IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

Crop Disease Prediction: A Deep Learning Based Plant Leaf Disease Detection System

Sai Kailas Kendale[1], Aashish Shankar Madhawai[2], Divya Arun Thakare[3], Shraddha Ananta Painjane[4], B. A. Abhale[5]
Jagdamba Education Society's S. N. D. College of Engineering & Research Center, Yeola, Nashik

Abstract : *This paper presents Crop Disease Prediction, a proposed image-based system that uses deep learning to identify possible crop diseases from plant-leaf images. Crop diseases damage leaves, reduce plant growth and cause production losses, while traditional identification depends on visual inspection by farmers or experts, which takes time and may be inaccurate when symptoms look similar or expert help is unavailable. The proposed system classifies a submitted leaf image into the disease categories present in its training data and returns a quick preliminary result that should be confirmed by an agricultural expert.*

The workflow comprises dataset collection from a reliable labelled source, data preparation (removal of unusable images, resizing and normalization), data augmentation such as flips and rotations, training of a convolutional neural network (CNN) or fine-tuning of a pretrained model, evaluation on separate test images, and display of the predicted class together with a note to consult an expert. Performance is to be measured using accuracy, precision, recall, F1-score and a confusion matrix.

The system is presented as a design and project plan. At the stage described here, no trained model, dataset size, accuracy or other experimental result is reported, and none should be inferred. Performance will depend on dataset quality and on field conditions such as lighting, angle and background, and the system should be tested on real-world images before any practical claim is made.

The contribution of this paper is a clearly structured architecture, an evaluation protocol and a statement of limitations for a leaf-image disease screening tool. The system is intended to give early, preliminary support to farmers and other users; it does not replace advice from agricultural experts or laboratory diagnosis.

Keywords: Crop Disease Prediction, Plant Leaf Disease Detection, Deep Learning, Convolutional Neural Network, Transfer Learning, Image Classification, Data Augmentation, Precision Agriculture.

I. INTRODUCTION

Agriculture is an important source of food and income, and healthy crops are essential for good yield. Plant pathogens and pests are a major threat to food production, and Savary et al. [1] estimated the substantial losses they cause to major food crops worldwide. Crop diseases can damage leaves, reduce plant growth and lead to large production losses, so early recognition of disease is valuable for farmers.

Traditional disease identification depends on visual inspection of leaves by farmers or experts. This is slow, depends on experience, and may be unreliable when different diseases produce similar symptoms or when experts are not available. Advances in deep learning for image recognition [2] have made it possible to learn disease patterns directly from leaf photographs, and studies have reported promising results for plant disease detection [3], [4], [5].

This project proposes an image-based system that uses deep learning to identify possible crop diseases from leaf images. A user uploads a photograph of a leaf; the system validates and preprocesses the image, a trained CNN extracts features and classifies the leaf, and the predicted class is displayed together with a reminder to consult an agricultural expert. The system can give an early indication of disease and help users take timely action, but final treatment decisions must be confirmed by an expert.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

It is important to state the status of the work. The system described in this paper is a proposed design with a planned evaluation. Reported accuracies of earlier studies [3], [4], [6] are not results of this project, and the performance of the proposed system can be claimed only after it has been trained and tested on suitable data.

Disease Burden, Detection Barriers and Image Variability:

Disease burden: Crop diseases reduce yield and quality, and their global impact on major food crops is considerable [1]. Early symptoms, such as small spots, discoloration or slight changes in leaf texture, can be difficult to identify, especially for non-experts, so that treatment may be delayed and avoidable losses may occur.

Detection barriers: Manual checking takes time and depends on the availability of experts. Different diseases can look alike, and the same disease can look different at different stages. A quick preliminary prediction from a photograph could help users decide when to seek timely advice, which motivates automated screening based on image classification [7], [8].

