Advanced Deep Learning Framework for Automated Lemon Leaf Canker Detection: A Comparative Study of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 for Precision Agriculture

By
Minhajul Abedin
201-15-3411

FINAL YEAR DESIGN PROJECT REPORT

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

Supervised by

Mr. Md. Abbas Ali Khan

Assistant Professor
Department of Computer Science and EngineeringDaffodil International University

Co-Supervised by

Md. Abdullah Al Kafi

Lecturer
Department of Computer Science and EngineeringDaffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY
Dhaka, Bangladesh

12 January, 2025

APPROVAL

This Project titled **"Advanced Deep Learning Framework for Automated Lemon Leaf Canker Detection: A Comparative Study of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 for Precision Agriculture"**, submitted by **Minhajul Abedin**, ID No: **201-15-3411** to the Department of Computer Science and Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of B.Sc. in Computer Science and Engineering and approved as to its style and contents. The presentation has been held on 12 January, 2025.

BOARD OF EXAMINERS

Hossain

Dr. Md. Fokhray Hossain **Chairman**
Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Siddiquee

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

Mr. Md Umaid Hasan **Internal Examiner**
Lecturer
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Nazibur Rahman **External Examiner**
Technical Lead - Database Administrator
Telenor - Grameen Phone Account

DECLARATION

We hereby declare that this project has been done by us under the supervision of **Mr.Md. Abbas Ali Khan, Assistant Professor**, Department of Computer Science and Engineering, Daffodil International University. We also declare that neither this project nor any part of this project has been submitted elsewhere for the award of any degree or diploma.

Supervised by:

Mr.Md. Abbas Ali Khan
Assistant Professor
Department of Computer Science and
EngineeringDaffodil International
University

Co-Supervised by:

Md. Abdullah Al Kafi
Lecturer
Designation
Department of Computer Science and
EngineeringDaffodil International
University

Submitted by:
Minhajul Abedin

Minhajul Abedin
ID: 201-15-3411
Department of Computer Science and
EngineeringDaffodil International
University

ACKNOWLEDGEMENTS

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

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

We are grateful and wish our profound indebtedness to **Mr.Md. Abbas Ali Khan** ,**Assistant Professor**, Department of Computer Science and Engineering, Daffodil International University, Dhaka, Bangladesh. Deep knowledge and keen interest of our supervisor in the field of **Precision Agriculture** and **Artificial intelligence (AI)** to carry out this project. His endless patience, scholarly guidance, continual encouragement, constant and energetic supervision, constructive criticism, valuable advice, reading many inferior drafts, and correcting them at all stages have made it possible to complete this project.

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

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

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

ABSTRACT

Lemon leaf canker, a bacterial disease affecting lemon cultivation globally, requires efficient detection methods to prevent economic losses and crop damage. This study evaluates the performance of four advanced YOLO models—YOLOv5, YOLOv8, YOLOv10, and YOLOv11—for automated detection of lemon leaf canker using a dataset of 884 annotated images representing healthy and diseased classes. The dataset was preprocessed and augmented to improve robustness under diverse conditions. The models were trained and evaluated on metrics such as precision, recall, mean Average Precision (mAP), and inference time. YOLOv11 emerged as the best-performing model, achieving a precision of 94%, a recall of 90.4%, and a mAP50-95 of 73.5%, indicating superior accuracy and computational efficiency. These results underscore YOLOv11's potential for real-time application in precision agriculture, enabling early disease detection and timely interventions. By reducing reliance on labor-intensive methods and excessive chemical treatments, the proposed approach supports sustainable farming practices, improves crop management, and mitigates environmental impact. This research demonstrates the capability of deep learning-based solutions like YOLOv11 to advance agricultural productivity and contribute to economic stability for farmers, while promoting environmentally conscious agricultural practices.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Lemon leaf canker is a destructive disease caused by bacterial and fungal pathogens, leading to lesions on leaves, reduced fruit yield, and economic losses. Traditional detection methods, such as manual inspections, are time-intensive and error-prone. This study explores the application of advanced deep learning models, like YOLO, to automate and improve the accuracy of lemon leaf canker detection, supporting efficient and sustainable agricultural practices.

1.1 Introduction

Lemon leaf canker, caused by various bacterial and fungal pathogens, is a devastating disease that affects lemon trees worldwide. The disease manifests as lesions on the leaves, which can lead to premature leaf drop, reduced photosynthesis, and ultimately, decreased fruit yield and quality. This not only causes significant financial losses for farmers but also impacts the agricultural industry's ability to meet market demands. Early detection of lemon leaf canker is critical for effective management and prevention, as it allows for prompt intervention, reducing the spread of the disease and minimizing crop damage. However, traditional methods of disease detection, such as manual inspection, are time-consuming, subjective, and prone to errors. These limitations highlight the need for automated detection systems that can provide accurate, rapid, and scalable solutions for monitoring plant health.

In recent years, the integration of artificial intelligence (AI) and deep learning in agricultural practices has gained significant attention. These technologies offer the potential to revolutionize how diseases in crops are detected and managed. Among the various AI techniques, deep learning models have shown exceptional promise in object detection tasks, particularly for identifying and classifying plant diseases. Object detection involves not only recognizing objects within an image but also determining their precise location and boundaries. This makes it an ideal approach for analyzing images of leaves and identifying symptoms of diseases like lemon leaf canker.

One of the most successful architectures for object detection is the YOLO (You Only Look Once) model, which has undergone significant advancements over the years. YOLO was introduced in 2015 as a real-time object detection system that directly predicts bounding boxes and class labels from images, making it faster and more efficient than previous models that relied on region proposals and classification steps. YOLO's key strength lies in its speed, accuracy, and ability to handle real-time video streams, which makes it suitable for applications in dynamic environments like farms.

Since its introduction, YOLO has evolved through several versions, each improving on the last in terms of speed, accuracy, and handling of more complex detection tasks. YOLOv2, released in 2016, introduced better resolution and faster inference, while YOLOv3 in 2018 made improvements to model performance and detection of smaller objects. The later versions, such as YOLOv4 and YOLOv5, further enhanced accuracy and introduced more efficient training techniques. YOLOv8, YOLOv10, and YOLOv11 have pushed the boundaries even further, incorporating state-of-the-art optimizations such as better use of transfer learning, fine-tuning for specific tasks, and improved handling of occlusions and small objects. These advancements have made YOLO models highly suitable for real-world applications in precision agriculture, where the goal is not only to detect plant diseases with high accuracy but also to do so quickly and efficiently at scale.

As the agricultural industry continues to adopt technology for disease management, YOLO's ability to perform accurate, automated detection of diseases like lemon leaf canker represents a significant step forward in precision agriculture. By leveraging deep learning models like YOLO, farmers can gain valuable insights into the health of their crops, enabling them to make data-driven decisions that can improve yields, reduce losses, and contribute to more sustainable farming practices.

1.2 Motivation

Early detection of lemon leaf canker is crucial for minimizing the negative impact of this disease on lemon crops. The disease can spread rapidly if left unchecked, leading to severe damage to the leaves, branches, and fruit. As a result, lemon trees may experience stunted growth, premature fruit drop, and reduced quality of harvest, ultimately causing significant economic losses for farmers. In some cases, if the disease is not detected early, it can devastate entire orchards, leading to long-term financial hardship for growers. Therefore, timely detection and intervention are essential for controlling the spread of lemon leaf canker and protecting the agricultural sector.

Traditional methods of detecting plant diseases, such as visual inspections by agricultural workers or experts, are labor-intensive and often ineffective in identifying early-stage infections. These methods are also prone to human error, as the symptoms of diseases like lemon leaf canker may resemble those of other leaf ailments, making it difficult to distinguish between them. Additionally, manual inspections can be time-consuming and impractical when dealing with large-scale farms, where the sheer volume of plants makes it impossible to inspect each one thoroughly.

This is where advanced deep learning models come into play. Deep learning, a subset of artificial intelligence (AI), has revolutionized various fields, including computer vision, by enabling machines to automatically learn patterns from large datasets and make accurate predictions. In the context of agriculture, deep learning models can be trained to identify subtle differences in leaf images, such as color changes, lesion patterns, and other indicators of disease, which might be difficult for the human eye to detect. These models, particularly convolutional neural networks (CNNs), have demonstrated high accuracy in classifying plant diseases from images, making them ideal for automating the detection process.

One of the most promising deep learning techniques for plant disease detection is object detection, which not only identifies the presence of a disease but also locates its position within an image. Advanced models like YOLO (You Only Look Once) offer significant advantages in this regard, as they are capable of performing real-time object detection with high accuracy. By leveraging YOLO-based models, farmers can quickly and accurately detect lemon leaf canker in images, enabling them to take immediate action to prevent the disease from spreading. This could include isolating affected plants, applying targeted treatments, or adjusting environmental conditions to mitigate further damage.

The use of deep learning models for early disease detection can significantly reduce the reliance on manual labor, minimize human error, and enable faster decision-making. By automating the detection process, farmers can monitor large areas efficiently and accurately, ensuring that potential outbreaks are addressed before they cause substantial harm. Ultimately, advanced deep learning models offer a powerful solution for mitigating agricultural losses due to lemon leaf canker, improving crop management, and contributing to more sustainable farming practices.

1.3 Objectives

The objective of this work is to design, implement, and evaluate a deep learning-based framework for the automated detection of lemon leaf canker, a disease that significantly impacts lemon cultivation. Specifically, this study aims to compare the performance of four state-of-the-art object detection models—YOLOv5, YOLOv8, YOLOv10, and YOLOv11—by analyzing their precision, recall, mean Average Precision (mAP), and computational efficiency. The research involves collecting and preprocessing a diverse dataset of lemon leaf images, applying advanced augmentation techniques, and optimizing the models for real-world deployment. By identifying the best-performing model, the study seeks to address the limitations of traditional detection methods, such as their time-intensive nature and susceptibility to error, while providing an efficient, accurate, and scalable solution for disease management. Ultimately, this work aims to contribute to the advancement of precision agriculture, promoting sustainable farming practices through timely disease detection and intervention.

1.4 Methodology

This study employs a systematic approach to design, train, and evaluate advanced deep learning models for the detection of lemon leaf canker. The methodology involves the following key steps

Dataset Collection and Preprocessing

A dataset comprising 2,124 images of lemon leaves, both healthy and diseased, was collected from various nursery gardens under diverse environmental conditions. Each image was manually annotated with two classes: 'canker' and 'healthy.' Preprocessing steps included auto-orientation, static cropping to focus on relevant leaf regions, and the application of data augmentation techniques such as rotation, brightness adjustment, noise addition, and hue variation to enhance model robustness and generalization.

Model Selection and Implementation

Four state-of-the-art YOLO models—YOLOv5, YOLOv8, YOLOv10, and YOLOv11—were selected for this study due to their proven efficiency in object detection tasks. Pre-trained weights were utilized for transfer learning, and each model was fine-tuned using the prepared dataset to optimize their performance for lemon leaf canker detection.

Training and Evaluation

The models were trained on 88% of the dataset, validated on 7%, and tested on the remaining 5%. Training involved the use of NVIDIA Tesla T4 GPUs and cloud-based platforms like Google Colab and Kaggle to ensure computational efficiency. Early stopping mechanisms were employed to prevent overfitting, and performance was evaluated using standard metrics such as precision, recall, mean Average Precision (mAP), and inference time per image.

Comparative Analysis

A detailed comparison was conducted to assess the strengths and limitations of each model in terms of accuracy, robustness, and computational efficiency. This involved analyzing the models' ability to detect subtle disease symptoms under various conditions and their effectiveness in handling class imbalances and environmental variability.

This methodology ensures a comprehensive evaluation of the YOLO models, providing actionable insights into their suitability for automated lemon leaf canker detection in real-world agricultural settings.

1.5 Project Outcome

The anticipated outcome of this research is to identify the best-performing YOLO model for the detection of lemon leaf canker, providing valuable insights into the model's potential for real-world applications in precision agriculture.

Specifically, the study expects to achieve the following outcomes:

Identification of the Best-Performing YOLO Model

Based on the comparative analysis of YOLOv5, YOLOv8, YOLOv10, and YOLOv11, the research aims to determine which version of YOLO delivers the highest accuracy, precision, and recall for detecting lemon leaf canker. YOLOv11 is expected to perform the best due to its latest enhancements in model architecture and optimization techniques, particularly in handling small objects and complex background variations.

Model Comparison in Terms of Computational Efficiency and Robustness

The study anticipates that the best-performing model will also offer a good balance between computational efficiency and robustness. While YOLOv11 is expected to have slightly higher computational demands, it should still be feasible for deployment in real-world agricultural settings, particularly with the availability of modern hardware like GPUs. The research will also highlight trade-offs between different YOLO versions, helping farmers choose models that fit their specific operational needs, balancing performance and resource requirements.

Application in Real-World Scenarios

The results of this study are expected to demonstrate the practical viability of YOLO models for early, automated detection of lemon leaf canker in field conditions. The best-performing model will provide an efficient tool for farmers to monitor the health of lemon orchards in real-time. By enabling faster identification of diseased leaves, the system will allow for timely intervention, such as applying targeted treatments or isolating affected plants to prevent further spread.

Contribution to Precision Agriculture

The findings are expected to contribute to the growing field of precision agriculture by offering a reliable, automated solution for plant disease detection. With the ability to handle large-scale data and provide real-time results, the selected YOLO model will support sustainable farming practices, optimize resource use, and reduce crop losses due to disease. Additionally, the research may encourage further exploration of deep learning-based solutions in other agricultural applications.

Overall, the anticipated outcomes of this research will not only help identify the most effective YOLO model for lemon leaf canker detection but will also demonstrate the practical benefits of advanced AI models in improving agricultural productivity and sustainability.

1.6 Organization of the Report

Chapter 1: Introduction

This chapter introduces the problem of lemon leaf canker, emphasizing its impact on global lemon cultivation and agricultural productivity. It outlines the motivation, objectives, and scope of the study, providing a roadmap for leveraging deep learning models like YOLO for automated disease detection.

Chapter 2: Background

The background provides an overview of lemon leaf canker, discussing its causes, symptoms, and traditional detection methods. It highlights the limitations of manual inspections and presents advanced deep learning models as a transformative solution. A literature review summarizes prior research on plant disease detection using YOLO models, identifying research gaps that this study addresses.

Chapter 3: Research Methodology

This chapter details the research process, including dataset collection from diverse nursery gardens, annotation techniques using Roboflow, and preprocessing steps to enhance image quality. It explains the implementation of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 models, along with training, validation, and performance evaluation methodologies using standard metrics such as precision, recall, and mAP.

Chapter 4: Implementation and Results

This chapter presents the experimental setups, training configurations, and results of the model evaluations. It provides a comparative analysis of the four YOLO models, highlighting YOLOv11 as the most effective for lemon leaf canker detection. Visualizations such as confusion matrices, precision-recall curves, and performance metrics are included to support the findings.

Chapter 5: Engineering Standards and Design Challenges

This chapter discusses the adherence to software, hardware, and communication standards throughout the study. It addresses design challenges such as class imbalances, environmental variability, and computational limitations, proposing strategies to optimize the deployment of AI-based solutions in real-world agricultural settings.

Chapter 6: Impact on Society, Environment, and Sustainability

This chapter explores the societal and environmental implications of automated disease detection. It emphasizes the reduction of pesticide usage, promotion of sustainable farming practices, and economic benefits for farmers. Ethical considerations, such as transparency and fairness, are also discussed.

Chapter 7: Conclusion and Future Work

The final chapter summarizes the study's contributions, emphasizing the potential of YOLOv11 for real-time, scalable agricultural applications. It suggests directions for future research, including expanding datasets, refining detection algorithms, and integrating the system into IoT-enabled devices for broader adoption.

References

Lists all scholarly articles, tools, and resources cited in the report.

Chapter 2

Background

The Background chapter of this report provides a comprehensive overview of lemon leaf canker as a critical issue in citrus farming. It outlines the causes and effects of the disease, including the lesions it causes on leaves, reduced photosynthesis, and the resulting economic losses for farmers. This section sets the stage by detailing the limitations of traditional disease detection methods and the need for advanced technologies like deep learning to automate and improve the accuracy of lemon leaf canker detection, thereby supporting sustainable agricultural practices.

