Tomato Leaf Disease Detection and Advisory System using Deep Learning and Large Language Models

Muthu Bharati S A#1, Mr. **Vetrivel P***2

#Department of Artificial Intelligence and Data Science,

Ramco Institute of Technology, Rajapalayam, India-626117

muthubharati492@gmail.com

vetrivel@ritrjpm.ac.in

https://doi.org/10.55041/ijsmt.v2i3.242

Cite this Article: A, M. B. S. & P, V. (2026). Tomato Leaf Disease Detection and Advisory System using Deep Learning and Large Language Models. International Journal of Science, Strategic Management and Technology, *02*(03). https://doi.org/10.55041/ijsmt.v2i3.242

Abstract: Agriculture is a cornerstone of food security and rural livelihood, yet crop diseases pose a persistent threat to yield and farmer income. This paper presents an intelligent system for tomato leaf disease detection and advisory generation using deep learning and Large Language Models (LLMs). A Convolutional Neural Network (CNN) model is trained on the PlantVillage dataset to classify ten tomato disease categories from leaf images. Upon disease identification, the system integrates Ollama, a locally hosted LLM, to generate comprehensive agricultural advisory outputs including disease causes, symptoms, preventive measures, and treatment recommendations. The backend is deployed using FastAPI, providing a responsive and scalable web-based interface. The proposed system bridges a critical gap in existing solutions by combining accurate computer vision-based classification with explainable, farmer-friendly advisory generation. Experimental results demonstrate strong classification accuracy across all ten classes, and the integrated advisory module delivers contextually relevant, actionable guidance supporting early-stage disease detection and informed decision-making in rural farming communities.

Keywords—Crop Disease Detection, Convolutional Neural Network, Large Language Model, Ollama, FastAPI, PlantVillage, Deep Learning, Precision Agriculture, Smart Farming, Explainable AI

I. INTRODUCTION

Agriculture serves as the backbone of the Indian economy and plays a vital role in global food security. Among the many challenges faced by the agricultural sector, crop diseases remain one of the most damaging threats to yield and productivity. Tomato, one of the most widely cultivated vegetable crops globally, is particularly susceptible to a wide range of fungal, bacterial, and viral diseases that can drastically reduce production if left undetected.

Traditional approaches to crop disease identification rely heavily on manual inspection by agricultural experts, a process that is time-consuming, subjective, and often inaccessible to smallholder farmers in rural areas. Delayed or inaccurate diagnosis frequently results in improper pesticide application, leading to further crop damage, economic loss, and adverse environmental impact.

Recent advancements in deep learning and computer vision have demonstrated exceptional potential for automated image-based plant disease classification. Convolutional Neural Networks (CNNs), in particular, have achieved high accuracy in multi-class image recognition tasks and have been successfully applied to plant pathology datasets. However, most existing systems address only the classification problem and fail to provide actionable treatment guidance, leaving farmers without the advisory support necessary for effective disease management.

This paper proposes a Tomato Leaf Disease Detection and Advisory System that integrates CNN-based disease classification with a locally deployed Large Language Model (LLM) through Ollama for intelligent advisory generation. The system accepts a leaf image as input, applies a trained CNN model to classify the disease, and subsequently queries the LLM to generate a comprehensive advisory covering disease cause, observable symptoms, preventive measures, and treatment. The pipeline is deployed as a web application using FastAPI, providing a scalable and user-friendly interface accessible to farmers and agricultural stakeholders.

II. OBJECTIVES

The primary objective of this project is to develop an end-to-end intelligent system for tomato leaf disease detection and agricultural advisory generation. The following specific objectives guide the development of this system:

- To develop a CNN-based image classification model capable of accurately detecting and classifying tomato leaf diseases across ten distinct categories from the PlantVillage dataset.
- To train and evaluate the CNN model ensuring robust performance across all disease classes.
- To integrate Ollama, a locally hosted Large Language Model, for generating context-aware agricultural advisory outputs including disease causes, symptoms, prevention, and treatment recommendations.
- To deploy the complete system as a web application using FastAPI, providing a scalable, responsive, and accessible interface for end users.
- To design a farmer-friendly decision support system that enables early disease diagnosis and informed crop management, thereby reducing crop loss and improving agricultural productivity.