Image variability: Real photographs differ in lighting, camera angle, distance and background. Studies have shown that such factors can strongly influence the performance of deep learning models for plant disease recognition [9], and that models trained on controlled images may not transfer directly to field conditions. The research problem is therefore to build and evaluate a leaf-disease classifier that is useful as a preliminary screening tool while being honest about these limits.

Deep Learning for Plant Disease Recognition:

Convolutional neural networks learn hierarchical visual features directly from images and have become the standard approach to image classification [2], [10]. For plant diseases, Mohanty et al. [3] trained deep networks on a large public image collection, Sladojevic et al. [5] recognized several disease classes from leaf images, and Ferentinos [4] compared deep architectures for disease detection and diagnosis.

Transfer learning reuses features learned on large general datasets [11], [12], [13]. Too et al. [6] compared fine-tuned deep learning models for plant disease identification, which indicates that fine-tuning a pretrained network is an effective option when the available plant-image dataset is limited. Recent work such as El-Behery et al. [14] combines CNN-based feature extraction with explainable, few-shot learning for crop disease detection in rice and sugarcane.

The proposed system uses a CNN or a fine-tuned pretrained model to classify leaf images. The choice between a small custom CNN and a pretrained backbone will be made during experimentation on the basis of validation results, and no specific model or accuracy is claimed in this paper.

Core Features and Techniques of the Proposed System:

- **Leaf Image Upload** – A simple interface through which a user submits a photograph of a crop leaf for preliminary analysis.
- **Image Validation and Preprocessing –** Checking that the input is a usable image, then resizing and normalizing it for consistent model input [15].
- **Data Augmentation –** Flips and rotations applied to training images, where appropriate, to improve training variety [16].
- CNN and **Transfer Learning** – A CNN or a fine-tuned pretrained model that learns disease patterns from labelled healthy and diseased leaves [12], [17], [18].
- **Classification Output** – The predicted class, with a confidence indication, shown only for classes that are present in the training data.
- **Expert Confirmation Note** – Every result is displayed with a reminder to confirm diagnosis and treatment with an agricultural expert.

II. RELATED WORK AND LITERATURE REVIEW

Early studies established that deep learning can recognize plant diseases from images. Mohanty et al. [3] trained deep networks on a large public collection of leaf images [19] and reported high accuracy on held-out images of the same collection, while also cautioning that performance on images taken under different conditions is substantially lower.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

Sladojevic et al. [5] used deep neural networks to recognize several plant disease classes by leaf image classification, and Ferentinos [4] evaluated several network architectures for detection and diagnosis.

Comparative and survey studies summarize the state of the field. Too et al. [6] compared fine-tuned deep learning models for plant disease identification, Kamilaris and Prenafeta-Boldú [7] surveyed deep learning in agriculture, and Liu and Wang [8] reviewed deep-learning approaches for plant disease and pest detection. These works show the breadth of methods in use while repeatedly pointing out the need for validation on real farm conditions.

Network architectures commonly used for image classification include VGG [20], GoogLeNet [21], ResNet [17] and DenseNet [22], together with efficient designs such as MobileNet [18] and EfficientNet [23] that are suitable for limited-resource deployment. Pretrained weights learned on ImageNet [11] can be adapted to a new task through transfer learning [12], [13], which is particularly relevant when the number of labelled plant images is limited.

Training practice also affects results. Data augmentation increases the variety of training images [16], batch normalization [24] and dropout [25] help stabilize training and reduce overfitting, and optimizers such as Adam [26] are widely used. Implementation can use mature libraries such as TensorFlow [27] with Keras [28] or PyTorch [29], and image preprocessing can use OpenCV [15].

Several studies emphasize limits to generalization. Barbedo [9] analysed factors that influence deep learning for plant disease recognition, including image background, symptom variability and the presence of multiple conditions, and Barbedo [30] examined how dataset size and variety affect deep learning and transfer learning. Geirhos et al. [31] describe shortcut learning, in which models exploit incidental cues, such as background, rather than the disease itself, which is a concern for leaf datasets collected under uniform conditions.