2.1 Introduction

Lemon leaf canker is a devastating plant disease that significantly threatens lemon cultivation worldwide. Caused by various bacterial and fungal pathogens, the disease manifests as lesions on leaves, leading to premature leaf drop, reduced photosynthesis, and ultimately decreased fruit yield and quality. These impacts result in substantial financial losses for farmers and challenge the agricultural industry's ability to meet market demands.

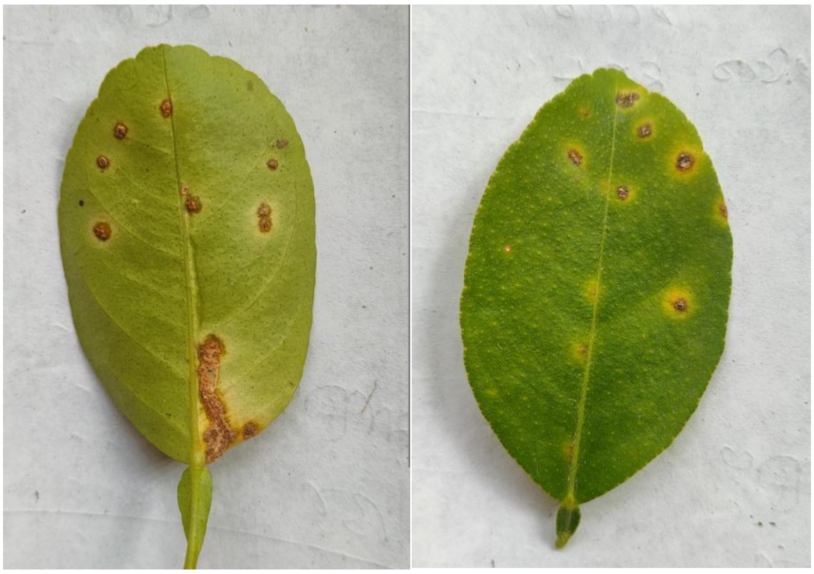

Figure 2.1: Sample Image

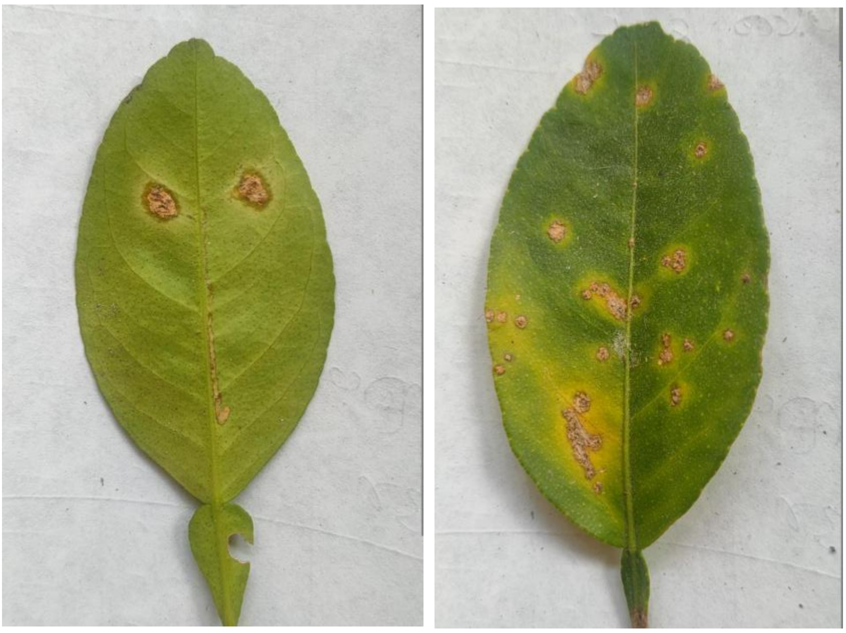

Figure 2.2: Sample Image

Traditional disease detection methods, primarily manual visual inspections, are inherently limited. These approaches are time-consuming, subjective, and prone to human error, especially when identifying early-stage infections. The complexity of detecting subtle disease symptoms necessitates a more advanced, reliable detection approach.

Object detection models, particularly YOLO (You Only Look Once), offer a transformative solution to these challenges. YOLO models represent a cutting-edge deep learning technique that can automatically identify and locate diseases within leaf images with remarkable speed and accuracy. Unlike traditional detection methods, YOLO can

- Rapidly analyze large volumes of images
- Provide real-time disease detection
- Identify precise locations of disease symptoms
- Distinguish between healthy and diseased leaf regions
- Work across varying imaging conditions

The evolution of YOLO models, from initial versions to recent iterations like YOLOv5, YOLOv8, YOLOv10, and YOLOv11, has progressively enhanced object detection capabilities. These advancements include improved accuracy, faster inference times, and better handling of complex detection tasks, making them particularly suitable for agricultural applications.

By leveraging YOLO's deep learning architecture, researchers can develop automated systems that;

- Detect lemon leaf canker at early stages
- Enable prompt intervention
- Minimize disease spread
- Reduce manual inspection efforts
- Support data-driven agricultural decision-making

The object detection approach allows for precise identification of disease symptoms, potentially catching infections before they become visually apparent to human observers. This early detection capability is crucial for implementing targeted interventions, such as isolating affected plants, applying specific treatments, or adjusting environmental conditions to prevent further disease progression.

Moreover, YOLO models offer scalability and efficiency, making them ideal for monitoring large agricultural areas. Farmers can potentially implement these technologies to continuously assess crop health, transforming disease management from a reactive to a proactive strategy.

As precision agriculture continues to evolve, YOLO-based object detection represents a significant technological breakthrough in plant disease management, promising to enhance crop health monitoring, reduce economic losses, and support more sustainable farming practices.

2.2 Literature Review

Table 2.1: Summary of Literature Reviewed.

Author (s)	**Year**	**Title**	**Methodology**	**Key Findings**
Qiu et al.	2022	An automatic identification system for citrus greening disease (Huanglongbing) using a YOLO convolutional neural network	YOLO-based CNN for citrus greening detection	Successfully developed an automated system for detecting Huanglongbing disease using YOLO architecture with high accuracy
Zhang et al.	2022	Automated identification of citrus diseases in	Deep learning methods for orchard	Created an automated system for identifying

		orchards using deep learning	disease detection	multiple citrus diseases in orchard settings
Krishna et al.	2023	On-spot Citrus Canker Disease Detection using YOLOv7	YOLOv7 implementation	Demonstrated effective real-time detection of citrus canker disease.
Pangaliman et al.	2024	Philippine Lime (Calamansi) Disease Detection and Classification Using YOLOv8 Model	YOLOv8 model	Successfully applied YOLOv8 to detect and classify diseases specific to Philippine lime crops.
Preanto et al.	2024	A Semantic Segmentation Approach on Sweet Orange Leaf Diseases Detection Utilizing YOLO	YOLO-based semantic segmentation	Developed a specialized approach combining YOLO with semantic segmentation for sweet orange leaf disease detection.
Dananjayan et al.	2022	Assessment of state-of-the-art deep learning based citrus disease detection techniques using annotated optical leaf images	Comparative analysis of deep learning techniques	Provided comprehensive evaluation of various deep learning methods for citrus disease detection.

Mo et al	2024	Lightweight citrus leaf disease detection model based on ARMS and cross-domain dynamic attention	ARMS with cross-domain dynamic attention	Created an efficient lightweight model for disease detection incorporating novel attention mechanisms.
Luo et al.	2023	Citrus Diseases and Pests Detection Model Based on Self-Attention YOLOV8	Self-attention enhanced YOLOv8	Integrated self-attention mechanisms into YOLOv8 for improved disease and pest detection.
da Silva et al.	2023	Using mobile edge AI to detect and map diseases in citrus orchards	Mobile edge AI implementation	Successfully deployed AI models on mobile edge devices for in-field disease detection and mapping.
Apacionado et al.	2023	Sooty Mold Detection on Citrus Tree Canopy Using Deep Learning Algorithms	Deep learning for sooty mold detection	Developed specialized algorithms for detecting sooty mold in citrus tree canopies.
Dai et al.	2024	YOLOv8-GABNet: An Enhanced Lightweight Network for the High-Precision Recognition of Citrus	YOLOv8 with GABNet enhancement	Created an improved lightweight network achieving high precision in disease and nutrient deficiency

		Diseases and Nutrient Deficiencies		detection.
Xiang et al.	2023	Real-Time Detection Algorithm for Kiwifruit Canker Based on a Lightweight and Efficient Generative Adversarial Network	GAN-based real-time detection	Developed efficient GAN-based algorithm for real-time kiwifruit canker detection.
Yao et al.	2024	Two-Stage Detection Algorithm for Plum Leaf Disease and Severity Assessment Based on Deep Learning	Two-stage deep learning approach	Implemented novel two-stage algorithm for both detection and severity assessment of plum leaf diseases.
Huang et al.	2024	YOLOv8-G: An Improved YOLOv8 Model for Major Disease Detection in Dragon Fruit Stems	Modified YOLOv8 architecture	Enhanced YOLOv8 specifically for dragon fruit stem disease detection.
Yue et al.	2023	Improved YOLOv8-Seg network for instance segmentatio	YOLOv8-Seg with improvements	Advanced instance segmentation capabilities for tomato plant

		n of healthy and diseased tomato plants in the growth stage		disease detection.
Rani et al.	2024	Infield disease detection in citrus plants: integrating semantic segmentation and dynamic deep learning object detection model for enhanced agricultural yield	Combined semantic segmentation and dynamic deep learning	Integrated multiple deep learning approaches for comprehensive in-field citrus disease detection.

2.2.1 Related Research

The application of deep learning techniques in plant disease detection has gained significant traction in recent years, with researchers exploring various approaches to enhance agricultural monitoring and disease management. Several studies have demonstrated the potential of YOLO models in detecting citrus and other plant diseases across different contexts.

[1] developed an automatic identification system for citrus greening disease using a YOLO convolutional neural network, highlighting the model's effectiveness in detecting critical citrus diseases. Similarly, [2] proposed an automated identification approach for citrus diseases in orchards, emphasizing the potential of deep learning in agricultural monitoring.

Specific studies have focused on canker detection and classification across different citrus varieties. [3] implemented YOLOv7 for on-spot citrus canker disease detection, while [4] [3] utilized YOLOv8 for Philippine lime (Calamansi) disease detection and classification. These studies underscore the versatility of YOLO models in addressing diverse plant disease challenges.

The research landscape reveals several key approaches:

Model Optimization: Researchers have been continuously improving YOLO architectures. [8] proposed a self-attention YOLO model for citrus disease detection, while [7] developed a lightweight citrus leaf disease detection model using dynamic attention mechanisms.

Technological Integration: [9] explored mobile edge AI for detecting and mapping diseases in citrus orchards, demonstrating the potential for real-world application of these technologies.

Comprehensive Detection: Some studies expanded beyond single disease detection. [11] introduced YOLOv8-GABNet for high-precision recognition of citrus diseases and nutrient deficiencies, showcasing the broader applicability of deep learning models.

Comparative studies have been crucial in understanding model performance. [6] assessed state-of-the-art deep learning techniques for citrus disease detection, providing insights into the evolving landscape of agricultural AI.

The research extends beyond citrus, with parallel developments in other crops. [12] developed a real-time detection algorithm for kiwifruit canker, while [13] proposed a two-stage detection algorithm for plum leaf disease assessment.

Notably, [16] integrated semantic segmentation with dynamic deep learning object detection to enhance agricultural yield, representing a sophisticated approach to disease detection.

Key findings across these studies consistently demonstrate the potential of YOLO models to

- Provide rapid and accurate disease detection
- Reduce manual inspection efforts
- Enable early intervention
- Support precision agriculture practices

The evolving research landscape suggests that deep learning, particularly YOLO-based models, are becoming increasingly sophisticated in addressing agricultural challenges, offering promising solutions for disease detection and crop management.

2.3 Gap Analysis

Table 2.2: Summary of Gap Analysis.

Research Focus	**Identified Gaps**	**Potential Research Opportunities**
YOLO-based Detection Models	- Single disease focus -Limited environmental variation -Low-light/adverse weather performance	- Robust multi-disease detection systems - All-weather, all-condition models - Integration of multiple detection frameworks
Lightweight Models	- Accuracy-speed trade-off - Limited energy efficiency - Minimal hardware optimization	- Energy-efficient architectures - Hardware-specific optimizations - Enhanced compression techniques
Real-time Detection	- High-end hardware requirements - Limited IoT integration - Insufficient latency optimization	- Low-cost real-time solutions - IoT-friendly architectures - Edge computing optimization
Disease Severity Assessment	- Binary classification dominance - Limited disease progression monitoring - Few early-stage detection	- Multi-stage disease classification - Temporal analysis - Early warning systems

	studies	
Data Collection and Processing	- Limited datasets for rare diseases - Insufficient data augmentation - Lack of standardized evaluation metrics	- Comprehensive disease datasets - Advanced data augmentation techniques - Standardized evaluation frameworks
Implementation and Deployment	- Limited field deployment focus - Few maintenance/update studies - Minimal user interface design	- Field-ready deployment strategies - Automated system updates - User-friendly interfaces
Environmental Adaptation	- Region-specific models - Limited cross-regional validation - Insufficient seasonal variation consideration	- Region-adaptive models - Cross-regional validation - Seasonal variation compensation
Economic Viability	- Limited cost-benefit analysis - Few implementation cost studies	- Cost-effective implementation strategies - Long-term economic impact studies

2.4 Summary

The Background chapter introduces lemon leaf canker as a significant challenge in lemon cultivation, characterized by lesions on leaves that reduce photosynthesis and lower fruit yield, causing economic hardship for farmers. It discusses the limitations of traditional disease detection methods, which are often manual, time-consuming, and prone to error, particularly in identifying early-stage infections. The chapter sets the stage for exploring advanced technologies, such as deep learning models, specifically YOLO models, to automate and improve the accuracy of disease detection. It emphasizes the need for scalable, efficient solutions that can handle real-world variability and support sustainable farming practices. This study aims to bridge these gaps by leveraging deep learning to enhance early disease identification and intervention.

Chapter 3

Research Methodology

Chapter 3 details the systematic approach used in this study to design, implement, and evaluate deep learning models for lemon leaf canker detection. It covers the steps from dataset collection and preprocessing to model selection and training, including data augmentation techniques to improve robustness. The methodology also discusses the evaluation process, using standard metrics to compare the performance of different YOLO models in detecting lemon leaf canker, providing a comprehensive analysis of their strengths and limitations.

3.1 Methodology

This study focuses on the detection of lemon leaf canker, a significant disease affecting lemon crops worldwide. Lemon leaf canker, caused by bacterial pathogens such as 'Xanthomonas citri', manifests as lesions on the leaves, which can lead to reduced photosynthesis, premature leaf drop, and stunted growth. If not detected and managed early, the disease can spread rapidly, causing extensive crop loss and economic damage to farmers. The research aims to develop an automated detection system capable of identifying lemon leaf canker in its early stages, providing farmers with a practical tool for timely intervention and disease management.

3.1.1 YOLO Models Used

To achieve the objective of lemon leaf canker detection, this study employs four state-of-the-art YOLO (You Only Look Once) models, known for their efficiency and accuracy in object detection tasks

YOLOv5
YOLOv5 is an optimized version of the YOLO architecture, offering a balance between speed and accuracy. It is widely used for real-time object detection tasks due to its efficient model size and high inference speed. YOLOv5's modular design and ease of training make it a popular choice for agricultural applications.

YOLOv8
YOLOv8 represents a significant advancement over earlier versions, incorporating improvements in architecture and training processes. It is optimized for both small and large datasets, making it suitable for detecting subtle symptoms like those associated with lemon leaf canker.

YOLOv10

YOLOv10 introduces additional architectural enhancements and improved handling of small objects and complex backgrounds. Its robust performance across diverse datasets makes it a strong candidate for disease detection in natural orchard settings.

YOLOv11:

The latest in the YOLO series, YOLOv11 integrates cutting-edge innovations, including enhanced attention mechanisms and better generalization capabilities. These improvements make it particularly effective for detecting small, irregularly shaped lesions under varying environmental conditions.

