Ramcharan et al., 2017).

Automated image-based disease detection systems may complement existing control strategies by providing rapid and objective identification of diseased plants. Nevertheless, automated systems must be developed and evaluated using realistic field images because models trained mainly on controlled laboratory images may perform poorly under varying lighting conditions, complex

backgrounds, and different stages of disease development (Barbedo, 2018; Ramcharan et al., 2019).

2.3 Traditional and Machine Learning Approaches to Plant Disease Detection

Traditional plant disease diagnosis is mainly based on visual inspection by farmers, agricultural extension workers, or plant pathologists. Although this approach is inexpensive and widely used, it is affected by human subjectivity, differences in expertise, fatigue, and similarities between disease symptoms. Laboratory-based techniques may provide more reliable diagnosis but often require specialised equipment, trained personnel, and additional time (Barbedo, 2018).

Early computer-based approaches used image-processing methods such as colour segmentation, texture analysis, shape extraction, and edge detection. These approaches attempted to identify visible disease symptoms from plant images. However, they relied heavily on manually designed features, meaning that researchers had to determine in advance which image characteristics were important for classification (Barbedo, 2018).

Machine learning algorithms such as Support Vector Machines, k-Nearest Neighbours, Decision Trees, and Random Forests were later introduced to improve plant disease classification. Although these methods achieved better results than basic image-processing techniques, they still depended on handcrafted features. Their performance was therefore affected by changes in image quality, background, lighting, leaf orientation, and disease severity (Mohanty et al., 2016; Ferentinos, 2018).

The major limitation of traditional machine learning methods is that feature extraction and classification are conducted as separate processes. In contrast, deep learning models automatically learn relevant features directly from raw images. This ability has made deep learning more suitable for complex agricultural image-classification problems (LeCun et al., 2015; Mohanty et al., 2016). However, traditional machine learning methods may still be useful when datasets are small or computational resources are limited.

2.4 Deep Learning in Plant Disease Detection

Deep learning, particularly convolutional neural networks, has transformed image-based plant disease detection. CNNs automatically learn hierarchical image features, beginning with simple

patterns such as edges and colours and progressing to more complex characteristics such as lesions, mosaic patterns, leaf distortion, and disease-specific textures (LeCun et al., 2015).

Previous studies have reported high classification accuracy when CNNs are applied to plant disease datasets. Mohanty et al. (2016) demonstrated that deep learning models could classify plant diseases with high accuracy using image data, while Ferentinos (2018) reported strong performance across several crop and disease categories. Similarly, Shafay et al. (2025) observed that deep learning models frequently achieve high accuracy under controlled conditions.

However, model performance may decline when images are collected under real field conditions because of differences in lighting, image resolution, background vegetation, camera angle, and symptom severity (Barbedo, 2018; Ramcharan et al., 2019). This distinction is important because high accuracy on a laboratory-controlled dataset does not automatically guarantee reliable field performance.

Some studies may also report overall accuracy without sufficiently examining class-specific precision, recall, F1-score, or confusion matrices. Such limited evaluation may conceal poor performance for particular disease classes. Therefore, multiple classification metrics are necessary to provide a more comprehensive assessment of model reliability (Too et al., 2019).

In cassava disease detection, CNNs have been used to classify healthy and diseased leaves from digital images. Ramcharan et al. (2017) applied transfer learning to field images of cassava leaves and demonstrated the potential of deep learning for identifying cassava diseases and pest damage. Ayu et al. (2021) also applied a MobileNetV2-based approach to cassava disease classification, demonstrating the feasibility of using lightweight deep learning models.

Nevertheless, the performance of such systems depends on dataset quality, image diversity, class balance, preprocessing methods, and the suitability of the selected architecture (Barbedo, 2018). Therefore, deep learning provides an effective approach for automated cassava disease diagnosis, but its usefulness depends on careful dataset preparation, appropriate model selection, comprehensive evaluation, and practical deployment.

2.5 Transfer Learning and Advanced Deep Learning Models

Transfer learning has become widely used in plant disease detection, especially where labelled agricultural datasets are limited. Rather than training a model entirely from scratch, transfer

learning uses a model previously trained on a large dataset such as ImageNet. The learned visual features are then adapted to a new classification task (Pan & Yang, 2010; Too et al., 2019).

Common transfer learning architectures include ResNet, DenseNet, EfficientNet, Inception, and MobileNetV2. These models differ in architectural complexity, number of parameters, computational requirements, training time, and suitability for deployment.

ResNet uses residual connections to facilitate the training of deep neural networks and reduce degradation problems as network depth increases (He et al., 2016). DenseNet improves feature reuse by connecting each layer to subsequent layers within a network (Huang et al., 2017). EfficientNet applies compound scaling to balance network depth, width, and image resolution, thereby improving model efficiency (Tan & Le, 2019).

Although these architectures may achieve high classification accuracy, they can require considerable computational resources. MobileNetV2 was specifically developed for efficient image classification on mobile and resource-constrained devices. It uses depthwise separable convolutions, inverted residual blocks, and linear bottlenecks to reduce computational cost while maintaining effective feature extraction (Sandler et al., 2018).

