



AI-Enabled Crop Disease Prediction and Management: A Deep Learning Framework for Sustainable Agriculture

Minal¹, Harshita², Ms. Mitu Sehgal³

Department of CSE-AI&DS,

Panipat Institute of Engineering & Technology, Samalkha, India^{1,2,3}

Mail ID: mitu.cse.ai-ds@piet.co.in

Abstract--Crop diseases can substantially reduce agricultural productivity and create economic losses, particularly when diagnosis is delayed or dependent on manual inspection. This paper presents an AI-enabled crop disease prediction and management framework based on Convolutional Neural Networks (CNNs) for image-based disease classification. The reported system uses a Kaggle-sourced dataset of 70,295 JPG images representing 38 crop-disease/healthy classes. The processing pipeline includes data cleaning, resizing, normalization, augmentation, feature learning, model training, validation, and application deployment. Transfer-learning architectures such as VGG16 and ResNet are considered to improve generalization when labeled data are limited. The framework is extended with a weather-aware risk module using temperature and humidity data, a bilingual interface, disease-specific recommendations, pesticide information, and downloadable reports. Experimental results reported in the project show 91% overall accuracy, with macro and weighted precision, recall, and F1-score also reported as 0.91. The system therefore combines visual diagnosis with contextual decision support to provide a practical precision-agriculture tool. The paper also identifies field-data diversity, model generalization, and multimodal sensing as important directions for future work.

Keywords: Crop disease prediction, convolutional neural network, deep learning, transfer learning, VGG16, ResNet, precision agriculture, weather-based risk assessment, sustainable agriculture.

INTRODUCTION

Agriculture is central to food security and economic activity, yet plant diseases remain a major constraint on crop productivity. Conventional diagnosis commonly depends on visual inspection or laboratory testing, which can be slow, costly, and difficult to access in remote areas. The project report identifies these limitations and motivates automated image-based diagnosis using artificial intelligence and deep learning [9]. CNNs are particularly appropriate because they learn hierarchical visual features such as color patterns, lesions, textures, and irregularities directly from plant images.

The proposed framework is designed not only to identify disease classes but also to support timely management. A farmer can upload a crop-leaf image and obtain a predicted disease, affected crop information, root-cause explanation, protective measures, treatment suggestions, and a weather-based risk indication [10]. The report further describes bilingual interaction, online pesticide information, and downloadable PDF reports as practical features. These design goals address accessibility as well as prediction [11].

The main contributions of this work are: (1) a CNN-based crop disease classification pipeline; (2) preprocessing and augmentation for improved robustness; (3) integration of real-time temperature and



humidity for disease-risk assessment; (4) an application layer for actionable recommendations and bilingual access; and (5) an evaluation using accuracy, precision, recall, F1-score, and a confusion matrix. The overall framework is intended to reduce delayed diagnosis and support more targeted crop management.

LITERATURE REVIEW

Research in automated plant-disease recognition has progressed from manual inspection and conventional image processing toward machine-learning and deep-learning approaches. Earlier classifiers such as decision trees, SVM, and k-nearest neighbors depended more heavily on engineered features. CNNs reduced this dependency by learning discriminative features directly from image data, enabling scalable classification of leaf symptoms. The reviewed literature also emphasizes benchmark datasets, environmental variables, IoT sensing, remote sensing, and precision-agriculture integration.

The project report summarizes studies covering CNN-based detection, transfer learning, classical machine learning, IoT-supported agriculture, disease and pest forecasting, and hybrid crop recommendation/disease detection. Table I consolidates the key outcomes reported in the source document. These studies indicate that image-based CNNs are effective for visual classification, while environmental and structured-data models can complement image analysis.

TABLE I. SUMMARY OF REPORTED STUDIES [1]–[8]

Ref.	Year	Approach	Reported outcome
[1]	2019	CNN-based crop disease detection	89% accuracy; rapid disease identification
[2]	2020	VGG16 transfer learning for paddy disease	90% accuracy; web/mobile deployment potential
[3]	2022	ML for crop disease through leaf images	SVM strong for classification; CNN/RF effective for image processing
[4]	2020	Adversarial image translation	Improved plant-disease recognition under image variation
[5]	2022	ML survey for crop diseases and pests	RNNs for meteorological forecasting; CNNs for images
[6]	2023	ML-based crop prediction	Up to 97.32% with Random Forest in reported study
[7]	2024	LightGBM + ResNet-50 framework	Crop recommendation plus disease detection
[8]	2024	PyTorch CNN plant-disease prediction	98.7% reported accuracy and robustness

Despite strong reported results, the literature and project report identify a recurring generalization problem. Many datasets contain isolated leaves captured under controlled backgrounds, whereas field images contain illumination changes, clutter, overlapping leaves, varying angles, and symptoms on stems



or fruits. Imbalanced disease classes and insufficient field-labeled data can further reduce reliability. These limitations motivate the proposed combination of augmentation, transfer learning, environmental context, and an accessible deployment layer.

