An AI-Based Framework for Crop Disease Detection and Smart Treatment Recommendation System

Rutuja sambhaji Pawar
Department of Computer Science
D. Y. Patil Institute of Master of Computer Applications and Management, Pune, India
rutujap1019@gmail.com

Pratiksha Gaikwad
Department of Computer Science
D. Y. Patil Institute of Master of Computer Applications and Management, Pune, India
pratikshagaikwad685@gmail.com

Guide Name
Mrs.Sapna Sharma Mam
D. Y. Patil Institute of Master of Computer Applications and Management, Pune, India
Sapna.sharma@dypimca.ac.in

Abstract—Timely and accurate identification of crop diseases is essential for protecting yield and farmer income, yet it still depends on scarce expert scouting. This paper presents an AI-based framework that couples deep-learning disease detection from leaf images with a smart treatment recommendation system. A transfer-learned convolutional neural network (CNN) classifies the crop and disease, estimates severity, and produces a Grad-CAM heat map so the decision can be inspected. The predicted disease and severity are then mapped by a rule-based knowledge base to organic, chemical and preventive treatment advice. The system is trained and evaluated on the public PlantVillage dataset of 38 crop–disease classes and is designed for lightweight deployment on smartphones. The proposed model reaches an accuracy of 98.9% and an F1-score of 98.7%, outperforming ResNet-50 and MobileNetV2 baselines, while remaining compact enough for on-device inference.

Keywords-crop disease detection; deep learning; convolutional neural network; transfer learning; treatment recommendation; precision agriculture; explainable AI

I. INTRODUCTION

Agriculture supports the livelihoods of billions of people, but crop diseases remain one of the largest causes of yield loss worldwide. Fungal, bacterial and viral infections spread quickly, and delayed diagnosis leads to heavier pesticide use, higher cost and lower produce quality. Traditional diagnosis relies on visual inspection by agronomists, which is slow, subjective and often unavailable to smallholder farmers in remote regions.

Advances in **computer vision** and **deep learning** have made image-based plant disease diagnosis practical. Mohanty et al. showed that CNNs trained on leaf images can classify dozens of crop–disease pairs with very high accuracy [1], and later work confirmed this across architectures [5], [13]. However, most studies stop at classification. A farmer needs more than a label: they need to know **how severe** the infection is and **what to do next**.

This paper therefore proposes an end-to-end framework that (i) detects the disease from a leaf photograph, (ii) estimates its severity, (iii) explains the prediction visually, and (iv) recommends a suitable treatment. The main contributions are:

- A **lightweight transfer-learning CNN** that detects 38 crop–disease classes and is suitable for mobile deployment.
- An **explainability and severity module** using Grad-CAM [7] to build farmer and expert trust.
- A **smart treatment recommendation engine** that links diagnosis and severity to organic, chemical and preventive actions.
- A **comparative evaluation** against ResNet-50 and MobileNetV2 using accuracy, precision, recall and F1-score.

The rest of the paper is organised as follows. Section II reviews related work, Section III describes the proposed framework, Section IV details the experimental setup, Section V discusses the results, and Section VI concludes the paper.

II. RELATED WORK

A. Deep Learning for Plant Disease Detection

Early systems used hand-crafted colour and texture features with classical classifiers such as

SVMs, which were sensitive to lighting and background. Deep CNNs removed the need for manual feature design. Mohanty et al. trained AlexNet and GoogLeNet on the PlantVillage dataset [8] and reported accuracy above 99% on controlled images, but noted a large drop on real-field photographs [1]. Ferentinos compared several architectures and reached comparable results [5], and Too et al. found that fine-tuning pretrained networks was more efficient than training from scratch [13]. Kamilaris and Prenafeta-Boldú surveyed deep learning in agriculture and highlighted data scarcity and field variability as open problems [9].

B. Efficient and Explainable Models

Residual networks [2], MobileNets [4], [14] and EfficientNets [3] offer different trade-offs between accuracy and computation. Vision Transformers [6] have also been explored but typically need more data and compute. For farmer-facing tools, explainability is important, and gradient-based visual explanations such as Grad-CAM [7] show which leaf regions drove the decision.