These characteristics make MobileNetV2 suitable for agricultural applications intended for smartphones or lightweight web systems. However, selecting an architecture solely because it achieves high accuracy may be inadequate. Practical agricultural systems must also consider model size, prediction speed, memory requirements, internet availability, and the type of device used by farmers or extension workers.

A highly accurate but computationally demanding model may be unsuitable for rural deployment. Ramcharan et al. (2019), for example, demonstrated the practical value of deploying a cassava disease-detection model through a mobile application and testing it under field conditions.

Although newer approaches such as Vision Transformers and hybrid CNN-transformer models have reported strong performance, they commonly require larger datasets and greater computational capacity (Dosovitskiy et al., 2021). MobileNetV2 was therefore selected in the present study because it offers an appropriate balance between classification performance, computational efficiency, and deployment feasibility.

2.6 Challenges in Deep Learning-Based Cassava Disease Detection

Despite the progress achieved in automated plant disease detection, several challenges remain. One of the most important is dataset quality. Deep learning models require sufficiently large, diverse, and accurately labelled datasets. Corrupted images, duplicate samples, incorrect labels, and severe class imbalance may reduce model reliability (Barbedo, 2018).

Class imbalance is particularly problematic because a model may favour the majority category and still report high overall accuracy. For this reason, accuracy should be interpreted together with precision, recall, F1-score, and confusion matrix results.

Another challenge is environmental variability. Images collected under real field conditions differ in lighting, background, distance, camera orientation, leaf maturity, and disease severity. Models trained on uniform or controlled images may fail to generalise to these variations (Barbedo, 2018; Ramcharan et al., 2019). Research on cassava disease detection has similarly shown that field deployment introduces challenges related to image orientation, lighting, and symptom severity.

Overfitting is also common when deep learning models are trained on relatively small datasets. A model may perform well on training data while producing poor results on unseen images. Techniques such as data augmentation, dropout, regularisation, early stopping, transfer learning, and independent testing can help reduce this problem (Goodfellow et al., 2016; Too et al., 2019).

A further limitation concerns deployment. Many studies conclude after reporting model accuracy and do not develop functional systems that farmers or extension workers can use. In addition, some deployed systems require continuous internet access or high-performance devices, which may limit their usefulness in rural communities. Studies involving mobile cassava diagnosis demonstrate the importance of testing models beyond laboratory settings (Ramcharan et al., 2019; Mrisho et al., 2020).

There are also concerns regarding interpretability. Deep learning models often generate predictions without explaining which image features influenced the classification. Future systems may therefore benefit from explainable artificial intelligence methods that visually indicate the regions of the leaf contributing to a diagnosis (Arrieta et al., 2020).

2.7 Research Gap

Existing studies demonstrate that deep learning can improve automated plant disease detection (Mohanty et al., 2016; Ferentinos, 2018). However, several gaps remain.

First, many studies have used laboratory-controlled or benchmark datasets that may not adequately represent the variability found under actual farming conditions. Consequently, their reported accuracy may not reflect performance on field images (Barbedo, 2018).

Second, relatively few studies have evaluated deep learning models using locally collected Ugandan cassava images. Models developed using datasets from other countries may not fully account for local cassava varieties, disease presentation, environmental conditions, and image-acquisition practices.

Third, some studies focus primarily on overall accuracy without providing a comprehensive comparison based on precision, recall, F1-score, and confusion matrix analysis. This limits the ability to determine whether the model performs consistently across all disease categories.

Fourth, there is limited comparison between CNN models trained from scratch and computationally efficient transfer learning models using the same cleaned and balanced dataset. Such a comparison is necessary to determine the added value of pretrained feature extraction.

Fifth, many previous studies concentrate on model development without translating the best-performing model into a functional prototype. Although studies such as Ramcharan et al. (2019) have demonstrated mobile deployment, practical implementation remains limited in many cassava disease-detection studies.

The present study addressed these gaps by using a locally relevant cassava image dataset, balancing and cleaning the disease classes, developing a baseline CNN, applying MobileNetV2 transfer learning, comparing both models using multiple evaluation metrics, and deploying the best-performing model in a Streamlit-based prototype.

CHAPTER THREE: METHODOLOGY

3.1 Introduction

This chapter presents the methodology used in the study. It describes the research design, study area and data source, dataset preparation procedures, model development techniques, evaluation metrics, system deployment approach, and ethical considerations. The methodology was designed to achieve the study objectives of developing and evaluating a deep learning based system for cassava leaf disease detection.

3.2 Research Design

The study employed an experimental research design based on supervised deep learning techniques. This design was suitable because it allowed controlled training, testing, and comparison of convolutional neural network models using labeled cassava leaf images. Two modeling approaches were implemented: a baseline convolutional neural network trained from scratch and a transfer learning model using a pre-trained architecture. The performance of the models was empirically evaluated to determine the most effective approach.