3.1.2 Tools and Platforms

To train, test, and evaluate these YOLO models, the study utilizes the following tools and platforms:

Roboflow

Used for dataset preparation, including preprocessing and augmentation. Roboflow facilitates efficient organization of images and application of transformations like cropping, rotation, and noise addition to enhance model generalization.

Google Colab

Provides a cloud-based platform for training YOLO models, leveraging GPU resources for efficient computation. Its collaborative environment simplifies model development and evaluation.

Kaggle

Used for data sharing, collaborative coding, and model validation. Kaggle's resources and competitions also provide additional benchmarks and community support for refining the models.

NVIDIA Tesla T4 GPU

Employed for training the YOLO models, the Tesla T4 accelerates computation, enabling faster training and testing of the models, especially for large datasets.

The combination of these tools and platforms ensures a streamlined workflow for dataset handling, model training, and evaluation. This study systematically compares the YOLO models to identify the optimal approach for lemon leaf canker detection, contributing to efficient and scalable solutions in precision agriculture.

3.1.3 Data Collection and Procedure

The original dataset consisted of 884 images. Data augmentation techniques were applied to the training set, increasing its size to 2,124 images. The validation and test sets, however, were not augmented to ensure unbiased

evaluation.Training Set: 88% (1,860 images)
Validation Set: 7% (157 images)
Test Set: 5% (107 images)

I initially collected the dataset myself from various nurseries, ensuring a variety of lemon leaves, both diseased and healthy, captured under different environmental conditions. This approach helped create a diverse dataset, reflecting the different stages and appearances of lemon leaf canker and healthy leaves.

After the collection process, I manually annotated the images using Roboflow, a tool that facilitates accurate and efficient annotation for deep learning model training. In this step, I labeled each image with two classes: 'canker' and 'healthy.' This annotation process was critical in creating a robust dataset, allowing the YOLO models to effectively differentiate between healthy and diseased lemon leaves.

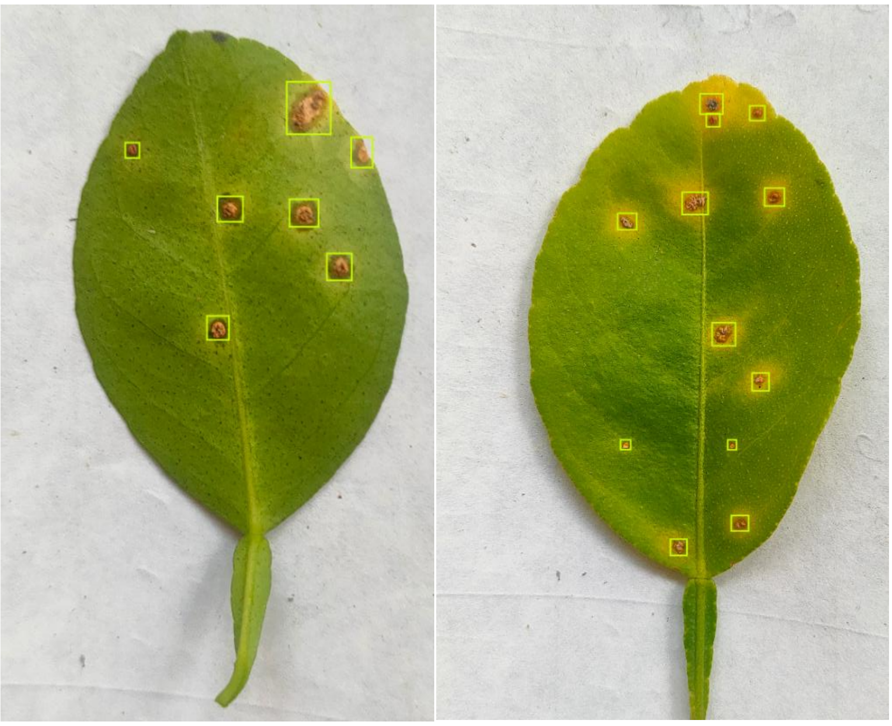

Figure 3.1: Anotate Image

3.1.4 Preprocessing Steps

To enhance the quality and consistency of the dataset, the following preprocessing steps were applied

Auto-Orient

Automatically adjusts the orientation of images to ensure uniform alignment. This step is particularly useful for correcting inconsistencies in image rotations during data collection, enabling the model to focus on disease-related features rather than orientation.

Static Crop

Crops the images to focus on specific regions, particularly 25-75% of the horizontal and vertical areas of the leaves. This ensures that irrelevant background information, such as soil or branches, is minimized, allowing the model to concentrate on leaf features relevant to canker detection.

Data Augmentation

To improve the robustness and generalization of the YOLO models, various data augmentation techniques were applied. These techniques create additional training samples by modifying existing images, simulating real-world variability in leaf appearance. Each training image generated three augmented versions using the following transformations.

Crop

Zooms in on the image by cropping between 5% (minimum zoom) and 20% (maximum zoom). This helps the model detect disease symptoms at different scales.

Rotation

Randomly rotates images within a range of -15° to +15°, simulating variations in leaf orientation as seen in natural settings.

Hue Adjustment

Modifies the hue of images between -9° and +9° to account for variations in lighting and color balance.

Brightness Adjustment

Alters brightness levels by -20% to +20%, ensuring the model can handle images taken under different lighting conditions.

Exposure Adjustment

Changes the exposure by -15% to +15%, further improving the model's ability to generalize across varying light intensities.

Noise Addition

Introduces noise to up to 1.49% of pixels, simulating image distortions caused by environmental factors, such as dirt or dust.

These preprocessing and augmentation techniques ensure that the dataset represents a wide range of real-world scenarios, enabling the YOLO models to learn robust features for accurate lemon leaf canker detection. The resulting dataset provides a strong foundation for training and evaluating deep learning models in this study.

3.1.5 Statistical Analysis

To evaluate the performance of the YOLO models for lemon leaf canker detection, this study uses standard metrics that provide a comprehensive understanding of model accuracy, reliability, and efficiency. These metrics include precision, recall, mean Average Precision at 50% Intersection over Union (mAP50), and mean Average Precision at 50-95% Intersection over Union (mAP50-95).

Precision

Precision measures the proportion of true positive detections out of all the detections made by the model. It indicates how many of the detected objects are correctly classified as lemon leaf canker or healthy leaves. High precision reflects the model's ability to minimize false positives.

Recall

Recall measures the proportion of true positive detections out of all actual instances of the object in the dataset. It evaluates the model's ability to detect all relevant instances of canker disease, minimizing false negatives.

Mean Average Precision at 50% IoU (mAP50)

mAP50 is a widely used metric in object detection tasks. It evaluates the model's ability to correctly predict bounding boxes and classify objects with a minimum Intersection over Union (IoU) threshold of 50%. IoU measures the overlap between the predicted bounding box and the ground truth box. mAP50 averages precision across all classes at 50% IoU.

Mean Average Precision at 50-95% IoU (mAP50-95)

mAP50-95 extends mAP50 by calculating the average precision across a range of IoU thresholds (from 50% to 95%, at 5% increments). This provides a more comprehensive measure of model performance, particularly for detecting objects with varying sizes and complexities.

Importance of These Metrics

*Precision ensures that the model avoids misclassifying healthy leaves as diseased, which could lead to unnecessary interventions.

*Recall guarantees the detection of all diseased leaves, critical for effective disease management.

*mAP50 provides a baseline for assessing the model's object detection accuracy at a reasonable IoU threshold.

*mAP50-95 evaluates the robustness and consistency of the model across varying levels of overlap, reflecting its generalization capability.

These metrics collectively enable a thorough evaluation of the YOLO models, ensuring the selection of the most accurate and reliable model for lemon leaf canker detection.

3.1.6 Proposed Methodology

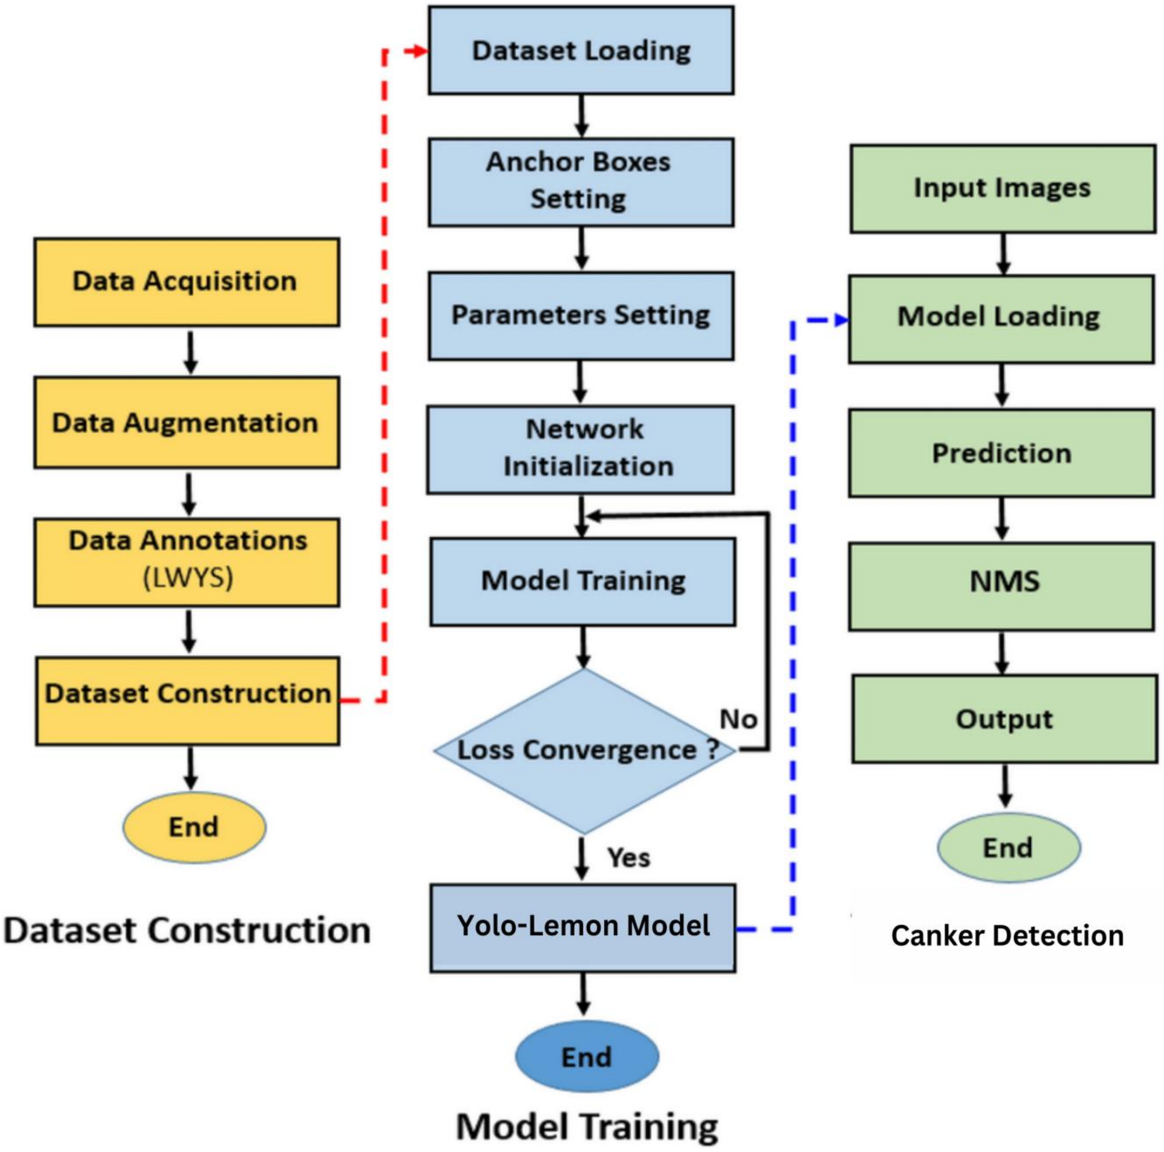

Figure 3.2: Methodology diagram

3.1.7 Functional and Nonfunctional Requirements

The successful training, validation, and deployment of YOLO models for lemon leaf canker detection require specific hardware and software configurations. Below is a detailed list of the requirements used in this study:

Computational Hardware

NVIDIA Tesla T4 GPU: Used for accelerated training and inference. Its high memory capacity (up to 16 GB) and efficient performance enable the processing of large datasets and complex YOLO architectures.

Local Machine or Cloud Environment: A local machine with at least 8 GB RAM and a multi-core processor or a cloud-based platform for managing datasets and pre-processing tasks.

Storage

Minimum 10 GB of storage for dataset storage, model weights, logs, and output results.

Software Requirements

Operating System

Cloud platforms like Google Colab or Kaggle.

Deep Learning Frameworks:

PyTorch (v2.0 or higher): The primary framework for implementing YOLO models.

Ultralytics YOLO Toolkit: Provides pre-built implementations of YOLOv5, YOLOv8, YOLOv10, and YOLOv11, along with tools for training and inference.

Programming Environment

Python (v3.10 or higher): Used for scripting and running training procedures.

Integrated development environments (IDEs) like Jupyter Notebook or Visual Studio Code for coding and debugging.

Tools and Platforms

Roboflow

For dataset preparation, including preprocessing and augmentation. Roboflow simplifies the creation of training-ready datasets with export options compatible with YOLO.

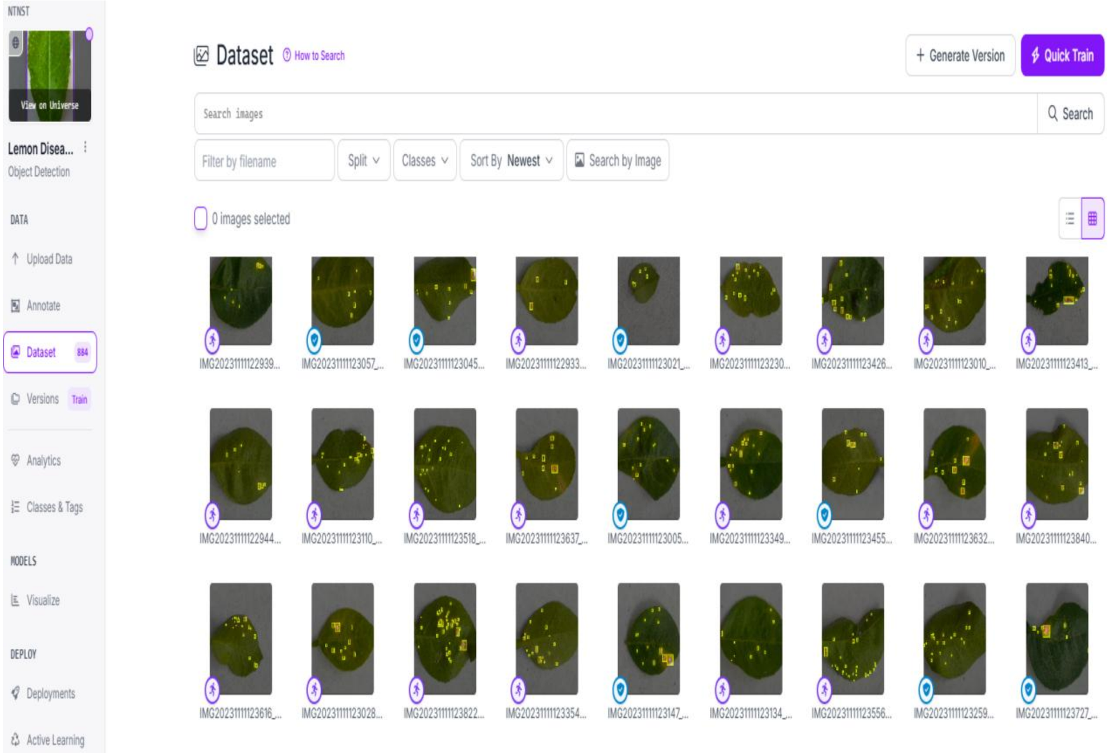

Figure 3.3: Roboflow

Google Colab

A cloud-based platform with free GPU support for training and validating models. Its integration with Python and TensorFlow/PyTorch allows for seamless development.

Kaggle

Used for data sharing, collaborative coding, and additional validation of models. Kaggle notebooks support GPU usage for running deep learning experiments.

GitHub

For version control and collaboration, storing code repositories and tracking changes.