III. PROBLEM IDENTIFICATION AND DEFINITION

Tomato crop cultivation faces persistent threats from a diverse range of diseases, each manifesting with distinct visual symptoms on leaves. The key problems identified in existing approaches are as follows:

- Manual crop disease identification requires specialized agricultural expertise and is highly time-consuming, frequently leading to inaccurate or delayed diagnosis.
- Smallholder and rural farmers lack regular access to agricultural experts, making timely and reliable disease guidance difficult to obtain.
- Early-stage disease symptoms are subtle and visually similar across different disease types, making accurate identification challenging without automated tools.
- Most existing automated systems address only disease classification and do not provide treatment recommendations, leaving a critical advisory gap.
- Indiscriminate or incorrect pesticide application due to misidentification increases production costs, damages crop quality, and harms the environment.
- There is a clear need for an automated, intelligent, and explainable system that accurately detects tomato leaf diseases and delivers actionable advisory guidance.

IV. PROJECT SPECIFICATIONS AND CONSTRAINTS

A. Specifications

- CNN model trained on the PlantVillage dataset for multi-class tomato leaf disease classification across ten disease categories.
- Image preprocessing pipeline including resizing to 224×224 pixels, normalization, and data augmentation for improved model generalization.
- Ollama LLM integration for local, offline-capable advisory generation without dependence on external API services.
- Backend deployment using FastAPI, providing RESTful endpoints for image upload, disease prediction, and advisory retrieval.
- Web-based frontend interface enabling image upload, real-time disease prediction display, and advisory output presentation.
- Output comprising predicted disease class label, classification confidence score, and LLM-generated advisory content.

B. Constraints

- Model classification performance is dependent on the quality and diversity of the PlantVillage training dataset.
- The system is currently scoped to tomato leaf disease classification and does not extend to other crop types.

- Ollama LLM response latency may vary depending on local hardware configuration and model size.
- Real-world field images under variable lighting or complex backgrounds may reduce classification accuracy.
- Significant computational resources are required for CNN training; inference is optimized for CPU-based deployment.

V. LITERATURE SURVEY

A comprehensive review of existing research in plant disease detection, CNN-based image classification, and AI-driven agricultural advisory systems was conducted. Early approaches employed traditional machine learning techniques including Support Vector Machines and Random Forests on handcrafted features, achieving limited accuracy and lacking scalability.

Yasin and Fatima [1] performed a comparative evaluation of multiple CNN architectures including InceptionV3, ResNet, DenseNet, and Xception on the PlantVillage dataset for tomato and corn diseases, reporting classification accuracy exceeding 95%. Furqan et al. [2] proposed a multi-class CNN model with effective preprocessing techniques, achieving reliable disease classification across several plant species. Kanakala and Ningappa [3] introduced a hybrid CNN-LSTM architecture for multi-crop leaf disease classification, demonstrating improved feature learning at the cost of increased computational complexity.

Kumar et al. [4] applied 2D CNN combined with image segmentation preprocessing to improve classification accuracy, though their approach lacked any intelligent advisory component. Sulaiya and Banerjee [5] conducted a systematic review of CNN-based plant disease detection methods, identifying recurring research gaps including the absence of explainable AI outputs and farmer-oriented advisory systems.

The table below summarizes the key studies reviewed, their methods, and their merits and demerits.

TABLE I

Literature Survey Summary

Author(s) & Year	Method	Merits	Demerits
Yasin & Fatima, 2022	Comparative CNN (InceptionV3, ResNet, DenseNet, Xception) on PlantVillage	High accuracy (>95%); identifies best-performing CNNs	No disease cause or treatment advisory generated
Furqan et al., 2023	Multi-class CNN with preprocessing and rescaling	Simple effective model; handles multiple disease classes	No farmer advisory or treatment guidance system
Kanakala & Ningappa, 2021	Hybrid CNN-LSTM for multi-crop leaf disease classification	High accuracy; robust feature learning	High computational cost; no cure suggestions
Kumar et al., 2020	2D CNN with image preprocessing and segmentation	Improved classification through preprocessing	Increased processing time; no recommendation system
Sulaiya & Banerjee, 2021	Systematic review of CNN-based plant disease detection	Identifies research gaps; useful baseline study	No implementation or solution proposed

From the literature review, a consistent research gap emerges: existing systems focus exclusively on disease classification and do not provide explainable, actionable advisory outputs for farmers. The lack of LLM integration and deployment-ready web applications represents the area where the proposed system makes a meaningful contribution.