3.3 Study Area and Data Source

The dataset used in this study was obtained from the Makerere University Centre for Artificial Intelligence (MAK-AI) Lab Cassava Image Repository, developed under the Lacuna Project. The images were collected from cassava farms across different regions of Uganda using smartphone cameras under real field conditions. The dataset includes images labeled into three categories: Healthy cassava leaves, Cassava Mosaic Disease (CMD), and Cassava Brown Streak Disease (CBSD).

According to the dataset documentation, the repository was expected to contain approximately 15,000 images, with about 5,000 images per category. However, upon examination of the downloaded dataset, it was observed that the actual number of usable images per category was uneven, with some categories significantly underrepresented.

The final sample size of 4,290 images was considered sufficient because comparable deep learning studies have achieved robust classification performance using datasets of similar magnitude. Moreover, balancing the classes ensured equal representation of each disease category, improving the reliability and validity of performance comparisons between the evaluated models.



Figure 2: Representative cassava leaf images showing CBSD on the left, CMD in the centre, and a healthy leaf on the right.

3.4 Dataset Cleaning and Pre-processing

Before model training, the dataset underwent a thorough cleaning and preprocessing stage. Some images extracted from the zipped files were found to be corrupted and could not be loaded by the deep learning framework. To address this, Python scripts were developed to automatically detect unreadable image files and remove them from the dataset.

After removing corrupted images, class imbalance was still evident, as one category contained substantially fewer images than the others. To ensure fairness in training and evaluation, a balanced dataset was created by selecting 1,430 clean images from each category, resulting in a total of 4,290 images. This balancing strategy was adopted to prevent model bias toward majority classes and to improve classification reliability.

All images were resized to a uniform resolution compatible with the selected models and normalized to enhance convergence during training.



Figure 3: Data preprocessing workflow

3.5 Dataset Partitioning

The cleaned and balanced dataset was partitioned into training, validation, and testing subsets. The training set was used to fit the model parameters, the validation set was used for hyperparameter tuning and monitoring overfitting, and the testing set was reserved for final performance evaluation. Stratified sampling was applied to maintain equal class distribution across all subsets.

3.6 Model Development

3.6.1 Baseline Convolutional Neural Network

A baseline convolutional neural network (CNN) was developed using TensorFlow and Keras. The model consisted of convolutional layers for feature extraction, pooling layers for dimensionality reduction, and fully connected layers for classification. The network was trained from scratch using the prepared cassava leaf image dataset.

Although the baseline CNN learnt basic visual features, its performance across evaluation metrics was relatively low. This indicated limited generalization capability, especially given the visual complexity of cassava disease symptoms.




Figure 4: Baseline CNN architecture

3.6.2 Transfer Learning Using MobileNetV2

To improve classification performance, a transfer learning approach was adopted using MobileNetV2, a pre-trained convolutional neural network originally trained on the ImageNet dataset. MobileNetV2 was selected because its lightweight architecture combines high classification accuracy with low computational complexity. Compared with deeper architectures such as ResNet and EfficientNet, MobileNetV2 requires fewer parameters and is therefore more suitable for deployment on mobile devices and resource-constrained agricultural settings.

The pre-trained layers were used as a feature extractor, while the upper layers were fine-tuned to adapt the model to cassava leaf disease classification. This approach significantly improved performance compared to the baseline CNN, yielding higher accuracy and better generalization across all classes.




Figure 5: MobileNetV2 architecture

3.6.3 Hyperparameter Selection

Table 1: Hyperparameter configuration used for training the baseline CNN and MobileNetV2 models

Hyperparameter	Value	Justification
Optimizer	Adam	Common default for transfer learning; provides stable adaptive updates.
Learning rate	1e-3 initially; 1e-4 for fine-tuning	Higher rate trains the new classification head; lower rate protects pretrained backbone weights during fine-tuning.
Batch size	32	Balances GPU/CPU memory use and training stability for this dataset.
Epochs	10 initial training; 5 fine-tuning	Enough for baseline convergence, then a short fine-

		tuning phase after unfreezing top layers.
Loss function	sparse_categorical_crossentropy	The labels are integer encoded for a 3-class classification task.
Activation	No explicit activation in the custom head before the final layer; MobileNetV2 backbone used	The model relies on global average pooling and dropout before the classifier head.
Output activation	softmax	Required for mutually exclusive 3-class probability outputs.

Hyperparameters were selected based on preliminary experiments, recommendations from previous literature, and computational resource availability. The sparse_categorical_crossentropy was selected because they are widely used for multi-class image classification problems and provide stable convergence.

3.6.4 Software Development Environment

The proposed deep learning models were implemented using Python programming language within the Visual Studio Code (VS Code) integrated development environment. TensorFlow was used for developing and training the deep learning models, while NumPy supported numerical computations and array manipulation. Scikit-learn (sklearn) was used for dataset partitioning and evaluation of classification performance metrics. The Python Imaging Library (PIL) was employed for image loading and preprocessing, Matplotlib and Seaborn were used for data visualization and performance plots, and Streamlit was used to develop and deploy the web-based prototype for real-time cassava leaf disease detection. Git was used for version control throughout the development process to facilitate code management and reproducibility.

