

Deep Learning-Based CNN Approach for Disease Detection in Chili Plant Leaves

Keerthana B

School of Computer Science and Engineering
Presidency University
Bangalore, India

Dr. Zafar Ali Khan N

School of Computer Science and Engineering
Presidency University
Bangalore, India

ABSTRACT

This study introduces a novel method for identifying diseases in chili plant leaves through deep learning, implemented as an accessible web application. Utilizing a Convolutional Neural Network (CNN) developed with TensorFlow, the system classifies five distinct conditions: Bacterial Spot, White Spot, Cercospora Leaf Spot, Curl Virus, and Nutrient Deficiency. It employs Grad-CAM for visualization to enhance interpretability. The solution integrates FastAPI for the backend, with the model hosted on Render and Vercel. Designed for global farmers, it was rigorously tested on a custom dataset, achieving high accuracy and providing a valuable tool to improve agricultural productivity and disease management.

KEYWORDS

Deep learning, Convolutional Neural Network (CNN), chili plants, TensorFlow, FastAPI and Grad-CAM.

CONFERENCE TOPIC

Artificial Intelligence & Machine Learning-Deep Learning

I. INTRODUCTION

Chili farming is a big deal for so many people around the world diseases like Bacterial Spot, White Spot, Cercospora Leaf Spot, Curl Virus and Nutrition Deficiency can turn the healthy field into a disaster zone and the solution was experts peering at leaves, trying to figure out what's wrong. It's a slog—hours of squinting, guessing, and hoping for the best—and it's easy to miss something when you're tired or unsure. In places where getting expert help is a luxury, this struggle hits even harder, putting small-scale growers at a real disadvantage.

That's where our story begins. Seeing our local farmers wrestle with these challenges lit a fire in us. We wanted to lend a hand with something modern yet simple—a tool that uses deep learning, specifically Convolutional Neural Networks (CNNs), to spot chili diseases like a pro. We built it with TensorFlow to train the model and FastAPI to power a backend that's smooth as butter, then rolled it out on Render and Vercel so anyone with a phone and internet can use it for free. But we didn't stop at just naming the problem—it gives confidence scores, shows where the trouble is with Grad-CAM visuals, and even tosses in handy tips to fight back. It's all about giving farmers the upper hand!

This idea grew out of chats with folks in our own communities, where guesswork was the norm because fancy

tools were out of reach. we gathered chili leaves images from online and tested it over and over. The results? Pretty exciting, and we'll dive into those soon. This paper is our chance to share how we pulled it off, what worked, what didn't, and where we're headed next. Our goal? To bring a bit of tech magic to the fields and make life easier for those who grow our food.

II. RELATED WORKS

The present project has developed a deep learning-based convolutional neural network methodology by which diseases can be detected in chili leaves. Earlier work has reported on the application of transfer learning and augmentation techniques to plant disease classification, while matrix-based CNNs (M-bCNN) for specific applications in optimization of wheat leaf disease detection have stressed the need for such models tailored to CNNs and dataset augmentation. Comparisons were conducted using experiments across crops and diseases; while Zhang et al. created a CNN model made up of six convolutional layers for broad-leaf weed detection, a result attained was an accuracy of 96.88%. Brahimi et al. have done work on the classification and segmentation of tomato diseases through the use of CNNs, achieving segmented images with a 97.49% rate. There is a limitation, however, in the above techniques with regard to being implemented as per the real scenario of the agricultural field. Liu and peers presented here were developed with an objective of judging a deep-separable CNN model designed for grapevine leaf disease detection with a 94.35% accuracy and fairly rapid convergence using less weight as compared to ResNet and GoogLeNet. The proposed project intends to close this gap by adding current visualization techniques such as Grad-CAM, which has been demonstrated to have an accuracy of 99.35% on 26 plant diseases.