PROPOSED FRAMEWORK

The proposed framework follows an end-to-end pipeline from image acquisition to prediction and management guidance. Fig. 1 summarizes the system architecture. Images may be captured using smartphones, cameras, or drones, while environmental information such as temperature and humidity can be obtained from external sources. After preprocessing, the image is passed to a CNN/transfer-learning model for disease classification. The application layer then combines the prediction with disease knowledge and weather information to produce risk and treatment guidance.




Fig. 1. System architecture for crop disease prediction (adapted from the project design).

The architecture is intended to support real-time use through a web/mobile interface. The source report describes a Streamlit-based interface, multilingual support, an OpenWeather API-based weather module, recommendation content, and PDF report generation. This makes the system a multi-component decision-support platform rather than an isolated classifier.

METHODOLOGY

A. Dataset and Data Preparation

The reported dataset is sourced from Kaggle and contains 70,295 JPG images. The class inventory contains 38 crop-disease or healthy categories spanning crops such as apple, blueberry, cherry, corn, grape, orange, peach, pepper, potato, raspberry, soybean, squash, strawberry, and tomato. The complete class counts in the project sum to 70,295 images. Table II gives a compact description of the dataset used in the implementation.

TABLE II. DATASET CHARACTERISTICS REPORTED IN THE PROJECT [18]

Property	Reported value
Source	Kaggle
Total images	70,295
Image format	.JPG
Classes	38
Data types	Healthy and diseased crop-leaf images
Preparation	Cleaning, resizing, normalization, augmentation
Splits	Training, validation, testing

Data cleaning is applied to remove duplicates, inconsistencies, and missing or unsuitable samples. Image resizing establishes a common input size, while normalization provides consistent numerical ranges. Augmentation operations reported in the project include rotation, horizontal/vertical flipping, brightness or color adjustment, and zooming [12]. These transformations increase sample diversity and help reduce



overfitting. The dataset is divided into training, validation, and testing subsets so that learning and generalization can be assessed separately.

B. CNN-Based Feature Learning

CNNs are selected as the principal image-classification mechanism because convolutional filters learn local visual structures and progressively combine them into higher-level representations. The project report describes a CNN containing convolutional, max-pooling, batch-normalization, dropout, flattening, and dense layers. The resulting representation is mapped to the disease classes through the final classification layer [13].

When data are limited or the visual variability is high, transfer learning is considered. The project evaluates architectures including VGG16, VGG19, and ResNet50. Pretrained networks provide reusable low- and mid-level visual features and can be fine-tuned for agricultural classes. This reduces the need to train a deep network entirely from scratch and is consistent with the transfer-learning direction identified in prior work [2], [5].

C. Preprocessing and Training Workflow




Fig. 2. Data flow and processing pipeline from acquisition to prediction.

As shown in Fig. 2, the workflow begins with data acquisition and preprocessing, followed by segmentation/feature learning and model training. After evaluation, the trained model is saved and loaded by the application. The interface accepts an uploaded image, applies the same preprocessing used during training, performs classification, and passes the predicted class to the recommendation and weather-risk modules.

The evaluation stage uses validation data to monitor validation loss and accuracy and to detect possible overfitting. The project report also describes Python, NumPy, Pandas, OpenCV, scikit-learn, TensorFlow/Keras, PyTorch, Jupyter, MLflow, Pillow/Albumentations, TensorFlow Lite, and TensorBoard as supporting tools [14].

SYSTEM IMPLEMENTATION

A. Weather-Aware Risk Assessment

A distinguishing component of the framework is the integration of environmental context. The system obtains temperature and humidity information and evaluates whether current conditions are favorable for disease propagation. The resulting risk indicator complements the visual prediction: the CNN answers what disease pattern is present in the uploaded image, while the weather module provides contextual information about possible outbreak conditions. The source implementation uses the OpenWeather API for weather data.



This design is important because disease expression and transmission are influenced by environmental conditions. The literature review in the project identifies weather, precipitation, IoT sensors, and structured agricultural data as useful complementary inputs. Therefore, multimodal integration is a logical extension of image-only classification [5], [6].