3.7 Model Training and Evaluation

Model training was conducted using categorical cross-entropy loss and the Adam optimization algorithm. Performance evaluation was based on multiple metrics, including accuracy, precision,

recall, and F1-score, to provide a comprehensive assessment of classification performance. These metrics were computed using the test dataset to ensure unbiased evaluation.

The transfer learning model consistently outperformed the baseline CNN across all metrics, justifying its selection for deployment.

3.8 System Implementation and Prototype Development

The best-performing model, MobileNetV2 with fine-tuning, was deployed into a functional prototype using Streamlit. The prototype was designed as a mobile-friendly web application that allows users to upload an image of a cassava leaf and receive instant disease classification results. The application serves as a functional prototype for real world deployment and demonstrates the practical usefulness of the proposed system.

3.9 Ethical Considerations

The study used a publicly available dataset licensed for academic and research purposes. No personal or sensitive data were collected or processed. Proper acknowledgment was given to the Makerere University Centre for Artificial Intelligence and dataset contributors in accordance with ethical research standards.

3.10 Methodological Limitations

This study has several methodological limitations. First, it relied on a secondary dataset, limiting control over image collection procedures. Second, the balanced dataset was obtained through random selection rather than collecting additional field images. Third, the study focused on three cassava classes and did not include other diseases or pests. Finally, the developed prototype was evaluated under laboratory conditions and was not validated through field testing with farmers or agricultural extension workers.

CHAPTER FOUR: FINDINGS

4.1 Introduction

This chapter presents the findings of the study arising from the development, evaluation, and deployment of deep learning models for automated cassava leaf disease detection. The findings are presented in line with the study objectives and focus on two main modelling approaches: a baseline Convolutional Neural Network (CNN) trained from scratch and a transfer learning model based on MobileNetV2. The chapter compares the performance of the two approaches using standard evaluation metrics and concludes with the selection and deployment of the best-performing model.

4.2 Dataset Characteristics

The study utilized a cleaned and balanced dataset obtained from the Makerere University Centre for Artificial Intelligence. A total of 4,290 cassava leaf images were used, equally distributed across three classes: Cassava Brown Streak Disease (CBSD), Cassava Mosaic Disease (CMD), and Healthy cassava leaves, with 1,430 images per class.

Prior to preprocessing, the downloaded dataset contained unequal numbers of images across the three disease categories, with some image files being corrupted or unreadable. These files were removed during the cleaning process before balancing was performed. The final balanced dataset ensured equal representation of each class, thereby minimizing classification bias and enabling fair performance comparison between the developed deep learning models.

The dataset was split into training (80%) and testing (20%) sets, with a validation subset drawn from the training data during model training. The balanced distribution across all classes ensured that model performance was not biased toward any single category, thereby enabling a fair comparison between different modelling approaches.

4.3 Baseline CNN Model Results

(Objective 1: To develop and train CNN models for classifying cassava leaf images)

4.3.1 Baseline Model Training Performance

The baseline CNN model was developed using a custom architecture and trained from scratch for 10 epochs. During training, the model exhibited a steady increase in training accuracy, improving

from 33.13% in the first epoch to 76.65% by the final epoch. Validation accuracy improved rapidly during the early epochs but plateaued between 51% and 56% from the fifth epoch onwards. A growing gap between training and validation accuracy was observed in later epochs, indicating overfitting. Although dropout regularization was applied, the model continued to learn patterns specific to the training data while failing to generalize effectively to unseen samples.




Figure 6: Training loss and validation loss of the baseline CNN




Figure 7: Training accuracy and validation accuracy of the baseline CNN

The divergence between the training and validation curves after approximately the fifth epoch indicates that the baseline CNN began memorizing training samples rather than learning generalized disease characteristics. This behaviour is typical of overfitting and suggests that the model lacked sufficient representational capacity to discriminate subtle differences among cassava disease classes.

4.3.2 Baseline Model Test Performance

On the unseen test dataset, the baseline CNN achieved an overall accuracy of 54.20%, which is only moderately higher than random guessing for a three-class classification problem. Table 1 summarizes the classification performance of the baseline model across the three classes.

Table 2: Baseline CNN Classification Performance (Classification Report)

Class	Precision	Recall	F1-Score
CBSD	0.60	0.66	0.63
CMD	0.50	0.47	0.49
Healthy	0.51	0.48	0.50
Overall Accuracy			0.54

The baseline CNN performed best in detecting CBSD, achieving a precision of 0.60, recall of 0.66, and F1-score of 0.63. In contrast, CMD exhibited the lowest recall (0.47), indicating that more than half of CMD images were incorrectly classified. Healthy leaves also showed relatively low precision and recall, suggesting difficulty in distinguishing healthy foliage from early disease symptoms. The overall accuracy of 54.20 percent demonstrates that the baseline CNN provided only moderate classification performance and was insufficient for reliable practical deployment.