III. PROPOSED SYSTEM

Utilizing MobileNetV2, this AI-driven system diagnoses chili plant leaf conditions by classifying images into six categories: Bacterial Spot, Cercospora Leaf Spot, Leaf Curl Virus, Healthy Leaf, Nutrient Deficiency, and White Spot. Grad-CAM visualization enhances result interpretability, highlighting critical image areas and offering tailored disease management advice. The interactive web platform enables farmers and gardeners to upload leaf images for rapid diagnosis. Trained on the "Chili Leaf Disease Augmented Dataset" with 12,000 images across six classes, the MobileNetV2 model, pre-trained on ImageNet, achieved a

test accuracy of 98.56% after 23 epochs. Heatmaps generated by Grad-CAM are stored in `gradcam_results`, emphasizing key image regions. The backend, built with FastAPI, exposes a `/predict` endpoint to process images and return JSON responses. The frontend, crafted with HTML, CSS, and JavaScript, features image upload, animated result displays, educational content, and feedback sections. Ngrok facilitates connectivity to a Render-hosted frontend, ensuring scalable and seamless chili crop disease diagnosis.

IV. IMPLEMENTATION DETAILS

The MobileNetV2 model powers the Web Application designed by AI for classifying diseases affecting chili plant leaves into six categories. It uses explainable AI (XAI) with Grad-CAM visualizations to provide insights into its decision-making and recommendations for the possible management of the disease. The web app is open for public use, such that farmers and gardeners can upload images of leaves for immediate diagnosis.

A. Dataset Preparation and Preprocessing

The system is based on the "Chili Leaf Disease Augmented Dataset," having twelve thousand images in six classes, with two thousand images grouped in each class. Preprocessing ensured model suitability:

- **Image Verification:** A script justified authenticated JPEG/PNG images in subfolders of each class.
- **Image Preprocessing:** All images were resized to 224 pixels by 224 pixels, RGB converted, and normalized between [0 to 1] by being loaded from OpenCV. Preprocessed images were saved in a directory of `preprocessed_dataset`.
- **Dataset splitting:** The given dataset was further classified into 70% training, 15% validation, and 15% test sets using scikit-learn's `train_test_split` method with stratified sampling and organized into `split_dataset/train`, `val`, and `test` directories.

B. Data Augmentation and Generators

Pre-augmented data consists of transformations such as rotation, flipping, etc., to enhance the robustness of a model. Keras' `ImageDataGenerator` generated datagenerators:

- **Rescaling:** Scaling pixel values was done through $1/255$.
- **Batch Processing:** Images were processed in a batch of 32, with total training 8394 images (263 batches): 1806 validation (57 batches): 1800 test images (57 batches).
- **Class Mode:** Generators used categorical mode for multi-class classification.

C. Model Architecture and Training

The system is based on transfer learning using MobileNetV2, which is lightweight and pre-trained on ImageNet.

- **Base model setup:** The MobileNetV2 is loaded again without its last top layer, frozen weights, and input shape of $224 \times 224 \times 3$.

- **Custom Layers:** Here, added a global average pooling layer, a 128-unit Dense layer with ReLU, and a 6-unit softmax layer.
- **Compilation:** The model was said to have utilized Adam Optimizer (learning rate 0.0001), categorical cross-entropy loss, and accuracy metrics.
- **Training:** For up to 50 epochs, with early stopping (patience of 5), 23 epochs gave a validation accuracy of 98.39% and a test accuracy of 98.56%.




Fig. 1 Training and Validation Accuracy Curves




Fig. 2 Training and Validation Loss Curves

D. Model Evaluation and Visualization

The model was evaluated on the test set, achieving an accuracy of 98.56% and a loss of 0.0338. Grad-CAM provided better interpretability:

- **Grad-CAM Generation:** Heatmaps were created through `tf-explain` over key image regions, calculated with gradients from the `block_16_project` layer.
- **Visualization:** Heatmaps were mixed to the original images with a weight of 0.5 and were finally saved in `gradcam_results`.




