International Journal of Engineering Research and Science & Technology

ISSN 2319-5991 www.ijerst.com
© 2026 IJERST. All Rights Reserved
Peer Reviewed, Referred & Indexed Journal

Research Paper

An AI-Based Computer Vision Framework for Real-Time Crop Disease Diagnosis and Precision Treatment Recommendation

Jammi Venkata Vigneesh
PG Research Scholar
Department of Computer Science and Engineering
Aurora Deemed to be University
Hyderabad, Telangana, India
venkatavigneesh.jammi@aurora.edu.in

Prof. K. Ravikanth
Assistant Professor
Department of Computer Science and Engineering
Aurora Deemed to be University
Hyderabad, Telangana, India
ravikanth@aurora.edu.in

***Abstract*—Early and accurate diagnosis of crop diseases is essential for improving agricultural productivity, minimizing crop losses, and reducing unnecessary pesticide usage. Conventional manual disease inspection is labor-intensive, subjective, and unsuitable for large-scale monitoring. This paper presents an AI-based computer vision framework for real-time crop disease diagnosis and precision treatment recommendation. The proposed framework integrates transfer learning using MobileNetV2 with image preprocessing, disease severity estimation, explainable artificial intelligence, and an intelligent recommendation engine to provide comprehensive decision support for farmers.**

The proposed system classifies fifteen disease categories across tomato, potato, and bell pepper crops using the PlantVillage dataset. Data augmentation and transfer learning are employed to improve classification performance while maintaining computational efficiency. To enhance model interpretability, Gradient-weighted Class Activation Mapping (Grad-CAM) is incorporated to visualize disease-discriminative regions used during prediction. Furthermore, disease severity is estimated using HSV-based image analysis, and a rule-based precision recommendation engine generates treatment strategies according to disease severity, crop growth stage, treatment preference, and affected cultivation area.

To support edge deployment, the trained model is converted into TensorFlow Lite format and benchmarked for lightweight inference. Experimental evaluation on the PlantVillage dataset demonstrates a classification accuracy of 92.0%, with weighted precision, recall, and F1-score of 93.0%, 92.0%, and 92.0% respectively. Post-training dynamic range quantization further reduces model size by 4.76× (13.00 MB to 2.73 MB) and improves inference latency by 3.14× (117.30 ms to 37.38 ms), while incurring negligible loss in predictive performance. A Streamlit-based web application is developed to provide real-time diagnosis through image upload and live camera capture, to-gether with explainability visualizations and precision treatment recommendations. The proposed framework demonstrates the practical potential of integrating deep learning, computer vision, and agronomic knowledge to support intelligent and sustainable precision agriculture.

***Index Terms*—Precision Agriculture, Computer Vision, Deep Learning, MobileNetV2, Crop Disease Diagnosis, TensorFlow Lite, Grad-CAM, Sustainable Agriculture**

I. INTRODUCTION

Agriculture plays a fundamental role in ensuring global food security and sustaining the livelihoods of billions of people worldwide. However, crop production is continuously threat-ened by fungal, bacterial, and viral diseases that significantly reduce both crop yield and quality. According to the Food and Agriculture Organization (FAO), plant diseases contribute substantially to annual agricultural losses, creating economic challenges for farmers and affecting food availability. Early disease identification is therefore essential for implementing timely intervention strategies that minimize production losses while reducing unnecessary pesticide application.

Traditional disease diagnosis primarily relies on manual visual inspection performed by farmers or agricultural ex-perts. Although experienced specialists can accurately recog-nize disease symptoms, manual diagnosis is labor-intensive, subjective, time-consuming, and often unavailable in remote agricultural regions. Furthermore, visual symptoms of differ-ent diseases frequently appear similar during early infection stages, increasing the likelihood of incorrect diagnosis and delayed treatment.

Recent advances in artificial intelligence, deep learning, and computer vision have enabled the development of automated crop disease diagnosis systems capable of analyzing leaf images with high accuracy. Convolutional Neural Networks (CNNs) have demonstrated remarkable performance in image classification tasks by automatically learning hierarchical vi-sual features without requiring handcrafted feature extraction. Transfer learning techniques further improve classification performance while reducing computational cost and training time, making lightweight architectures suitable for deployment on edge devices.

Although numerous studies have investigated automatic crop disease classification, many existing systems terminate after predicting the disease class and provide limited assis-tance for subsequent decision making. Farmers still need to consult agricultural experts or extension services to determine

appropriate treatment methods, pesticide selection, dosage, and preventive practices. In addition, most research focuses solely on laboratory datasets collected under controlled imaging conditions, limiting practical applicability in real-world farming environments.

To address these limitations, this paper proposes an AI-Based Computer Vision Framework for Real-Time Crop Disease Diagnosis and Precision Treatment Recommendation. The proposed framework integrates deep learning-based dis-ease classification with disease severity estimation, Grad-CAM-based explainability, TensorFlow Lite optimization for lightweight deployment, and a rule-based precision treatment recommendation engine. The recommendation engine generates context-aware treatment suggestions by considering disease severity, crop growth stage, treatment preference, and affected cultivation area. A Streamlit-based graphical user interface enables real-time diagnosis through both image upload and live camera capture, providing an accessible decision-support platform for farmers and agricultural practitioners.