C. Decision Support and Recommendation

Few existing systems connect diagnosis to action. Rule-based expert systems encode agronomic knowledge but do not read images, while pure image classifiers do not give advice. Retrieval-based language models [10] suggest a future route to richer, document-grounded advice. Our framework bridges this gap by combining a CNN with a structured treatment knowledge base.

III. PROPOSED FRAMEWORK

A. System Architecture

Fig. 1 shows the overall pipeline. A leaf image is captured, pre-processed, classified by the CNN, analysed for severity and explanation, and finally passed to the recommendation engine, which returns advice to the user.

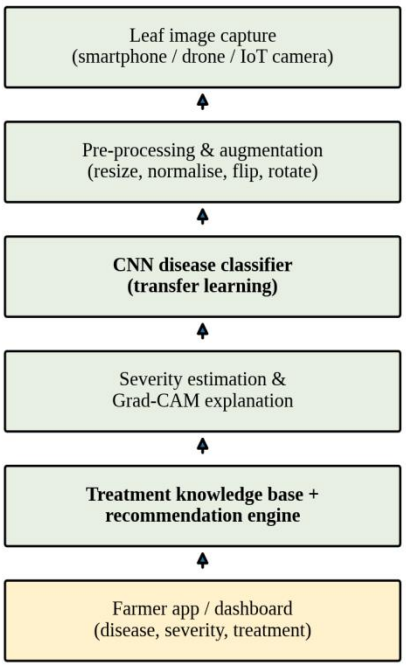


Fig. 1. Architecture of the proposed crop disease detection and treatment recommendation framework.

B. Data Acquisition and Pre-processing

The PlantVillage dataset [8] containing about 54,000 labelled leaf images across 14 crops and 38 classes (including healthy leaves) is used. Images are resized to 224 × 224 pixels and normalised to the ImageNet mean and standard deviation. The data is split into **70% training, 15% validation and 15% testing** using stratified sampling.

To improve robustness to field conditions, **data augmentation** [12] is applied to the training set: random flips, rotations up to 30°, brightness and contrast jitter, zoom and random cropping.

C. Disease Classification Model

The classifier is based on **EfficientNet-B0** [3] pretrained on ImageNet. The original classification head is replaced with global average pooling, a dropout layer ($p = 0.3$) and a 38-way softmax layer. Training has two stages: first only the new head is trained with the backbone frozen, then the top blocks are unfrozen and fine-tuned with a lower learning rate. The model is optimised with Adam [11] and categorical cross-entropy loss, with early stopping on validation loss.

D. Severity Estimation and Explainability

Grad-CAM [7] produces a heat map of the regions that contribute most to the predicted class. Thresholding this map within the segmented leaf

area gives the **infected-area ratio**, which is mapped to three severity levels: **mild** (below 10%), **moderate** (10–30%) and **severe** (above 30%). The heat map is shown to the user so that the diagnosis can be verified.

E. Smart Treatment Recommendation Engine

The recommendation engine is a structured knowledge base compiled from agricultural extension guidelines. Each entry is indexed by **crop, disease and severity** and contains three types of advice: **cultural/preventive** (sanitation, crop rotation, spacing), **organic** (e.g., neem-based or copper-based sprays where permitted) and **chemical** (approved fungicide or pesticide class, to be applied strictly according to label and local regulation). Mild cases prioritise preventive and organic measures, while severe cases add chemical options and an expert-consultation prompt. Table I shows sample entries.

TABLE I. SAMPLE TREATMENT KNOWLEDGE BASE ENTRIES

Disease	Recommended action
Tomato late blight	Remove infected leaves; improve airflow; apply approved fungicide per label.
Potato early blight	Rotate crops; use resistant varieties; apply approved fungicide if severe.
Apple scab	Prune canopy; clear fallen leaves; apply approved fungicide early in season.