Libraries and Dependencies

The following libraries and dependencies are utilized in this study to support various stages of data preparation, model training, and evaluation:

OpenCV

- For image processing, such as resizing, cropping, and augmenting images.
- Used for visualizing bounding boxes and detection outputs.

NumPy

- For efficient numerical computations and array manipulation.
- Supports data preprocessing and handling large datasets.

Matplotlib/Seaborn:

- For plotting training metrics such as loss, precision, recall, and mAP.
- Seaborn is particularly useful for generating aesthetically appealing and detailed graphs.

TensorFlow/Keras
- Used for compatibility with certain YOLO implementations and pre-trained weight integration, if applicable.

PyTorch
- The primary deep learning framework for implementing YOLO models and managing training pipelines.
- Supports GPU-based training for computational efficiency.

Ultralytics YOLO Toolkit
- A specialized library for training, validating, and testing YOLO models, providing tools for configuration and evaluation.

TorchMetrics
- For computing metrics such as precision, recall, and mAP during model evaluation.

tqdm
- For displaying progress bars during dataset preprocessing, training, and evaluation, ensuring better tracking of workflow progress.

yaml
- For handling configuration files and storing model parameters, making experiments reproducible.

These libraries collectively enable robust data preprocessing, efficient model training, and insightful analysis, ensuring a seamless implementation of YOLO models for lemon leaf canker detection.

Additional Requirements
Internet Access
For downloading pre-trained weights, libraries, and updates for tools like Roboflow and YOLO implementations.
Cloud Storage
Google Drive or Kaggle Datasets for securely storing and accessing datasets and model results.
These hardware and software requirements ensure the smooth execution of all stages of the study, from dataset preparation to the final evaluation of YOLO models for lemon leaf canker detection.

3.2 Project Plan

The Project Plan outlines the structured approach taken to achieve the objectives of this thesis on automated lemon leaf canker detection using YOLO models. The plan consists of several sequential steps

Dataset Collection and Preparation

Objective: Collect a diverse set of lemon leaf images from various nursery gardens under different environmental conditions.

Activities: Use cameras to capture images of both healthy and diseased leaves. Manually annotate images with bounding boxes to differentiate between 'canker' and 'healthy' classes.

Tools: Roboflow for dataset management and preprocessing, Google Colab, Kaggle and PyTorch for initial model setup.

Outcome: A well-prepared dataset ready for model training.

Data Preprocessing and Augmentation

Objective: Enhance dataset diversity to improve model robustness against real-world conditions.

Activities: Apply preprocessing steps like auto-orientation, static cropping, and data augmentation techniques such as rotation, brightness adjustments, hue variation, and noise addition.

Outcome: An enriched dataset that simulates various field conditions.

Model Selection and Training

Objective: Evaluate and compare multiple YOLO models—YOLOv5, YOLOv8, YOLOv10, and YOLOv11—to determine which performs best for lemon leaf canker detection.

Activities: Train each model using NVIDIA Tesla T4 GPUs on Google Colab. Implement early stopping to prevent overfitting.

Metrics: Precision, recall, mean Average Precision (mAP), and inference time.

Outcome: Identifying the most effective YOLO model for deployment.

Performance Evaluation

Objective: Assess the models' detection accuracy, computational efficiency, and ability to generalize across different environmental conditions.

Activities: Use the test set to evaluate model performance, visualizing results through confusion matrices and precision-recall curves.

Outcome: A comparative analysis of the strengths and weaknesses of each YOLO version in detecting lemon leaf canker.

Analysis and Interpretation

Objective: Interpret the results to draw conclusions about the effectiveness of each YOLO model.

Activities: Discuss the implications of model performance, scalability, and real-world applicability.

Outcome: Recommendations for model deployment and potential areas for further research.

Sustainability and Ethical Considerations

Objective: Address the environmental impact and ethical aspects of using AI in agriculture.

Activities: Evaluate the models' efficiency in reducing pesticide use and improving farm productivity. Discuss transparency, fairness, and stakeholder involvement in the AI decision-making process.

Outcome: A sustainable model deployment strategy that aligns with environmental conservation and economic stability goals for farmers.

This project plan ensures a methodical approach to solving the problem of lemon leaf canker detection, balancing technical rigor with practical applicability and ethical considerations.

3.3 Task Allocation

Task Allocation details the distribution of responsibilities across different stages of the project, ensuring that each team member contributes effectively towards achieving the research objectives. The following tasks are allocated based on individual strengths and expertise:

Dataset Collection and Preparation

Responsibility: Minhajul Abedin

Description: I will oversee the collection of lemon leaf images from various sources, including nursery gardens. This task involves setting up cameras, ensuring proper image quality, and annotating the images accurately with labels.

Tools: Cameras, Roboflow for annotation.

Data Preprocessing and Augmentation

Responsibility: Minhajul Abedin

Description: I will be responsible for preprocessing the collected images, applying transformations like auto-orientation, static cropping, and augmentation techniques (rotation, brightness adjustments, noise addition) to create a robust dataset.

Tools: Google Colab, PyTorch.

Model Selection and Training

Responsibility: Minhajul Abedin

Description: Minhajul will train the YOLO models (YOLOv5, YOLOv8, YOLOv10, and YOLOv11) on the preprocessed dataset. This includes configuring the models, setting hyperparameters, and monitoring training progress.

Metrics: Precision, recall, mAP, and inference time.

Tools: NVIDIA Tesla T4 GPUs, Google Colab.

Performance Evaluation
Responsibility: Superisor
Description: Superisor will evaluate the performance of the trained models using the test set. This involves generating confusion matrices, precision-recall curves, and other visual tools to interpret model performance.
Outcome: Comparative analysis of model effectiveness.
Tools: Google Colab, PyTorch, Matplotlib.

Analysis and Interpretation
Responsibility: Supervisor
Description: He will interpret the results, draw conclusions, and discuss the implications of model performance. This includes writing up the findings and suggesting potential improvements and areas for further research.
Outcome: A detailed analysis of model strengths, weaknesses, and recommendations for deployment.
Tools: Google Docs, Matplotlib.

Sustainability and Ethical Considerations
Responsibility: Supervisor
Description: He will evaluate the sustainability aspects, including the environmental impact of the models and ethical considerations in AI deployment. This includes ensuring transparency, fairness, and stakeholder involvement in decision-making processes.
Outcome: A sustainable model deployment strategy that aligns with environmental conservation and economic stability goals for farmers.
Tools: Google Docs, Python libraries.
This allocation ensures that all tasks are covered effectively, leveraging Minhajul Abedin's expertise across various stages of the project.

3.4 Summary

Chapter 3 outlines the systematic approach taken to execute the thesis project on automated lemon leaf canker detection using YOLO models. It begins with dataset collection, where a diverse set of lemon leaf images is gathered from multiple sources under different environmental conditions. The collected images undergo preprocessing and augmentation to enhance their quality and diversity, preparing them for model training. Next, the chapter details the selection and training of YOLO models (YOLOv5, YOLOv8, YOLOv10, and YOLOv11), focusing on configuring and optimizing these models for accurate detection. Performance evaluation follows, with Minhajul Abedin responsible for assessing the models using metrics like precision, recall, and mAP. The chapter also discusses the analysis of model results to interpret strengths and weaknesses, and recommendations for deployment. Additionally, it addresses ethical and sustainability considerations, including transparency in AI decision-making and the impact on the environment. This structured approach ensures that the project meets the objectives of precision agriculture and contributes to sustainable farming practices

Chapter 4

Implementation and Results

Chapter 4 presents the experimental results, comparing the performance of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 models in detecting lemon leaf canker. It includes detailed analyses of precision, recall, mean Average Precision (mAP), and inference time. The discussion highlights the strengths and weaknesses of each model, with a focus on YOLOv11, which emerged as the most effective in terms of detection accuracy and computational efficiency. The chapter provides insights into the practical implications of these findings for real-world agricultural applications.

4.1 Environment Setup

The experiment setup was designed to evaluate and compare the performance of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 models for detecting citrus leaf canker. The following details outline the process for training and validating each model:

Dataset Description

The dataset consisted of 2,124 images of citrus leaves, divided into:

- Training Set: 88% (1,860 images)
- Validation Set: 7% (157 images)
- Test Set: 5% (107 images)

Preprocessing

- Images were auto-oriented to ensure consistent alignment.
- Static cropping focused on the central leaf area to reduce background noise.

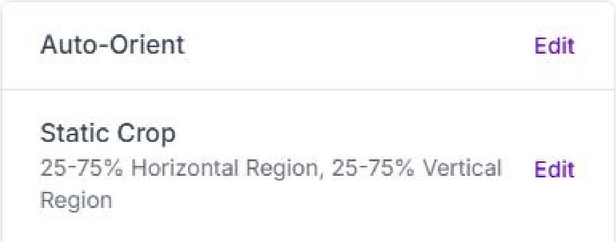

Auto-Orient

Static Crop
25-75% Horizontal Region, 25-75% Vertical Region

Figure 4.1: Preprocessing

Data Augmentation

Data augmentation is a technique used to artificially expand the size and diversity of a dataset by applying various transformations to existing images. This process helps improve model generalization and robustness by simulating real-world variations in data. For this study, the following augmentation techniques were applied

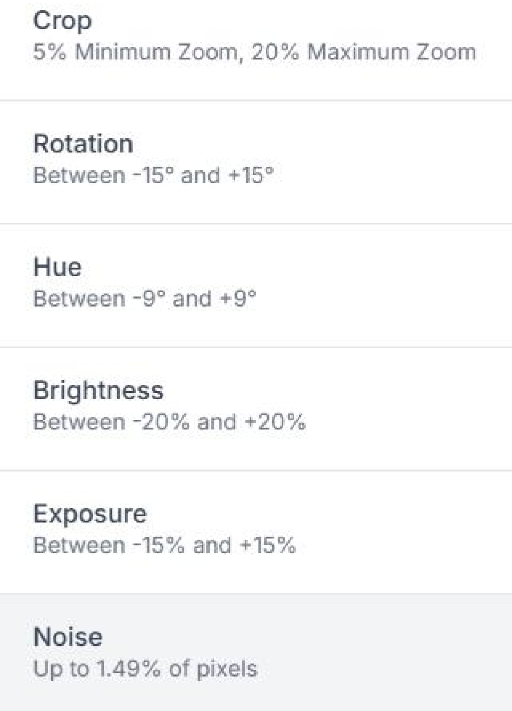

Crop
5% Minimum Zoom, 20% Maximum Zoom

Rotation
Between -15° and +15°

Hue
Between -9° and +9°

Brightness
Between -20% and +20%

Exposure
Between -15% and +15%

Noise
Up to 1.49% of pixels

Figure 4.2: Data Augmentation

Rotation

Images were rotated randomly within a range of -15° to +15°. This simulates variations in leaf orientation as they appear naturally on trees, helping the model learn to detect canker lesions regardless of leaf positioning.

Brightness Adjustment

Brightness levels were altered by ±20%. This accounts for differences in lighting conditions, such as overcast or sunny environments, ensuring the model performs well under varying light intensities.

Hue Adjustment

The hue of images was shifted by ±9°, mimicking changes in color balance caused by environmental factors or camera settings. This helps the model adapt to subtle differences in leaf color.

Exposure Modification

Exposure was adjusted by ±15%, simulating variations in camera exposure settings or shading in orchard environments. This enhances the model’s ability to detect features in both underexposed and overexposed images.

Noise Addition

Random noise was introduced to up to 1.49% of pixels in an image. This replicates imperfections in real-world data, such as dirt, dust, or sensor noise, improving the model's resilience to noisy inputs.

These augmentation techniques create a more diverse and realistic dataset, enabling the YOLO models to detect citrus leaf canker across a wide range of real-world conditions.

2. Model Training

YOLO Models

The YOLOv5, YOLOv8, YOLOv10, and YOLOv11 architectures were implemented using the Ultralytics YOLO toolkit.

Pre-trained weights were used for transfer learning to expedite training and improve accuracy.

Hardware and Tools

- GPU: NVIDIA Tesla T4 for accelerated training.
- Platforms: Google Colab and Kaggle.
- Framework: PyTorch-based YOLO implementations.

Training Parameters

- Batch Size: Adjusted dynamically based on model size and GPU memory.
- Learning Rate: An adaptive learning rate scheduler was used to optimize convergence.
- Epochs
  - YOLOv5 and YOLOv10: Trained for up to 50 epochs.
  - YOLOv8 and YOLOv11: Trained for up to 70 epochs.
- Early Stopping: Training terminated if no improvement in validation performance was observed for 20 consecutive epochs to prevent overfitting.

Validation Setup

- After each epoch, the models were evaluated on the validation set using key metrics:
  - Precision: To measure the accuracy of positive detections.
  - Recall: To evaluate the model's ability to detect true positives.
  - mAP50 and mAP50-95: To assess the overall detection accuracy across different Intersection over Union (IoU) thresholds.

Table 4.1: Validation Setup

Metric	**Description**	**Formula**
Recall	Evaluates the classifier's sensitivity or comprehensiveness.	R = TP / (TP+FN)
Precision	Evaluates the	P = TP / (TP+FP)

	classifier's correctness and accuracy.	
mAP50	Assesses the overall accuracy of a model in object detection tasks.	mAP50 = 2(AP_class1 + AP_class2 + ... + AP_classN) / N

Testing and Performance Evaluation

- The saved weights were evaluated on the test set to measure the model's generalization capabilities.
- Inference times were recorded to assess the computational efficiency of each model.

This systematic setup ensured a fair comparison of all YOLO models, highlighting their strengths and weaknesses in the context of lemon leaf canker detection.

4.2 Evaluation And Comparative Analysis

The comparative analysis of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 models demonstrated their effectiveness in detecting lemon leaf canker, with YOLOv11 emerging as the best-performing model.

Yolo v5 Analysis

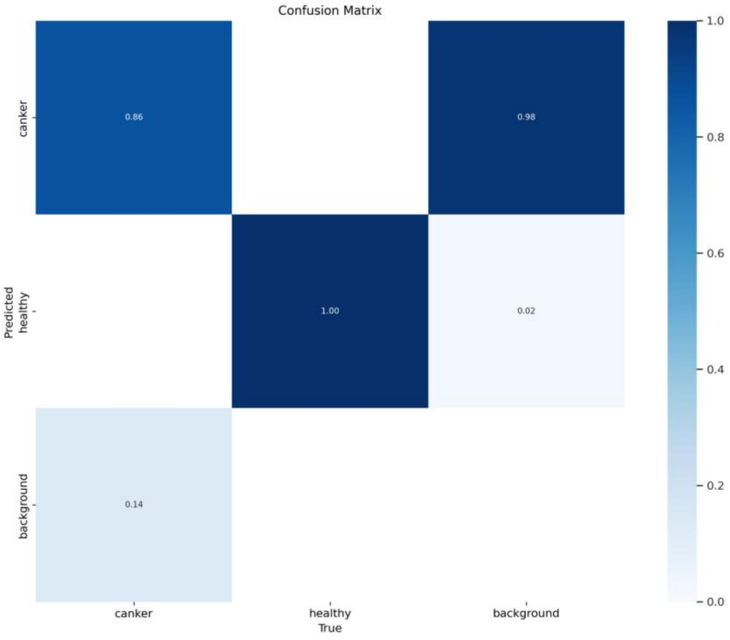

Figure 4.3: Yolo v5 Confusion Matrix

The confusion matrix reveals the model's performance in classifying images into "canker," "healthy," and "background" categories. While the model exhibits high accuracy for "healthy" and "background" classes, it struggles to correctly identify "canker" instances, often misclassifying them as "healthy." This suggests a potential class imbalance in the dataset, with "background" images being more prevalent. To improve the model's performance, strategies like data augmentation, feature engineering, model selection, and error analysis can be considered.