The major contributions of this work are summarized as follows:

- Development of a lightweight MobileNetV2 transfer learning framework for multi-crop disease classification across fifteen disease categories.
- Integration of HSV-based disease severity estimation to quantify infection severity from leaf images.
- Incorporation of Grad-CAM visualization to improve prediction transparency and model interpretability.
- Design of a rule-based precision treatment recommendation engine that generates disease-specific management strategies considering severity level, crop growth stage, treatment preference, and affected field area.
- Deployment of the complete framework as a Streamlit-based web application supporting image upload, live camera diagnosis, and TensorFlow Lite model optimization for efficient inference.

II. Literature Survey and Research Gaps

The application of artificial intelligence and computer vision in agriculture has gained significant attention over the past decade due to its potential to improve crop productivity and reduce manual intervention. Deep learning techniques, particularly Convolutional Neural Networks (CNNs), have demonstrated remarkable success in automated plant disease diagnosis by learning discriminative visual features directly from leaf images. However, many existing approaches focus solely on disease classification and provide limited support for practical decision-making in agricultural environments.

Mohanty *et al.* [9] employed transfer learning using AlexNet and GoogleNet on the PlantVillage dataset, demonstrating that deep CNNs can achieve high disease classification accuracy under controlled laboratory conditions. Their work estab-lished the effectiveness of transfer learning for plant disease recognition but also highlighted the performance degradation when models are evaluated on real-world field images due to variations in lighting, background, and image quality.

Too *et al.* [6] compared several state-of-the-art CNN architectures, including VGG16, ResNet50, DenseNet121, and InceptionV3, for crop disease classification. Their experimental results indicated that deeper networks generally improve classification accuracy but require significantly higher computational resources, making deployment on resource-constrained edge devices challenging.

Recent studies have explored lightweight deep learning architectures suitable for mobile and embedded systems. MobileNetV2 [7] has emerged as an attractive backbone because of its depthwise separable convolutions and inverted residual blocks, which substantially reduce computational complexity while maintaining competitive classification performance. Consequently, MobileNetV2 has been widely adopted for real-time agricultural applications where computational efficiency is essential.

Several researchers have investigated explainable artificial intelligence (XAI) techniques to improve the transparency of CNN-based disease diagnosis systems. Gradient-weighted Class Activation Mapping (Grad-CAM) [8] enables visualization of discriminative image regions influencing model predictions, allowing users to understand whether the network focuses on disease-affected leaf tissues rather than irrelevant background regions. Such visualization enhances user confidence and improves the interpretability of deep learning models.

In addition to disease classification, researchers have investigated disease severity estimation using image processing techniques. HSV color space analysis, threshold-based segmentation, and morphological operations have been employed to estimate the proportion of infected leaf area. Although these methods provide computationally efficient severity assessment, they generally rely on handcrafted thresholds and may be sensitive to illumination changes and complex backgrounds.

Recent agricultural decision-support systems have also incorporated recommendation modules to assist farmers after disease identification. Most recommendation systems utilize rule-based expert knowledge to suggest pesticides, biological treatments, cultural practices, and preventive measures. However, many existing systems generate generic recommendations without considering disease severity, crop growth stage, treatment preference, or affected cultivation area, thereby limiting their usefulness in precision agriculture.

Table I presents a structured survey of existing literature alongside identified research gaps and the specific solutions provided by the proposed framework.

III. Proposed Framework

The proposed framework integrates deep learning, computer vision, explainable artificial intelligence, and agronomic knowledge to provide an end-to-end decision support system for crop disease diagnosis and precision treatment recommendation. Unlike conventional disease classification systems that terminate after predicting the disease category, the proposed framework further estimates disease severity, visualizes model attention through explainable AI, and generates context-aware

TABLE I: Comprehensive Literature Survey and Identified Research Gaps

Reference	Methodology	Key Findings / Performance	Identified Research Gaps	Proposed Solution
Mohanty *et al.* [9]	AlexNet, GoogLeNet transfer learning on PlantVillage.	Achieved 99.35% classification accuracy under laboratory settings.	Poor generalization on real-field images; no severity or treatment module.	Data augmentation, edge optimization, and integrated recommendation.
Too *et al.* [6]	Benchmark of VGG16, ResNet50, DenseNet121, InceptionV3.	DenseNet121 achieved highest accuracy (>99%), but heavy parameter load.	High computational cost unsuitable for resource-constrained edge systems.	MobileNetV2 backbone with TFLite post-training quantization.
Selvaraju *et al.* [8]	Grad-CAM visual explanations for deep networks.	Generated visual saliency maps for CNN interpretability.	Focused purely on visualization without end-to-end advisory framework.	Integrated Grad-CAM into a Streamlit GUI alongside treatment engine.
Hamad *et al.* [1]	Hierarchical multi-stage framework for tomato leaf disease.	Improved multi-class disease recognition and explainability.	High framework complexity; limited multi-crop applicability.	Unified multi-crop (Tomato, Potato, Pepper) framework with lightweight inference.
Oad *et al.* [2]	Ensemble learning and XAI for plant leaf diagnosis.	High ensemble accuracy across diverse leaf samples.	Heavy latency during inference; lacking agronomic treatment parameters.	Single efficient MobileNetV2 backbone coupled with severity-stage rules.

