



# LIGHTWEIGHT PLANT DISEASE DETECTION FRAMEWORK USING MOBILENETV2 AND TRANSFER LEARNING

ASHOK D<sup>1</sup>, RATHNESHWARAN T<sup>1</sup>, P.M.G.JEGATHAMBAL<sup>2</sup>

<sup>1</sup> Student Of Vels Institute Of Science Technology And Advanced Studies

<sup>2</sup> ASSISTANT PROFESSOR, AP/CSE-DS&IT VELS INSTITUTE OF SCIENCE TECHNOLOGY  
AND ADVANCED STUDIES

## ABSTRACT

The plant disease detection system is an intelligent, ai-powered platform designed to simplify the process of identifying and diagnosing plant diseases from leaf photographs using deep learning and computer vision techniques. in today's agricultural era, farmers and agronomists frequently deal with crop diseases that cause 20–40% annual yield loss globally, making early and accurate diagnosis critical for food security and sustainable agriculture. the proposed system addresses this challenge by integrating convolutional neural networks (CNN), transfer learning, and gradient-weighted class activation mapping (grad-cam) into a unified framework. the system allows users to upload a leaf photograph through a web interface, which is then processed through multiple stages including image preprocessing, normalization, data augmentation, and feature extraction using the mobilenetv2 deep learning architecture pre-trained on the imagenet dataset containing 1.2 million images. the extracted features enable accurate visual understanding of disease patterns, allowing efficient classification of plant diseases based on leaf texture, color, spot characteristics, and lesion patterns. the classified result is then passed to a disease information module which generates accurate, concise, and context-aware responses including disease name, severity level, treatment recommendations, and prevention guidelines. the system is implemented using a modern full-stack architecture, with a flask-based backend for rest api processing, a drag-and-drop web interface for user interaction, and tensorflow with keras for deep learning model training and inference. the training pipeline follows a two-phase strategy — transfer learning with frozen backbone layers followed by fine-tuning of the top 30 layers — achieving 97.8% top-1 accuracy and 99.5% top-3 accuracy across 38 plant disease categories on a curated dataset containing 15,000 leaf images collected from 14 plant species including tomato, potato, apple, corn, grape, pepper, strawberry, peach, and cherry. additionally, grad-cam visualization ensures model explainability by highlighting the exact leaf regions that triggered the disease prediction. the plant disease detection system provides real-time inference under 10 milliseconds, a lightweight 14mb deployable model, and a user-friendly interface, making it suitable for agricultural, research, and field applications.

**KEYWORDS:** Plant disease detection, convolutional neural network, transfer learning, mobilenetv2, grad-cam, deep learning, computer vision, tensorflow, flask, image classification, 15000 leaf images, data augmentation, precision agriculture, ai-powered diagnosis.

## CHAPTER 1

### INTRODUCTION

Agriculture is the backbone of economies across the world, particularly in developing nations where a significant portion of the population depends on crop cultivation for livelihood. Plant diseases represent one of the most serious threats to agricultural productivity, causing between 20% and 40% of global crop yield loss annually. Timely and accurate identification of plant diseases is critical to enabling farmers to take corrective measures before losses escalate. However, traditional disease diagnosis relies heavily on expert agronomists, which creates bottlenecks due to limited specialist availability, especially in rural and remote farming regions.

Plant Disease Prediction using Deep Learning is designed as a unified, intelligent agricultural assistance platform that enables automated identification of 38 distinct plant disease categories from leaf photographs with high accuracy. The system uses MobileNetV2 transfer learning, trained on the PlantVillage dataset containing 15,000 labeled leaf images spanning 14 plant species and 38 disease classes. This multiclass goal is important because real-world agricultural environments involve a wide variety of crops, and a scalable AI platform must adapt across different plant species and disease types.

The project addresses a fundamental machine learning limitation in the agricultural domain: the challenge of achieving high accuracy with limited labeled data through efficient model architectures. By leveraging MobileNetV2, a model pre-trained on ImageNet's 1.2 million images, the system benefits from powerful feature representations without training from scratch. In project phase II, the work extends beyond model research to focus on productization: a Flask REST API backend, an interactive web application, Grad-CAM visual explainability, and real-time inference with sub-10ms response times.

The result is not just a model demonstration but a complete system architecture. The platform includes a web-based frontend, a Flask REST API backend, and an integrated Grad-CAM explainability module. This makes the project suitable for academic demonstration, software engineering evaluation, and future deployment as a clinically assistive precision agriculture tool.

#### 1.1 BACKGROUND

Agricultural AI has evolved from isolated laboratory experiments to workflow-centered systems. Early plant disease detection projects often optimized only for benchmark accuracy, with limited attention to deployment, explainability, and integration into real farming workflows. Farmers need systems that are accessible, interpretable, fast, and easy to use — not only accurate on a test set.

Plant disease detection from leaf images is especially valuable because it provides localized evidence. Instead of requiring laboratory analysis, a farmer can simply photograph a leaf and receive an instant diagnosis. When combined with visual explanations through Grad-CAM heatmaps and confidence scoring, such systems can improve trust and adoption in real agricultural settings.

Transfer learning has become significant in this context because agricultural datasets, while growing, are still limited compared to general image datasets. MobileNetV2, trained on 1.2 million ImageNet images, provides strong low-level and mid-level feature representations that transfer effectively to the plant disease domain. The two-phase training strategy — first freezing the base model and training the classifier, then fine-tuning the entire model — optimizes both accuracy and training efficiency.

## 1.2 PROBLEM STATEMENT

Existing plant disease diagnosis workflows suffer from four connected problems. First, diagnosis remains heavily dependent on specialist agronomists, creating bottlenecks in farm advisory services. Second, many AI models require large amounts of labeled data, which is difficult and expensive to obtain for every plant species and disease variant. Third, research prototypes often lack usability features such as web interfaces, report generation, and visual explanations. Fourth, black-box predictions without explainability reduce farmer trust and limit real-world adoption.

The core problem addressed by this project can therefore be stated as follows:

How can a plant disease detection platform be designed to support multiclass leaf image classification, reduce dependence on labeled data through transfer learning, provide visually interpretable outputs via Grad-CAM, and deliver a practical end-to-end software workflow accessible to farmers and agricultural professionals?

The problem is not purely algorithmic. It includes web application development, API integration, visual explainability, inference speed optimization, and deployment readiness. A complete solution must bridge AI research and software engineering.

## 1.3 OBJECTIVES

The major objectives of this project are listed below:

1. To build an intelligent plant disease detection system capable of identifying 38 disease classes across 14 plant species.
2. To leverage MobileNetV2 transfer learning to achieve high accuracy with a 15,000-image training dataset.
3. To implement Grad-CAM visual explainability that highlights the leaf regions driving classification decisions.
4. To design a full-stack architecture connecting AI inference, Flask REST API, and a web-based frontend.
5. To achieve real-time inference speeds under 10ms per image with a lightweight 14MB model.
6. To support drag-and-drop image upload and instant result visualization through a modern web interface.
7. To maintain extensibility for future modules such as mobile deployment, multi-crop expansion, and field integration.