Interpretation and evaluation are also discussed in the literature. El-Behery et al. [14] propose an explainable few-shot framework for rice and sugarcane, and gradient-based visual explanations [32] can show which image regions influence a prediction. Classification should be reported with appropriate measures [33], and predicted confidence values should not be treated as probabilities of correctness without calibration [34].

A. RESEARCH GAP AND MOTIVATION

Existing studies show that deep learning can classify plant diseases, but they also show that results depend on selected crops, controlled image conditions and dataset composition, and that broader field testing is still needed. Lighting, background and similar symptoms can affect predictions [9]. The motivation of this work is to define a simple, clearly scoped leaf-image screening system for early support, with an evaluation protocol that includes separate test images and standard metrics, and with an explicit expert-confirmation step. This is a design objective; the effectiveness of the system must be established through experiments and is not assumed here.

PROBLEM STATEMENT:

Crop diseases are often identified by manually examining leaves. This process can be slow and may be inaccurate when symptoms look similar or expert help is unavailable. Early symptoms can be difficult to identify, especially for non-experts, and different lighting, angles and backgrounds may affect prediction accuracy. The project therefore aims to develop an image-based crop disease prediction system that classifies leaf images and gives a quick preliminary result, which can help users seek timely advice and reduce avoidable losses.

OBJECTIVES:

I. Collect leaf images: prepare a suitable dataset of healthy and diseased crop leaves from a reliable labelled source.
II. Preprocess images: resize and normalize images for consistent model input, and remove unusable images.
III. Train a model: use a cnn or a transfer-learning model to learn disease patterns from the training images.
IV. Predict disease: classify a submitted leaf image into the supported classes and display the result with a note to confirm treatment with an expert.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

V. Evaluate results: measure performance on separate test images using accuracy, precision, recall, f1-score and a confusion matrix.
VI. Provide a simple interface: display understandable results in a simple user interface after image upload.
VII. Report limits honestly: state the supported classes, the dataset conditions and the need for real-world testing.

PREDICTION WORKFLOW AND USER INTERACTION

the system organizes a prediction request into connected stages so that every result is traceable from the uploaded image to the displayed output.

- the user uploads a leaf image through the interface.
- the system validates the image and rejects unusable input.
- the image is preprocessed (resized and normalized) to match the model input.
- the trained CNN extracts features and classifies the leaf into a supported class with a confidence indication.
- the user reviews the result, optionally the result is stored, and an agricultural expert is consulted for confirmation.

SIGNIFICANCE OF THE STUDY:

This work addresses the need for quick, accessible preliminary screening of crop diseases. By allowing a user to photograph a leaf and receive an immediate indication, the system could reduce dependence on immediate manual checking and support earlier action. The clear statement of scope, supported classes and expert-confirmation step is intended to prevent over-reliance on the result. The work is a basis for further development and testing and does not claim validated field performance.

III. SYSTEM ARCHITECTURE AND MODULE DESCRIPTION

The proposed architecture follows the flow shown in Fig. 1. The user uploads a leaf image; the image is validated and preprocessed; the trained CNN model, which takes a crop-leaf photograph as input, performs feature extraction and disease classification; and a confidence and class output is produced. The output is a predicted disease class and applies only to classes that are present in the training data. The prediction is displayed to the user, the result may optionally be stored, the user reviews it, and an agricultural expert is consulted.

The input module accepts a leaf photograph and performs validation, for example checking file type and readability and rejecting corrupt or unsuitable images. Preprocessing resizes the image to the input size of the network and normalizes pixel values so that the model receives consistent input [15]. No specific resolution or normalization scheme is claimed at this stage, and it will be fixed according to the chosen model.

The model module contains the trained CNN, either a custom network or a pretrained backbone fine-tuned on the leaf dataset [12], [6]. Convolutional layers perform feature extraction, and a final classification layer produces scores for the supported classes. The module is limited to the classes found in the training data; an image of an unsupported plant or disease should not be presented as a reliable result.