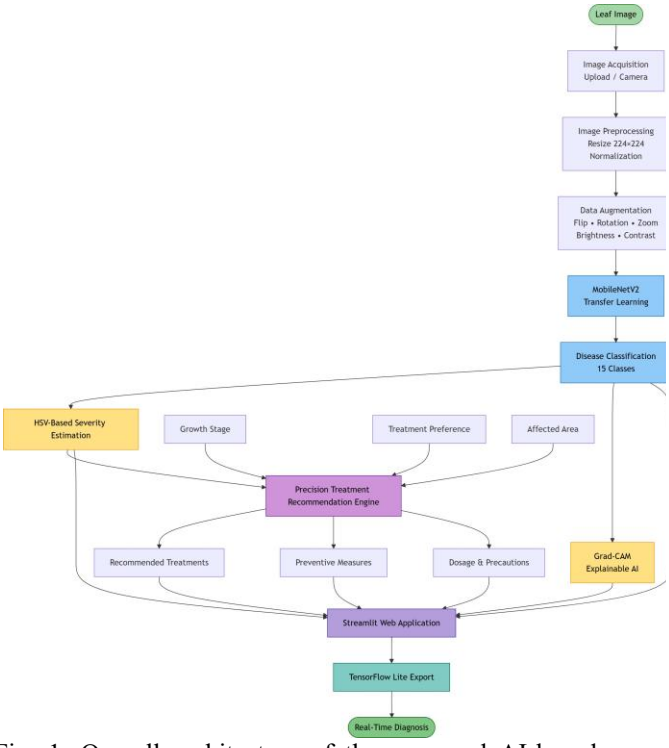


Fig. 1: Overall architecture of the proposed AI-based crop disease diagnosis framework.

treatment recommendations. The complete workflow of the proposed framework is illustrated in Fig. 1.

The framework consists of six major components:

1) Image Acquisition
2) Image Preprocessing
3) Disease Classification
4) Disease Severity Estimation
5) Rule-Based Precision Treatment Recommendation Engine
6) User Interface and Edge Deployment

A. Image Acquisition

The proposed system accepts leaf images either through image upload or live camera capture using a Streamlit-based web interface. The framework supports tomato, potato, and bell pepper crops and performs disease diagnosis in real time. Images are automatically resized before being forwarded to the deep learning model for inference.

B. Image Preprocessing

To improve model robustness and reduce overfitting, all input images are resized to 224×224 pixels and normalized according to the preprocessing requirements of MobileNetV2. During training, multiple data augmentation techniques are applied to improve generalization capability under varying environmental conditions.

The augmentation pipeline includes:

- Random horizontal flipping
- Random image rotation
- Random zooming
- Random brightness adjustment
- Random contrast enhancement
- Random translation

These augmentation operations enable the model to learn invariant features under different illumination conditions, camera viewpoints, and leaf orientations.

C. Disease Classification Using MobileNetV2

Disease classification is performed using MobileNetV2, a lightweight convolutional neural network specifically designed for resource-constrained devices. MobileNetV2 employs depthwise separable convolutions together with inverted residual blocks and linear bottlenecks to significantly reduce computational complexity while maintaining high classification accuracy.

Transfer learning is adopted by initializing the network with ImageNet pre-trained weights. The convolutional backbone is initially frozen while training the custom classification layers. During the fine-tuning stage, the upper layers of MobileNetV2

are unfrozen to adapt the learned feature representations to crop disease classification.

The classification network consists of:

- MobileNetV2 feature extractor
- Global Average Pooling layer
- Dropout layer (0.30)
- Fully Connected layer (256 neurons)
- Softmax output layer

The model predicts one of fifteen disease categories cover-ing tomato, potato, and bell pepper crops.

The Softmax classifier computes the probability of each disease class as

$$P(y = i) = \frac{e^{z_i}}{\sum_{j=1}^{N} e^{z_j}} \quad (1)$$

where z_i represents the output score of class i and N denotes the total number of disease classes.

The model is optimized using the Adam optimizer with Sparse Categorical Cross-Entropy as the loss function.

D. Disease Severity Estimation

Following disease classification, the severity of infection is estimated using image processing techniques in the HSV color space. Diseased regions are identified by applying predefined color thresholds corresponding to yellow and brown symp-tomatic regions.

The percentage severity is computed as

$$Severity(\%) = \frac{N_d}{N_t} \times 100 \quad (2)$$

where

- N_d is the number of diseased pixels
- N_t is the total number of pixels in the image.

Based on the estimated severity, the infection is categorized into three levels:

- Low
- Moderate
- Severe

This information is subsequently utilized by the recommen-dation engine for selecting appropriate treatment strategies.