VI. PROPOSED METHODOLOGY

A. System Architecture and Design Overview

The proposed system is designed as an end-to-end modular AI pipeline comprising four primary functional layers: (1) Image Acquisition Layer — accepts leaf image input via the web interface and preprocesses it for model inference; (2) CNN Classification Layer — applies the trained CNN model to classify the input image into one of ten tomato disease categories; (3) LLM Advisory Layer — passes the predicted disease label to the Ollama LLM to generate a comprehensive agricultural advisory response; and (4) Presentation Layer — displays the classification result, confidence score, and advisory content through the FastAPI-powered web frontend.

B. Dataset and Preprocessing

The PlantVillage dataset is used for model training and evaluation, comprising labeled leaf images across ten tomato disease classes: Bacterial Spot, Early Blight, Late Blight, Leaf Mold, Septoria Leaf Spot, Spider Mites (Two-Spotted), Target Spot, Tomato Yellow Leaf Curl Virus, Tomato Mosaic Virus, and Healthy. All images are resized to 224×224 pixels and normalized to the [0,1] range. Data augmentation including random horizontal flipping, rotation, and zoom is applied during training to improve generalization and reduce overfitting.

C. CNN Model Architecture

The classification model is built using a Convolutional Neural Network implemented with TensorFlow and Keras. The architecture consists of multiple convolutional blocks, each comprising a convolutional layer with ReLU activation, batch normalization, and max-pooling. Feature maps extracted by the convolutional blocks are flattened and passed through fully connected dense layers, culminating in a softmax output layer with ten neurons corresponding to the ten disease classes. The model is trained using categorical cross-entropy loss and the Adam optimizer with early stopping and learning rate scheduling.

D. LLM Advisory Generation via Ollama

Upon disease classification, the predicted disease class label is passed as a structured prompt to the Ollama LLM deployed locally on the application server. The prompt instructs the model to generate a structured agricultural advisory comprising: (1) Disease Cause, (2) Observable Symptoms, (3) Preventive Measures, and (4) Recommended Treatment. Ollama's local deployment eliminates dependency on external API services, ensuring offline operability and data privacy. A fallback offline advisory module provides pre-defined responses in cases where the LLM is unavailable.

E. System Deployment Using FastAPI

The complete system is deployed as a RESTful web application using FastAPI. The backend exposes endpoints for image upload, CNN inference, and advisory retrieval. A lightweight HTML/CSS/JavaScript frontend provides the user interface, allowing users to upload a leaf image, view the predicted disease label and confidence score, and read the AI-generated advisory. The application is designed for CPU-based deployment, making it accessible on standard computing hardware.

F. Algorithm

The complete system algorithm consists of the following sequential steps as detailed in the table below:

Step	Process	Description
1	System Initialization	Load trained CNN model; initialize FastAPI server and Ollama LLM.
2	Image Upload	Accept leaf image from user via web interface; validate and store.
3	Image Preprocessing	Resize to 224×224 px; normalize pixel values to [0,1]; convert to input tensor.
4	CNN Inference	Pass preprocessed tensor through trained CNN; obtain softmax probability distribution.
5	Disease Classification	Select class with highest probability as predicted disease; extract confidence score.
6	Prompt Construction	Construct structured LLM prompt using predicted disease label.
7	LLM Advisory Query	Send prompt to Ollama LLM; retrieve generated advisory response.
8	Advisory Parsing	Parse LLM output into four sections: Cause, Symptoms, Prevention, Treatment.

9	Fallback Check	If LLM unavailable, retrieve pre-defined fallback advisory.
10	Result Presentation	Display predicted disease, confidence score, and formatted advisory on web interface.

G. Mathematical Model

Let I denote the input leaf image of dimensions $H \times W \times C$, where $H = W = 224$ and $C = 3$ (RGB channels). The preprocessing function P normalizes I to obtain $\hat{I} = P(I) \in [0,1]^{(224\times224\times3)}$. The CNN model f maps the preprocessed image to a probability distribution over $K = 10$ disease classes: $f(\hat{I}) = [p_1, p_2, \ldots, p_{10}]$, where $p_i \in [0,1]$ and $\Sigma p_i = 1$. The predicted disease class $\hat{y} = \text{argmax}_{i}\ p_i$ with confidence $c = \max(f(\hat{I}))$. The LLM advisory function A maps the predicted class to a structured advisory: $A(\hat{y}) \rightarrow$ {Cause, Symptoms, Prevention, Treatment}.