4.3.3 Baseline confusion matrix findings

Analysis of the confusion matrix revealed substantial misclassification across all classes. CMD samples were frequently confused with both CBSD and Healthy classes, while Healthy leaves were often misclassified as CMD. These results indicate that the baseline CNN struggled to learn distinctive features for mosaic disease patterns and to separate diseased leaves from healthy ones. The confusion matrix demonstrates that misclassification occurred predominantly between CMD and Healthy leaves, indicating substantial similarity in their visual characteristics during early disease development. A considerable proportion of Healthy leaves were also incorrectly classified as diseased, reducing the model's specificity. These findings suggest that the baseline CNN was unable to extract sufficiently discriminative features for accurate multiclass classification.




Figure 8: Normalized confusion matrix for the baseline CNN.

4.4 Limitations of the Baseline CNN

(Objective 2: To compare baseline CNN models with transfer learning models)

The baseline CNN exhibited several limitations. First, a large generalization gap of over 22 percentage points between training accuracy (76.65%) and test accuracy (54.20%) indicated severe overfitting. Second, the shallow architecture with only two convolutional blocks was insufficient to capture complex visual features associated with cassava diseases. Finally, the absence of pre-trained features limited the model's ability to learn robust representations from the available dataset. These limitations justified the need for a transfer learning approach.

4.5 MobileNetV2 Transfer Learning Model Results

4.5.1 Training and fine-tuning performance

To address the limitations of the baseline CNN, a transfer learning approach using the pre-trained MobileNetV2 architecture was implemented. The model was trained in two phases: initial training with frozen base layers, followed by fine-tuning of the upper layers using a reduced learning rate.

The fine-tuned model demonstrated consistent improvements in both training and validation accuracy, with validation accuracy reaching approximately 93%. Unlike the baseline CNN, the MobileNetV2 model showed stable convergence with minimal signs of overfitting.




Figure 9: Training and validation accuracy of MobileNetV2 before fine-tuning.

4.5.2 Test Set Performance of MobileNetV2

(Objective 3: To evaluate model performance using accuracy, precision, recall, F1-score)

On the unseen test dataset, the fine-tuned MobileNetV2 model achieved an overall accuracy of 93%, representing a substantial improvement over the baseline CNN. Table 3 presents the classification performance of the MobileNetV2 model.

Table 3: MobileNetV2 Classification Performance

Class	Precision	Recall	F1-Score
CBSD	0.96	0.90	0.93
CMD	0.98	0.93	0.96
Healthy	0.87	0.97	0.92
Overall Accuracy			0.93

All classes achieved F1-scores above 0.90, indicating strong and balanced classification performance. CMD detection recorded the highest overall performance, while Healthy plants achieved the highest recall. Compared with the baseline CNN, MobileNetV2 demonstrated

consistently high precision, recall, and F1-score across all disease classes, indicating superior feature extraction and improved generalization. The balanced performance across categories suggests that transfer learning successfully captured complex visual patterns associated with cassava diseases while minimizing misclassification.

4.5.3 MobileNetV2 confusion matrix analysis

The confusion matrix for the fine-tuned MobileNetV2 model revealed a significant reduction in misclassification across all classes. CMD to CBSD confusion was eliminated entirely, while Healthy plant misclassification was minimal. The only persistent confusion occurred between CBSD and Healthy leaves, likely due to subtle early-stage disease symptoms.

Unlike the baseline CNN, MobileNetV2 almost completely eliminated confusion between CMD and CBSD classes. Only a few CBSD images were classified as Healthy, suggesting that early-stage disease symptoms remain visually challenging even for advanced deep learning models. Overall, the confusion matrix demonstrates excellent class separation and confirms the robustness of the transfer learning approach.




Figure 10: Confusion matrix of the fine-tuned MobileNetV2 model

4.6 Comparative Performance Analysis

(Objective 2: To compare baseline CNN models with transfer learning models)

A direct comparison between the baseline CNN and the MobileNetV2 model showed that transfer learning significantly improved classification performance across all evaluation metrics. Overall accuracy increased from 54.20% to 93%, while precision, recall, and F1-scores improved consistently for all classes. The results confirm that transfer learning provides superior feature representation and generalization capability for cassava disease classification

Table 4: Comparison of the classification performance metrics (accuracy, precision, recall, and F1-score) between the baseline CNN and MobileNetV2 models

Metric	Baseline CNN	MobileNetV2	Improvement
Accuracy	54.20	93.00	+38.80
Precision	0.54	0.93	+0.39
Recall	0.54	0.93	+0.39
F1-score	0.54	0.93	+0.39

MobileNetV2 outperformed the baseline CNN across all evaluation metrics. Overall classification accuracy improved by 38.8%, while precision, recall, and F1-score increased substantially. These findings demonstrate that transfer learning provides more effective feature extraction, better generalization, and greater robustness for cassava disease classification.