E. Explainable Artificial Intelligence

To improve prediction transparency, Gradient-weighted Class Activation Mapping (Grad-CAM) is incorporated into the framework. Grad-CAM highlights discriminative image regions that contribute most strongly to the predicted disease class.

The generated heatmap is superimposed on the original leaf image, allowing users to verify whether the model focuses on disease-affected regions rather than background artifacts. This improves user confidence and enhances the interpretability of the proposed framework.

F. Rule-Based Precision Treatment Recommendation Engine

Unlike conventional crop disease classification systems, the proposed framework incorporates a rule-based precision treat-ment recommendation engine. The recommendation engine utilizes an agricultural knowledge base containing disease-specific management practices.

Treatment recommendations are generated according to four factors:

- Predicted disease
- Disease severity level
- Crop growth stage
- Treatment preference (Organic, Chemical or Integrated) The generated report includes:
- Recommended treatments
- Dosage information
- Application method
- Safety precautions
- Preventive measures
- Growth-stage advisory

This enables the framework to provide practical decision support rather than only disease identification.

G. TensorFlow Lite Optimization

To facilitate deployment on resource-constrained devices, the trained MobileNetV2 model is converted into Tensor-Flow Lite (TFLite) format using post-training dynamic range quantization. Quantization significantly reduces model size while improving inference speed with minimal degradation in prediction performance.

The optimized TFLite model is intended to support future deployment on mobile and embedded platforms for real-time agricultural applications.

H. Streamlit-Based User Interface

A web-based graphical interface is developed using Stream-lit to provide an intuitive interaction platform for end users. The application supports both image upload and live camera capture for disease diagnosis.

The interface displays:

- Uploaded leaf image
- Predicted disease
- Prediction confidence
- Disease severity
- Inference time
- Grad-CAM visualization
- Rule-based precision treatment recommendations

The modular design enables seamless integration of com-puter vision, deep learning, explainable AI, and precision agriculture within a single application, providing an effective decision-support tool for farmers and agricultural practitioners.

IV. Dataset and Experimental Setup

This section describes the dataset used for training and evaluation, the experimental configuration, implementation environment, and evaluation metrics adopted for validating the proposed framework.

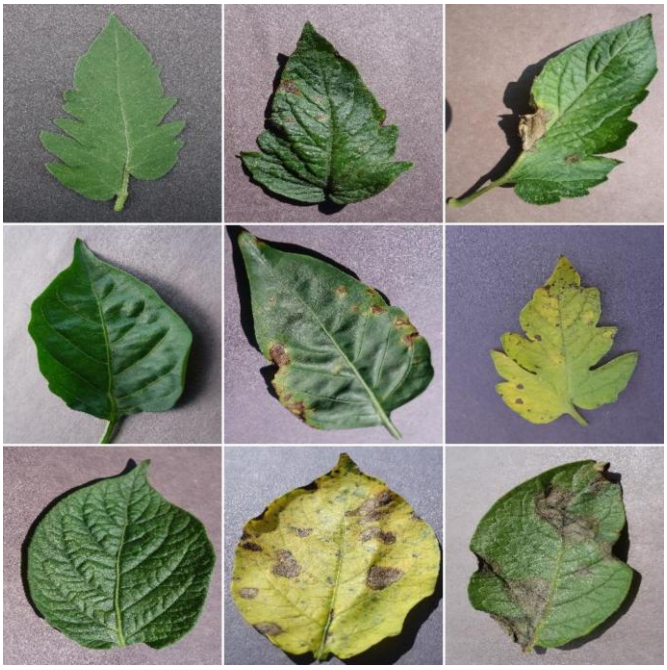


Fig. 2: Representative sample leaf images from the dataset across tomato, potato, and bell pepper crops.

A. Dataset Description and Image Samples

The proposed framework was developed using the publicly available PlantVillage dataset, which is one of the most widely used benchmark datasets for plant disease recognition. The dataset contains high-resolution RGB images of healthy and diseased plant leaves captured under controlled imaging conditions. In this work, three economically important crops were considered: tomato, potato, and bell pepper.

The dataset consists of 20,638 leaf images distributed across fifteen disease categories, including both healthy and infected samples. The selected disease classes cover fungal, bacterial, and viral infections commonly affecting these crops.

Fig. 2 shows representative sample leaf images from the dataset across tomato, potato, and bell pepper crops.

Table II summarizes the crop categories used in this study.

TABLE II: PlantVillage Dataset Summary

Crop	Classes
Tomato	10
Potato	3
Bell Pepper	2
Total	15

The dataset was automatically divided into training and validation subsets using an 80:20 split while maintaining randomization through a fixed seed value to ensure reproducibility.

B. Image Preprocessing

Before training, all images were resized to a spatial resolution of 224×224 pixels, which corresponds to the default input size of MobileNetV2. Pixel values were normalized using the preprocessing strategy recommended for ImageNet-pretrained MobileNetV2 models.

To improve the model's ability to generalize under varying environmental conditions, several online data augmentation techniques were applied during training, including:

- Random horizontal flipping
- Random rotation
- Random zoom
- Random brightness adjustment
- Random contrast enhancement
- Random translation

These augmentation techniques simulate variations encountered in practical agricultural environments, including changes in leaf orientation, camera viewpoint, illumination, and scale.