B. Recommendation and User Interface

After classification, the application provides disease-specific information including the affected crop, likely cause, preventive measures, and treatment guidance. The project also describes pesticide suggestions with online purchase links and downloadable PDF reports. A bilingual English/Hindi interface is included to improve accessibility for users with different language preferences.



Fig. 3. Example prediction screen showing disease diagnosis, weather risk, and recommendations.

Fig. 3 illustrates the integrated output rather than only a class label. For the example shown in the project, the application predicts Cedar Apple Rust for an apple image and displays a low outbreak-risk indication for Delhi, together with cause, affected crop, protective advice, pesticide information, and report generation. Such contextual output can support faster action, although recommendations should be treated as decision support and validated with local agricultural expertise.

C. Model Evaluation

Performance is evaluated using accuracy, precision, recall, F1-score, and a confusion matrix [15]. Accuracy measures the proportion of correctly classified samples. Precision measures the correctness of positive predictions, recall measures the proportion of actual positives detected, and F1-score balances precision and recall. The confusion matrix provides class-wise insight into correct predictions and misclassifications.




Fig. 4. Training/validation accuracy and loss curves reported in the project.

According to Fig. 4, training accuracy increases across epochs while validation accuracy improves and then stabilizes at approximately 90%. Training and validation losses decrease sharply during the initial learning period. The reported behavior indicates that the model learned useful discriminative features without a pronounced divergence between training and validation performance.



RESULTS AND DISCUSSION

A. Overall Classification Performance

The project reports an overall classification accuracy of 0.91 (91%). Macro-average precision, recall, and F1-score are each reported as 0.91, and the weighted averages are also reported as 0.91. These results as shown in Table III indicate balanced aggregate performance across the evaluated classes, although class-level variation remains visible in the detailed report.

TABLE III. REPORTED MODEL PERFORMANCE

Metric	Reported value	Interpretation
Accuracy	0.91	91% of evaluated samples classified correctly
Macro precision	0.91	Equal-weight average precision across classes
Macro recall	0.91	Equal-weight average recall across classes
Macro F1	0.91	Balanced class-level precision/recall measure
Weighted precision	0.91	Support-weighted precision
Weighted recall	0.91	Support-weighted recall
Weighted F1	0.91	Support-weighted F1




Fig. 5. Precision, recall, and F1-score across crop-disease classes.

Fig. 5 shows generally high precision, recall, and F1 values across the evaluated categories, while some classes are more difficult. For example, the reported precision/recall/F1 values are 0.93/0.86/0.89 for Apple_Apple_scab, 0.98/0.95/0.96 for Blueberry_healthy, and 0.75/0.68/0.71 for Pepper_bell_Bacterial_spot. These variations demonstrate why class-wise metrics are important in addition to overall accuracy.

B. Confusion-Matrix Analysis




Fig. 6. Confusion matrix for the crop-disease classification task.



The confusion matrix in Fig. 6 shows a strong diagonal pattern, indicating that most samples are assigned to their correct classes. Off-diagonal entries represent class confusion. Such errors may arise when visually similar diseases share symptoms or when image conditions differ from those represented during training. The project therefore identifies field-realistic data collection and domain variation as important improvements for future versions.

C. Practical Application

The reported application accepts JPG, JPEG, and PNG images and presents the uploaded image before classification. The output combines the disease name with risk assessment and recommendations. The interface therefore converts model output into an actionable workflow. The project also provides a Hindi interface, which can improve usability for users who prefer regional-language interaction.

From a sustainability perspective, targeted disease identification can reduce indiscriminate pesticide application by associating treatment guidance with a predicted disease. The report positions this capability as a means of supporting precision agriculture, reducing unnecessary chemical use, and improving resource efficiency. However, the paper does not claim a measured reduction in pesticide use or crop loss; these benefits remain expected applications requiring field validation.

LIMITATIONS AND FUTURE SCOPE

The principal limitation is the gap between controlled image datasets and real farm environments. The project report notes that lighting, background clutter, leaf overlap, camera angle, crop growth stage, and symptoms occurring on stems or fruits can reduce generalization. Dataset bias, class imbalance, and limited field-representative images are additional concerns. Consequently, the reported 91% accuracy should not be interpreted as universal field accuracy without independent validation.

Future development can proceed in several directions. First, IoT sensors can provide continuous temperature, humidity, soil-moisture, and other environmental measurements. Second, satellite and drone imagery can extend disease monitoring from individual leaves to field-scale surveillance. Third, advanced transfer learning, ensemble methods, and generative augmentation can improve robustness to diverse crops and environments. Fourth, multimodal learning can combine image features with weather, soil, crop-history, and geographic information. Finally, lightweight models and TensorFlow Lite-style deployment can support resource-constrained mobile and edge devices.