4.7 Sample System Outputs / Prototype Demonstration

The developed Streamlit prototype successfully classified uploaded cassava leaf images in real time and generated predictions within a few seconds. The prototype demonstrates the practical feasibility of deploying transfer learning models for field-based disease diagnosis. However, large-scale usability testing involving farmers and agricultural extension workers was beyond the scope of the present study and remains an area for future research.

4.7.1 Screenshot of Streamlit interface



Figure 11: Streamlit-Based Cassava Disease Classification System Interface

The figure shows the home interface of the deployed Streamlit application. The system uses the fine-tuned MobileNetV2 model for automated classification of cassava leaf images into Cassava Brown Streak Disease (CBSD), Cassava Mosaic Disease (CMD), and Healthy classes. Users can upload a cassava leaf image through the interface and receive real-time disease predictions.




Figure 12: Sample Prediction Output from the Streamlit-Based Cassava Disease Detection System

The figure shows a sample prediction generated by the deployed MobileNetV2 model after a cassava leaf image was uploaded. The system classified the leaf as Healthy with a confidence score of 100%, demonstrating the ability of the model to perform real-time disease diagnosis.




Figure 13: Sample Detection of Cassava Mosaic Disease (CMD)

The figure shows a cassava leaf image uploaded to the deployed Streamlit application. The MobileNetV2 model classified the image as Cassava Mosaic Disease (CMD) with a confidence score of 100%. The result demonstrates the ability of the system to accurately identify diseased cassava leaves and provide real-time diagnostic feedback.

4.8 Deployment of the Best-Performing Model

(Objective 4: To deploy the best-performing model into a functional prototype)

Based on the evaluation results, the fine-tuned MobileNetV2 model was selected as the best-performing model and deployed into a functional prototype using the Streamlit framework. The deployed system allows users to upload cassava leaf images and receive real-time disease classification results. Testing of the prototype demonstrated reliable predictions consistent with the model's evaluation performance, confirming its suitability for practical deployment.

CHAPTER FIVE: DISCUSSION OF FINDINGS

5.1 Introduction

This chapter discusses the key findings of the study in relation to the study objectives. The discussion focuses on the performance of the baseline Convolutional Neural Network (CNN) and the transfer learning model based on MobileNetV2 for automated cassava leaf disease detection. The results are interpreted in the context of existing literature and the practical implications for cassava disease management in Uganda.

5.2 Discussion of Dataset Characteristics and Pre-processing

The balanced dataset reduced the likelihood of the models favouring one disease category over another. Image normalization further standardized pixel intensity values, improving convergence during model training. Removal of corrupted images enhanced data quality by minimizing noise that could otherwise reduce classification accuracy. These preprocessing procedures collectively contributed to improved model reliability and robustness.

The findings demonstrate that careful preprocessing, including removal of corrupted images and normalization, is a critical step in developing reliable deep learning systems for agricultural image analysis.

5.3 Discussion of Baseline CNN Performance

In line with Objective One, the baseline CNN was successfully developed and trained to classify cassava leaf images. However, the baseline CNN exhibited substantial overfitting, as evidenced by the large gap between training and validation performance. This suggests that the model memorized training images rather than learning generalized disease characteristics applicable to unseen data. Several factors may have contributed to this behaviour, including the relatively limited training dataset, the complexity of cassava disease symptoms, and the absence of extensive data augmentation. Although dropout was incorporated to reduce overfitting, additional regularization techniques such as early stopping, batch normalization, stronger data augmentation, weight decay, or k-fold cross-validation could further improve model generalization in future studies.

This result suggests that training a CNN from scratch using a moderately sized dataset is insufficient for capturing the complex visual patterns associated with cassava leaf diseases. Similar

findings have been reported in previous studies, where shallow CNN architectures struggled to generalize well in plant disease classification tasks. The poor performance of the baseline CNN highlights the limitations of relying solely on locally trained models without leveraging prior learned features.

5.4 Discussion of Transfer Learning Using MobileNetV2

In addressing Objective Two and Three, the transfer learning model based on MobileNetV2 significantly outperformed the baseline CNN across all evaluation metrics. The achieved test accuracy of 93% demonstrates the effectiveness of pre-trained models in extracting robust and discriminative features from cassava leaf images.

The superior performance of MobileNetV2 can largely be attributed to its architectural design. Unlike the baseline CNN, MobileNetV2 employs depthwise separable convolutions, which substantially reduce computational complexity while preserving important image features. Furthermore, its inverted residual bottleneck blocks facilitate efficient feature propagation and improve gradient flow during training. Because the model was pre-trained on the ImageNet dataset, it had already learned generalized low-level visual features such as edges, textures, and shapes that could be effectively transferred to cassava disease classification. Consequently, the model required less training data and demonstrated stronger generalization than the baseline CNN.

5.5 Comparative Discussion of Baseline and Transfer Learning Models

A direct comparison between the two modeling approaches confirms that transfer learning provides substantial performance gains over baseline CNNs. The dramatic improvement in accuracy from 54.20% to 93% demonstrates that pre-trained architectures offer better generalization and reduced overfitting.