The output module converts model scores into a predicted class and a confidence indication that is shown to the user together with a note to confirm treatment with an expert. Raw softmax scores are not calibrated probabilities [34], so the interface should present confidence cautiously, and low-confidence predictions should be marked as uncertain.

The storage and review module optionally stores the uploaded image, predicted class and confidence for later review and for building a record that could be used to improve the model. The privacy and handling of stored images, particularly any location or farmer information, should be defined before real deployment, and storage is optional in the proposed design.

The user-interface module provides a simple and understandable screen for upload and result display. The interface should use plain language, show the predicted class and the expert-confirmation note clearly, and avoid presenting a prediction as a diagnosis. Usability and accessibility have not been evaluated at this stage.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

The architecture is designed to be simple, to keep a clear boundary between prediction and expert advice, and to allow the model to be replaced as better data or networks become available. Its real-world usefulness will require testing on images captured in field conditions and evaluation with intended users.

PROPOSED METHODOLOGY

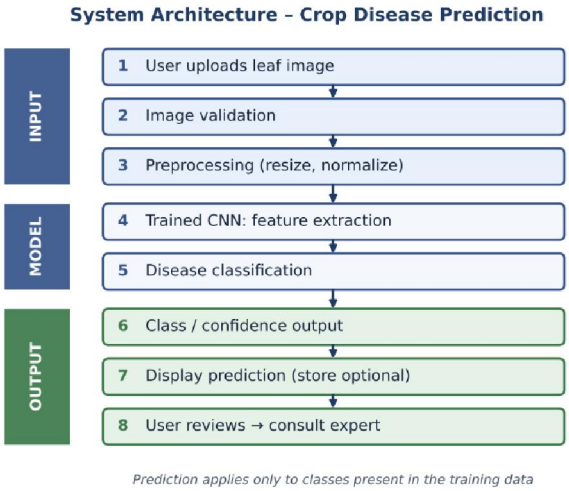


Fig. 1 System Architecture

- Proposed Methodology for the System

The methodology consists of six stages: dataset collection, data preparation, data augmentation, model training, model evaluation and prediction output. Labelled healthy and diseased leaf images are gathered from a reliable dataset; public collections such as PlantVillage [19], [3] are examples of candidate sources, and the dataset to be used is to be finalized and documented.

The six stages are followed by expert confirmation of the output and, as future work, field testing (Fig. 2). Data are divided into training, validation and test subsets so that final metrics are computed on images that the model has not seen during training or tuning.

- IMPLEMENTATION METHOD AND STUDY DESIGN
- 1. DATASET AND MODEL DEVELOPMENT:
- Dataset Preparation: Collect labelled healthy and diseased leaves, remove unusable images, resize and normalize, and keep an independent held-out test set. Dataset size, crop and class names remain to be specified.

IV. PROPOSED EXPERIMENTAL PROCEDURE

The planned experimental procedure has four phases. Phase A prepares the data: images are checked, labels are verified, unusable images are removed, and the dataset is divided into training, validation and test subsets. Phase B develops the model: a baseline CNN and one or more pretrained networks are trained and compared on the validation set. Phase C evaluates the selected model on the held-out test set. Phase D, proposed as future work, evaluates the model on field images captured under real conditions.

To avoid optimistic estimates, images of the same leaf or plant should not appear in both training and test subsets, and class balance should be examined, since highly imbalanced classes can make accuracy misleading. Where the dataset has uniform backgrounds, results may not hold on field photographs [9], [31], and this limitation should be reported with the results.

Training details are chosen on the validation set. Candidates include transfer learning with a pretrained backbone [17], [18], [23] with initial layers frozen and later layers fine-tuned, augmentation by flips and rotations [16], regularization

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

by dropout [25] or batch normalization [24], and optimization with Adam [26]. All hyperparameters, such as learning rate, batch size and number of epochs, are not fixed here and will be reported with the experiments.