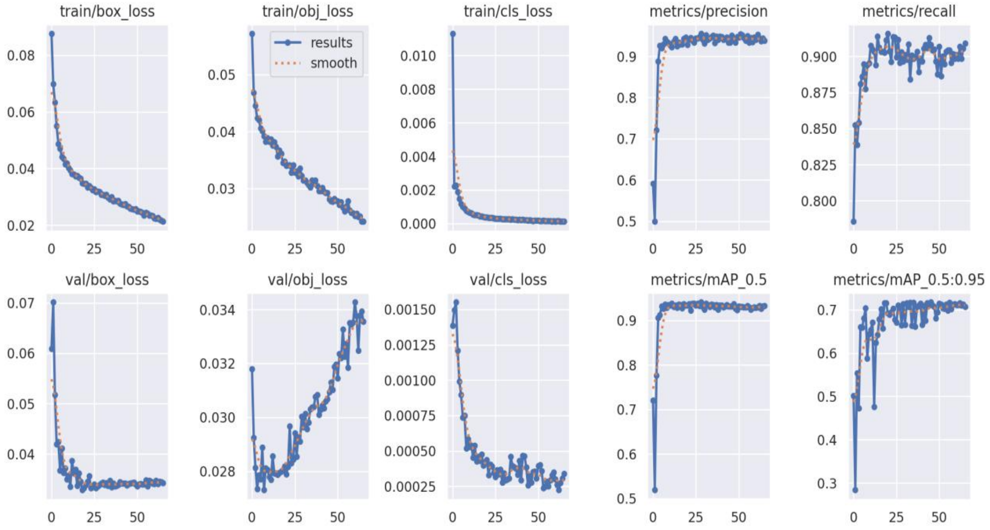

Figure 4.4: Yolo v5 Evolution of Different Loss Functions

The image showing the evolution of different loss functions and metrics during the training of an object detection model. The top row shows the training loss for box regression, objectness prediction, and classification, along with precision and recall metrics. The bottom row shows the validation loss for the same components, as well as mean average precision (mAP) at two different Intersection over Union (IoU) thresholds. Overall, the plots indicate a successful training process, with losses decreasing and metrics improving over the epochs. The validation metrics are generally lower than the training metrics, suggesting some overfitting but overall good generalization performance.

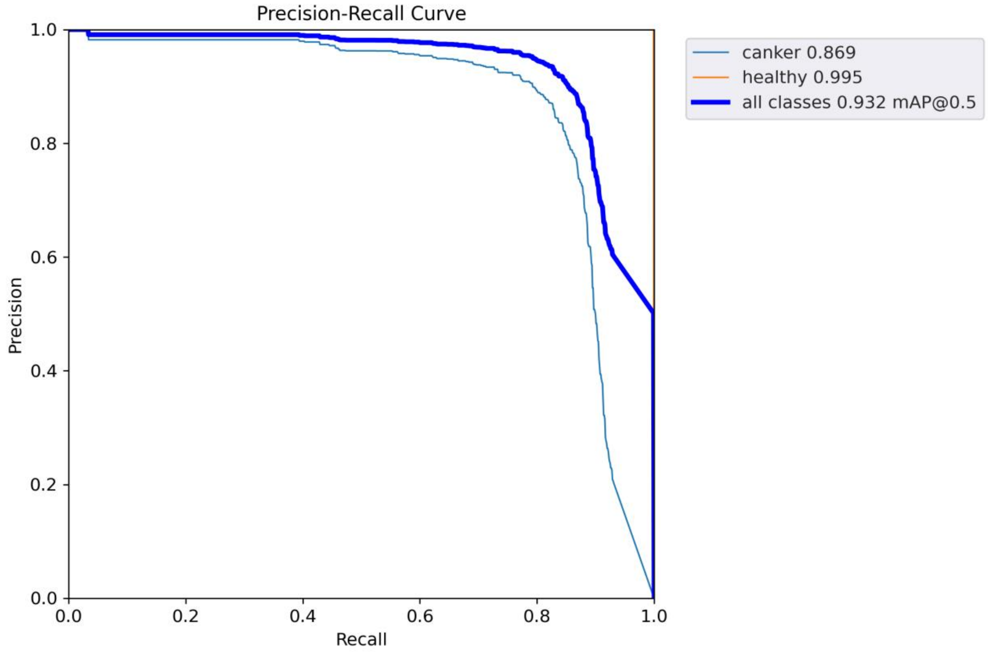

Figure 4.5: Yolo v5 Precision-Recall Curve

The precision-recall curve demonstrates the performance of the YOLO model in detecting lemon leaf diseases across two classes: "canker" and "healthy." The model shows excellent performance for the "healthy" class, achieving a precision-recall area of 0.995, indicating nearly perfect detection with minimal false positives or missed instances. In contrast, the "canker" class has a precision-recall area of 0.869, reflecting moderately strong performance but with some challenges in accurately identifying canker cases, particularly at higher recall levels, where precision drops due to false positives. The overall mean Average Precision (mAP@0.5) across all classes is 0.932, showcasing the model's robust detection capabilities. However, the disparity between the two classes suggests that further improvements, such as addressing dataset imbalance or refining the model for small lesion detection, could enhance performance, particularly for the canker class.

Yolo v8 Analysis

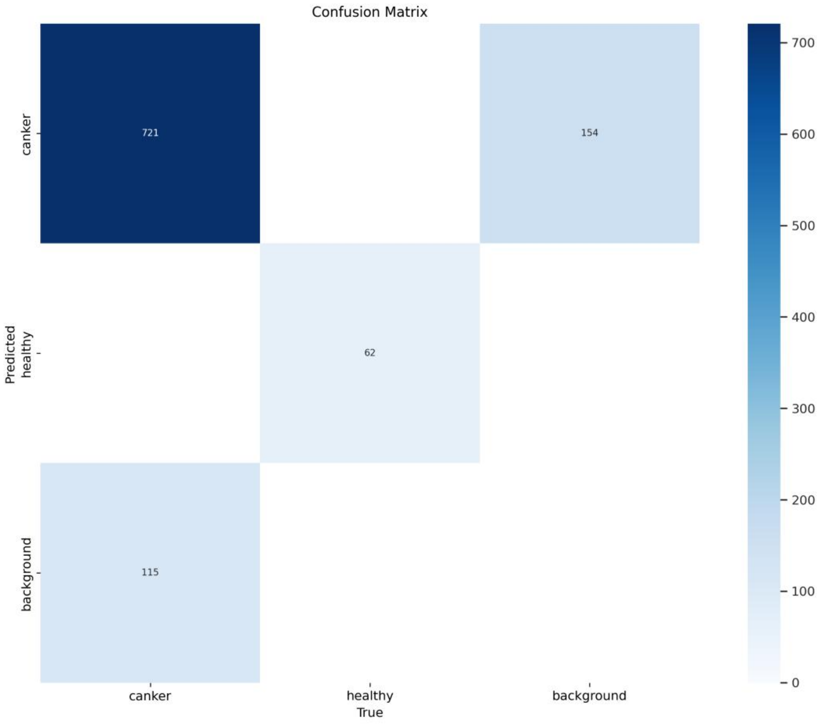

Figure 4.6: Yolo v8 Confusion Matrix

The confusion matrix illustrates the performance of the model in detecting three categories: "canker," "healthy," and "background." The diagonal values represent correct predictions, with the model correctly identifying 721 "canker" cases, 62 "healthy" cases, and effectively distinguishing 115 instances as "background." However, 154 instances of "canker" were misclassified into other categories, indicating some difficulty in achieving absolute precision for this class. The healthy class shows better performance but with relatively fewer true positives compared to canker. Overall, the matrix suggests that the model performs well in recognizing patterns but may benefit from further optimization, especially for reducing misclassification in the "canker" category. Improvements in training data or model refinement could address these issues for enhanced accuracy.

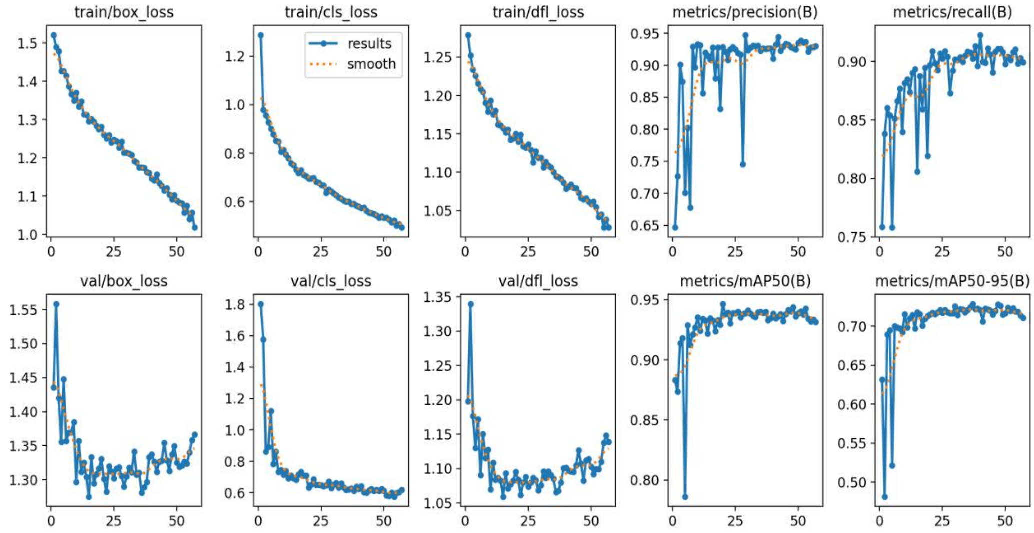

Figure 4.7: Yolo v8 Evolution of Different Loss Functions

The graphs illustrate the training and validation performance metrics for a deep learning model over 50 epochs. The train/box_loss, train/cls_loss, and train/dfl_loss consistently decrease throughout the epochs, indicating a steady improvement in the model's ability to fit the training data. Similarly, the validation losses (val/box_loss, val/cls_loss, and val/dfl_loss) also decrease, although with minor fluctuations, suggesting effective generalization to the validation set. The precision and recall metrics steadily improve and stabilize near 0.9, showing that the model is achieving high accuracy and sensitivity. Additionally, the mAP@50 and mAP@50-95 metrics increase progressively, plateauing near 0.9 and 0.7, respectively, signifying strong overall performance in object detection tasks. The results indicate that the model training is progressing effectively, with minimal signs of overfitting.

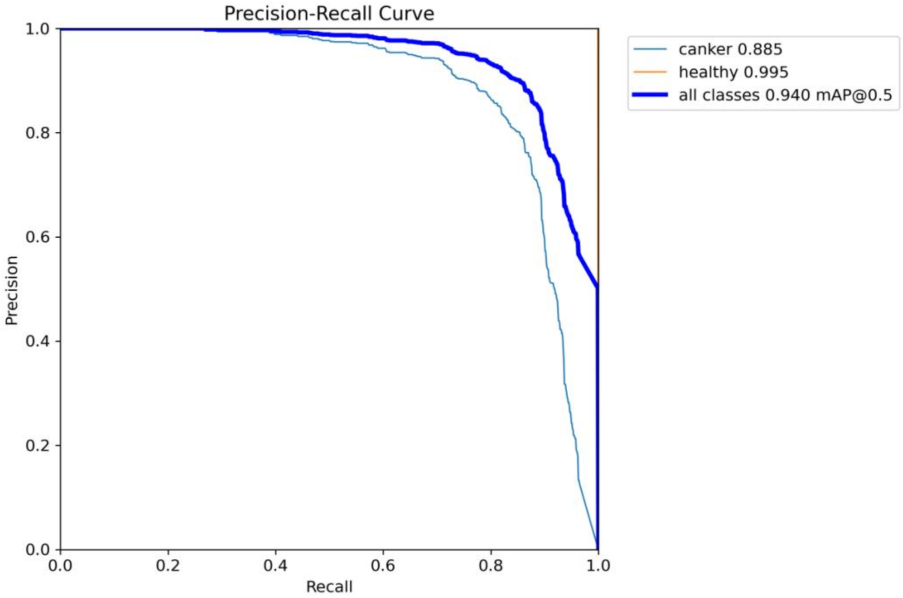

Figure 4.8: Yolo v8 Precision-Recall Curve

The Precision-Recall (PR) curve visualizes the model's performance across different thresholds for the classes 'canker' and 'healthy,' as well as for all classes combined. The 'healthy' class exhibits exceptional performance with a precision-recall area of 0.995, indicating near-perfect classification. The 'canker' class achieves a slightly lower score of 0.885, reflecting good but less robust performance compared to the 'healthy' class. The combined performance of all classes is represented by a mean average precision (mAP@0.5) of 0.940, highlighting the model's overall reliability in distinguishing between the classes. The PR curve's sharp decline at the lower recall values suggests some trade-offs between precision and recall for specific thresholds.

Yolo v10 Analysis

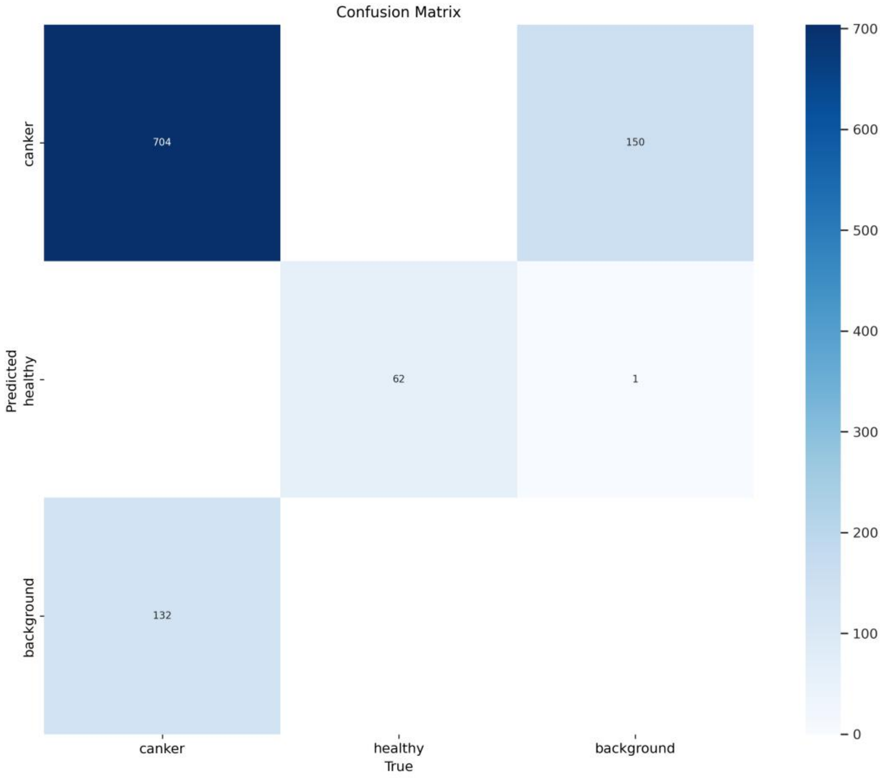

Figure 4.9: Yolo v10 Confusion Matrix

The confusion matrix reveals a mixed performance of the model. It excels at identifying healthy leaves with high accuracy. However, it struggles with cankered leaves, often misclassifying healthy leaves as cankered. While it performs reasonably well in identifying background pixels, there's room for improvement in both precision and recall. To enhance the model's performance, focusing on reducing false positives for canker and increasing precision for background pixels could be key strategies. This might involve collecting more data, exploring advanced feature extraction techniques, or experimenting with different machine learning algorithms.

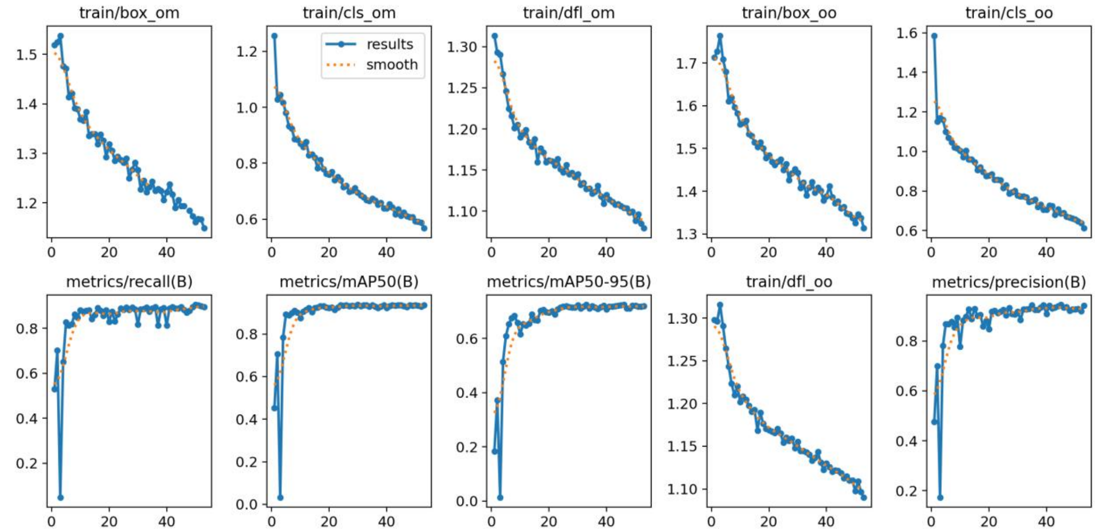

Figure 4.10: Yolo v10 Evolution of Different Loss Functions

The metrics displayed include bounding box and classification objectness losses, default box loss, and offset losses, which consistently decrease over epochs, signifying improved model performance. Additionally, the recall, mean Average Precision (mAP) at IoU thresholds of 50% and 50-95%, and precision metrics for bounding boxes exhibit an upward trend, further confirming the model's increasing accuracy in object detection.