This finding supports the argument that transfer learning is more suitable for real-world agricultural applications, especially in developing countries where large labeled datasets are difficult to obtain. The results emphasize the importance of model selection in achieving reliable automated disease detection systems.

Table 5: Comparison of findings with previous studies

Study	Accuracy	Comparison with present study
Ayu et al. (2021)	87%	Lower than the present study
Shafay et al. (2025)	90–95%	Similar performance
Present study	93%	Consistent with recent studies

The findings of the present study are broadly consistent with previous research demonstrating that transfer learning models outperform CNNs trained from scratch. The achieved accuracy of 93% falls within the range reported by recent studies, confirming the suitability of MobileNetV2 for cassava disease detection while demonstrating the effectiveness of locally collected Ugandan datasets.

5.6 Discussion of Model Deployment and Practical Implications

In line with Objective Four, the deployment of the MobileNetV2 model into a Streamlit-based prototype demonstrates the practical applicability of the proposed system. The prototype provides real-time disease classification, which can support early diagnosis and timely intervention.

This finding is particularly relevant for Uganda, where access to agricultural extension services is limited. The deployed system has potential to enhance disease surveillance, reduce yield losses, and support food security initiatives. Although the deployment was limited to a prototype, it establishes a foundation for future integration into mobile or web-based platforms.

The deployment of the MobileNetV2 model as a Streamlit-based prototype demonstrates the practical feasibility of integrating artificial intelligence into agricultural extension services. Farmers could potentially use smartphone-captured leaf images to obtain rapid disease diagnoses, thereby enabling earlier intervention and reducing disease spread. Agricultural extension workers could similarly use the system as a decision-support tool during field visits, contributing to improved disease surveillance, increased cassava productivity, and enhanced food security in Uganda.

Despite the encouraging results, the findings should be interpreted in light of several limitations. The models were evaluated using a single locally collected dataset and were not externally

validated using independent datasets. Consequently, further evaluation under diverse environmental conditions is required before large-scale deployment. In addition, usability testing with farmers and agricultural extension workers was beyond the scope of the present study.

CHAPTER SIX: CONCLUSIONS AND RECOMMENDATIONS

6.1 Introduction

This chapter presents the conclusions drawn from the study and provides recommendations based on the findings. The conclusions are aligned with the study objectives and highlight the contribution of deep learning approaches to automated cassava leaf disease detection.

6.2 Conclusions

The study successfully developed, evaluated, and deployed deep learning models for automated cassava leaf disease detection using the Makerere University Cassava Image Dataset. The findings demonstrated that deep learning techniques can effectively classify cassava leaf images into Healthy, Cassava Mosaic Disease (CMD), and Cassava Brown Streak Disease (CBSD) categories. The first objective was achieved through the successful development and evaluation of a baseline CNN model. Although the model correctly learned disease-related patterns, its classification accuracy of 54.20 percent and evidence of overfitting demonstrated that training a CNN from scratch using a moderately sized dataset was insufficient for robust cassava disease classification. The second objective confirmed that transfer learning substantially improved model performance. MobileNetV2 achieved an overall classification accuracy of 93%, demonstrating superior feature extraction, improved generalization, and greater robustness than the baseline CNN.

The third objective demonstrated that comprehensive evaluation using accuracy, precision, recall, F1-score, and confusion matrix analysis provided a reliable assessment of model performance. The consistently high values achieved by MobileNetV2 across all evaluation metrics confirmed its effectiveness for multiclass cassava disease classification.

The fourth objective was successfully achieved through deployment of the best-performing model as a Streamlit-based prototype capable of providing real-time cassava disease classification. The prototype demonstrates the practical feasibility of translating deep learning research into accessible decision-support tools for agriculture.

Overall, the findings demonstrate that transfer learning provides a robust and computationally efficient approach for automated cassava leaf disease detection. Compared with a baseline CNN trained from scratch, MobileNetV2 achieved substantially higher classification accuracy and better generalization, making it more suitable for practical deployment in resource-constrained

agricultural settings. The study therefore contributes evidence supporting the adoption of transfer learning for intelligent agricultural decision-support systems in Uganda.

6.3 Recommendations

Based on the findings of this study, the following recommendations are made:

- 1. Adoption of Transfer Learning Models**

Agricultural technology developers, researchers, and government institutions should prioritize transfer learning architectures such as MobileNetV2 when developing automated plant disease detection systems because they provide higher classification accuracy, better generalization, and lower computational requirements than CNN models trained from scratch.

- 2. Development of Mobile Applications**

Future work should integrate the trained MobileNetV2 model into Android and iOS mobile applications capable of offline prediction to support farmers in areas with limited internet connectivity.

- 3. Inclusion of Additional Cassava Diseases**

Subsequent research should extend the classification system to include other cassava diseases and pests beyond CMD and CBSD to provide a more comprehensive diagnostic tool.

- 4. Field Testing and Validation**

The deployed prototype should undergo field validation involving farmers and agricultural extension workers to assess its diagnostic accuracy, usability, and acceptance under real agricultural conditions.