Evaluation uses accuracy, precision, recall, F1-score and a confusion matrix on the held-out test set [33]. Per-class precision and recall are reported in addition to overall accuracy so that poorly recognized diseases are visible. Where possible, gradient-based visual explanations [32] will be inspected to check whether the model attends to the diseased region of the leaf rather than to background. No performance figure is claimed in this paper until these experiments have been carried out.

- **Split Strategy:** Keep training, validation and test subsets separate, avoid leakage of the same plant across subsets, and report class distribution.
- **Metrics:** Report accuracy, precision, recall, F1-score and a confusion matrix, with per-class values, computed only on the held-out test set.

V. FUNCTIONAL ANALYSIS:

Module Testing: The validation, preprocessing, model inference and output modules are tested individually, including behaviour on corrupt, very small or non-leaf images.

Workflow Testing: The complete path from image upload to result display and expert-consultation note is tested, including optional storage and the handling of low-confidence predictions.

SYSTEM IMPLEMENTATION AND PROCESSING:

Upload Image: The user selects and uploads a photograph of a crop leaf.

Validate and Preprocess: The system validates the image, then resizes and normalizes it.

Classify: The trained CNN extracts features and classifies the leaf into a supported class.

Display Prediction: The predicted class and confidence are shown, with a note that the result is only preliminary.

Store Result: The result may optionally be stored for later review.

Consult Expert: The user reviews the result and consults an agricultural expert before taking treatment decisions.

VI. DISCUSSION, LIMITATIONS AND THREATS TO VALIDITY

This paper reports a design and an evaluation plan and does not report experimental results. Reported accuracies in the literature [3], [4], [6] describe other models and datasets and must not be read as performance of the proposed system. The actual performance will depend on dataset quality, the number and balance of classes and field conditions.

Several limitations are known in advance. The system recognizes only the classes present in its training data, and an image of an unsupported disease may be assigned to a wrong class with apparent confidence [34]. Lighting, angle, background and similar symptoms can affect predictions [9], and models trained on images with uniform backgrounds may rely on background cues [31]. Multiple simultaneous diseases, nutrient deficiency or pest damage may not be handled correctly.

The dataset has not yet been fixed or documented in this paper, and dataset size and variety are known to influence results [30]. Public datasets collected under controlled conditions may not represent farm conditions, so real-world images should be collected and added in later stages. The usability of the interface has not been evaluated.

Threats to validity include data leakage between training and test subsets, class imbalance, small test sets, overfitting to the dataset background and optimistic reporting of accuracy. These threats should be reduced by independent test data, per-class metrics, visual inspection of model attention [32] and honest reporting of failures.

Scope Statement: The system provides preliminary screening only. Its predictions are not a diagnosis and should not replace advice from agricultural experts; no result of this system has yet been experimentally validated.

KEY METHODS AND TOOLS CONSIDERED: CNN, TRANSFER LEARNING, DATA AUGMENTATION, TENSORFLOW/KERAS OR PYTORCH, AND OPENCV.

- **EVALUATION METRICS**
- **CLASSIFICATION METRICS:**

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

Accuracy: The proportion of test images classified correctly, reported together with class balance.
Precision: The proportion of images predicted as a disease that truly have that disease, reported per class.
Recall: The proportion of images of a disease that are correctly identified, reported per class.
F1-Score: The harmonic mean of precision and recall, reported per class and as a macro average.
Confusion Matrix: A table of predicted versus true classes showing which diseases are confused with each other.
Generalization Check: The change in metrics when the model is tested on images from different conditions.

- **PROPOSED FIELD VALIDATION CRITERIA:**

Performance on leaf photographs captured by users with ordinary cameras under varied lighting.
Performance on different backgrounds, angles and distances.
Behaviour on unsupported diseases and non-leaf images, including how often low confidence is shown.
Agreement of predictions with diagnoses made by agricultural experts.