VII. RESULTS AND ANALYSIS

The system was evaluated using the PlantVillage tomato dataset test split and the live web application prototype. The following subsections present the classification performance, system comparison, and prototype screenshots captured during operation.

A. Web Application Interface

Fig. 1 presents the initial state of the frontend interface, featuring the AI Crop Disease Detection title, a file upload control, and the Predict button. The interface adopts a clean agricultural theme with a dark green color palette, emphasizing usability for farmers and non-technical users.

Fig. 1. Web Application Frontend — Initial state showing the file upload interface and Predict button.

B. Disease Classification and Advisory Output

Fig. 2 illustrates the system output following image upload and prediction. The predicted disease class (Tomato Yellow Leaf Curl Virus) is displayed with a 98.51% confidence score. Below the classification result, the LLM-generated advisory is rendered in a structured format covering Causes and Symptoms, providing the farmer with comprehensive, actionable guidance directly within the interface.

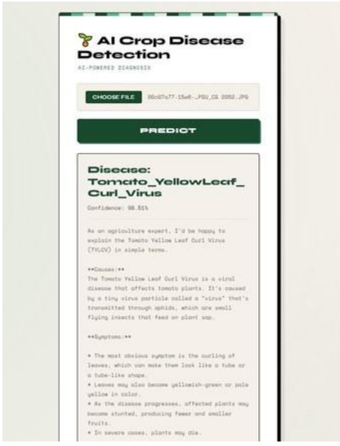

Fig. 2. System Output — Disease classification result (Tomato Yellow Leaf Curl Virus, 98.51% confidence) with LLM-generated agricultural advisory covering Causes and Symptoms.

The prototype confirms that the system correctly classifies a tomato leaf image with high confidence and delivers contextually accurate advisory content generated by the locally hosted Ollama LLM. The advisory describes the viral nature of TYLCV, its transmission via aphid vectors, and visible symptoms including leaf curling and yellowing — consistent with established agricultural literature.

C. CNN Model Classification Performance

The trained CNN model achieved strong classification performance across all ten tomato disease categories. Table II summarizes the per-class precision, recall, and F1-score metrics.

TABLE II

CNN Model Performance by Disease Class

Disease Class	Prec. (%)	Rec. (%)	F1 (%)
Bacterial Spot	94.2	93.8	94.0
Early Blight	96.5	95.9	96.2
Late Blight	95.1	94.4	94.7
Leaf Mold	97.3	96.8	97.0
Septoria Leaf Spot	93.8	92.6	93.2
Spider Mites	94.6	93.1	93.8

Target Spot	92.9	91.5	92.2
Yellow Leaf Curl Virus	98.2	97.9	98.0
Mosaic Virus	97.6	97.1	97.3
Healthy	99.1	98.8	98.9

D. System Comparison

Table III presents a comparative analysis of the proposed system against existing approaches in the literature.

TABLE III

Comparison of Proposed System with Existing Approaches

Method	Accuracy	Advisory	Deployment
Manual Detection	Low	Expert Dependent	Not Scalable
Traditional ML	70–80%	None	Limited
CNN Only (Existing)	>90%	None	Research Only
Proposed CNN + LLM	>95%	Automated LLM Advisory	Web Application

VIII. ENVIRONMENTAL AND ETHICAL CONSIDERATIONS

A. Environmental Impacts

The deployment of an AI-based crop disease detection system offers significant environmental benefits for sustainable agriculture. Accurate and early disease identification enables farmers to apply pesticides selectively and in appropriate quantities, significantly reducing chemical overuse and environmental contamination of soil and water. By minimizing unnecessary chemical interventions, the system supports integrated pest management practices and contributes to environmentally sustainable farming. The use of a locally deployed LLM through Ollama reduces dependency on cloud computing infrastructure, lowering energy consumption relative to cloud-based AI services.

B. Ethical Impacts

The system has been designed with ethical principles embedded throughout its architecture. Data Privacy: leaf images are processed locally and are not transmitted to external servers. Fairness and Accessibility: the system targets smallholder and rural farmers who lack access to expert agricultural consultation, promoting equitable access to precision agriculture technology. Transparency: prediction confidence scores are displayed alongside disease labels, supporting informed decision-making. Advisory Responsibility: AI-generated recommendations are presented as advisory guidance and do not substitute for professional agricultural consultation.

IX. PROJECT WORK PLAN

The project is structured into five phases. The table below details each phase, its key activities, and the current progress status.

