JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal

ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

DISEASE DETECTION IN SUGARCANE AND TOMATO CROPS USING CONVOLUTIONAL NEURAL NETWORKS: A DEEP LEARNING APPROACH FOR PRECISION AGRICULTURE

Zahinul Haque Chowdhury[*1], Kanza Saleem[*2], Sikandar Ali Mandan[3], Ayaz Ali Mandan[4]

[*1]*Department of Computer Science, Purdue University Fort Wayne, Fort Wayne, Indiana, USA.*
[*2]*Institute of Agricultural Sciences, University of Punjab, Lahore, Punjab, Pakistan.*
[3]*Faculty of Artificial Intelligence, Department of Smart Engineering & Advanced Technology, Universiti Teknologi Malaysia, Kuala Lumpur*
[4]*Universiti Teknologi Malaysia (UTM), Kuala Lumpur 54100, Malaysia,*
[4]*Benazir Bhutto Shaheed University of Technology and Skill Development, Khairpur Mirs, Sindh, Pakistan*

[*1]zahinulhaquechowdhury@gmail.com, [*2]kanzasaleem222@gmail.com , [3]sikandermandan@gmail.com, [4]ayazmandan@gmail.com

DOI: https://doi.org/10.5281/zenodo.23099656

Keywords
Convolutional neural network, deep learning, plant disease detection, image classification, artificial intelligence, crop health.

Article History
Received: *28 June 2026*
Accepted: *15 August 2026*
Published: *31 August 2026*

Copyright @Author
Corresponding Author: *
Zahinul Haque Chowdhury
Kanza Saleem

Abstract

Plant diseases are a persistent constraint on crop productivity, quality, and farm profitability, while timely diagnosis remains difficult when disease symptoms are visually similar or when expert plant-pathology support is limited. This study developed and evaluated a convolutional neural network (CNN) framework for image-based recognition of healthy and diseased leaves of sugarcane and tomato. A dataset comprising 10,650 images representing 11 classes was assembled from publicly available agricultural image resources. The dataset included five sugarcane classes (healthy, red rot, rust, smut, and mosaic) and six tomato classes (healthy, early blight, late blight, bacterial spot, leaf mold, and Septoria leaf spot). Images were standardized to 224 × 224 pixels, pixel values were normalized, noise reduction was applied, and training images were augmented using rotation, flipping, zooming, shifting, brightness adjustment, and cropping. A supervised CNN was implemented in TensorFlow/Keras using four convolutional stages, ReLU activation, max pooling, a dense layer, dropout regularization, and a Softmax output layer. The model was trained with Adam optimization at a learning rate of 0.001, a batch size of 32, and a maximum of 50 epochs with early stopping. Using a 70% training, 15% validation, and 15% testing partition, the reported model achieved 98.6% training accuracy and 97.8% test accuracy, with overall precision, recall, and F1-score of 97.5%, 97.3%, and 97.4%, respectively. The results indicate that the proposed image-classification framework can distinguish the studied crop-disease categories with high overall performance under the evaluated dataset conditions. However, because the images were primarily obtained from public datasets and the study does not report independent field validation or class-wise confusion statistics, the findings should be interpreted as

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal

ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

evidence of classification performance rather than proof of field-ready diagnostic capability. Further work should incorporate geographically diverse field images, independent external testing, explainable AI, disease-severity estimation, and lightweight deployment on mobile or edge devices.

1. Introduction

Agricultural productivity is increasingly dependent on the ability to monitor crop health rapidly and accurately. Plant diseases can reduce photosynthetic capacity, marketable yield, product quality, and stand establishment, while delayed recognition can increase the spatial and economic extent of an outbreak. Disease diagnosis is therefore not only a biological problem but also an information-management problem: useful disease-management decisions depend on obtaining reliable information at the right time and at an operationally practical cost. Traditional diagnosis based on visual symptoms remains important, but it is influenced by observer experience, symptom overlap, crop developmental stage, lighting, and environmental conditions. Computer vision offers a complementary route by converting visible disease symptoms into quantitative image features that can be processed consistently (Strange & Scott, 2005; Bock et al., 2010; Mahlein, 2016).