- **ETHICAL AND RESPONSIBLE USE CONSIDERATIONS**

Advisory Use Only: Present predictions as preliminary indications and not as diagnoses or treatment recommendations.

Data Handling: Obtain permission and protect privacy when storing user images or related information.
Honest Reporting: Report limitations, supported classes and failures, and avoid claiming accuracy that has not been measured.

EXPECTED OUTCOMES

CLASSIFICATION OF SUPPORTED LEAF IMAGES INTO HEALTHY OR DISEASE CATEGORIES.
EVALUATION METRICS SUCH AS ACCURACY, PRECISION, RECALL AND F1-SCORE.
A QUICK PRELIMINARY PREDICTION AFTER IMAGE UPLOAD, IN A SIMPLE INTERFACE.
SUPPORT FOR EARLY ACTION, WITH DIAGNOSIS AND TREATMENT CONFIRMED BY EXPERTS.

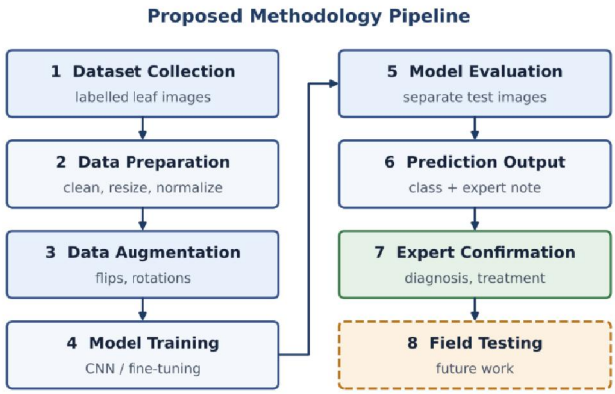


Fig. 2 Proposed Methodology Pipeline

Dataset Collection
→ Labelled healthy and diseased leaf images from a reliable dataset
Data Preparation
→ Remove unusable images; resize and normalize
Data Augmentation
→ Flips and rotations applied to training images where appropriate
Model Training
→ Train a CNN or fine-tune a pretrained model on training images
Model Evaluation
→ Test on separate images; compute accuracy, precision, recall and F1-score

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

Prediction Output

→ Show predicted class with a note to confirm treatment with an expert

Expert Confirmation

→ Agricultural expert confirms diagnosis and treatment

Field Testing (future)

→ Evaluate on real-world images captured under farm conditions

THEORETICAL FRAMEWORK

The theoretical framework of the system rests on representation learning with convolutional networks, transfer learning, data augmentation and regularization, and standard classification evaluation. These concepts explain how a leaf image can be mapped to a disease class and how the quality of that mapping should be measured.

· **Convolutional Representation Learning:**

CNNs learn hierarchical features from pixels, from edges and textures to disease-specific patterns [2], [10], and have been applied to leaf disease classification [3], [5].

· **Transfer Learning:**

Features learned on large datasets are reused and fine-tuned on a smaller target dataset [12], [13], [11], which has been compared across architectures for plant disease identification [6].

· **Augmentation and Regularization:**

Augmentation [16], dropout [25] and batch normalization [24] reduce overfitting and increase robustness to variations in orientation and appearance.

· **Classification Evaluation:**

Accuracy, precision, recall, F1-score and confusion matrices [33] describe different aspects of performance, and confidence values require calibration to be interpreted as probabilities [34].

· **Synthesis of Frameworks:**

Together, these concepts support a system that learns from labelled leaves, generalizes as far as the data allow, is evaluated on independent images and presents its output as preliminary advice [14], [32].

VII. IMPLEMENTATION CONSIDERATIONS AND FUTURE ENHANCEMENTS

Before implementation, the dataset should be selected and documented, including source, number of images, crops, disease classes and capture conditions. The data split should be fixed before training so that the test set remains untouched until final evaluation.