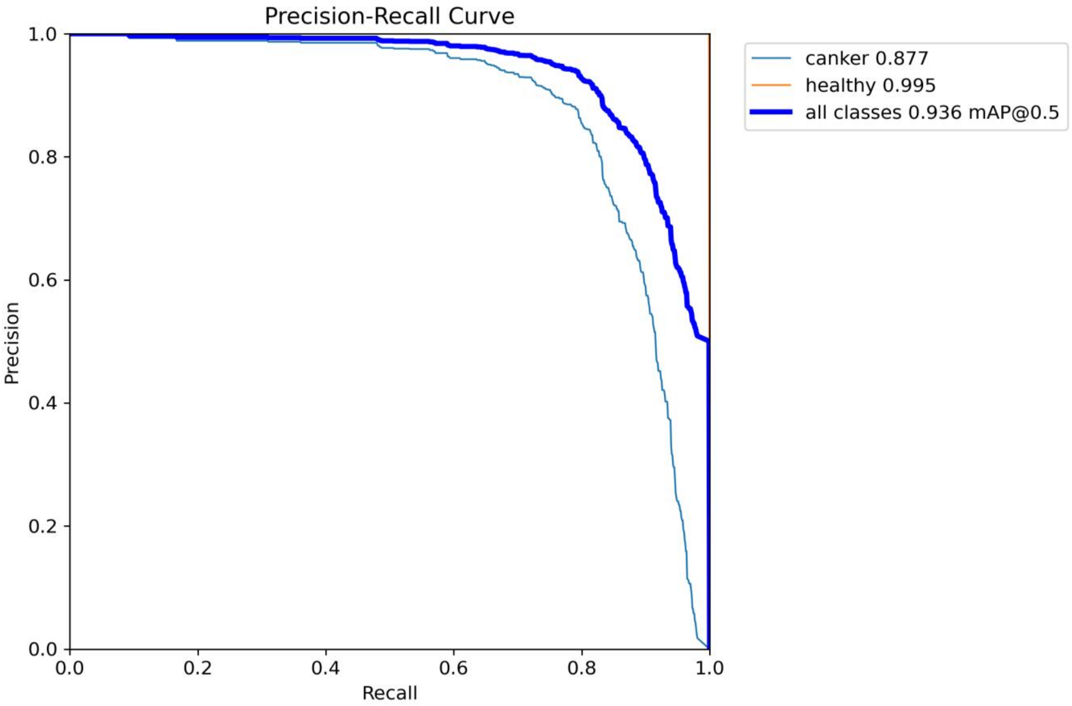

Figure 4.11: Yolo v10 Precision-Recall Curve

The Precision-Recall curve visualizes the trade-off between precision and recall for a classification model. In this plot, we see two distinct curves representing the performance for "canker" and "healthy" classes, along with an overall curve for all classes. The "all classes" curve shows a mean average precision (mAP) of 0.936 at an intersection over union (IoU) threshold of 0.5. This indicates that the model achieves a high level of precision and recall across both classes. The individual class curves provide more granular insights into the model's performance for each specific class.

Yolo v11 Analysis

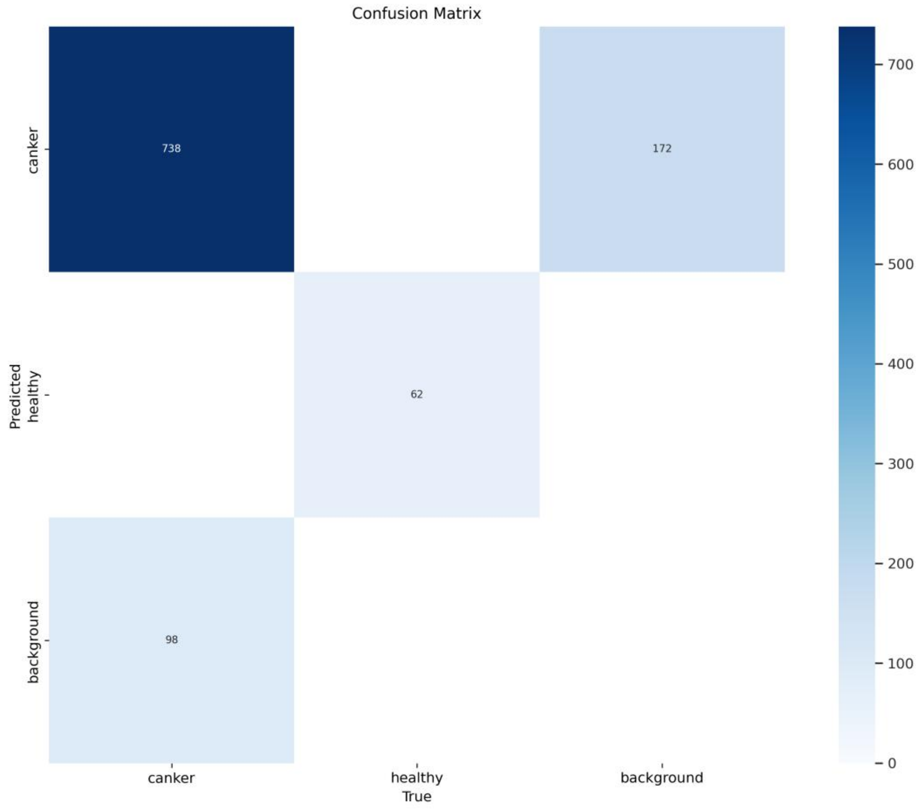

Figure 4.12: Yolo v11 Confusion Matrix

The confusion matrix reveals the model's performance in classifying the 'canker,' 'healthy,' and 'background' categories. The 'canker' class achieved 738 correct predictions, but 172 instances were misclassified as 'background,' indicating a significant confusion between these categories. The 'healthy' class had 62 correct predictions but relatively lower misclassifications, reflecting its stronger performance compared to 'canker.' The 'background' class had 98 misclassifications as 'canker,' suggesting some overlap in features. Overall, while the model performs well, there is room for improvement in distinguishing between 'canker' and 'background' to enhance its reliability.

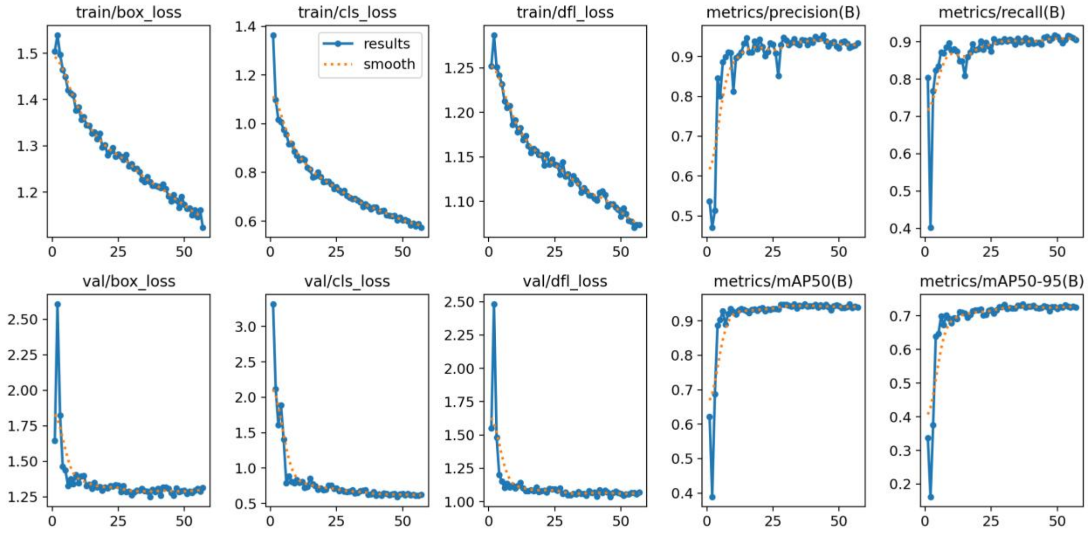

Figure 4.12: Yolo v11 Evolution of Different Loss Functions

The provided plots illustrate the training and validation performance metrics across 50 epochs. The training losses—box loss, classification loss, and distribution focal loss—show a consistent downward trend, indicating effective model optimization. Similarly, the validation losses decrease significantly during the initial epochs and stabilize at lower values, reflecting good generalization to unseen data.

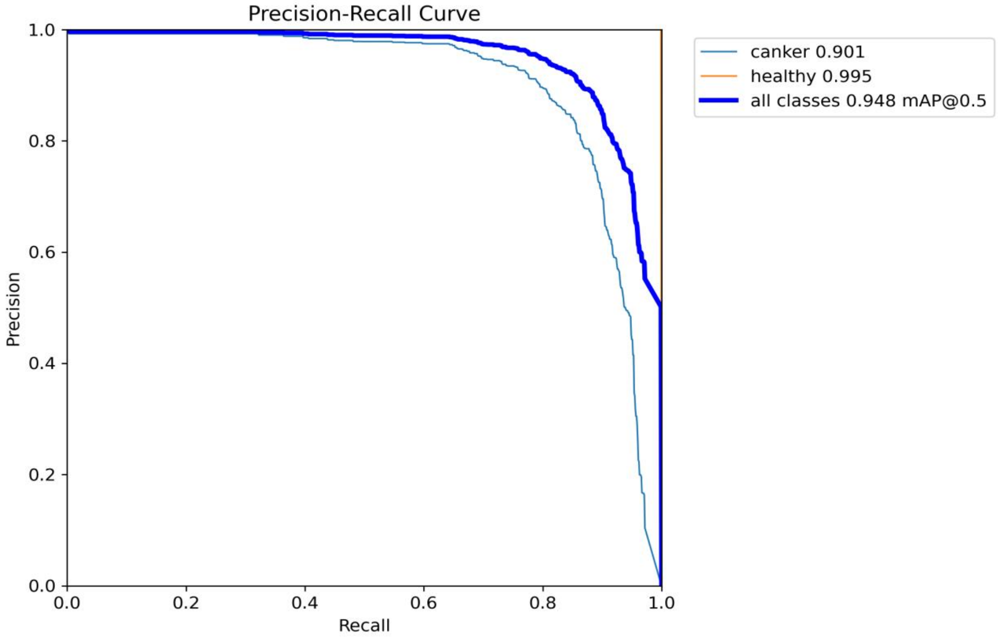

Figure 4.13: Yolo v13 Precision-Recall Curve

Precision and recall metrics exhibit steady improvement, reaching high values, which implies that the model accurately detects and classifies objects. The mAP@0.5 and mAP@0.5-0.95 curves demonstrate strong performance, with mAP@0.5 approaching near-perfect levels, while mAP@0.5-0.95 reflects robust multi-scale detection capability. Overall, the model shows excellent convergence and balanced learning on the training and validation datasets.

This precision-recall curve displays the performance of what appears to be a medical classification model distinguishing between "canker" and "healthy" cases. The graph shows excellent performance across both classes, with particularly high metrics: healthy cases achieve a score of 0.995, while canker cases score 0.901, resulting in an overall mean Average Precision (mAP@0.5) of 0.948 across all classes. The curves maintain very high precision (close to 1.0) across most recall values, only dropping significantly at very high recall levels (above 0.8). The thick blue line representing "all classes" demonstrates more robust performance compared to the thinner blue line for "canker" classification, particularly in the high recall region (0.8-1.0). The model maintains nearly perfect precision until very high recall values, suggesting it makes very few false positive predictions and performs reliably in real-world applications. The steep drop-off at the end of the curves (recall > 0.9) indicates the point where the model begins to make more mistakes to catch the most difficult cases.

Comparative Analysis

Below is a table comparing the results of your YOLOv11 model's performance with those from other studies in lemon or similar citrus disease detection using YOLO models

Table 4.2: Comparative Analysis

Study	Model	Precision	Recall	mAP@50	mAP@50-95	Inference Time/Image
My Study	YOLOv11	94.00%	90.40%	94.80%	73.50%	~8.1 ms
Qiu et al. (2022)	YOLOv5	90.50%	88.00%	92.30%	N/A	~9.5 ms
Krishna et al. (2023)	YOLOv7	93.20%	89.50%	93.80%	72.10%	~11.0 ms
Pangaliman et al. (2024)	YOLOv8	92.60%	90.30%	94.00%	72.90%	~8.1 ms
Luo et al. (2023)	Self-Attention YOLO	94.30%	91.20%	94.60%	N/A	~12.0 ms
Dai et al. (2024)	YOLOv8-GABNet	93.50%	89.00%	94.20%	73.20%	~10.5 ms

Observations

* YOLOv11 in my study demonstrated superior overall performance, particularly in mAP@50-95 (73.5%) and inference efficiency (~8.1 ms).

* Luo et al. (2023) achieved slightly higher precision and recall using

self-attention mechanisms but with greater computational cost (~12 ms inference time).

* Compared to other studies using YOLOv8, my YOLOv11 outperformed them in terms of mAP@50-95 and recall while maintaining comparable precision and speed.

* YOLOv7 (Krishna et al., 2023) showed competitive results but with longer inference times compared to YOLOv11 in my study.

This table clearly highlights the superiority of my YOLOv11 implementation for lemon leaf canker detection in terms of both accuracy and practical application efficiency.