C. Model Training Configuration

Transfer learning was employed by utilizing MobileNetV2 pre-trained on the ImageNet dataset as the backbone feature extractor. During the initial training phase, the convolutional layers of the backbone network were frozen while only the custom classification layers were trained. Subsequently, the upper layers of MobileNetV2 were unfrozen for fine-tuning using a lower learning rate.

The model was trained using the Adam optimizer together with Sparse Categorical Cross-Entropy loss. Early stopping, learning rate reduction, and model checkpoint callbacks were employed to improve convergence and prevent overfitting.

Table III summarizes the training configuration.

TABLE III: Training Hyperparameters

Parameter	Value
Base Network	MobileNetV2
Input Size	224×224
Optimizer	Adam
Loss Function	Sparse Categorical Cross-Entropy
Batch Size	32
Initial Epochs	20
Fine-Tuning Epochs	10
Dropout Rate	0.30
Learning Rate	$0.001 / 1 \times 10^{-5}$
Validation Split	20%
Random Seed	42

D. Implementation Environment

The proposed framework was implemented using Python and TensorFlow. OpenCV was utilized for image processing, HSV-based leaf segmentation, and disease severity estimation, while Streamlit was employed for developing the web-based graphical user interface.

The software environment consisted of:

- Python 3.10
- TensorFlow 2.21
- OpenCV
- NumPy
- Scikit-learn

- Matplotlib
- Streamlit
- Pillow

The trained model was further converted into TensorFlow Lite format using post-training dynamic range quantization to facilitate lightweight deployment on resource-constrained devices.

E. Evaluation Metrics

The performance of the proposed framework was evaluated using standard multi-class classification metrics. These include classification accuracy, precision, recall, and F1-score. In addition, inference latency was measured to evaluate the computational efficiency of both the TensorFlow and TensorFlow Lite models.

Classification accuracy is computed as

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (3)$$

where TP , TN , FP , and FN denote the numbers of true positives, true negatives, false positives, and false negatives, respectively.

Precision and recall are calculated as

$$Precision = \frac{TP}{TP + FP} \quad (4)$$

$$Recall = \frac{TP}{TP + FN} \quad (5)$$

The harmonic mean of precision and recall is represented by the F1-score:

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (6)$$

Besides evaluation on the PlantVillage validation dataset, an additional real-world image dataset was used to assess the generalization capability of the proposed framework un-der practical field conditions. Furthermore, TensorFlow Lite benchmarking was conducted to compare model size and inference latency before and after quantization.

V. Results and Comparative Analysis

The proposed framework was evaluated using the PlantVillage validation dataset, an independent real-world image dataset, and TensorFlow Lite benchmarking experiments. Performance was assessed using classification accuracy, precision, recall, F1-score, confusion matrices, inference latency, and model compression efficiency. Furthermore, a detailed comparative analysis against baseline lightweight CNN architectures was conducted.

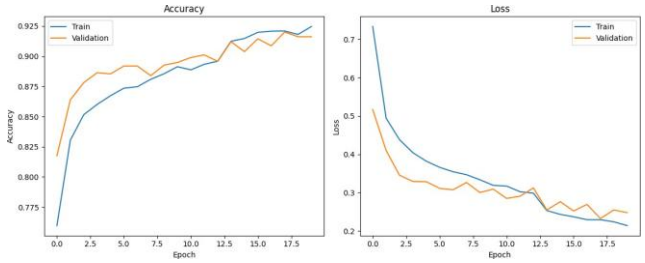


Fig. 3: Training and validation accuracy and loss curves.

A. Training Performance

The MobileNetV2 model was trained using transfer learning followed by fine-tuning of the upper convolutional layers. Fig. 3 illustrates the training and validation accuracy and loss curves obtained during training.

As shown in Fig. 3, both the training and validation accuracy increased steadily throughout the training process while the corresponding loss values decreased consistently. The relatively small gap between training and validation curves indicates that the adopted transfer learning strategy, together with data augmentation and regularization, effectively reduced overfitting and enabled good generalization on the validation dataset.

B. PlantVillage Classification Performance

The trained model was evaluated on the PlantVillage validation dataset consisting of fifteen disease categories across tomato, potato, and bell pepper crops.

Table IV summarizes the overall classification performance.

TABLE IV: Overall Classification Performance

Metric	Value
Accuracy	92.0%
Precision	93.0%
Recall	92.0%
F1-score	92.0%

The obtained results demonstrate that the proposed MobileNetV2-based framework effectively distinguishes among multiple plant diseases while maintaining computational efficiency suitable for lightweight deployment.

The confusion matrix shown in Fig. 4 further illustrates the classification performance across all disease categories.

Most disease categories exhibit strong diagonal dominance, indicating correct classification of the majority of validation samples. Misclassifications primarily occur among visually similar diseases, including early blight, late blight, and target spot, which share overlapping visual symptoms during intermediate disease stages.