Sugarcane and tomato are particularly relevant targets for automated disease recognition because both crops are widely cultivated and affected by multiple diseases with visually recognizable foliar symptoms. In sugarcane, red rot, smut, rust, and mosaic can affect plant vigor and productivity, whereas tomato production is challenged by diseases such as early blight, late blight, bacterial spot, leaf mold, and Septoria leaf spot. The diversity of symptom patterns creates a multi-class classification problem in which a useful model must distinguish both healthy and diseased leaves and, within diseased leaves, separate visually related disease categories.

The development of artificial intelligence has changed image analysis by allowing models to learn discriminative representations directly from pixels. Convolutional neural networks are especially suited to this task because convolutional filters learn local patterns such as edges, textures, color transitions, lesions, and other spatial structures, while deeper layers combine these low-level features into more complex representations (LeCun et al., 2015; Krizhevsky et al., 2012; Simonyan & Zisserman, 2015; He et al., 2016). Unlike conventional machine-learning pipelines that often depend on manually designed color, shape, or texture descriptors, CNNs can jointly optimize feature extraction and classification.

The feasibility of this approach for plant pathology has been demonstrated by several landmark studies. Mohanty et al. (2016) trained deep models on 54,306 PlantVillage images and reported very high accuracy under the dataset conditions, while also demonstrating a substantial performance reduction on images acquired under different conditions. Ferentinos (2018) evaluated deep learning models using 87,848 images across 58 plant-disease classes and reported 99.53% accuracy on an unseen test set. Too et al. (2019) compared fine-tuned architectures and found that DenseNet-based transfer learning performed strongly on PlantVillage data. These studies collectively establish both the promise of CNNs and the importance of evaluating their generalization beyond controlled image collections.

Data quality is central to this problem. Public datasets have accelerated research by providing thousands of labeled images, but many images are acquired against relatively clean backgrounds and under controlled conditions. In operational agriculture, leaves may overlap, illumination can change rapidly, symptoms can occur at different severity levels, and multiple stresses can appear simultaneously. Reviews of plant-disease recognition have therefore emphasized domain shift, dataset bias, class imbalance, field variability, and the need for independent validation as major unresolved challenges (Abade et al., 2021; Liu & Wang, 2021; Salman et al., 2023). Image preprocessing

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

and augmentation are practical strategies for improving robustness. Resizing creates a consistent input representation; pixel normalization helps stabilize optimization; and noise reduction can reduce irrelevant variation. Augmentation techniques such as rotation, reflection, cropping, zooming, translation, and brightness changes expose the network to plausible variations and can reduce overfitting when the original dataset is limited (Shorten & Khoshgoftaar, 2019). Nevertheless, augmentation cannot fully substitute for real field diversity, because synthetic transformations do not reproduce all combinations of weather, camera devices, background vegetation, disease stage, and mixed stresses.

Model architecture has also evolved rapidly. Classical CNNs such as AlexNet and VGG demonstrated the value of deeper hierarchical representations, while Inception, ResNet, DenseNet, MobileNet, and EfficientNet improved accuracy, optimization, parameter efficiency, or computational cost (Krizhevsky et al., 2012; Simonyan & Zisserman, 2015; Szegedy et al., 2015; He et al., 2016; Huang et al., 2017; Howard et al., 2017; Tan & Le, 2019). More recently, vision transformers and hybrid computer-vision systems have expanded the range of models considered for plant monitoring, although CNNs remain attractive where data, computation, and deployment simplicity are important (Dosovitskiy et al., 2021; Singh et al., 2023). For precision agriculture, classification accuracy alone is insufficient to establish practical usefulness. A model can obtain high aggregate accuracy while performing poorly for a minority class or confusing diseases with similar symptoms. Consequently, accuracy should be complemented by precision, recall, F1-score, confusion matrices, and, where appropriate, receiver-operating-characteristic analysis. Class-wise evaluation is particularly important when the consequences of false negatives differ among diseases (Powers, 2011; Sokolova & Lapalme, 2009). Explainability is another emerging requirement because farmers and plant pathologists may need to know which image regions influenced a prediction before acting on an automated diagnosis (Selvaraju et al., 2017).