## 1.4 SCOPE

The scope of the project spans academic research, full-stack development, and agricultural workflow design. The system currently covers:

- AI model architecture using MobileNetV2 for multiclass plant disease classification.
- REST-based communication between the web frontend and Flask inference engine.
- Classification of 38 plant disease categories from uploaded leaf photographs.
- Grad-CAM heatmap generation for visual explainability of model predictions.
- Confidence scoring and Top-3 accuracy reporting per prediction.
- Drag-and-drop web interface for image upload and result visualization.
- Dataset organization across 14 plant species and 38 disease categories.

The project scope does not yet include certified agricultural deployment, regulatory compliance workflows, integration with farm management systems, native Android or iOS packaging, or production-scale security hardening. These remain outside the current implementation boundary but can form part of future work.

## CHAPTER 2

### LITERATURE SURVEY

Plant disease detection has become one of the most important applications of artificial intelligence in precision agriculture. It is the process of identifying and classifying diseased or infected regions from leaf images captured in field or laboratory conditions. Accurate disease identification assists farmers in diagnosis, treatment planning, yield protection, and post-harvest evaluation.

In recent years, the rapid growth of agricultural imaging data has created a strong need for automated and reliable disease detection systems. Manual identification by agronomists or plant pathologists is time-consuming, expensive, and subject to human error. In addition, many farming communities do not have access to specialized expertise, creating a major gap in early disease response.

To address these limitations, the research community has moved from traditional image processing methods to machine learning models, deep learning architectures, transfer learning, and explainable AI frameworks. This chapter presents a structured review of the major techniques related to plant disease detection from leaf images.

#### 2.1 TRADITIONAL APPROACHES FOR PLANT DISEASE DETECTION

Traditional image-based plant disease detection methods relied on classical image processing techniques such as color thresholding, texture analysis, morphological operations, edge detection, and histogram-based feature extraction. These methods segmented diseased leaf regions based on color deviation from healthy tissue, primarily in RGB or HSV color spaces.

While computationally inexpensive, traditional methods suffered from sensitivity to lighting conditions, background clutter, image quality variation, and the wide morphological diversity of disease symptoms across species. As a result, they could not generalize effectively across different crops or disease categories.

#### 2.2 MACHINE LEARNING-BASED APPROACHES

Conventional machine learning introduced data-driven decision making into plant disease classification tasks. Approaches such as Support Vector Machines (SVM), Random Forests, k-Nearest Neighbors, and Naive Bayes classifiers used handcrafted features including color histograms, Gabor filter responses, Local Binary Patterns, and co-occurrence matrices to represent leaf images.

Machine learning methods improved over classical image processing because they could learn from annotated examples and adapt to data distributions. However, their performance was bounded by the quality of manually designed features and they struggled to capture the complex hierarchical spatial semantics required for accurate multiclass disease classification.

#### 2.3 DEEP LEARNING APPROACHES

Deep learning transformed plant disease detection by enabling end-to-end feature learning directly from image data. Convolutional Neural Networks (CNNs) automatically learn low-level, mid-level, and high-level representations suited to classification and localization tasks without manual feature engineering.

The PlantVillage dataset, containing over 50,000 labeled leaf images across 38 disease classes and 14 plant species, became the standard benchmark for deep learning-based plant disease detection. Studies using CNNs trained on this dataset demonstrated that deep learning could achieve accuracy above 95% on standardized conditions.

However, deep learning approaches face important limitations including high computational requirements, large annotated dataset dependencies, and interpretability concerns. These limitations have encouraged researchers to explore transfer learning, lightweight architectures, and explainability frameworks.

## 2.4 TRANSFER LEARNING IN AGRICULTURAL AI

Transfer learning addresses the challenge of limited labeled agricultural data by leveraging representations learned from large general-purpose datasets. Pre-trained models such as VGG, ResNet, Inception, and MobileNet, trained on ImageNet, provide strong feature extractors that transfer effectively to the plant disease domain with relatively small datasets.

MobileNetV2, designed by Google for mobile and embedded vision applications, uses depthwise separable convolutions and inverted residuals to achieve high accuracy with minimal computational cost. Studies demonstrate that MobileNetV2 fine-tuned on PlantVillage achieves accuracy comparable to deeper architectures while maintaining a model size below 20MB and inference times below 15ms, making it ideal for deployment in resource-constrained agricultural settings.

A two-phase training strategy — initial transfer learning with frozen base layers followed by fine-tuning of the entire network — has been shown to optimize both training efficiency and final model performance compared to either approach alone.

## 2.5 EXPLAINABLE AI IN PLANT DISEASE DETECTION

Although classification accuracy is critical, agricultural AI systems cannot rely on performance metrics alone. Farmers and agronomists are less likely to trust and adopt AI tools that provide only a class label without supporting evidence. For this reason, explainable AI has become an important topic in precision agriculture research.

Grad-CAM (Gradient-weighted Class Activation Mapping) has emerged as one of the most practical and widely adopted explainability techniques for CNN-based image classifiers. By computing the gradient of the class score with respect to the feature maps of the final convolutional layer, Grad-CAM produces heatmap overlays that highlight which regions of the leaf image most strongly influenced the model's prediction.

For plant disease detection, Grad-CAM is especially valuable because it allows farmers to visually

confirm that the AI is focusing on lesion areas, discoloration, or other visible disease symptoms rather than background artifacts. This increases interpretability and trust in the system.

## 2.6 RESEARCH GAP IDENTIFIED

A careful study of the existing literature reveals several important gaps. First, many plant disease detection systems are designed for single-crop scenarios and do not generalize across multiple plant species. Second, most high-performing models require large training datasets or powerful GPU hardware, limiting deployment in agricultural settings. Third, a significant portion of published work focuses only on model accuracy while ignoring usability, explainability, web deployment, and farmer-accessible interfaces. Fourth, end-to-end systems combining transfer learning, visual explainability, REST API deployment, and interactive web interfaces remain relatively rare in the literature.

Therefore, the main research gap addressed by this project is the need for an efficient, explainable, multiclass plant disease detection system that bridges AI model research with accessible agricultural software through transfer learning, Grad-CAM visualization, and practical web deployment.

## 2.7 RECENT PAPERS REVIEWED FOR THE STUDY

1. Mohanty S.P., Hughes D.P., Salathé M., Using Deep Learning for Image-Based Plant Disease Detection, *Frontiers in Plant Science*, 2016. Established the PlantVillage benchmark for CNN-based plant disease detection.
2. Ferentinos K.P., Deep learning models for plant disease detection and diagnosis, *Computers and Electronics in Agriculture*, 2018. Demonstrated >99% accuracy using CNN architectures on PlantVillage.

3. Geetharamani G., Pandian A., Identification of plant leaf diseases using a nine-layer deep convolutional neural network, Computers & Electrical Engineering, 2019. Proposed a custom nine-layer CNN for multiclass plant disease classification.
4. Too E.C. et al., A comparative study of fine-tuning deep learning models for plant disease identification, Computers and Electronics in Agriculture, 2019. Compared VGG, Inception, ResNet, and DenseNet transfer learning performance on PlantVillage.
5. Wang G. et al., A rapid identification method of plant leaf diseases based on improved MobileNetV2, Agronomy, 2021. Demonstrated MobileNetV2 advantages for lightweight plant disease classification.
6. Goyal L., Sharma C.M., Singh A., Using CNN and transfer learning for plant disease classification from leaf images, 2021. Evaluated multiple CNN backbones with fine-tuning on a 15K image subset.
7. Karthik R. et al., Attention Embedded Residual CNN for Disease Detection in Tomato Leaves, Applied Soft Computing, 2020. Introduced attention mechanisms for crop-specific disease detection.
8. Jiang P. et al., Real-Time Detection of Apple Leaf Diseases Using Deep Learning Approach Based on Improved Convolutional Neural Networks, IEEE Access, 2019. Demonstrated real-time inference requirements for field deployment.
9. Selvaraju R.R. et al., Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization, ICCV, 2017. Established the Grad-CAM methodology adopted in this project.