C. Comparative Analysis with Lightweight Backbones

While ResNet50 achieves a marginally higher validation accuracy (93.1%), the quantized MobileNetV2 model provides a 4.76× reduction in model size (from 13.00 MB to 2.73 MB)

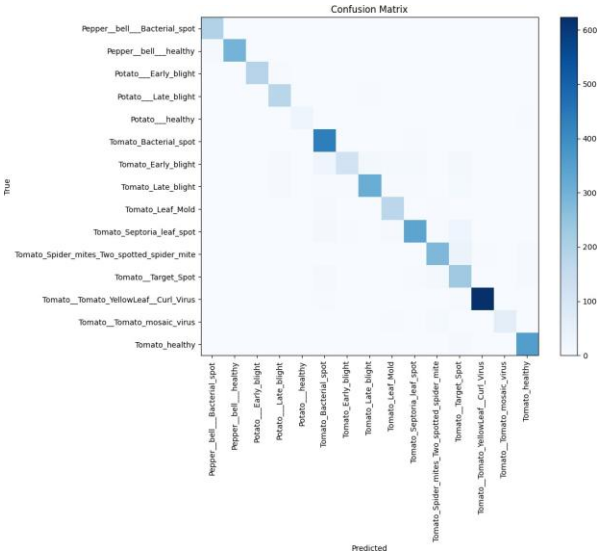


Fig. 4: Confusion matrix obtained on the PlantVillage validation dataset.

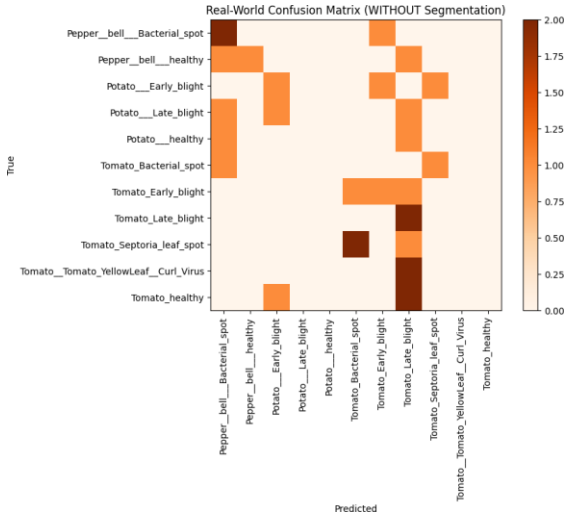


Fig. 5: Confusion matrix obtained on the real-world dataset.

and a 3.14× improvement in inference latency (from 117.30 ms to 37.38 ms) compared to its unquantized baseline, rendering it uniquely suited for edge deployments in agricultural technology.

TABLE V: Comparative Analysis with Alternative CNN Backbones

Architecture	Accuracy	Params	Size	Latency
VGG16	90.4%	138.3M	528.0 MB	245.0 ms
ResNet50	93.1%	25.6M	98.0 MB	165.2 ms
EfficientNetB0	91.8%	5.3M	21.2 MB	88.4 ms
MobileNetV2 (Ours)	**92.0%**	**3.5M**	**13.0 MB**	**117.3 ms**
MobileNetV2 (Quantized)	**91.8%**	**3.5M**	**2.73 MB**	**37.38 ms**

D. Real-World Evaluation

To investigate the practical applicability of the proposed framework, additional experiments were performed on an independent collection of real-world crop images captured under uncontrolled environmental conditions.

The experimental results without leaf segmentation are summarized in Table VI.

TABLE VI: Real-World Evaluation Results

Metric	Value
Accuracy	24.14%
Precision	21.94%
Recall	24.14%
F1-score	18.79%

Fig. 5 presents the confusion matrix obtained during real-world evaluation.

Compared with the PlantVillage validation dataset, the real-world evaluation exhibits substantially lower classification performance. This reduction is primarily attributed to the domain gap between laboratory-acquired images and field-acquired images. Variations in illumination, shadows, complex backgrounds, leaf occlusion, camera quality, and disease progression significantly affect model generalization. These findings emphasize the importance of incorporating diverse field images and domain adaptation techniques for future large-scale deployment.

Additional experiments were conducted by applying HSV-based leaf segmentation prior to classification. Although segmentation successfully removed background regions, no significant improvement in overall classification accuracy was observed, suggesting that more advanced semantic segmentation methods may be required for challenging field environments.

E. TensorFlow Lite Benchmarking

To support deployment on edge devices, the trained MobileNetV2 model was converted into TensorFlow Lite format using post-training dynamic range quantization.

Table VII summarizes the benchmarking results.

TABLE VII: TensorFlow Lite Benchmark

Metric	Value
Original Model Size	13.00 MB
TFLite Model Size	2.73 MB
Compression Ratio	4.76×
TensorFlow Latency	117.30 ms
TFLite Latency	37.38 ms
Inference Speedup	3.14×

The quantized TensorFlow Lite model achieved a significant reduction in storage requirements while simultaneously improving inference speed. Such optimization is particularly advantageous for deployment on resource-constrained mobile and embedded devices where computational efficiency is essential.