Against this background, the present study focuses on a practical CNN framework for simultaneous recognition of selected sugarcane and tomato leaf conditions. The study uses 10,650 labeled images covering 11 classes, applies standardized preprocessing and augmentation, and evaluates the resulting model using multiple performance measures. The specific objectives were to: (i) construct a multi-class image dataset for sugarcane and tomato disease recognition; (ii) standardize and augment the images before training; (iii) develop a four-stage CNN classifier using TensorFlow/Keras; and (iv) quantify overall classification performance using accuracy, precision, recall, and F1-score. The study is intended as a computational proof-of-concept and as a basis for future field-oriented disease-monitoring systems.

2. Literature Review and Research Gap

Early image-based plant disease studies commonly relied on handcrafted descriptors derived from color, texture, shape, and lesion geometry, followed by classifiers such as support vector machines, random forests, or artificial neural networks. These pipelines can be effective when symptom characteristics are stable, but their performance may decline when disease appearance varies with illumination, cultivar, disease stage, or background. The transition to deep learning reduced the need for manual feature engineering by allowing hierarchical representations to be learned directly from image data (LeCun et al., 2015; Liu & Wang, 2021).

The PlantVillage initiative was particularly influential because it created an open repository of labeled plant-health images that could be used to benchmark automated diagnosis. Mohanty et al. (2016) demonstrated the potential of CNNs on this resource, but their cross-domain experiment also showed that strong performance on controlled images does not automatically translate into equivalent performance on independently sourced images.

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

This observation is important for the present study because the dataset described here also contains public-domain agricultural images rather than a prospective field-acquisition campaign.

Transfer learning has become a major strategy for plant disease classification because pretrained networks can provide useful visual representations when agricultural datasets are smaller than general computer-vision datasets. Architectures such as VGG, Inception, ResNet, and DenseNet have been applied to plant disease images, and comparative work has shown that architecture choice, fine-tuning strategy, and dataset characteristics can materially affect performance (Too et al., 2019). Lightweight networks such as MobileNet and EfficientNet are particularly relevant for eventual smartphone or edge deployment because they seek favorable accuracy-to-computation ratios (Howard et al., 2017; Tan & Le, 2019).

Tomato disease recognition has received substantial attention because multiple fungal, bacterial, and viral conditions produce leaf symptoms that can overlap. CNN-based studies have reported strong classification results for tomato disease categories, including approaches using relatively compact architectures and transfer learning. Such work supports the feasibility of automated tomato diagnosis but also reinforces the need for evaluation under uncontrolled conditions (Fuentes et al., 2017; Brahimi et al., 2017; Too et al., 2019). Sugarcane disease recognition has received comparatively less attention in the mainstream computer-vision literature, creating an opportunity for multi-crop frameworks that address disease classes across economically important crops.

Recent reviews emphasize that future plant-disease systems should move beyond single-image classification toward field robustness, disease localization, severity estimation, multimodal sensing, explainability, and integration with mobile or Internet-of-Things platforms (Abade et al., 2021; Liu & Wang, 2021; Salman et al., 2023; Singh et al., 2024). Object detection and segmentation are particularly useful when diseased tissue occupies only part of a leaf or when multiple leaves occur within the same image. Semantic segmentation can additionally estimate lesion area and therefore support severity assessment, while object detection can identify affected leaves in complex scenes. The research gap addressed here is therefore not simply the absence of CNNs in plant disease detection; that gap has largely closed. Instead, the contribution is a unified experimental framework that combines sugarcane and tomato classes within one supervised classifier and reports multiple aggregate performance measures. The main remaining limitation is external validity: the current evidence does not establish performance on independently collected field imagery. The study therefore provides a reproducible baseline rather than a final field-deployment solution.

3. Research Objectives

- Develop a CNN-based framework for classifying healthy and diseased sugarcane and tomato leaves.
- Construct and organize a multi-class image dataset representing five sugarcane classes and six tomato classes.
- Apply standardized image preprocessing and augmentation to improve input consistency and reduce overfitting.
- Train the CNN using supervised learning with Adam optimization and early stopping.
- Evaluate model performance using accuracy, precision, recall, and F1-score, while identifying the principal limitations relevant to field deployment.