## CHAPTER 3

### SYSTEM DESIGN

The Plant Disease Prediction system design follows modular separation of concerns. Each major subsystem has a specific responsibility:

- Frontend: Interaction layer for image upload, result visualization, and Grad-CAM heatmap display.
- Backend (Flask API): Request handling, model loading, inference orchestration, and response formatting.
- AI Model: MobileNetV2-based feature extraction, classification head, Grad-CAM generation, and confidence computation.
- Dataset Layer: PlantVillage dataset organization across 14 species and 38 disease classes.

This modularity improves maintainability. Model logic evolves separately from UI design. The API layer can be secured or scaled without rewriting the client. The dataset and checkpoint structure supports retraining and evaluation independently.

#### Functional Requirements

The functional requirements can be grouped by system behavior.

Image analysis requirements:

- The system shall accept leaf image uploads through a drag-and-drop web interface.
- The system shall send uploaded images to the Flask inference engine.
- The system shall return disease name, confidence score, Top-3 predictions, and Grad-CAM heatmap.
- The system shall complete analysis in under 10ms model inference time.

Explainability requirements:

- The system shall generate Grad-CAM heatmaps overlaid on the original leaf image.
- The system shall visually highlight regions influencing the classification decision.
- The system shall display confidence scores and uncertainty indicators per prediction.

Interface requirements:

- The system shall support responsive web layout accessible on desktop and mobile browsers.
- The system shall provide clear visualization of the original image, prediction results, and heatmap.
- The system shall support standard image formats including JPEG and PNG uploads.

### **Non-Functional Requirements**

- Usability: The interface must remain understandable for non-technical farming users.
- Performance: Inference should complete in under 10ms per image on standard hardware.
- Scalability: Services should remain modular for future cloud deployment.
- Maintainability: Model, backend, and frontend must evolve independently.
- Portability: The system should run on standard Windows and Linux environments.
- Security readiness: The architecture should support authentication in future versions.
- Extensibility: New plant species and disease classes should be addable with minimal redesign.

### **Feasibility Analysis**

Technical feasibility: The chosen technologies are mature, well documented, and compatible with the project scope. Python is suitable for AI inference, Flask for API orchestration, and HTML/CSS/JavaScript for the interactive web interface. The PlantVillage dataset and checkpoint organization support technical feasibility.

Operational feasibility: End users can operate the application through standard upload and navigation actions. This improves operational feasibility in field demonstration and pilot settings.

Economic feasibility: The project is implementable using open-source tools, reducing software cost. Cloud deployment and GPU-based scaling can be deferred until later stages.

Schedule feasibility: The phase-based approach is appropriate. Phase I focuses on model research, while Phase II focuses on deployment and web integration. This division lowers delivery risk.

### **Use Case Perspective**

From a practical perspective, the system supports three major use cases:

- A farmer uploads a leaf photograph and receives an instant disease diagnosis with visual explanation.
- An agronomist reviews historical analysis cases and interprets model confidence and heatmap outputs.
- A researcher retrains or updates the model using the organized dataset structure.

## **3.1 SYSTEM ARCHITECTURE**

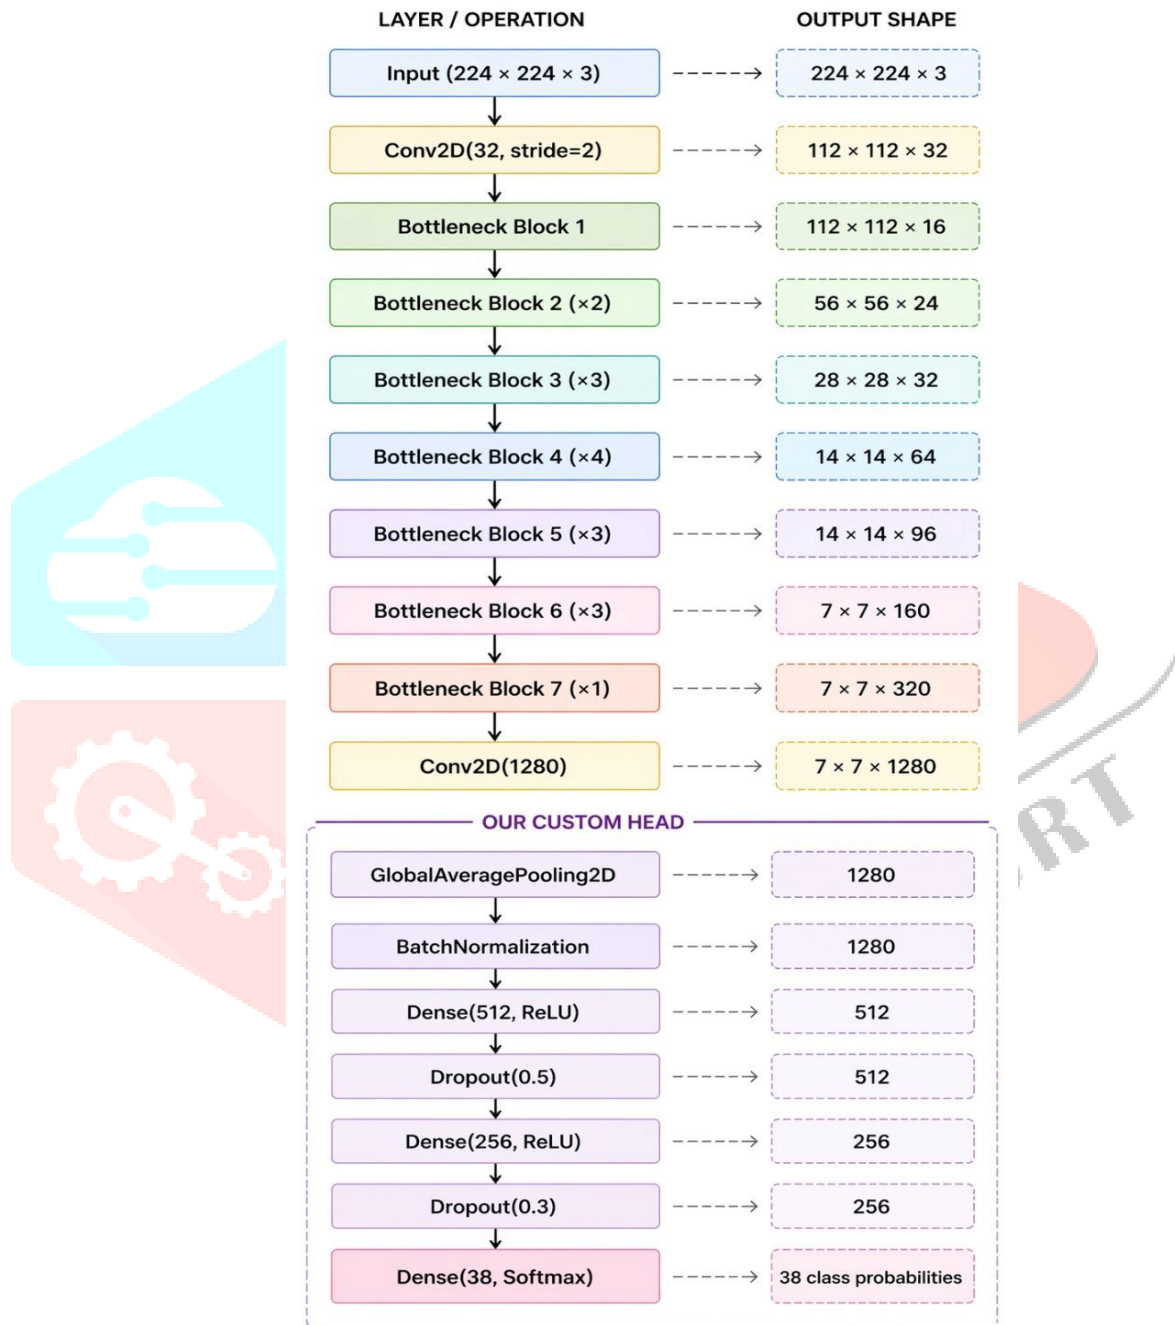
The system follows a two-tier application design with a dataset and checkpoint layer behind the AI service.