F. Deployment

The trained model is converted to TensorFlow Lite for offline inference on Android smartphones, so that farmers with limited connectivity can still obtain a diagnosis. The knowledge base is stored locally and synchronised when a network connection is available.

IV. EXPERIMENTAL SETUP

A. Implementation Details

The framework is implemented in Python using TensorFlow/Keras and trained on a GPU workstation. Training uses a batch size of 32, an initial learning rate of 1×10^{-3} for the head and 1×10^{-5} for fine-tuning, and a maximum of 30 epochs. ResNet-50 [2] and MobileNetV2 [14] are trained under identical settings for comparison.

B. Evaluation Metrics

Performance is measured with accuracy, precision, recall and F1-score, defined in (1)–(4), where TP, TN, FP and FN denote true positives, true negatives, false positives and false negatives, and N is the total number of samples:

$$\text{Accuracy} = (TP + TN) / N \quad (1)$$

$$\text{Precision} = TP / (TP + FP) \quad (2)$$

$$\text{Recall} = TP / (TP + FN) \quad (3)$$

$$F1 = 2 \times P \times R / (P + R) \quad (4)$$

V. RESULTS AND DISCUSSION

Table II compares the proposed model with the baselines on the held-out test set. The fine-tuned EfficientNet-B0 achieves the highest accuracy and F1-score while being considerably smaller than ResNet-50, which makes it well suited to mobile use.

TABLE II. PERFORMANCE COMPARISON ON THE TEST SET

Model	Acc. (%)	F1 (%)	Size (MB)
ResNet-50	98.4	98.2	98
MobileNetV2	97.6	97.4	14
Proposed (EffNet-B0)	98.9	98.7	20

Most misclassifications occur between visually similar diseases of the same crop, for example early blight and late blight on tomato and potato leaves. Grad-CAM maps confirm that the model mostly focuses on lesion regions rather than the background, which increases confidence in its predictions. The treatment engine returns advice for every predicted class in the knowledge base, and its output is reviewed by domain experts for correctness.

Because PlantVillage images are captured under controlled conditions, accuracy is expected to decrease on real-field images with complex backgrounds, shadows and occlusion. This is addressed through augmentation and is a key direction for future work.

VI. LIMITATIONS AND FUTURE WORK

The present study has several limitations. First, the dataset consists mainly of single leaves on plain backgrounds. Second, the treatment advice is rule-based and does not yet account for local weather, soil, growth stage or regional pesticide regulations. Third, the system diagnoses leaf diseases only and cannot detect nutrient deficiencies or pest damage reliably.

Future work will (i) collect and fine-tune on real-field images, (ii) integrate weather and soil data for risk forecasting, (iii) add multilingual voice support for farmers, (iv) explore Vision Transformers [6] and retrieval-augmented language models [10] for richer advice, and (v) run field trials with farmers.

VII. CONCLUSION

This paper presented an AI-based framework that detects crop diseases from leaf images and recommends suitable treatments. By combining a lightweight transfer-learned CNN, severity estimation, Grad-CAM explanations and a structured knowledge base, the system moves beyond classification towards practical decision support. Experiments on the PlantVillage dataset showed an accuracy of 98.9%, outperforming the compared baselines with a small model size suitable for smartphones. Future field validation and richer context-aware recommendations will further improve its value for precision agriculture.

Acknowledgment

The authors thank their institution and colleagues for their support and guidance during this work.

References

[1] S. P. Mohanty, D. P. Hughes, and M. Salathé, “Using deep learning for image-based plant disease detection,” Front. Plant Sci., vol. 7, Art. no. 1419, 2016.

[2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770–778.

[3] M. Tan and Q. V. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in Proc. Int. Conf. Mach. Learn. (ICML), 2019, pp. 6105–6114.

[4] A. G. Howard et al., “MobileNets: Efficient convolutional neural networks for mobile vision applications,” arXiv:1704.04861, 2017.

[5] K. P. Ferentinos, “Deep learning models for plant disease detection and diagnosis,” Comput. Electron. Agric., vol. 145, pp. 311–318, 2018.