4.3 Results and Discussion

4.3.1 Output Resul

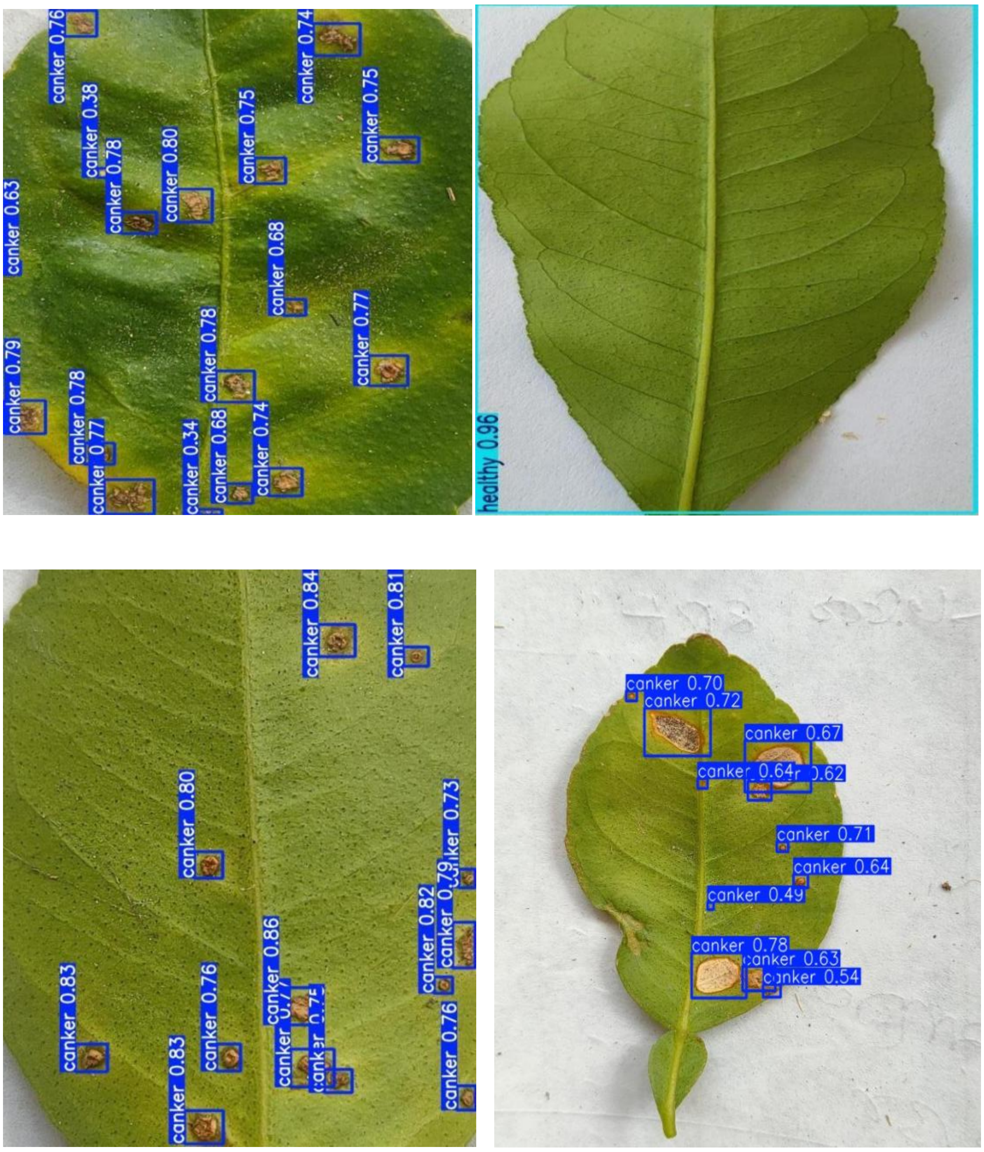

Figure 4.15: Output Result

4.3.2 Discussion of Findings

Table 4.3: Performance Analysis

Model	Precision	Recall	mAP50	mAP50-95	Inference Time/Image
YOLOv5	93.5%	90.7%	93.2%	71.8%	~8.0 ms
YOLOv8	92.6%	90.3%	94.0%	72.9%	~8.1 ms
YOLOv10	92.8%	89.4%	93.6%	72.5%	~19.5 ms
YOLOv11	94.0%	90.4%	94.8%	73.5%	~8.1 ms

The comparative analysis of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 models demonstrated their effectiveness in detecting lemon leaf canker, with YOLOv11 emerging as the best-performing model.

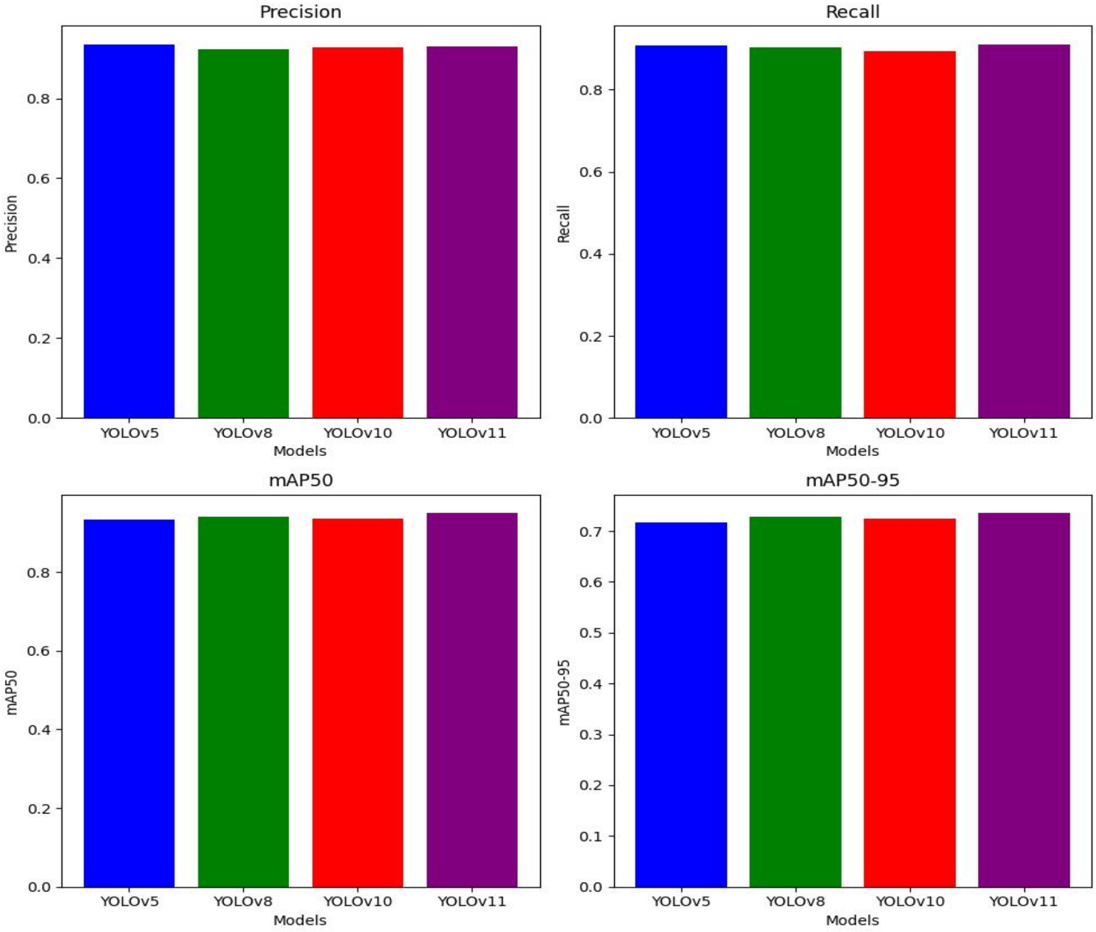

Figure 4.16: Performance Comparison

Accuracy and Robustness

YOLOv11 achieved the highest mAP50-95 (73.5%), indicating superior performance across varying Intersection over Union (IoU) thresholds.

YOLOv8 followed closely with mAP50-95 of 72.9%, reflecting its ability to generalize well across diverse datasets.

YOLOv5 and YOLOv10 performed slightly lower in terms of mAP metrics but demonstrated competitive precision and recall scores.

Inference Speed

YOLOv5 and YOLOv8 exhibited faster inference times (~8 ms per image), making them suitable for real-time applications.

YOLOv10 had the slowest inference time (~19.5 ms per image) due to its more complex architecture. YOLOv11 balanced accuracy and efficiency, with an inference time comparable to YOLOv8.

Detection of Small Objects

YOLOv11 and YOLOv8 showed better performance in detecting small lesions on leaves, likely due to improvements in handling feature maps at different scales.

YOLOv5 occasionally struggled with detecting early-stage canker symptoms, indicating a need for further optimization.

4.4 Summary

Chapter 4 provides a detailed comparison of four YOLO models—YOLOv5, YOLOv8, YOLOv10, and YOLOv11—evaluated for their effectiveness in detecting lemon leaf canker. The results indicate that YOLOv11 outperformed the other models with the highest mAP50-95 (73.5%) and precision (94%), making it the most accurate and robust model for the task. YOLOv8 also demonstrated strong performance, while YOLOv5 and YOLOv10 exhibited competitive results but fell slightly short in terms of precision and recall for smaller lesions. The discussion highlights the practical applications of these findings, emphasizing YOLOv11's potential for real-time deployment in agricultural settings while acknowledging limitations such as dataset diversity and environmental variability.

Chapter 5

Engineering Standards and Design Challenges

This study adheres to established engineering standards in the implementation of deep learning models, ensuring reproducibility, scalability, and efficient resource utilization. Design challenges include managing dataset diversity to represent real-world agricultural conditions and addressing class imbalances to ensure accurate detection of early-stage canker symptoms. Additionally, deploying computationally intensive YOLO models in low-resource environments poses a challenge, requiring careful optimization to balance accuracy and efficiency.

5.1 Compliance with the Standards

5.1.1 Software Standards

The thesis adheres to widely recognized software development standards to ensure the reliability, scalability, and maintainability of the deep learning framework used for lemon leaf canker detection. The implementation leverages established libraries and tools, including PyTorch for model training, Roboflow for dataset management, and Google Colab for computational support, all of which follow industry-standard practices. Code development adheres to modularity and version control principles, with repositories maintained on GitHub for collaboration and reproducibility. Standardized data preprocessing techniques and metric evaluations (e.g., mAP, precision, and recall) ensure consistency across experiments and compliance with software engineering best practices in AI model development.

5.1.2 Hardware Standards

The thesis employs state-of-the-art hardware standards to ensure efficient training and evaluation of the deep learning models for lemon leaf canker detection. The computational setup includes the use of NVIDIA Tesla T4 GPUs, which are industry-standard for AI and deep learning tasks, offering high memory capacity and processing speed to handle large datasets and complex YOLO architectures. Cloud platforms like Google Colab and Kaggle provide accessible and scalable environments, meeting modern hardware standards for research while ensuring resource optimization. The hardware configuration ensures compatibility with the chosen software tools and supports the execution of computationally intensive tasks like training advanced YOLO models.

5.1.3 Communication Standards

The thesis adheres to communication standards by employing clear, consistent, and structured methods for documenting and sharing research findings. Visual tools such as precision-recall curves, confusion matrices, and performance metric graphs are used to effectively convey model evaluation results. The study also integrates standardized reporting formats for data and results, ensuring compatibility and transparency for future research and collaborations. Additionally, platforms like GitHub and Kaggle facilitate seamless communication and sharing of code, datasets, and insights with the research community, adhering to established norms in academic and technical communication.

5.2 Impact on Society, Environment and Sustainability

5.2.1 Impact on Life

Early detection of plant diseases has a profound impact on society, particularly in the agricultural sector. By identifying diseases at an early stage, farmers can take timely actions to mitigate their spread, leading to significant reductions in crop losses. This directly translates to increased agricultural productivity and economic stability for farmers, many of whom rely on their crops as their primary source of income.

Moreover, early detection reduces the excessive use of chemical pesticides, promoting environmentally sustainable farming practices. Targeted treatments not only preserve the ecosystem but also lower production costs for farmers, improving overall profitability.

On a broader scale, effective plant disease management supports global food security by ensuring stable crop yields. This is particularly critical in addressing the challenges posed by a growing global population and climate change, which threaten agricultural outputs. Thus, leveraging technology for the early detection of plant diseases contributes significantly to reducing hunger and supporting the livelihood of farming communities worldwide.

5.2.2 Impact on Society & Environment

Timely disease management in agriculture offers significant environmental benefits by minimizing the reliance on chemical pesticides. Early detection of plant diseases allows farmers to address specific problem areas rather than applying pesticides broadly. This targeted approach reduces the volume of harmful chemicals released into the soil, water, and air, thereby preserving the natural ecosystem.

Furthermore, decreased pesticide usage helps maintain biodiversity, protecting beneficial insects, microorganisms, and other wildlife that are often unintentionally harmed by excessive chemical application. It also lowers the risk of pesticide residues entering the food chain, promoting healthier agricultural products.

Additionally, reduced chemical usage contributes to improved soil health, as excessive pesticide use often depletes essential nutrients and disrupts soil microbial balance. By preserving soil fertility and preventing contamination of water sources, timely disease management supports sustainable farming practices that are essential for long-term environmental conservation.

5.2.3 Ethical Aspects

Ethical considerations are paramount in the development and deployment of AI-based models for automated lemon leaf canker detection. This research prioritizes two key aspects:

Transparency in AI-Based Decisions

It is critical that the best-performing YOLO model (YOLOv11) provides interpretable results to farmers and stakeholders. Transparency can be achieved by offering visualizations of the model's predictions, such as bounding boxes indicating diseased and healthy regions, and explaining the reasoning behind these predictions. This transparency fosters trust and encourages adoption in agricultural practices.

Fairness in Dataset Representation

The dataset used in this study consists of 2,124 lemon leaf images collected from diverse nursery gardens. Efforts were made to include both diseased (canker) and healthy samples under varying environmental conditions to ensure fair representation. However, bias could still arise if certain conditions (e.g., regional or seasonal variations) are underrepresented. To address this, the dataset must be continually updated to include images from different regions and conditions, ensuring that the model performs well across all farming contexts without favoring specific groups.

By addressing these ethical considerations, this study ensures that the AI-based solution is fair, transparent, and suitable for real-world agricultural use.

5.2.4 Sustainability Plan

To sustain the implementation of the best-performing YOLOv11 model for lemon leaf canker detection in real-world agricultural systems, the following strategies are proposed

Integration with Farming Tools

Embed the model into existing agricultural apps or hardware tools like drones and IoT-enabled devices. These tools can capture real-time images of crops and provide instant feedback to farmers, enabling quick intervention.

Cloud and Edge Deployment

Deploy the model on lightweight edge devices, such as smartphones or portable AI systems, allowing farmers to use the model offline in remote areas. Alternatively, cloud-based platforms can host the model for larger-scale use, enabling multi-farm analysis.

Farmer Training and Awareness

Conduct training programs and workshops for farmers and agricultural workers to educate them on the use and benefits of AI-based disease detection. Highlighting the cost savings and yield improvements associated with early detection can encourage adoption.

Continuous Model Optimization

Regular updates to the model using new, diverse datasets will maintain its accuracy and reliability over time. Collaborative efforts with agricultural research organizations can support data collection and validation.

Subsidies and Partnerships

Collaborate with government agencies and NGOs to provide financial support, such as subsidies, to small-scale farmers for accessing AI-based tools. Partnerships with agricultural companies can further enhance the model's reach and usability.

Environmental Monitoring

Use the model to track disease trends across regions, enabling proactive measures to prevent outbreaks. This monitoring can also help optimize pesticide usage, reducing environmental impact and promoting sustainable farming practices.

By implementing these strategies, the YOLOv11 model can become a robust and sustainable solution for lemon leaf canker detection, benefiting both agriculture and the environment in the long term.

5.3 Project Management and Financial Analysis

In terms of financial considerations, the project does not have external funding or financial backing at this stage. All resources used, including cloud-based platforms like Google Colab and Kaggle, offer free-tier access or are funded through educational or personal accounts. While the study does not involve significant financial expenditure, there may be potential costs associated with scaling the project, such as for premium cloud services, additional computational resources, or future model deployment in large-scale field applications. If the project were to expand or transition to commercial applications, funding considerations could include the costs of infrastructure, maintenance, and the development of real-time detection systems for farmers.

5.4 Complex Engineering Problem

5.4.1 Complex Problem Solving

In this section, provide a mapping with problem solving categories. For each mapping add subsections to put rationale (Use Table 5.1). For P1, you need to put another mapping with Knowledge profile and rational thereof.

Table 5.1: Mapping with complex problem solving.

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

EP1: **Depth of Knowledge**

The study demonstrates advanced knowledge in deep learning, specifically in the application of YOLO models for agricultural disease detection. It involves systematic methodologies and thorough analysis of model performance.

EP2: Range of Conflicting Requirements

Balances the need for high accuracy, computational efficiency, and scalability in the models. Conflicting requirements like precision in detection versus inference speed were addressed through optimization and model comparison.

EP3: Depth of Analysis

Provides an in-depth evaluation of YOLOv5, YOLOv8, YOLOv10, and YOLOv11 using metrics like mAP, precision, recall, and inference time. The study includes robust discussions on model strengths, limitations, and potential improvements.

EP4: Familiarity with Issues

Demonstrates strong familiarity with issues in agricultural disease detection, such as dataset imbalance, environmental variability, and real-world deployment challenges, with solutions proposed in the methodology and analysis.

EP5: Extent of Applicable Codes

Adheres to software and engineering standards, leveraging frameworks like PyTorch, Roboflow, and Google Colab while ensuring reproducibility and compliance with data handling and computational requirements.

EP7: Interdependence

Highlights interdependence between various components, such as dataset preparation, model optimization, and evaluation, to create a cohesive framework for lemon leaf canker detection.

Mapping with Knowledge Profile for EP1

Table 5.2: Mapping with knowledge Profile.

K3 Engineering Fundamentals	K4 Specialist Knowledge	K5 Engineering Design	K6 Engineering Practice	K8 Research Literature
✓	✓	✓	✓	✓

K3: Engineering Fundamentals

Utilizes foundational principles of image processing and machine learning to clean, annotate, and prepare data for training deep learning models.