4. Materials and Methods

4.1 Study Design

The study used an experimental quantitative design based on supervised image classification. The workflow consisted of dataset assembly, image quality screening, preprocessing, data augmentation, CNN development, model training, validation, and independent testing. The supplied manuscript specifies a 70% training, 15% validation, and 15% test

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

partition and reports that cross-validation was used as an additional stability check. Because the source manuscript does not provide the number of cross-validation folds or fold-specific results, these details are not inferred in the present revision.

4.2 Dataset and Class Structure

The dataset contained 10,650 images distributed across 11 classes (Table 1). Images represented healthy and diseased leaves of sugarcane and tomato and were obtained from publicly available agricultural image resources, including PlantVillage and other open-source collections. The original manuscript states that images were screened for duplication and blurring before partitioning. The class distribution reported in the source is reproduced below. The distribution is reasonably broad but not perfectly equal across classes. Therefore, the dataset should be described as multi-class and relatively balanced rather than mathematically balanced. This distinction is important because aggregate accuracy can be influenced by the larger classes.

Table 1. Distribution of healthy and diseased sugarcane and tomato leaf images across the 11 classes used for dataset development.

Crop	Class	Images
Sugarcane	Healthy	1,000
Sugarcane	Red rot	900
Sugarcane	Rust	850
Sugarcane	Smut	900
Sugarcane	Mosaic	850
Tomato	Healthy	1,200
Tomato	Early blight	1,100
Tomato	Late blight	1,050
Tomato	Bacterial spot	950
Tomato	Leaf mold	900
Tomato	Septoria leaf spot	950
Total	11 classes	10,650

4.3 Dataset Partitioning and Leakage Control

Images were divided into training (70%), validation (15%), and testing (15%) subsets. The test set was reserved for final performance assessment. To reduce the risk of optimistic performance, generate augmented variants only from training images after the split, not before. Where multiple images originate from the same physical leaf, plant, or image sequence, group-wise partitioning is preferable because near-duplicate samples can otherwise leak visual information across subsets. The original manuscript does not provide leaf-level or plant-level identifiers, so the extent to which this type of grouping was possible cannot be established retrospectively.

4.4 Image Preprocessing

All images were resized to 224 × 224 pixels to create a uniform CNN input (Table 2). RGB values were normalized to the 0–1 range. The authors used median filtering to reduce noise and state that they minimized image backgrounds so the network could focus on disease-relevant visual characteristics. Preprocessing is consistent with established computer-vision practice, but background suppression should be applied using the same procedure to training, validation, and test images to prevent distributional inconsistencies.

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

Table 2. Summary of image preprocessing procedures, including resizing, RGB normalization, noise reduction, and background minimization, applied consistently across the training, validation, and test datasets.

Preprocessing step	Implementation	Purpose
Resizing	224 × 224 pixels	Standardize CNN input dimensions
RGB normalization	Pixel values scaled to 0–1	Improve numerical stability and optimization
Median filtering	Noise reduction	Suppress isolated image noise
Background minimization	Reduce irrelevant background regions	Increase focus on leaf/symptom features

4.5. Data Augmentation

Training images were augmented using random rotation (±20°), horizontal and vertical flipping, zooming, width/height shifting, brightness adjustment, and random cropping (Table 3). Augmentation was intended to expose the CNN to plausible changes in leaf orientation, scale, framing, and illumination. The use of augmentation is supported by the broader deep-learning literature because it can increase effective sample diversity and reduce overfitting; however, transformations should remain biologically plausible because unrealistic transformations can alter disease morphology or color cues (Shorten & Khoshgoftaar, 2019).

Table 3. Summary of data augmentation techniques and transformation ranges applied to training images of sugarcane and tomato leaves to account for variations in orientation, scale, framing, and illumination.

Technique	Purpose
Rotation (±20°)	Increase orientation invariance
Horizontal/vertical flipping	Increase spatial diversity
Zoom	Simulate camera-distance variation
Width/height shifting	Improve tolerance to framing changes
Brightness adjustment	Improve illumination robustness
Random cropping	Increase positional and scale variation

4.6 CNN Architecture

The proposed classifier consisted of four convolutional stages followed by dense classification layers. Convolutional layers used 3 × 3 kernels with ReLU activation and max pooling (Table 4). The reported filter progression was 32, 64, 128, and 256 filters. The final representation was flattened and passed to a dense layer with 512 neurons, followed by dropout (0.5) and an 11-class Softmax output. The architecture therefore follows the general CNN principle of progressively learning higher-level spatial representations while reducing feature-map dimensions through pooling.