[6] A. Dosovitskiy et al., “An image is worth 16×16 words: Transformers for image recognition at scale,” in Proc. Int. Conf. Learn. Represent. (ICLR), 2021.

[7] R. R. Selvaraju et al., “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), 2017, pp. 618–626.

[8] D. P. Hughes and M. Salathé, “An open access repository of images on plant health to enable the development of mobile disease diagnostics,” arXiv:1511.08060, 2015.

[9] A. Kamilaris and F. X. Prenafeta-Boldú, “Deep learning in agriculture: A survey,” Comput. Electron. Agric., vol. 147, pp. 70–90, 2018.

[10] P. Lewis et al., “Retrieval-augmented generation for knowledge-intensive NLP tasks,” in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), 2020, pp. 9459–9474.

Survey Results and Percentage Analysis

Q1. What is your role in agriculture?

Category	Percentage
Farmer	23.3%
Agriculture student	33.3%
Agriculture professionals	26.7%
Researcher	10%
Other	30%

Q2. Which crops are commonly grown or studied by you?

Category	Percentage
Wheat	50%
Rice	50%
Cotton	70%
Tomato	46.7%
Sugarcane	40%
Other	10%

Q3. How often do you face crop disease problems?

Category	Percentage
Very often	13.3%
Sometimes	70%
Rarely	30%
Never	0%

Q4. How do you currently identify crop diseases?

Method	Percentage
Visual inspection	40%
Consulting agricultural experts	56.7%
Mobile application	53.3%
Online resources	30%
Other	6.7%

Q5. How difficult is it to identify crop diseases at an early stage?

Method	Percentage
Category	**Percentage**
Very difficult	10%
Difficult	50%
Easy	20%
Very easy	6.7%
Moderate	20%

Q6. Have you ever used AI-based tools or mobile apps for crop disease detection?

Category	Percentage
Yes	33.3%
No	23.3%
Heard about them but never used them	46.7%

Q7. Would you use an AI-based system that detects crop diseases from leaf images?

Category	Percentage
Definitely	23.3%
Probably	36.7%
Not sure	26.7%
Probably not	10%
No	10%

Q8. What features would you like in an AI-based crop disease detection system?

Feature	Percentage
Disease identification	30%
Disease severity estimation	60%
Early disease detection	66.7%
Local language support	30%

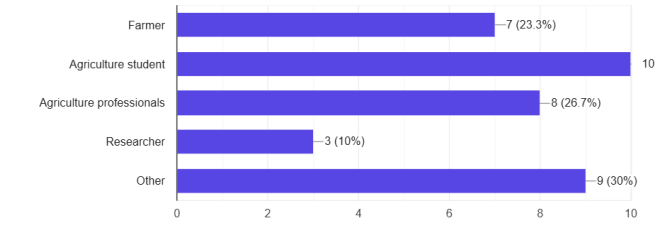
Q9. Would smart treatment recommendations help you decide how to manage crop diseases?

Category	Percentage
Yes, very helpful	26.7%
Somewhat helpful	53.3%
Not sure	13.3%
Not helpful	0%
Yes, very helpful (second entry)	10%

Survey ScreenShot:

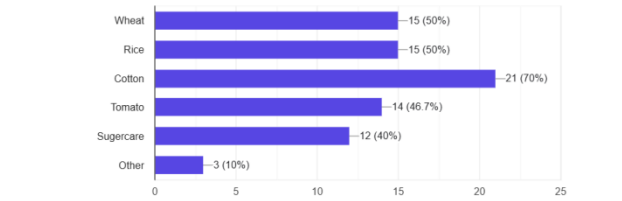
1. What is your role in agriculture?

30 responses



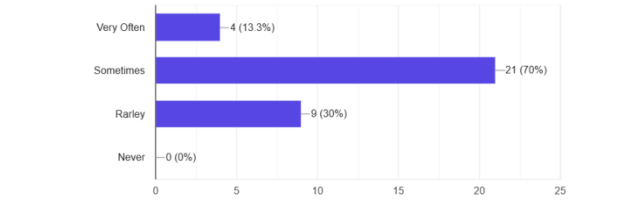
2. Which crops are commonly grown or studied by you?