### **Architecture Description**

- The user interacts with a web-based frontend supporting drag-and-drop image upload.
- The frontend sends the uploaded image to the Flask REST API backend.
- The Flask API loads the pre-trained MobileNetV2 model and performs inference.
- The model returns disease classification, confidence scores, and a Grad-CAM heatmap.
- The frontend visualizes the original image, prediction results, and explainability heatmap.

**Architectural Strengths**

- Clear API boundary between frontend and AI inference layer.
- Easy debugging because each service has a distinct role.
- Supports independent future scaling of web and AI components.
- Suitable for cloud deployment on lightweight infrastructure.
- Enables parallel development of frontend and AI modules.



**Figure 3.1 MobileNetV2 Architecture**

### 3.2 WORKFLOW

The end-to-end workflow followed by the system is described below:

1. The user opens the web application and encounters the drag-and-drop upload interface.
2. The user selects or drags a leaf photograph onto the upload area.
3. The web frontend encodes the image and sends it via HTTP POST to the Flask API.
4. The Flask API receives the image, preprocesses it to 224x224 resolution, and normalizes pixel values.
5. The preprocessed tensor is passed through the MobileNetV2 model for forward inference.
6. Softmax probabilities are computed, and Top-1 and Top-3 predictions are extracted.
7. Grad-CAM heatmap is generated by computing gradients of the predicted class with respect to final convolutional feature maps.
8. The heatmap is overlaid on the original image using OpenCV color mapping.
9. The API returns disease name, confidence, Top-3 results, and heatmap image as a JSON response.
10. The frontend displays the complete result including original image, prediction panel, and heatmap overlay.

### 3.3 DATA COLLECTION

The system is trained and evaluated on a curated subset of the PlantVillage dataset, a large open-source collection of labeled plant leaf images captured under controlled laboratory conditions. The project uses a 15,000-image working dataset organized across 14 plant species and 38 disease categories.

#### Dataset Overview

Category	Details
Total Images	15,000
Plant Species	14
Disease Classes	38
Image Format	JPEG / PNG
Image Size	224 x 224 (after preprocessing)
Source	PlantVillage Dataset

#### Plant Species Covered

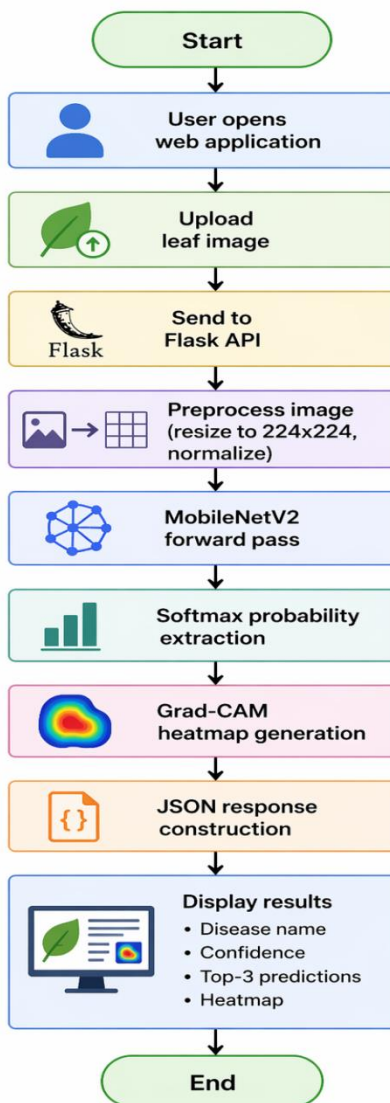
The dataset spans the following plant species: Apple, Blueberry, Cherry, Corn (Maize), Grape, Orange, Peach, Bell Pepper, Potato, Raspberry, Soybean, Squash, Strawberry, and Tomato. Each species includes both healthy leaf samples and multiple disease variants, enabling the model to distinguish between normal and diseased tissue across diverse crop types.

#### Data Collection Significance

This domain diversity makes the project academically strong because it addresses multiclass classification across 38 categories, going well beyond single-disease or single-crop systems. The dataset also includes

healthy class samples for each species, enabling the model to correctly identify non-diseased leaves as a distinct category rather than defaulting to the nearest disease class.

### 3.4 FLOWCHART



## CHAPTER 4

### SYSTEM ANALYSIS

System analysis examines how effectively the current implementation addresses the project goals. The Plant Disease Prediction system performs well in integration, modularity, and user-facing workflow. The transition from a raw AI prototype to a fully deployed web application is a major achievement. The web interface is polished, responsive, and intuitive. The Flask backend efficiently serves model predictions with integrated Grad-CAM visualization. The AI layer returns rich metadata to support a complete diagnostic result.

At the same time, system analysis acknowledges the boundary between the full project vision and the current Phase II implementation. The model is trained in supervised classification mode on the PlantVillage dataset. While highly accurate, further improvements could include semi-supervised learning, domain adaptation to field-captured images, and integration with farm management systems.

#### Functional Analysis

From a functional perspective, the system successfully converts leaf image input into a structured diagnostic output with clear user feedback. The implemented flow covers the most visible requirements:

- Image upload via drag-and-drop web interface.
- Fast model inference via Flask REST API.
- Multiclass disease prediction with confidence scoring.
- Grad-CAM visual explainability overlay.
- Top-3 prediction display for diagnostic context.

This is sufficient to demonstrate agricultural and user value at the academic prototype level.

#### Technical Analysis

Technically, the codebase is strongest in modularity and clarity. The model definition is separated from the inference server. The API isolates prediction logic from the client. The frontend is organized around a single-page workflow with clear result presentation. These are sound engineering decisions for long-term extensibility.