Table 4. Architecture and configuration of the proposed convolutional neural network (CNN) classifier for sugarcane and tomato leaf disease classification.

Stage	Configuration	Role
Input	224 × 224 × 3	RGB leaf image
Conv block 1	32 filters; 3 × 3; ReLU + max pooling	Low-level edges and textures
Conv block 2	64 filters; 3 × 3; ReLU + max	Intermediate patterns

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

	pooling	
Conv block 3	128 filters; 3 × 3; ReLU + max pooling	Higher-level symptom features
Conv block 4	256 filters; 3 × 3; ReLU	Deep visual representation
Dense	512 neurons	Integrated feature representation
Dropout	0.5	Regularization
Output	11-way Softmax	Multi-class prediction

4.7 Model Training

The model was implemented in Python using TensorFlow/Keras. Adam optimization was used with an initial learning rate of 0.001, batch size of 32, categorical cross-entropy loss, and a maximum of 50 epochs (Table 5). Early stopping was applied when validation loss ceased to improve, thereby limiting unnecessary training and helping control overfitting. The manuscript states that training was accelerated using an NVIDIA GPU. The exact GPU model, CUDA version, TensorFlow version, and random seed were not reported and should be added in a future reproducibility revision if available.

Table 5. Hyperparameters and training configuration of the proposed CNN model, including optimizer, learning rate, batch size, loss function, maximum epochs, and early-stopping strategy.

Parameter	Value
Framework	TensorFlow/Keras
Input size	224 × 224 × 3
Optimizer	Adam
Learning rate	0.001
Batch size	32
Maximum epochs	50
Loss	Categorical cross-entropy
Regularization	Dropout = 0.5
Early stopping	Applied
Hardware	NVIDIA GPU; exact model not reported

4.8 Performance Metrics

Model performance was assessed using accuracy, precision, recall, and F1-score. Accuracy represents the proportion of all predictions that were correct. Precision measures the proportion of predicted positives that are correct, recall measures the proportion of actual positives that are identified, and the F1-score is the harmonic mean of precision and recall. For multi-class classification, the reported aggregate values should be accompanied by class-wise metrics and a confusion matrix in future experiments because macro-, weighted-, and micro-averaging can produce different summaries (Sokolova & Lapalme, 2009; Powers, 2011).

Accuracy = (TP + TN) / (TP + TN + FP + FN)

Precision = TP / (TP + FP)

Recall = TP / (TP + FN)

F1-score = 2 × (Precision × Recall) / (Precision + Recall)

5. Results

5.1 Dataset Composition

The final dataset comprised 10,650 images distributed among 11 disease/health categories (Table 6). Sugarcane contributed 4,500 images, whereas tomato contributed 6,150 images. The five sugarcane categories were healthy, red rot, rust, smut, and mosaic. The six tomato categories were healthy, early blight, late blight, bacterial spot, leaf mold, and Septoria leaf spot. The inclusion of healthy leaves enabled the classifier to address the practical first-level distinction between apparently healthy and symptomatic material.

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

Table 6. Class-wise distribution of the 10,650 images comprising five sugarcane and six tomato health/disease categories used for CNN-based classification.

Crop	No. of classes	Images	Share of dataset
Sugarcane	5	4,500	42.25%
Tomato	6	6,150	57.75%
Total	11	10,650	100%

5.2 Overall CNN Performance

The reported training accuracy was 98.6%, while test accuracy reached 97.8%. The corresponding aggregate precision, recall, and F1-score were 97.5%, 97.3%, and 97.4%, respectively (Table 7). The relatively small difference between training and test accuracy (0.8 percentage points) is consistent with limited apparent overfitting under the reported evaluation setup. Nevertheless, this interpretation is conditional because class-wise results, fold-level variability, and an independent external dataset were not reported.

Table 7. Comparative training and test performance of the proposed CNN classifier based on accuracy, precision, recall, and F1-score.