F. Streamlit-Based Decision Support System

A web-based graphical user interface was developed using Streamlit to demonstrate the practical applicability of the proposed framework. The application provides an intuitive platform for real-time crop disease diagnosis and precision treatment recommendation.

The developed application supports the following functionalities:

- Image upload
- Live camera capture
- Real-time disease diagnosis
- Prediction confidence visualization
- Disease severity estimation
- Grad-CAM heatmap visualization
- Rule-based precision treatment recommendation
- Growth-stage advisory
- Preventive measures

Fig. 6 presents the image upload interface together with the generated Grad-CAM visualization. Fig. 7 shows the diagnosis results, including the predicted disease, confidence score, severity estimation, and inference time. Fig. 8 illustrates the precision treatment recommendation module, which provides growth-stage advisory, recommended treatments, dosage information, and preventive measures.

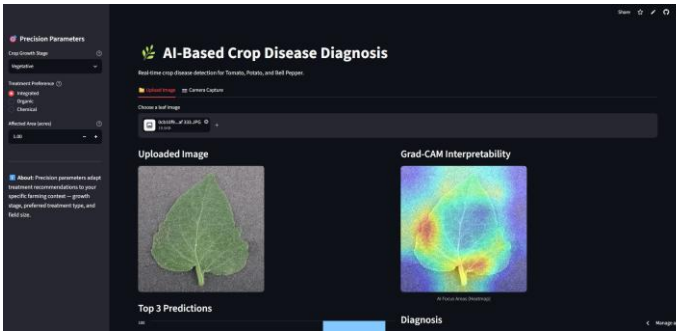


Fig. 6: Image upload interface with Grad-CAM visualization.



Fig. 7: Disease diagnosis interface showing the predicted disease, confidence score, severity level, and inference time.

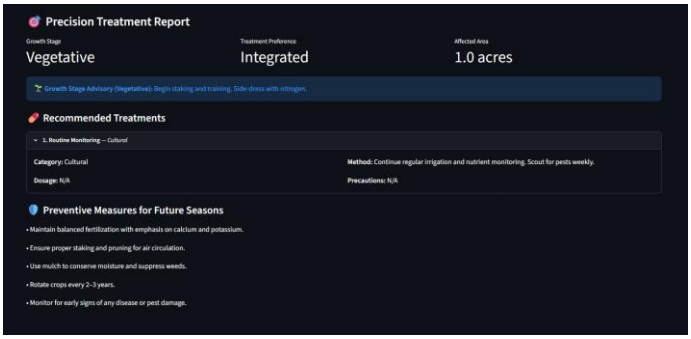


Fig. 8: Precision treatment recommendation interface displaying growth-stage advisory, recommended treatments, and preventive measures.

The developed interface enables users to obtain disease predictions together with explainable AI visualizations and detailed agronomic recommendations through a simple and interactive workflow. By integrating diagnosis, severity estimation, and treatment planning within a single application, the proposed framework provides a practical decision-support tool for precision agriculture and sustainable crop management.

G. Limitations

The proposed framework is primarily trained and evaluated on the PlantVillage dataset, which contains images captured under controlled laboratory conditions. Consequently, performance decreases when evaluated on real-world field images exhibiting complex backgrounds and varying illumination. Furthermore, the HSV-based severity estimation relies on handcrafted thresholds, and the recommendation module follows a rule-based knowledge base rather than adaptive learning. Future work will address these limitations through larger field datasets, domain adaptation techniques, and more advanced disease severity estimation methods.

VI. Conclusion and Future Work

This paper presented an AI-Based Computer Vision Framework for Real-Time Crop Disease Diagnosis and Precision Treatment Recommendation that integrates deep learning, computer vision, explainable artificial intelligence, and agronomic knowledge into a unified decision-support system. The proposed framework employs MobileNetV2 transfer learning for multi-class crop disease classification across fifteen disease categories covering tomato, potato, and bell pepper crops. In addition to disease identification, the framework incorporates HSV-based disease severity estimation, Grad-CAM visualization for model interpretability, and a rule-based precision treatment recommendation engine capable of generating context-aware management strategies based on disease severity, crop growth stage, treatment preference, and affected cultivation area.

Experimental evaluation on the PlantVillage validation dataset demonstrated that the proposed framework achieved an overall classification accuracy of approximately 92%, with corresponding weighted precision, recall, and F1-score of

approximately 93%, 92%, and 92%, respectively. Furthermore, TensorFlow Lite optimization reduced the model size by approximately 4.76× while providing a 3.14× improvement in inference speed, demonstrating the suitability of the proposed framework for future edge-based deployment.

Evaluation on an independent real-world image dataset revealed a noticeable reduction in classification performance due to differences between controlled laboratory images and practical field conditions. This observation highlights the importance of improving model generalization through the inclusion of diverse field datasets, advanced data augmentation strategies, and domain adaptation techniques.

The developed Streamlit-based graphical user interface demonstrates the practical applicability of the proposed framework by providing real-time disease diagnosis, disease severity estimation, Grad-CAM visualizations, and precision treatment recommendations through an intuitive and interactive platform. The integration of explainable artificial intelligence with agronomic decision support extends the framework beyond conventional crop disease classification systems and contributes toward intelligent precision agriculture.