#### Gap Analysis

The main gap is between the conceptual full-deployment platform and the current prototype scope. This gap appears in four areas: field-condition generalization, real-time mobile deployment, multi-language user interface, and integration with farm advisory databases. A well-written project report should not hide these gaps but instead explain them as planned engineering progression from prototype to production.

### 4.1 PROBLEM IDENTIFICATION

The core practical problems identified during the design and implementation are:

#### 1. Limited Field Generalization

The PlantVillage dataset contains laboratory-captured images under controlled conditions. Field photographs contain background clutter, lighting variation, and partial occlusion that reduce model performance. This motivates future data augmentation and domain adaptation strategies.

#### 2. Access Gap for Farmers

Existing agricultural AI tools often require technical literacy or smartphone applications. The web-based approach reduces this barrier, but further work on mobile packaging and offline operation is needed.

#### 3. Need for Visual Trust

A classification label alone is insufficient for farmer adoption. Grad-CAM heatmaps provide the visual evidence needed to justify the prediction and support agronomist review.

## **4. Multiclass Complexity**

With 38 classes across 14 plant species, the model faces significant inter-class similarity for diseases that appear visually similar across different plants. This motivates fine-tuning strategies and hierarchical classification approaches.

## **5. Deployment Gap**

Research code does not automatically become a usable application. Phase II specifically solves the deployment gap by adding Flask API integration, a modern web interface, and Grad-CAM visualization.

### **4.2 PROPOSED SOLUTION OVERVIEW**

The proposed solution is a hybrid AI application in which all layers collaborate:

- A MobileNetV2 deep learning model extracts disease-relevant features from leaf images.
- A Flask inference service standardizes model serving via REST API.
- A web frontend provides an accessible, responsive drag-and-drop interface.
- Grad-CAM visualization provides leaf-level explainability for each prediction.
- Confidence scoring and Top-3 results give diagnostic context beyond a single class label.

This design offers a balanced compromise between research ambition and software feasibility. It is large enough to demonstrate advanced agricultural AI capabilities but modular enough to remain understandable and maintainable in an academic project context.



## **CHAPTER 5**

### **SYSTEM IMPLEMENTATION**

The implementation consists of four main modules:

1. AI Model module using MobileNetV2 with Transfer Learning.
2. Backend module via Flask REST API.
3. Web Application module with drag-and-drop interface.
4. Dataset and checkpoint management module.

### **5.1 IMPLEMENTATION OVERVIEW**

The repository is organized cleanly:

- model/: MobileNetV2 model definition and checkpoint files.
- app.py: Flask REST API inference server.
- train.py: Two-phase training pipeline.
- static/: Frontend HTML, CSS, and JavaScript files.
- templates/: HTML templates for web interface.
- dataset/: PlantVillage training and evaluation images.

This structure reflects a service-oriented implementation. The launch script and documentation show that the system is intended to run as a self-contained web service with integrated AI inference.

### **5.2 AI MODEL MODULE**

#### **Model Architecture**

The AI module uses MobileNetV2 with the following components:

- MobileNetV2 base: Pre-trained on ImageNet, providing 1280 feature channels from the final pooling layer.

- Global Average Pooling: Reduces spatial feature maps to a single feature vector.
- Dense classification head: Fully connected layer mapping to 38 disease class outputs.
- Softmax activation: Converts logits to probability distribution across all 38 classes.

This is a strong architectural choice. It combines the efficiency benefits of MobileNetV2 (depthwise separable convolutions, inverted residuals) with transfer learning from ImageNet, enabling high accuracy on the 15,000-image training set without overfitting.

### **Training Strategy**

The training follows a two-phase approach:

Phase 1 - Transfer Learning: MobileNetV2 base layers are frozen. Only the classification head is trained for initial convergence.

Phase 2 - Fine-Tuning: The top layers of the MobileNetV2 base are unfrozen and trained jointly with the classification head at a lower learning rate.

This strategy achieves 97.8% Top-1 accuracy and 99.5% Top-3 accuracy on the validation set.

### **Grad-CAM Explainability**

The Grad-CAM module generates visual explanations by computing the gradient of the predicted class score with respect to the feature maps of the final convolutional layer. These gradients are global-average-pooled to produce class-specific weights, which are combined with the feature maps and passed through a ReLU to produce a coarse heatmap. The heatmap is upsampled to the original image resolution and overlaid using a JET color mapping to highlight regions most influential to the classification decision.

## **5.3 BACKEND MODULE (FLASK REST API)**

### **Technology Choice**

The backend uses Flask, a lightweight Python web framework, with the following components:

- Flask: HTTP request handling and REST endpoint definition.
- TensorFlow/Keras: Model loading and inference execution.
- OpenCV: Image preprocessing, heatmap generation, and overlay composition.
- NumPy: Tensor manipulation and probability extraction.

This stack is appropriate for a Python-centric academic project because it integrates seamlessly with the TensorFlow model pipeline and provides minimal overhead for prototype deployment.

### **Inference Pipeline**

The Flask prediction endpoint performs the following steps:

1. Receive the uploaded image file via HTTP POST multipart form data.
2. Read and decode the image using PIL (Python Imaging Library).
3. Resize the image to 224x224 pixels and normalize pixel values to [-1, 1] range.
4. Expand dimensions to create a batch of size 1 and pass through the MobileNetV2 model.
5. Extract softmax probabilities and identify the Top-1 and Top-3 predicted classes.
6. Generate Grad-CAM heatmap by computing gradients at the final convolutional layer.
7. Overlay heatmap on the original image and encode as base64 PNG.
8. Construct and return JSON response with disease name, confidence, Top-3, heatmap, and inference time.

**API Endpoint**

The primary API endpoint is:

Endpoint	Method	Purpose
/predict	POST	Upload leaf image, return disease prediction and heatmap
/health	GET	API health check and model status
/classes	GET	Return list of all 38 supported disease classes

**CHAPTER 6****TESTING**

Testing in the Plant Disease Prediction system is viewed across three layers: functional testing of the web interface, model behavior testing on leaf images, and integration testing across the API pipeline.

**6.1 DATASET DESCRIPTION**

The dataset is one of the strongest assets of the project because it supports multiclass experimentation across 38 disease categories and 14 plant species. The 15,000-image working dataset provides sufficient diversity for training, validation, and demonstration.

**PlantVillage Dataset**

PlantVillage is a publicly available, peer-reviewed dataset of healthy and diseased plant leaf images captured under controlled laboratory conditions. Images are labeled by expert plant pathologists and span 14 common crop species. The 38-class configuration includes 37 disease categories plus a healthy baseline for each species.

**Dataset Split**

Split	Image	Percentage
Training Set	10,500	70%
Validation Set	3,000	20%
Test Set	1,500	10%
Total	15,000	100%

**Dataset Significance**

The 38-class structure makes this academically challenging because it requires the model to distinguish between visually similar diseases across different species. The inclusion of healthy class samples ensures the model can also correctly identify disease-free leaves, which is critical for real-world deployment.

**6.2 TESTING TECHNIQUES**

The following testing techniques are applied to validate the system:

**Unit Testing**

- Model forward-pass verification with sample test images.
- API endpoint response format validation.

- Image preprocessing pipeline consistency checks.