Performance metric	Reported value (%)	Interpretation
Training accuracy	98.6	Overall correctness on training data
Test accuracy	97.8	Overall correctness on held-out test data
Precision	97.5	High overall proportion of correct positive predictions
Recall	97.3	High overall identification of actual classes
F1-score	97.4	Balanced summary of precision and recall

5.3 Error Pattern

The source manuscript reports that the principal classification errors occurred between visually similar tomato diseases, particularly early blight and late blight. This observation is biologically plausible because visually related foliar lesions can share color, shape, and distribution characteristics. However, the supplied study does not provide a numerical confusion matrix or class-wise precision/recall values; therefore, the magnitude of these class-specific errors cannot be quantified from the available results. Overall, the reported metrics indicate strong discrimination across the 11 evaluated categories. The combination of high test accuracy with precision, recall, and F1-score values above 97% suggests that the model was not relying solely on a small number of dominant classes under the reported test conditions. Even so, performance should be considered dataset-specific until independently reproduced on field-acquired images.

6. Discussion

The present study demonstrates the feasibility of a compact CNN architecture for multi-class recognition of selected sugarcane and tomato leaf conditions. The 97.8% test accuracy and closely aligned precision (97.5%), recall (97.3%), and F1-score (97.4%) indicate strong overall classification performance under the evaluated dataset conditions. These findings are directionally consistent with the wider literature, in which CNNs have repeatedly shown high performance on curated plant-disease datasets (Mohanty et al., 2016; Ferentinos, 2018; Too et al., 2019; Liu & Wang, 2021). The high performance should nevertheless be interpreted in the context of dataset construction. PlantVillage and related

JOURNAL FOR SOCIAL SCIENCE STUDIES
An International Peer-Reviewed Open Access Journal
ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

public datasets have been extremely valuable for advancing plant-disease recognition, but many images are acquired under relatively controlled conditions. Mohanty et al. (2016) demonstrated this distinction particularly clearly: high performance on the public dataset was accompanied by much lower accuracy when the model was evaluated on images collected under different conditions. Thus, a 97.8% test accuracy in the present study indicates that the model learned useful visual patterns within the evaluated dataset; it does not establish that the same accuracy will be obtained in commercial fields.

The comparatively small training-to-test performance gap is encouraging from a modeling perspective. Data augmentation, dropout, validation monitoring, and early stopping can all help reduce overfitting. The present pipeline incorporates these measures, consistent with evidence that augmentation can improve generalization when training data are limited (Shorten & Khoshgoftaar, 2019). However, the effect of each component was not isolated through an ablation study. Consequently, the study cannot determine how much of the final performance gain was attributable to preprocessing, augmentation, dropout, architecture, or other implementation choices. The reported confusion between tomato early blight and late blight is important because visually similar diseases are a known challenge in image-based diagnosis. A model may correctly recognize that a leaf is abnormal while still confusing two diseases with overlapping symptoms. This is why class-wise confusion matrices and per-class F1-scores are more informative than a single global accuracy value for deployment decisions (Sokolova & Lapalme, 2009; Powers, 2011). Future work should report these statistics and, where possible, incorporate expert-confirmed diagnoses rather than relying exclusively on image labels.

The architecture used in this study is deliberately simpler than many modern pretrained networks. Four convolutional blocks followed by dense classification provide a transparent baseline and can be less computationally demanding than very deep architectures. Nevertheless, transfer learning from ImageNet-pretrained models such as ResNet, DenseNet, EfficientNet, or MobileNet may offer advantages when field datasets are limited (He et al., 2016; Huang et al., 2017; Howard et al., 2017; Tan & Le, 2019). Comparative experiments would determine whether the present architecture provides an appropriate accuracy–complexity trade-off for the intended deployment environment. Mobile and edge deployment is a logical future direction. Lightweight architectures have been designed specifically to reduce computation while maintaining useful accuracy (Howard et al., 2017; Sandler et al., 2018; Tan & Le, 2019). A field system could combine a smartphone camera with an optimized classifier to provide rapid preliminary screening. However, deployment should be treated as decision support rather than an unquestioned replacement for plant-pathology expertise, particularly for diseases with similar visual symptoms or when symptoms are caused by abiotic stress.