The trained model should be saved with its preprocessing settings and class list so that predictions are reproducible, and the interface should read the class list from the model rather than from hard-coded text. Low-confidence outputs should be flagged to the user as uncertain.

If images are stored, the storage should be protected and the user informed. For deployment on mobile phones, a lightweight network such as MobileNet [18] may be considered, and inference time and memory use should be measured on the target device.

Model updates should be versioned and re-evaluated whenever new data or classes are added, since the system can recognize only the classes on which it was trained [7], [8].

This paper presented Crop Disease Prediction, a proposed deep-learning-based system for identifying possible crop diseases from leaf images. The system addresses the slow and sometimes inaccurate manual identification of crop diseases by offering a quick preliminary indication from a photograph.

The proposed workflow includes dataset preparation, preprocessing, augmentation, model training, evaluation and prediction, followed by expert confirmation. The architecture and evaluation plan follow established practice for leaf-image classification [3], [4], [6] and a clear statement of scope.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

At this stage the work is a design and plan. No trained model, accuracy or field result is reported, and actual performance will depend on dataset quality and field conditions. The system should be tested on real-world images before practical claims are made.

The system gives preliminary support to farmers and other users, and its predictions should not replace advice from agricultural experts. The contribution of this paper is a structured architecture, an evaluation protocol and an explicit statement of limitations that can guide implementation and further study.

FUTURE SCOPE: ADVANCEMENTS IN THE CROP DISEASE PREDICTION SYSTEM

A. Dataset Construction and Documentation

A documented dataset covering target crops and diseases should be built, preferably including field images with varied backgrounds and lighting.

B. Model Training and Comparison

Custom and pretrained CNNs should be trained and compared on validation data, and the best model should be reported with its full training settings.

C. Independent Evaluation

The selected model should be evaluated on a held-out test set with accuracy, precision, recall, F1-score and a confusion matrix, including per-class results.

D. Field Testing

The system should be tested on real-world images taken by farmers under farm conditions, with comparison against expert diagnoses.

E. Explainability and Confidence

Visual explanations [32] and calibrated confidence [34] could help users and experts judge when a prediction can be trusted.

F. Mobile and Offline Deployment

A lightweight model could be deployed in a mobile application that works with limited connectivity, with inference time and memory use measured.

G. Extension to More Crops and Conditions

More crops, diseases and growth stages could be added, together with handling of multiple diseases, pests and nutrient deficiency, if suitable data become available.

Further development and independent testing could help establish the usefulness of the system in practical farming. Evaluation findings could guide improvements to the model, the interface and the way results are communicated.

Future research could compare architectures, augmentation strategies and training regimes to determine which give the most robust results on field images, and how much labelled data is needed for new crops.

Additional work could examine few-shot and explainable approaches [14] to reduce the labelling effort for new diseases and to make predictions easier to verify.

Over time, the system could be integrated with advisory content prepared by agricultural experts, provided each addition is separately tested and its status is reported honestly.

REFERENCES

[1] S. Savary, L. Willocquet, S. J. Pethybridge, P. Esker, N. McRoberts, and A. Nelson, "The global burden of pathogens and pests on major food crops," Nature Ecology & Evolution, vol. 3, pp. 430–439, 2019. https://doi.org/10.1038/s41559-018-0793-y

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015. https://doi.org/10.1038/nature14539

[3] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, Art. no. 1419, 2016. https://doi.org/10.3389/fpls.2016.01419