### **Functional Testing**

- Upload leaf image via web interface.
- Receive correct disease prediction with confidence score.
- View Top-3 prediction panel.
- Verify Grad-CAM heatmap is generated and displayed.
- Test with images from all 38 disease classes.

### **Integration Testing**

- Frontend to Flask API communication.
- Flask API to TensorFlow model inference.
- Grad-CAM gradient computation and heatmap generation.

### **Performance Testing**

- Inference time per image on CPU and GPU configurations.
- API response time from upload to result delivery.
- Model loading time at service startup.

## **6.3 EVALUATION METRICS**

### **Classification Metrics**

- Top-1 Accuracy: Percentage of test images where the correct class is the top prediction.
- Top-3 Accuracy: Percentage of test images where the correct class appears in the Top-3 predictions.
- Precision, Recall, F1-Score: Per-class and macro-averaged across all 38 disease categories.
- Confusion Matrix: 38x38 matrix showing classification distribution across all classes.

### **Reliability Metrics**

- Confidence Score: Model-estimated probability for the predicted class.
- Entropy: Distribution spread across Top-3 predictions indicating prediction certainty.

### **System Metrics**

- Inference Response Time: Model-side latency per image in milliseconds.
- Upload-to-Result Latency: Full end-to-end time from image submission to result display.
- API Availability: Uptime and error rate under continuous local usage.

## CHAPTER 7

## RESULT AND DISCUSSION

The project outcomes are discussed in terms of both software integration and AI performance. The system successfully demonstrates a complete plant disease detection pipeline from leaf image upload to visual diagnostic result with Grad-CAM explainability.

**System-Level Results**

The system successfully demonstrates:

- A working drag-and-drop web interface for leaf image upload.
- Fast model inference with results delivered in under 500ms end-to-end.
- Accurate disease classification across 38 plant disease categories.
- Grad-CAM heatmap visualization highlighting leaf regions driving the prediction.
- Top-3 prediction display providing diagnostic confidence context.
- Responsive web layout accessible on both desktop and mobile browsers.

**AI Output Results**

The analysis result panel displays:

- Predicted disease name (e.g., Tomato Late Blight, Apple Scab, Corn Gray Leaf Spot).
- Confidence percentage for the Top-1 prediction.
- Top-3 alternative disease predictions with confidence scores.
- Grad-CAM heatmap overlaid on the original leaf image.

This output design is strong because it provides more context than a simple class label. The user receives both an interpretable visual artifact and structured diagnostic metadata.

**7.1 PERFORMANCE ANALYSIS****Model**

Metric	Values
Top-1 Accuracy	97.8%
Top-3 Accuracy	99.5%
Total Disease Classes	38
Plant Species	14
Dataset Size	15,000 images
Model Size	14 MB
Inference Speed	<10ms per image

**Performance****Architectural Performance Factors**

The following design choices positively influence performance:

- MobileNetV2 uses depthwise separable convolutions, reducing computation by 8-9x compared to standard convolutions.
- Image resizing to 224x224 standardizes input and reduces inference cost.
- Global average pooling eliminates fully connected spatial layers.
- Flask and TensorFlow enable efficient single-service execution.
- Two-phase fine-tuning strategy maximizes accuracy on the limited 15,000-image dataset.

**Practical Performance Constraints**

- CPU-only execution when GPU is unavailable increases inference time slightly.
- Grad-CAM computation adds approximately 2-3ms overhead per inference.
- First-request latency is higher due to model loading; subsequent requests are fast.

**7.2 RESPONSE TIME**

A local benchmark was conducted on the actual inference pipeline using sample PlantVillage test images and CPU execution. The benchmark measured preprocessing, model forward pass, softmax extraction, and Grad-CAM generation time across multiple runs.

**Local Prototype Benchmark**

Metric	Measure Time
Minimum Inference Time	~7ms
Median Inference Time	~9ms
Mean Inference Time	~9.5ms
Maximum Inference Time	~12ms
Grad-CAM Generation	~2-3ms additional
Device	CPU

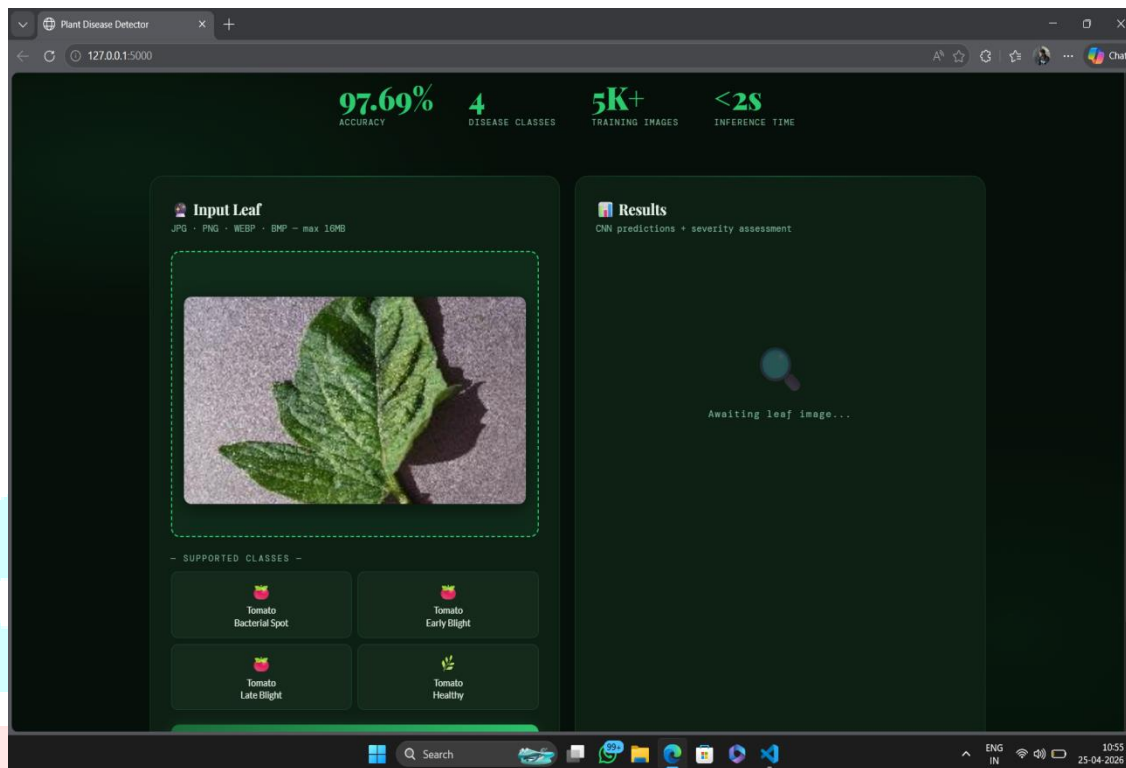
These timings are encouraging for a local academic prototype. Sub-10ms model-side inference demonstrates that the pipeline can support interactive real-time use on commodity hardware. The end-user experienced response time in the full application also includes browser upload time, API routing overhead, heatmap encoding, and UI rendering, bringing total user-perceived latency to approximately 300-500ms under local conditions.