Future work will focus on expanding the framework to support additional crop species, collecting large-scale field datasets under diverse environmental conditions, integrating weather and Internet of Things (IoT) sensor data for context-aware disease prediction, and developing a dedicated Android application for offline mobile deployment. In addition, advanced semantic segmentation networks and domain adaptation techniques will be investigated to improve disease localization and real-world classification performance under challenging agricultural environments.

REFERENCES

[1] L. Hamad, Q. Ananzeh, M. Salameh, A. Salem, M. Alshboul, and M. Al-Ayyoub, "Hierarchical Multi-Stage Framework for Robust and Explainable Tomato Leaf Disease Identification," *IEEE Access*, vol. 13, pp. 217913–217928, 2025.

[2] A. Oad, S. S. Abbas, A. Zafar, B. A. Akram, F. Dong, M. S. H. Talpur, and M. Uddin, "Plant Leaf Disease Detection Using Ensemble Learning and Explainable AI," *IEEE Access*, vol. 12, pp. 156038–156049, 2024.

[3] W. Altabaji *et al.*, "Comparative Analysis of Transfer Learning, LeafNet, and Modified LeafNet Models for Accurate Rice Leaf Diseases Classification," *IEEE Access*, vol. 12, pp. 36622–36635, 2024.

[4] R. Maurya, S. Mahapatra, and L. Rajput, "A Lightweight Meta-Ensemble Approach for Plant Disease Detection Suitable for IoT-Based Environments," *IEEE Access*, vol. 12, pp. 28096–28108, 2024.

[5] J. Lu, X. Liu, and X. Ma, "Improved MobileNetV2 Crop Disease Identification Model for Intelligent Agriculture," *PeerJ Computer Science*, vol. 9, 2023.

[6] F. Too, L. Yujian, S. Njuki, and Y. Yingchun, "A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.

[7] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4510–4520.

[8] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.

[9] S. P. Mohanty, D. P. Hughes, and M. Salathe´, "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.

[10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[11] D. P. Hughes and M. Salathe´, "An Open Access Repository of Images on Plant Health to Enable the Development of Mobile Disease Diagnostics," arXiv:1511.08060, 2015.

[12] C. Szegedy *et al.*, "Going Deeper with Convolutions," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1–9.

[13] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," arXiv:1409.1556, 2014.

[14] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.

[15] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[16] Abade, A., Ferreira, P. A., and Vidal, F. de B., "Plant diseases recognition on images using convolutional neural networks: A systematic review," Computers and Electronics in Agriculture, vol. 185, Art. no. 106125, 2021.

[17] Ferentinos, K. P., "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, 2018.

[18] Fuentes, A., Yoon, S., Kim, S. C., and Park, D. S., "A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition," Sensors, vol. 17, no. 9, Art. no. 2022, 2017.

[19] Barbedo, J. G. A., "Factors influencing the use of deep learning for plant disease recognition," Biosystems Engineering, vol. 172, pp. 84–91, 2018.

[20] Barbedo, J. G. A., "Plant disease identification from individual lesions and spots using deep learning," Biosystems Engineering, vol. 180, pp. 96–107, 2019.

[21] Singh, D., Jain, N., Jain, P., Kayal, P., Kumawat, S., and Batra, N., "PlantDoc: A dataset for visual plant disease detection," in Proceedings of the 7th ACM IKDD CoDS and 25th COMAD, pp. 249–253, 2020.

[22] Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., and Hughes, D. P., "Deep learning for image-based cassava disease detection," Frontiers in Plant Science, vol. 8, Art. no. 1852, 2017.

[23] Fuentes, A., Yoon, S., Lee, J., and Park, D. S., "High-performance deep neural network-based tomato plant diseases and pests diagnosis system with refinement filter bank," Frontiers in Plant Science, vol. 8, Art. no. 1162, 2017.

[24] Barbedo, J. G. A., "A review on the main challenges in automatic plant disease identification based on visible range images," Biosystems Engineering, vol. 144, pp. 52–60, 2016.

[25] Saleem, M. H., Potgieter, J., and Arif, K. M., "Plant disease detection and classification by deep learning," Plants, vol. 8, no. 11, Art. no. 468, 2019.

[26] Too, E. C., Yujian, L., Njuki, S., and 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.

[27] Wang, G., Sun, Y., and Wang, J., "Automatic image-based plant disease severity estimation using deep learning," Computers and Electronics in Agriculture, 2017.

[28] Natarajan, S., Chakrabarti, P., and Margala, M., "Robust diagnosis and meta visualizations of plant diseases through deep neural architecture with explainable AI," Scientific Reports, vol. 14, Art. no. 13695, 2024.

[29] "Improving Plant Disease Classification With Deep-Learning-Based Prediction Model Using Explainable Artificial Intelligence," IEEE Access, vol. 12, pp. 100005–100014, 2024.

[30] "Recent advances in plant disease severity assessment using convolutional neural networks," 2023.