30 responses



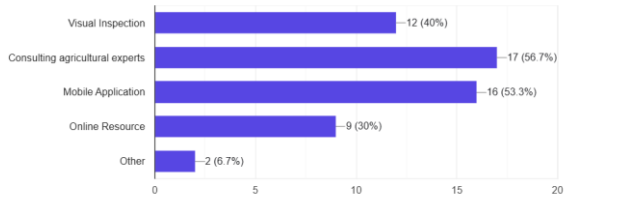
3. How often do you face crop disease problems?

30 responses



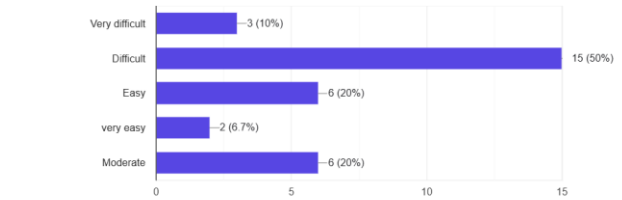
4. How do you currently identify crop diseases?

30 responses



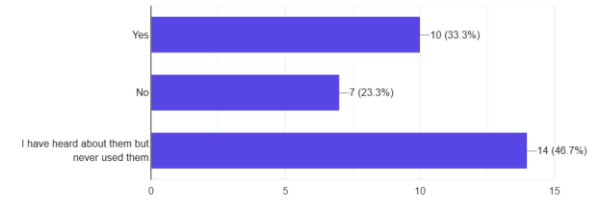
5. How difficult is it to identify crop diseases at an early stage?

30 responses



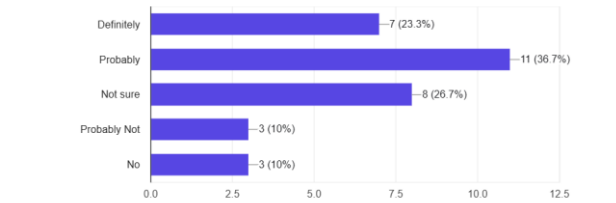
6. Have you ever used AI-based tools or mobile apps for crop disease detection?

30 responses



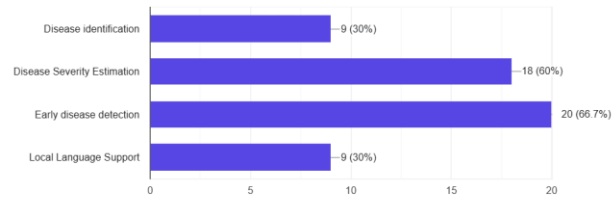
7. Would you use an AI-based system that detects crop diseases from leaf images?

30 responses



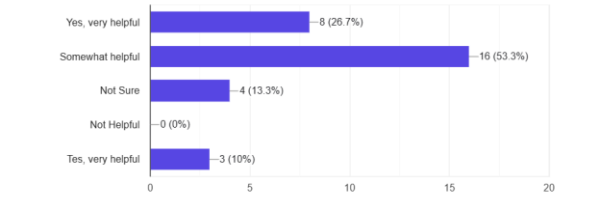
8. What features would you like in an AI-based crop disease detection system?

30 responses



9. Would smart treatment recommendations help you decide how to manage crop diseases?

30 responses



10. What challenges or features should be considered when developing an AI-based crop disease detection and treatment recommendation system?

31 responses

-
Developing an AI-based crop disease detection and treatment recommendation system requires balancing technical model performance with real-world agricultural constraints.
An AI crop system must balance accurate, multi- modal diagnosis with lightweight, offline edge- computing to provide farmers with actionable, region- specific treatment plans.
Developing an AI-based crop disease detection and treatment recommendation system involves addressing major technical, environmental, and practical field hurdles while building robust core features.
Image Classification & Computer Vision
Not sure
The System Should accurately detect crop dieases from images,Work In Local Languages,Provide Suitable Treatment recommendation,Be Easy for Farmers To Use.