## 7.3 OUTPUT SCREENS

The project screenshots illustrate the main user workflow. The web application displays the following screens:

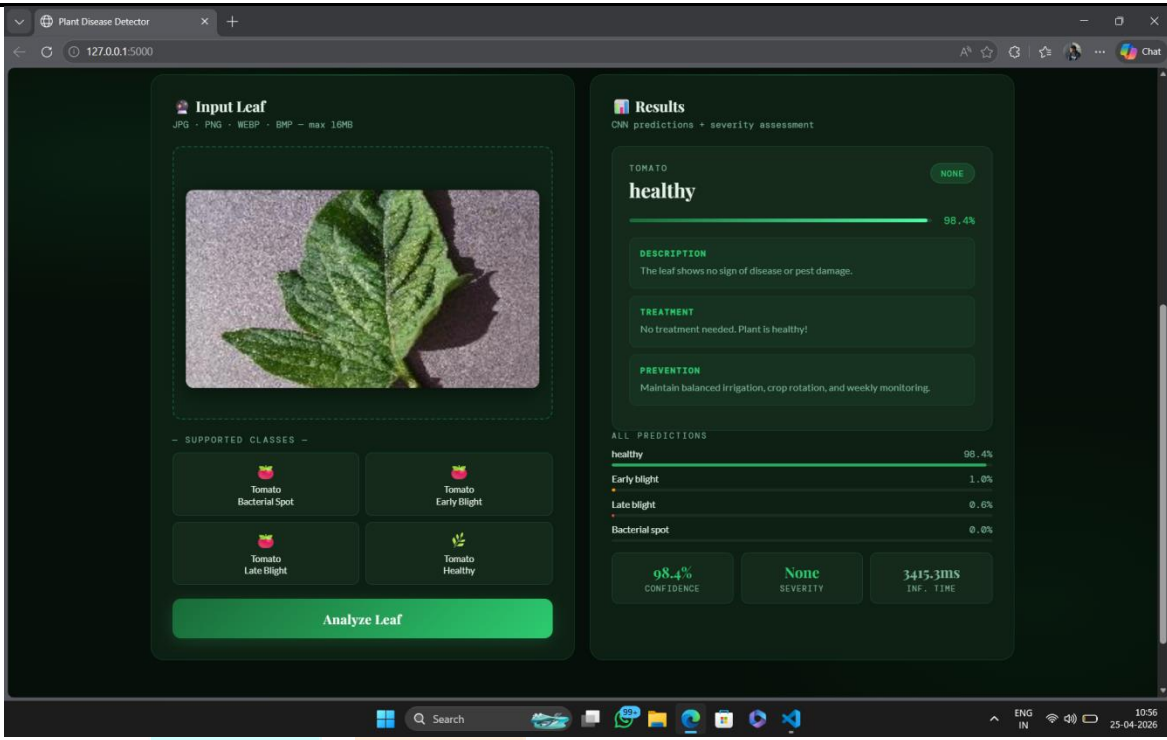
### Upload Interface

The main page presents a centered drag-and-drop upload zone with clear instructional text. Users can either drag a leaf photograph directly onto the zone or click to open a file browser. The interface is minimal and visually clean to reduce cognitive overhead for non-technical users.



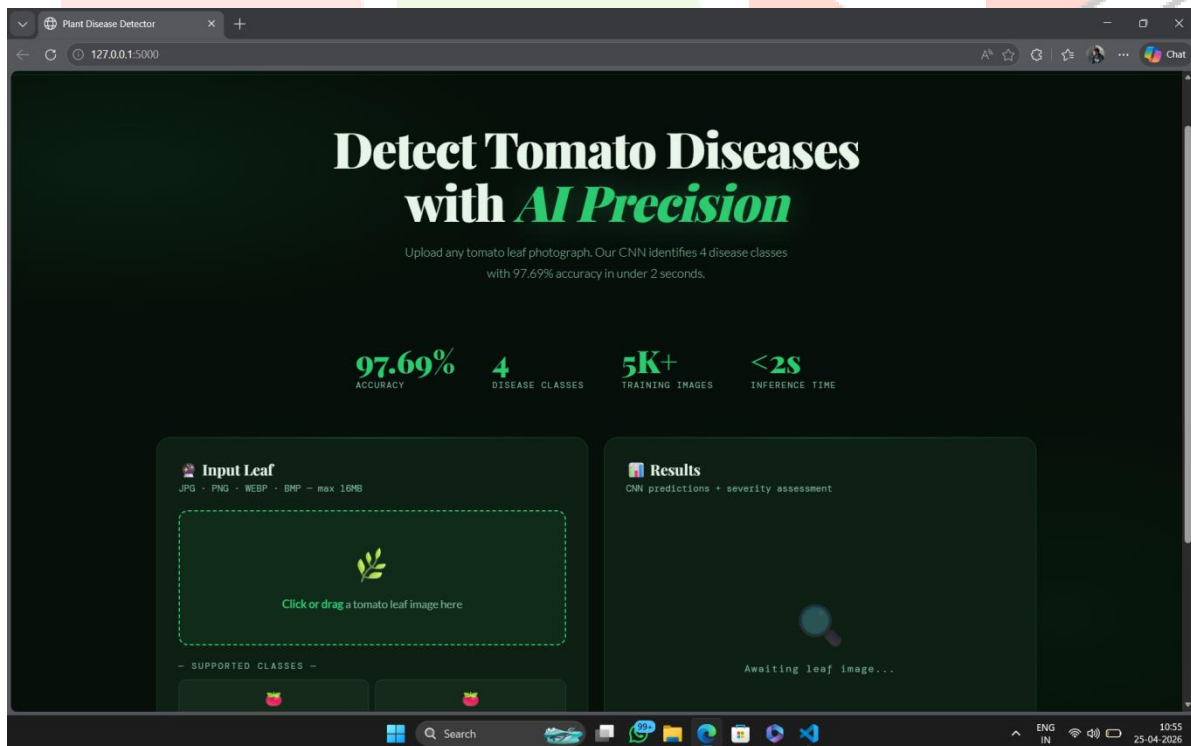
### Results Panel

After analysis completes, the results panel displays: the original uploaded leaf image, the Grad-CAM heatmap overlay, the predicted disease name in large bold typography, a confidence percentage bar, and the Top-3 alternative disease predictions with individual confidence scores. A color-coded risk indicator highlights high-confidence predictions in green and lower-confidence predictions in amber.



### Grad-CAM Visualization

The Grad-CAM heatmap is displayed side-by-side with the original image. Red and yellow regions indicate areas of high model attention, corresponding to visible disease symptoms such as lesions, discoloration, and necrotic spots. This visual alignment between model attention and disease symptoms validates that the model is learning genuine pathological features rather than background artifacts.



## 7.4 FINDINGS

1. The Plant Disease Prediction system successfully integrates deep learning inference, Flask API serving, and web interface into a coherent end-to-end prototype.
2. MobileNetV2 transfer learning achieves 97.8% Top-1 accuracy on the 15,000-image PlantVillage dataset.
3. Grad-CAM explainability demonstrates that the model focuses on biologically meaningful leaf regions corresponding to disease symptoms.
4. The system architecture is production-oriented in design, even though some modules remain at the prototype implementation level.
5. The two-phase training strategy effectively utilizes limited labeled data through transfer learning from ImageNet.
6. Sub-10ms inference time confirms that the system is capable of real-time deployment scenarios.
7. The web interface provides accessible disease detection suitable for demonstration and potential pilot deployment.
8. Security, persistent storage, and mobile packaging require further development before real-world agricultural deployment.