The project also identifies automated pest detection, disease-resistant crop selection, precision pesticide application, cross-domain agricultural monitoring, and global deployment as longer-term opportunities. These directions should be evaluated using geographically diverse field datasets, independent test sets, uncertainty estimates, and expert-validated treatment recommendations.

CONCLUSION

This paper presented an AI-enabled crop disease prediction and management framework centered on CNN-based image classification. The reported implementation combines a 70,295-image, 38-class dataset with preprocessing, augmentation, deep feature learning, transfer-learning options, weather-based risk assessment, recommendation support, bilingual interaction, and downloadable reports. The reported experimental results show 91% overall accuracy and 0.91 macro and weighted averages for precision, recall, and F1-score. These findings demonstrate the feasibility of integrating image intelligence with environmental context in a practical agricultural decision-support system.



The main value of the framework is its end-to-end orientation: it moves from image acquisition and disease recognition to contextual risk assessment and actionable information. Nevertheless, real-world agricultural deployment requires stronger field validation, broader datasets, and careful handling of uncertainty. Future work should therefore focus on multimodal sensing, field-scale monitoring, model compression, explainability, and expert-validated recommendations.

REFERENCES

- [1] N. Patil, V. Wankhedkar, R. Ali, and D. Nayak, "Crop Disease Detection using Deep Convolutional Neural Networks," Dept. of Information Technology, Vidyalankar Institute of Technology, Mumbai, India.
- [2] S. S. Rautaray, M. Pandey, M. K. Gourisaria, R. Sharma, and S. Das, "Paddy Crop Disease Prediction – A Transfer Learning Technique," *International Journal of Recent Technology and Engineering (IJRTE)*, vol. 8, no. 6, Mar. 2020.
- [3] S. Nandhini and K. Ashokkumar, "Machine Learning Technique for Crop Disease Prediction Through Crop Leaf Image," *Applied Mathematics & Information Sciences*, vol. 16, no. 2, pp. 149–158, Mar. 2022, doi: 10.18576/amis/160202.
- [4] H. Nazki, S. Yoon, A. Fuentes, and D. S. Park, "Unsupervised image translation using adversarial networks for improved plant disease recognition," Elsevier B.V., 2020.
- [5] T. Domingues, T. Brandão, and J. C. Ferreira, "Machine Learning for Detection and Prediction of Crop Diseases and Pests: A Comprehensive Survey," *Agriculture*, vol. 12, no. 9, p. 1350, Sep. 2022, doi: 10.3390/agriculture12091350.
- [6] E. Elbasi et al., "Crop Prediction Model Using Machine Learning Algorithms," *Applied Sciences*, vol. 13, no. 16, p. 9288, Aug. 2023, doi: 10.3390/app13169288.
- [7] K. S. T., K. P. Revathi, and B. Arivuselvam, "Machine Learning Approaches for Crop Prediction and Disease Detection System with Personalized Recommendations," *International Journal of Gender, Science and Technology*, vol. 13, no. 2, Jul. 2024.
- [8] P. Boopathirajan, P. Rathidevi, S. R. Sanjay, and K. Gokulnath, "Predictive Modeling for Crop Disease Outbreaks," *International Journal*, vol. 7, no. 2, paper no. 55.
- [9] S. Zhang, S. Zhang, C. Zhang, X. Wang, and Y. Shi, "Cucumber leaf disease identification with global pooling dilated convolutional neural network," Elsevier B.V., 2019.
- [10] U. P. Singh, S. S. Chouhan, S. Jain, and S. Jain, "Multilayer Convolution Neural Network for the Classification of Mango Leaves Infected by Anthracnose Disease," 2019.
- [11] V. Singh, "Sunflower leaf diseases detection using image segmentation based on particle swarm optimization," Elsevier, 2019.
- [12] S. Mishra, R. Sachan, and D. Rajpal, "Deep Convolutional Neural Network based Detection System for Real-time Corn Plant Disease Recognition," Elsevier B.V., 2020.
- [13] P. Sharma, Y. P. S. Berwal, and W. Ghai, "Performance analysis of deep learning CNN models for disease detection in plants using image segmentation," Elsevier B.V., 2019.
- [14] M. Agarwal, A. Singh, S. Arjaria, A. Sinha, and S. Gupta, "Tomato Leaf Disease Detection using Convolution Neural Network," Elsevier, 2019.
- [15] A. Khamparia et al., "Seasonal Crops Disease Prediction and Classification Using Deep Convolutional Encoder Network," 2019.