Explainability is another important consideration. Saliency-based methods such as Grad-CAM can visualize image regions contributing to a CNN prediction (Selvaraju et al., 2017). Applying explainability to the present model could help determine whether predictions are driven by disease lesions, leaf margins, background artifacts, or other unintended cues. This would be particularly valuable for public datasets, where models can inadvertently learn correlations associated with image acquisition conditions rather than biological disease features. The study also has implications for precision agriculture. Automated image analysis can support rapid scouting, prioritize samples for expert inspection, and potentially reduce unnecessary broad-spectrum pesticide application when integrated with appropriate disease-management protocols. Computer vision is increasingly being positioned as one component of precision plant-health monitoring rather than a stand-alone diagnostic

ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

endpoint (Mahlein, 2016; Singh et al., 2023; Salman et al., 2023).

Finally, future research should move from classification toward more realistic field tasks. These include disease localization in cluttered scenes, segmentation of lesions, disease-severity estimation, recognition of multiple simultaneous stresses, and multimodal integration with weather, spectral, or sensor information. Recent developments in transformer-based vision models provide additional alternatives, but their usefulness in agriculture must be judged against dataset size, computational requirements, interpretability, and field robustness rather than benchmark accuracy alone (Dosovitskiy et al., 2021; Singh et al., 2024).

7. Study Limitations and Future Research

- The dataset was assembled from public/open agricultural image resources; independent field-collected validation was not reported.
- The manuscript does not provide the exact source and number of images contributed by each database, so complete dataset provenance cannot be independently reconstructed from the supplied material.
- Class-wise precision, recall, F1-score, and a numerical confusion matrix were not reported, limiting detailed assessment of disease-specific errors.
- Cross-validation is mentioned, but the number of folds and fold-level performance are not provided.
- The exact hardware/software environment, random seed, and complete preprocessing implementation are not specified.
- No ablation study was conducted to isolate the contributions of augmentation, dropout, preprocessing, or architecture.
- The current classifier predicts one of 11 predefined classes; it does not estimate disease severity, localize lesions, or explicitly distinguish biotic disease from abiotic stress.
- Future work should use geographically diverse field images, expert-verified labels, independent external test sets, class-wise statistical reporting, explainable AI, and lightweight deployment experiments.

8. Conclusion

This study developed a CNN-based framework for image classification of selected sugarcane and tomato leaf conditions. Using 10,650 images across 11 classes, standardized preprocessing, augmentation, a four-stage CNN, Adam optimization, and early stopping, the reported model achieved 97.8% test accuracy, with aggregate precision, recall, and F1-score of 97.5%, 97.3%, and 97.4%, respectively. These results support the feasibility of automated image-based recognition under the evaluated dataset conditions. The findings should not, however, be interpreted as evidence that the system is already validated for unrestricted field diagnosis. The next stage of development should prioritize independent field validation, class-wise error analysis, disease-severity estimation, explainability, and resource-efficient deployment. Such validation would determine whether the demonstrated computational performance can translate into a reliable precision-agriculture tool.

Funding: Not Applicable

Clinical trial no: Not Applicable

Competing Interest declaration: There are no Competing Interests

Author Contribution declaration

Zahinul Haque Chowdhury: Conceptualization, methodology, data curation, formal analysis, investigation, and writing—original draft. Kanza Saleem: Methodology, investigation, data analysis, and writing—review and editing. Sikandar Ali Mandan: Data collection, investigation, validation, and visualization. Ayaz Ali Mandan: Data curation, validation, interpretation of results, and writing—review and editing. All authors contributed to the manuscript, critically reviewed the final version, and approved the submitted manuscript.

ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

Acknowledgement: We are thankful to all the authors who participated in the current work.

Ethics approval and consent to participate: Not Applicable

Consent for publication: Not applicable.

Availability of data and materials: The data are available on request from the corresponding author.

REFERENCES

Abade, A. S., Ferreira, P. A., & Vidal, F. de B. (2021). Plant diseases recognition on images using convolutional neural networks: A systematic review. Computers and Electronics in Agriculture, 185, 106125. https://doi.org/10.1016/j.compag.2021.106125

Bock, C. H., Poole, G. H., Parker, P. E., & Gottwald, T. R. (2010). Plant disease severity estimated visually, by digital photography and image analysis, and by hyperspectral imaging. Critical Reviews in Plant Sciences, 29(2), 59–107.