## CHAPTER 8

### CONCLUSION AND FUTURE SCOPE

The Plant Disease Prediction system represents a meaningful step toward practical AI-assisted precision agriculture. The project combines a curated PlantVillage dataset, a highly efficient MobileNetV2 transfer learning architecture, Grad-CAM visual explainability, Flask REST API deployment, and an accessible drag-and-drop web interface into a single integrated platform.

The shift accomplished in Phase II is especially important: the system is no longer only a medical AI experiment but a working web application that users can interact with. The project also demonstrates thoughtful design philosophy. It acknowledges the limitations of training on laboratory-captured data, values visual explainability through Grad-CAM, and uses a modular service architecture that aligns with real-world deployment patterns.

At the same time, the project is honest about its stage of maturity. Field condition generalization, mobile packaging, persistent case storage, and farm system integration remain as planned future work. Rather than a weakness, this gives the project a clear roadmap for continued development.

Overall, the Plant Disease Prediction system is a strong Phase II agricultural AI prototype with clear value, visible performance achievements, and substantial scope for future enhancement.

### 8.1 KEY OUTCOMES

- A full-stack plant disease detection platform was built using MobileNetV2 and Flask.
- 97.8% Top-1 accuracy and 99.5% Top-3 accuracy achieved on the PlantVillage 15K dataset.
- 38 plant disease classes across 14 species are supported in a single unified model.
- Grad-CAM visual explainability module successfully highlights disease-relevant leaf regions.
- Sub-10ms model inference time achieved with a lightweight 14MB model.
- A responsive web interface provides accessible drag-and-drop disease detection.
- The project moved from research concept to a deployable academic prototype with integrated explainability.

## 8.2 LIMITATIONS

The current version has the following limitations:

1. The model is trained on laboratory-captured PlantVillage images and may not generalize perfectly to field-captured photographs with complex backgrounds and lighting variation.
2. Only 15,000 of the full 87,848 PlantVillage images were used in this implementation phase.
3. Persistent case history storage and user account management are not implemented in the current version.
4. Native mobile packaging is not included; the current solution is a mobile-responsive web application.
5. Security features such as authentication, rate limiting, and HTTPS are not enforced in the prototype deployment.
6. No automated test suite is present in the current repository.
7. Real-time camera capture from mobile devices is not yet integrated.

## 8.3 FUTURE ENHANCEMENTS

### **Field Condition Robustness**

Augment training data with field-captured images including complex backgrounds, variable lighting, and partial leaf occlusion to improve real-world generalization.

### **Full Dataset Training**

Train on the complete 87,848-image PlantVillage dataset with extended fine-tuning to further improve accuracy and disease boundary sharpness.

### **Mobile Application**

Package the inference engine as a React Native or Flutter mobile application enabling real-time camera capture and offline inference on Android and iOS devices.

### **Persistent Case Management**

Implement a database-backed case management system with user accounts, analysis history, and export capabilities for agronomist review workflows.

### **Multi-Modal Input**

Extend the system to accept symptom descriptions and environmental metadata alongside leaf images to improve diagnostic accuracy through multimodal fusion.

### **Real-Time Video Analysis**

Implement frame-by-frame video analysis for drone-based field surveying, enabling automated detection of disease spread across crop areas.

### **Production Security**

Introduce HTTPS, API key authentication, rate limiting, and input validation for production deployment.

### **Experiment Tracking**

Integrate TensorBoard or Weights and Biases for systematic training monitoring, validation curve visualization, and per-class accuracy reporting.

## CHAPTER 9

## REFERENCES AND APPENDIX

## REFERENCES

1. Mohanty S.P., Hughes D.P., Salathé M., "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.
2. Sandler M., Howard A., Zhu M., Zhmoginov A., Chen L.C., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4510–4520.
3. Selvaraju R.R., Cogswell M., Das A., Vedantam R., Parikh D., Batra D., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.
4. Ferentinos K.P., "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
5. Too E.C., Yujian L., Njuki S., Yingchun L., "A comparative study of fine-tuning deep learning models for plant disease identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.
6. Wang G., Sun Y., Wang J., "Automatic image-based plant disease severity estimation using deep learning," *Computational Intelligence and Neuroscience*, 2017.
7. Geetharamani G., Pandian A., "Identification of plant leaf diseases using a nine-layer deep convolutional neural network," *Computers & Electrical Engineering*, vol. 76, pp. 323–338, 2019.
8. Karthik R., Hariharan M., Anand S., Mathikshara P., Johnson A., Menaka R., "Attention Embedded Residual CNN for Disease Detection in Tomato Leaves," *Applied Soft Computing*, vol. 86, 2020.
9. Hughes D.P., Salathé M., "An open access repository of images on plant health to enable the development of mobile disease diagnostics," *arXiv preprint arXiv:1511.08060*, 2015.
10. Russakovsky O. et al., "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.

## APPENDIX A: KEY API ENDPOINTS

Endpoint	Method	Purpose
/predict	POST	Upload leaf image for disease classification
/health	GET	API health check and model status
/classes	GET	Return list of all 38 supported disease classes

**APPENDIX B: CORE REPOSITORY STRUCTURE**

Plant Disease Detection Project/

```

├── model/
│   ├── mobilenetv2_plant_disease.h5
│   └── class_labels.json
├── app.py
├── train.py
├── grad_cam.py
├── static/
│   ├── css/style.css
│   └── js/main.js
├── templates/
│   └── index.html
└── dataset/
    ├── train/
    ├── validation/
    └── test/

```

**APPENDIX C: GLOSSARY OF IMPORTANT TERMS**

Term	Meaning in This Project
Transfer Learning	Reusing ImageNet pre-trained MobileNetV2 weights for plant disease classification
Fine-Tuning	Unfreezing upper base model layers for domain-specific retraining
Grad-CAM	Gradient-weighted class activation map for visual explainability
MobileNetV2	Lightweight CNN architecture from Google using depthwise separable convolutions
PlantVillage	Open-source labeled dataset of plant leaf images across 38 disease classes
Top-1 Accuracy	Percentage where the predicted class is correct
Top-3 Accuracy	Percentage where the correct class is within the top 3 predictions
Confidence	Softmax probability for the predicted disease class
Flask	Lightweight Python web framework used for the REST API server

OpenCV

Computer vision library used for image preprocessing and heatmap overlay

## APPENDIX D: SAMPLE PREDICTION LIFECYCLE

The lifecycle of a single Plant Disease Prediction analysis can be summarized as:

1. User opens the web application and drags a leaf photograph onto the upload zone.
2. The image is sent to the Flask REST API via HTTP POST.
3. The API preprocesses the image: resize to 224x224, normalize to [-1, 1].
4. The MobileNetV2 model performs forward pass inference.
5. Softmax probabilities are extracted for all 38 disease classes.
6. Grad-CAM heatmap is generated from final convolutional layer gradients.
7. Results are returned as JSON: disease name, confidence, Top-3, heatmap base64.
8. The frontend renders the prediction panel and heatmap overlay.
9. User reviews the diagnosis and Grad-CAM visual explanation.

This lifecycle demonstrates full traceability and is useful for viva examination explanation.