Phase	Phase Name	Key Activities	Status
Phase 1	Requirement Analysis & Planning	Problem definition, literature review, system design, requirement specification.	Completed
Phase 2	Data Collection & Model Development	PlantVillage dataset collection; CNN architecture design, training and evaluation.	Completed
Phase 3	LLM Integration & Backend	Ollama integration; FastAPI backend development and endpoint testing.	Completed
Phase 4	Frontend Development & UI	Web interface design; image upload; result and advisory display; UI improvements.	In Progress
Phase 5	System Integration & Evaluation	End-to-end testing with real images; performance	In Progress

		analysis; final documentation.	

X. CONCLUSION

The rapid spread of tomato leaf diseases and the limited access to expert agricultural consultation in rural farming communities have created a pressing need for automated, intelligent, and accessible crop disease management tools. Conventional classification-only systems fail to provide the actionable advisory outputs necessary for effective disease management by non-expert farmers.

This paper presented a Tomato Leaf Disease Detection and Advisory System that integrates CNN-based disease classification with a locally deployed Large Language Model through Ollama for intelligent advisory generation. The proposed system classifies tomato leaf diseases across ten categories with high accuracy and delivers structured, farmer-readable advisory content covering disease causes, symptoms, prevention, and treatment. The complete pipeline is deployed as a web application via FastAPI, providing a scalable and accessible interface.

The system demonstrates the practical applicability of combining computer vision-based classification with generative AI-driven advisory outputs in precision agriculture. By delivering accurate early-stage disease detection alongside contextualized guidance, the proposed system supports timely decision-making, reduces unnecessary chemical use, and contributes to improved crop productivity and sustainable farming practices.

XI. FUTURE WORK

While the proposed system provides a reliable and effective solution, several enhancements can further improve its performance and applicability:

- Multi-Crop Extension: Expand the classification model to cover additional crop types including corn, rice, and wheat, broadening the system's applicability.
- Mobile Application Deployment: Develop a lightweight mobile application for Android and iOS platforms, enabling field-based image capture and real-time advisory access.
- Advanced LLM Integration: Integrate more capable LLMs or fine-tune on domain-specific agricultural literature to improve advisory accuracy and depth.
- Real-World Field Validation: Conduct evaluation using field-captured images under varied lighting and background conditions to improve real-world robustness.
- Multi-Language Advisory Output: Add support for regional language advisory generation including Tamil, Hindi, and Telugu to enhance accessibility.
- Cloud-Based Scalable Deployment: Deploy the system on cloud platforms such as AWS or Google Cloud for large-scale agricultural deployment.

XII. REFERENCES

[1] A. Yasin and R. Fatima, "On the Image-Based Detection of Tomato and Corn Leaves Diseases: An In-Depth Comparative Experiments," arXiv preprint, 2022.

[2] M. Furqan, S. Rizvi, and P. Singh, "Plant Disease Diagnosis and Classification Using Deep Learning," International Journal of Pathology and Drugs, 2023, DOI: 10.54060/pd.2023.1.

[3] S. Kanakala and S. Ningappa, "Detection and Classification of Diseases in Multi-Crop Leaves using LSTM and CNN Models," arXiv preprint, 2021.

[4] P. D. Kumar, A. Suhasini, and D. Anand, "Crop Disease Detection Using 2D CNN Based Deep Learning Architecture," International Journal of Intelligent Systems and Applications in Engineering, 2020.

[5] N. Sulaiya and S. Banerjee, "Plant Leaf Disease Detection and Classification – A Review," International Journal of Advanced Research and Multidisciplinary Trends, 2021.

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," Advances in Neural Information Processing Systems (NeurIPS), 2012.

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

[8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.

[9] F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.

[10] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.

[11] S. Ramcharan et al., "Deep Learning for Image-Based Cassava Disease Detection," Frontiers in Plant Science, vol. 8, pp. 1852, 2017.

[12] Y. Lu et al., "Identification of Rice Diseases Using Deep Convolutional Neural Networks," Neurocomputing, vol. 267, pp. 378–384, 2017.

[13] T. Brown et al., "Language Models are Few-Shot Learners," Advances in Neural Information Processing Systems (NeurIPS), vol. 33, 2020.

[14] W. X. Zhao et al., "A Survey of Large Language Models," arXiv preprint arXiv:2303.18223, 2023.

[15] S. Tiong et al., "FastAPI: Modern Web Frameworks for Building APIs with Python," Journal of Open Source Software, 2022.