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

[4] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, 2018. https://doi.org/10.1016/j.compag.2018.01.009
[5] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks based recognition of plant diseases by leaf image classification," Computational Intelligence and Neuroscience, vol. 2016, Art. no. 3289801, 2016. https://doi.org/10.1155/2016/3289801
[6] E. C. Too, L. Yujian, S. Njuki, and L. 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. https://doi.org/10.1016/j.compag.2018.03.032
[7] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018. https://doi.org/10.1016/j.compag.2018.02.016
[8] J. Liu and X. Wang, "Plant diseases and pests detection based on deep learning: a review," Plant Methods, vol. 17, Art. no. 22, 2021. https://doi.org/10.1186/s13007-021-00722-9
[9] J. G. A. Barbedo, "Factors influencing the use of deep learning for plant disease recognition," Biosystems Engineering, vol. 172, pp. 84–91, 2018. https://doi.org/10.1016/j.biosystemseng.2018.05.013
[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems (NeurIPS), vol. 25, 2012. https://papers.nips.cc/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html
[11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in Proc. IEEE CVPR, 2009, pp. 248–255. https://doi.org/10.1109/CVPR.2009.5206848
[12] S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359, 2010. https://doi.org/10.1109/TKDE.2009.191
[13] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in Advances in Neural Information Processing Systems (NeurIPS), vol. 27, 2014. https://arxiv.org/abs/1411.1792
[14] H. El-Behery, A.-F. Attia, and N. G. Rezk, "An explainable deep learning framework for few shot crop disease detection in rice and sugarcane using CNN based feature extraction," Scientific Reports, vol. 16, Art. no. 8272, 2026. https://doi.org/10.1038/s41598-026-37501-2
[15] G. Bradski, "The OpenCV library," Dr. Dobb's Journal of Software Tools, 2000. https://opencv.org/
[16] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," Journal of Big Data, vol. 6, Art. no. 60, 2019. https://doi.org/10.1186/s40537-019-0197-0
[17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE CVPR, 2016, pp. 770–778. https://doi.org/10.1109/CVPR.2016.90
[18] A. G. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861, 2017. https://arxiv.org/abs/1704.04861
[19] 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. https://arxiv.org/abs/1511.08060
[20] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. ICLR, 2015. https://arxiv.org/abs/1409.1556
[21] C. Szegedy et al., "Going deeper with convolutions," in Proc. IEEE CVPR, 2015, pp. 1–9. https://doi.org/10.1109/CVPR.2015.7298594
[22] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE CVPR, 2017, pp. 4700–4708. https://doi.org/10.1109/CVPR.2017.243
[23] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. ICML, PMLR vol. 97, 2019, pp. 6105–6114. https://proceedings.mlr.press/v97/tan19a.html
[24] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in Proc. ICML, PMLR vol. 37, 2015, pp. 448–456. https://arxiv.org/abs/1502.03167

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

[25] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," Journal of Machine Learning Research, vol. 15, pp. 1929–1958, 2014. https://jmlr.org/papers/v15/srivastava14a.html

[26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. ICLR, 2015. https://arxiv.org/abs/1412.6980

[27] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in Proc. USENIX OSDI, 2016, pp. 265–283. https://arxiv.org/abs/1605.08695

[28] F. Chollet et al., "Keras." 2015. [Online]. Available: https://keras.io

[29] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Systems (NeurIPS), vol. 32, 2019. https://arxiv.org/abs/1912.01703

[30] J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," Computers and Electronics in Agriculture, vol. 153, pp. 46–53, 2018. https://doi.org/10.1016/j.compag.2018.08.013

[31] R. Geirhos et al., "Shortcut learning in deep neural networks," Nature Machine Intelligence, vol. 2, pp. 665–673, 2020. https://doi.org/10.1038/s42256-020-00257-z

[32] 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 ICCV, 2017, pp. 618–626. https://doi.org/10.1109/ICCV.2017.74

[33] M. Sokolova and G. Lapalme, "A systematic analysis of performance measures for classification tasks," Information Processing & Management, vol. 45, no. 4, pp. 427–437, 2009. https://doi.org/10.1016/j.ipm.2009.03.002

[34] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in Proc. ICML, PMLR vol. 70, 2017, pp. 1321–1330. https://arxiv.org/abs/1706.04599