K4: Specialist Knowledge

Involves advanced knowledge of object detection architectures, specifically YOLO versions, for real-time and accurate plant disease detection.

K5: Engineering Design

Includes designing experiments to compare YOLOv5, YOLOv8, YOLOv10, and YOLOv11, optimizing model parameters for the best balance of accuracy and efficiency.

K6: Engineering Practice

Applies industry-standard tools like PyTorch, Roboflow, and Google Colab for developing, training, and validating deep learning models.

K8: Research Literature

Involves synthesizing existing studies on YOLO-based detection models and identifying research gaps to advance knowledge in precision agriculture.

5.4.2 Engineering Activities

Table 5.3: Mapping with complex engineering activities.

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

EA1: Range of Resources

The study requires a diverse set of resources, including computational power (NVIDIA Tesla T4 GPUs, cloud platforms like Google Colab and Kaggle), dataset management tools (Roboflow), and standard deep learning frameworks (PyTorch). This range ensures that the models can be trained effectively and evaluated across various conditions, which is essential for handling the complex task of lemon leaf canker detection.

EA2: Level of Interaction

The research involves interaction at multiple levels, including dataset preprocessing, model training, evaluation, and interpretation of results. This level of interaction ensures that different stages of the project are tightly integrated, allowing for iterative improvements and validation of model performance against real-world agricultural challenges.

EA3: Innovation

The study introduces innovation by comparing multiple versions of YOLO models (YOLOv5, YOLOv8, YOLOv10, and YOLOv11) to determine which is best suited for detecting lemon leaf canker. This approach pushes the boundaries of existing knowledge by systematically exploring the strengths and limitations of these models in an agricultural context.

EA4: Consequences for Society and Environment

The implications of this study extend beyond technical performance to its impact on society and the environment. By providing accurate, real-time detection of lemon leaf canker, the research helps farmers take timely action to mitigate crop losses, which in turn supports sustainable farming practices and reduces the need for excessive chemical interventions that could harm the environment.

EA5: Familiarity

The study builds on existing knowledge in deep learning and precision agriculture, integrating well-known methodologies and tools like YOLO models and data augmentation techniques. This familiarity ensures that the research aligns with current best practices in the field, providing a solid foundation for both methodological rigor and practical applicability.

5.5 Summary

Chapter 5 discusses the adherence to engineering standards and the associated design challenges faced in implementing deep learning models for lemon leaf canker detection. The study aligns with established software standards by using recognized libraries and cloud platforms, ensuring reliability, scalability, and reproducibility. Hardware standards are maintained through the use of NVIDIA Tesla T4 GPUs and cloud computing, facilitating efficient training and evaluation of YOLO models. The communication standards are upheld through structured documentation and sharing of research findings via visual tools and platforms like GitHub and Kaggle. Design challenges include managing dataset diversity to reflect real-world conditions, addressing class imbalances for accurate detection, and optimizing YOLO models for deployment in low-resource environments. The sustainability plan outlines strategies for integrating the model into agricultural tools and ensuring long-term utility and environmental conservation.

Chapter 6

Conclusion

The conclusion of this report highlights the effectiveness of advanced YOLO models, particularly YOLOv11, for automated detection of lemon leaf canker. By leveraging deep learning techniques, the study addresses key challenges such as dataset diversity, class imbalance, and environmental variability, providing a scalable solution for real-time deployment. The findings underscore the importance of integrating AI into precision agriculture, promoting sustainable farming practices and contributing to food security and economic stability for farmers.

6.1 Summary

This research conducted a comparative analysis of four state-of-the-art YOLO models—YOLOv5, YOLOv8, YOLOv10, and YOLOv11—for automated lemon leaf canker detection. The results revealed that YOLOv11 outperformed the other models in terms of overall precision (94%), recall (90.4%), and mAP50 (94.8%), making it the most effective model for detecting lemon leaf canker. It demonstrated superior performance in identifying the 'canker' class, achieving a mAP50 of 90.1%, which is a significant improvement over the other models. The 'healthy' class consistently showed near-perfect detection across all models, with YOLOv11 achieving a mAP50 of 94.8%.

Additionally, YOLOv11 exhibited balanced computational efficiency, requiring only 67.7 GFLOPs and completing validation with an average inference time of 8.1 ms per image. This makes it a practical choice for real-time deployment in agricultural environments. While YOLOv5 and YOLOv8 demonstrated competitive performance, particularly for the 'healthy' class, they fell short in precision and recall for the 'canker' class. YOLOv10, while slightly better than YOLOv8 for the 'canker' class, also could not surpass YOLOv11 in overall performance.

6.2 Limitation

Despite the promising results, the study faced several limitations

Dataset Size

While 884 images provided a reasonable starting point, a larger dataset with more diverse samples would further improve model generalization and robustness. Limited data for certain conditions, such as early-stage symptoms or leaves with overlapping lesions, may have impacted performance.

Class Imbalance

The dataset had an unequal distribution of healthy and diseased samples. This imbalance can bias the models, making them favor the majority class. Techniques like oversampling or synthetic data generation could mitigate this issue.

Environmental Variability

The dataset did not fully capture the diversity of real-world conditions, such as extreme lighting, occlusions, or heavily cluttered backgrounds. Models might require fine-tuning or domain adaptation for deployment in highly variable environments.

Computational Requirements:While YOLOv11 offered the best performance, its higher computational demands might pose challenges for deployment on low-resource devices, limiting its accessibility for some farmers.

Single Disease Focus

The study focused exclusively on lemon leaf canker, limiting its applicability to other diseases. Expanding the scope to multi-disease detection could enhance its utility in agricultural monitoring systems.

6.3 Future Work

The results of this study have significant implications for the development of automated systems for agricultural disease management. Deploying the YOLOv11 model in agricultural settings could enable real-time monitoring and early detection of lemon leaf canker, potentially minimizing crop losses and improving yield.

Future research could expand upon this work by

Expanding the Dataset

Incorporating more diverse leaf images from different environmental conditions and geographic locations to improve model robustness.

Testing Additional Models

Exploring other state-of-the-art object detection models or novel architectures to further improve detection accuracy.

Optimizing Deployment

Developing lightweight versions of YOLOv11 for deployment on edge devices, ensuring real-time performance in resource-constrained environments.

Multi-Disease Detection

Extending the framework to detect multiple plant diseases concurrently, making it more versatile for agricultural applications.

These advancements could pave the way for a comprehensive automated solution that supports farmers in managing crop health more efficiently.

References

[1] R. Z. Qiu and S. P. Chen and M. X. Chi and R. B. Wang and T. Huang and G. C. Fan and Q. Y. Weng. An automatic identification system for citrus greening disease (Huanglongbing) using a YOLO convolutional neural network. Frontiers in Plant Science, 13(1):1002606, 2022.

[2] X. Zhang and Y. Xun and Y. Chen. Automated identification of citrus diseases in orchards using deep learning. Biosystems Engineering, 223(1):249-258, 2022.

[3] A. S. Krishna and M. Sunag and N. S. Rani and B. R. Pushpa. On-spot Citrus Canker Disease Detection using YOLOv7. International Conference on Applied Artificial Intelligence and Computing, 2(1):944-949, 2023.

[4] M. M. S. Pangaliman and K. M. Manalo and P. C. Naval Jr. Philippine Lime (Calamansi) Disease Detection and Classification Using YOLOv8 Model. Asian Conference on Intelligent Information and Database Systems, 1(1):106-117, 2024.

[5] S. A. Preanto and M. T. Ahad and Y. R. Emon and S. Mustofa and M. Alamin. A Semantic Segmentation Approach on Sweet Orange Leaf Diseases Detection Utilizing YOLO. arXiv preprint, 2409.06671(1):1-15, 2024.

[6] S. Dananjayan and Y. Tang and J. Zhuang and C. Hou and S. Luo. Assessment of state-of-the-art deep learning based citrus disease detection techniques using annotated optical leaf images. Computers and Electronics in Agriculture, 193(1):106658, 2022.

[7] H. Mo and L. Wei. Lightweight citrus leaf disease detection model based on ARMS and cross-domain dynamic attention. Journal of King Saud University-Computer and Information Sciences, 36(7):102133, 2024.

[8] D. Luo and Y. Xue and X. Deng and B. Yang and H. Chen and Z. Mo. Citrus Diseases and Pests Detection Model Based on Self-Attention YOLOV8. IEEE Access, 11(1):1-12, 2023.

[9] J. C. da Silva and M. C. Silva and E. J. Luz and S. Delabrida and R. A. Oliveira. Using mobile edge AI to detect and map diseases in citrus orchards. Sensors, 23(4):2165, 2023.

[10] B. V. Apacionado and T. Ahamed. Sooty Mold Detection on Citrus Tree Canopy Using Deep Learning Algorithms. Sensors, 23(20):8519, 2023.

[11] Q. Dai and Y. Xiao and S. Lv and S. Song and X. Xue and S. Liang and Z. Li. YOLOv8-GABNet: An Enhanced Lightweight Network for the High-Precision Recognition of Citrus Diseases and Nutrient Deficiencies. Agriculture, 14(11):1964, 2024.

[12] Y. Xiang and J. Yao and Y. Yang and K. Yao and C. Wu and X. Yue and G. Gong. Real-Time Detection Algorithm for Kiwifruit Canker Based on a Lightweight and Efficient Generative Adversarial Network. Plants, 12(17):3053, 2023.

[13] C. Yao and Z. Yang and P. Li and Y. Liang and Y. Fan and J. Luo and J. Mu. Two-Stage Detection Algorithm for Plum Leaf Disease and Severity Assessment Based on Deep Learning. Agronomy, 14(7):1589, 2024.

[14] L. Huang and M. Chen and Z. Peng. YOLOv8-G: An Improved YOLOv8 Model for Major Disease Detection in Dragon Fruit Stems. Sensors, 24(15):1-15, 2024.

[15] X. Yue and K. Qi and X. Na and Y. Zhang and Y. Liu and C. Liu. Improved YOLOv8-Seg network for instance segmentation of healthy and diseased tomato plants in the growth stage. Agriculture, 13(8):1643, 2023.

[16] N. S. Rani and A. S. Krishna and M. Sunag and M. A. Sangamesha and B. R. Pushpa. Infield disease detection in citrus plants: integrating semantic segmentation and dynamic deep learning object detection model for enhanced agricultural yield. Neural Computing and Applications, 1(1):1-26, 2024.

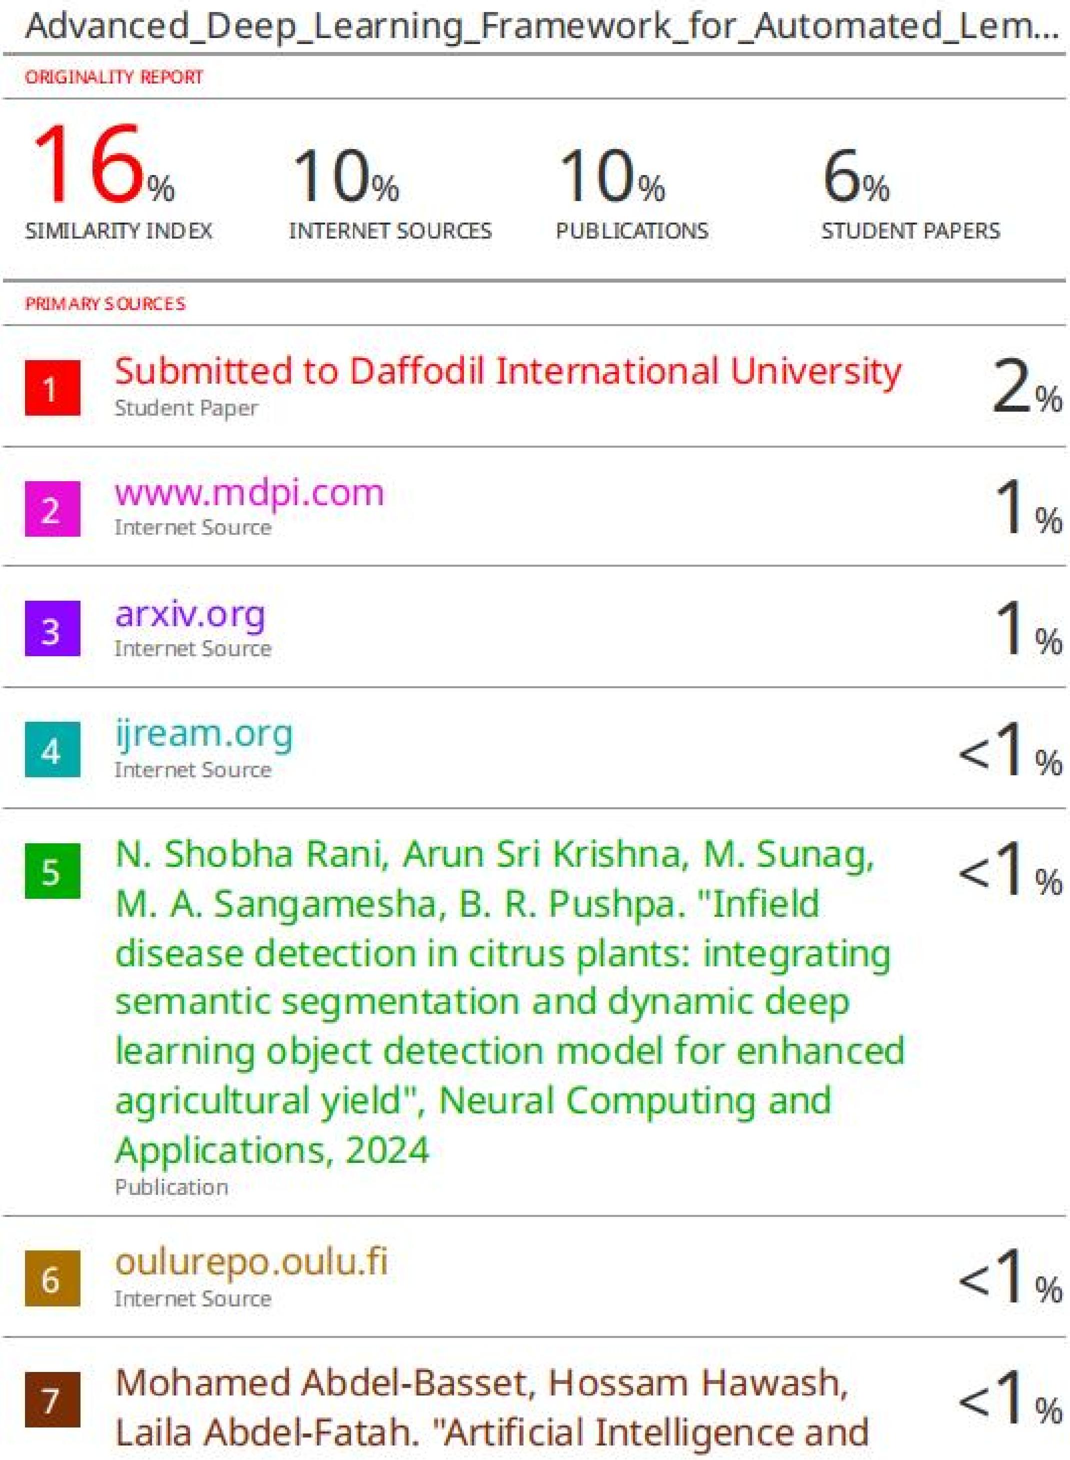

Advanced_Deep_Learning_Framework_for_Automated_Lem...

ORIGINALITY REPORT

16%	10%	10%	6%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Daffodil International University Student Paper	2%
2	www.mdpi.com Internet Source	1%
3	arxiv.org Internet Source	1%
4	ijream.org Internet Source	<1%
5	N. Shobha Rani, Arun Sri Krishna, M. Sunag, M. A. Sangamesha, B. R. Pushpa. "Infield disease detection in citrus plants: integrating semantic segmentation and dynamic deep learning object detection model for enhanced agricultural yield", Neural Computing and Applications, 2024 Publication	<1%
6	oulurepo.oulu.fi Internet Source	<1%
7	Mohamed Abdel-Basset, Hossam Hawash, Laila Abdel-Fatah. "Artificial Intelligence and	<1%