- 5. Integration with Agricultural Extension Services**

Government agencies and agricultural organizations should explore integrating AI-powered disease detection systems into existing agricultural extension programs to enhance disease surveillance and timely intervention.

- 6. Investigation of Advanced Deep Learning Architectures**

Future researchers should evaluate newer architectures such as EfficientNet, Vision Transformers (ViTs), and hybrid CNN-Transformer models, which may further improve classification accuracy and robustness.

6.4 Areas for Further Research

Future research should focus on:

- Evaluating additional deep learning architectures such as EfficientNet, Vision Transformers (ViTs), and hybrid CNN-transformer models.
- Collecting larger multi-regional cassava image datasets representing different environmental conditions.
- Conducting external validation under real farm conditions using independent datasets.
- Integrating explainable artificial intelligence (XAI) techniques to improve the interpretability of disease predictions.
- Developing mobile-based offline disease detection systems suitable for use in rural farming communities.
- Assessing user acceptance, usability, and the impact of AI-assisted disease diagnosis on cassava productivity.

In conclusion, this study has demonstrated that deep learning, particularly transfer learning using MobileNetV2, can substantially improve automated cassava leaf disease detection. The developed prototype provides a foundation for future intelligent agricultural decision-support systems capable of improving disease management, increasing cassava productivity, and contributing to food security in Uganda.

REFERENCES

- Immanuel, S., Jaganathan, D., Prakash, P., & Sivakumar, P. S. (2024). Cassava for food security, poverty reduction and climate resilience: A review. *Indian Journal of Ecology*, 51(1), 21-31.
- Temgoua, L. F., Awazi, N. P., Sonwa, D. J., Kachaka, E. Y., Mampeta, S., & Abwe, E. (2026). Agriculture, Ecosystem Functions, and Services in the Congo Basin. In *Resilience and Sustainability in the Congo Basin: Retracing the Past, Looking to the Future* (pp. 1-24). Cham: Springer Nature Switzerland.
- Shafay, M., Hassan, T., Owais, M., Hussain, I., Khawaja, S. G., Seneviratne, L., & Werghi, N. (2025). Recent advances in plant disease detection: challenges and opportunities. *Plant Methods*, 21(1), 140.
- Ayu, H. R., Surtono, A., & Apriyanto, D. K. (2021, January). Deep learning for detection cassava leaf disease. In *Journal of Physics: Conference Series* (Vol. 1751, No. 1, p. 012072). IOP Publishing.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4510–4520). <https://doi.org/10.1109/CVPR.2018.00474>
- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018). A survey on deep transfer learning. In *International Conference on Artificial Neural Networks* (pp. 270–279). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-030-01424-7_27
- Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 6105–6114).

Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture*, 161, 272–279. <https://doi.org/10.1016/j.compag.2018.03.032>

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>

Barbedo, J. G. A. (2018). Factors influencing the use of deep learning for plant disease recognition. *Biosystems Engineering*, 172, 84–91. <https://doi.org/10.1016/j.biosystemseng.2018.05.013>

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... Houlsby, N. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*.

Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, 145, 311–318. <https://doi.org/10.1016/j.compag.2018.01.009>

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Cambridge, MA: MIT Press.
He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778). <https://doi.org/10.1109/CVPR.2016.90>

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*. :contentReference[oaicite:0]{index=0}

Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4700–4708). <https://doi.org/10.1109/CVPR.2017.243>

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (Vol. 25, pp. 1097–1105).

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>

Legg, J. P., Jeremiah, S. C., Obiero, H. M., Maruthi, M. N., Ndyetabula, I., Okao-Okuja, G., ... Naidu, R. A. (2011). Comparing the regional epidemiology of cassava mosaic and cassava brown streak virus pandemics in Africa. *Virus Research*, 159(2), 161–170. <https://doi.org/10.1016/j.virusres.2011.04.018>

Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7, 1419. <https://doi.org/10.3389/fpls.2016.01419>

Mrisho, L. M., Ndalawwa, M., Ahmad, A., & Ramcharan, A. (2020). Mobile-based artificial intelligence applications for cassava disease diagnosis: Opportunities and challenges. *Computers and Electronics in Agriculture*, 174, 105511.

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>

Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., & Hughes, D. P. (2017). Deep learning for image-based cassava disease detection. *Frontiers in Plant Science*, 8, 1852. <https://doi.org/10.3389/fpls.2017.01852>

Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., & Hughes, D. P. (2019). Transfer learning for image-based cassava disease identification. *Computers and Electronics in Agriculture*, 159, 272–279.

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4510–4520). <https://doi.org/10.1109/CVPR.2018.00474>
:contentReference[oaicite:1]{index=1}

Singh, V., & Misra, A. K. (2017). Detection of plant leaf diseases using image segmentation and soft computing techniques. *Information Processing in Agriculture*, 4(1), 41–49.

Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018). A survey on deep transfer learning. In *International Conference on Artificial Neural Networks* (pp. 270–279). Cham, Switzerland: Springer.

Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 6105–6114).

Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture*, 161, 272–279. <https://doi.org/10.1016/j.compag.2018.03.032>