Boulent, J., Foucher, S., Théau, J., & St-Charles, P. L. (2019). Convolutional neural networks for the automatic identification of plant diseases. Frontiers in Plant Science, 10, 941. https://doi.org/10.3389/fpls.2019.00941

Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for tomato diseases: Classification and symptoms visualization. Applied Artificial Intelligence, 31(4), 299–315.

Dosovitskiy, A., et al. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. International Conference on Learning Representations.

Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318. https://doi.org/10.1016/j.compag.2018.01.009

Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. Sensors, 17(9), 2022.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778.

Howard, A. G., et al. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv:1704.04861.

Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 4700–4708.

Hughes, D. P., & Salathé, M. (2015). An open access repository of images on plant health to enable the development of mobile disease diagnostics. arXiv:1511.08060.

Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. International Conference on Learning Representations.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521, 436–444.

Liu, J., & Wang, X. (2021). Plant diseases and pests detection based on deep learning: A review. Plant Methods, 17, 22. https://doi.org/10.1186/s13007-021-00722-9

ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

Mahlein, A.-K. (2016). Plant disease detection by imaging sensors—Parallels and specific demands for precision agriculture and plant phenotyping. Plant Disease, 100(2), 241–251.

Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419. https://doi.org/10.3389/fpls.2016.01419

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. IEEE Transactions on Knowledge and Data Engineering, 22(10), 1345–1359.

Picon, A., et al. (2019). Deep convolutional neural networks for mobile capture device-based crop disease classification in the wild. Computers and Electronics in Agriculture, 161, 280–290.

Powers, D. M. W. (2011). Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation. Journal of Machine Learning Technologies, 2(1), 37–63.

Ramcharan, A., et al. (2017). Deep learning for image-based cassava disease detection. Frontiers in Plant Science, 8, 1852.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention, 234–241.

Salman, H., Muhammad, A., Piran, M. J., & Han, D. (2023). Crop-saving with AI: Latest trends in deep learning techniques for plant pathology. Frontiers in Plant Science, 14, 1224709. https://doi.org/10.3389/fpls.2023.1224709

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 4510–4520.

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. Proceedings of the IEEE International Conference on Computer Vision, 618–626.

Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. Journal of Big Data, 6, 60. https://doi.org/10.1186/s40537-019-0197-0

Singh, A., Ganapathysubramanian, B., Singh, A. K., & Sarkar, S. (2016). Machine learning for high-throughput stress phenotyping in plants. Trends in Plant Science, 21(2), 110–124.

Singh, V., Kumar, A., & Sharma, R. (2022). Deep learning approaches for sugarcane disease detection using convolutional neural networks. Computers and Electronics in Agriculture, 196, 106875.

Singh, A., et al. (2023). Computer vision and deep learning approaches for plant disease monitoring: Recent advances and challenges. Plant Methods, 19.

Singh, S., et al. (2024). Next generation of computer vision for plant disease monitoring in precision agriculture: A contemporary survey, taxonomy, experiments, and future direction. Information Sciences.

Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations.

Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. Information Processing & Management, 45(4), 427–437.

Strange, R. N., & Scott, P. R. (2005). Plant disease: A threat to global food security. Annual Review of Phytopathology, 43, 83–116.

ISSN (Online): 3006-2071 | ISSN (Print): 3006-2063
Volume 4 • Issue 4 • August 2026

Szegedy, C., et al. (2015). Going deeper with convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1–9.

Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. Proceedings of the 36th International Conference on Machine Learning, 6105–6114.

Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. Computers and Electronics in Agriculture, 161, 272–279. https://doi.org/10.1016/j.compag.2018.03.032

Van Rijsbergen, C. J. (1979). Information Retrieval (2nd ed.). Butterworths.

Zoph, B., Vasudevan, V., Shlens, J., & Le, Q. V. (2018). Learning transferable architectures for scalable image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 8697–8710.

Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. Computational Intelligence and Neuroscience, 2016, 3289801. https://doi.org/10.1155/2016/3289801

JOURNAL FOR SOCIAL SCIENCE STUDIES
ADVANCING KNOWLEDGE, INSPIRING SOCIETY