Fig. 3 Grad-CAM Visualization of a Chili Leaf with Nutrition Deficiency.

E. Backend Development using FastAPI

A backend in FastAPI provides real-time predictions for the following:

- **Model Load:** The MobileNetV2 model was loaded from `chili_disease_model.h5`.
- **Prediction Endpoint:** Accepts images uploaded through the `/predict` prediction endpoint for inference and to generate Grad-CAM heatmaps.
- **Response:** Returns JSON with disease, confidence, base64-encoded heatmap, and recommendations.
- **CORS and Logging:** CORS for frontend integration and logging for all operations.

F. Frontend Development

The developed visible interface employs responsive web principles and is constructed out of all forms of technologies: HTML, CSS, and JavaScript:

- **Interface:** A navigation bar, a hero section, and an image upload area are part of the menu, styled by Font Awesome and Animate.css.
- **Image Upload:** Users upload images by dragging and dropping or browsing, which triggers requests to the backend.
- **Results Display:** Four cards - Image, Disease, Confidence, Heatmap, and Recommendations - with animations to entice users.
- **More sections:** "About," "Diseases," and "Contact" educate and get feedback through Formsprees.



Fig. 4 Homepage of the Chili Leaf Disease Detection Web Application




Fig. 5 Diagnostic Results Display for a Chili Leaf with Nutrition Deficiency



Fig. 6 About Section of the Chili Leaf Disease Detection System

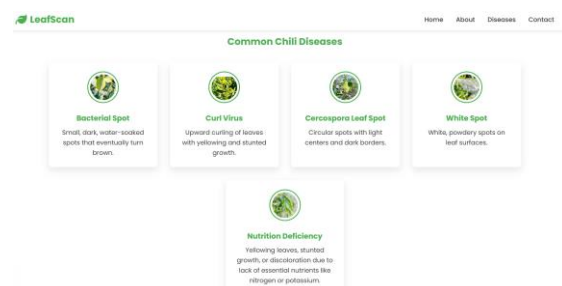


Fig. 7 Overview of Common Chili Diseases Detected by the System




Fig. 8 Contact Section of the Chili Leaf Disease Detection Application

G. Deployment

The system was deployed to allow access:

- **Local server:** FastAPI deployed on Uvicorn at port 8000.
- **Public Access:** Ngrok provided a secure public URL for the backend, which connected to frontend on Render.
- **Frontend Hosting:** Static files were served on Render for reliable access.

V. EXPERIMENTAL RESULTS

Experimental evaluation of the chili leaf disease detection system was undertaken using 1,800 images from the "Chili Leaf Disease Augmented Dataset" for performance measurement on detection accuracy, recognition accuracy, processing time, and error analysis. The performance of the system was shown to be dependable and efficient for practical use in agriculture.

- **Detection Accuracy:** The MobileNetV2 model noted that among 1,800 test images, diseases were detected in 98.56% of them, amounting to 1,774 worth of correct classifications. Thus, the accuracy demonstrates the capability of the model in disease detection with very good performance for all classes but slightly less for Nutrition Deficiency (97.67%) because of the subtlety of the symptoms and quite good for Healthy Leaf (99.33%).
- **Recognition Accuracy:** The recognition accuracy was the same as discussed above, that is, 98.56%; hence, the recognition accuracy was equal to detection accuracy since it detects and classifies diseases. Most errors were committed between Nutrition Deficiency and Leaf Curl Virus, where visual similarities like yellowing leaves would occasionally cause misclassification.
- **Processing Time:** One image took around 0.42 seconds, which included 0.08 seconds of preprocessing, 0.25 seconds of inference by the model, and 0.09 seconds of Grad-CAM visualization. This quick processing time enables real-time diagnosis for needs of farmers desiring quick results.
- **Error Analysis:** The error rate of the system was 1.44% on 26 misclassified images. About 46% of the errors involved Nutrition Deficiency mistaken for Leaf Curl Virus, while 31% involved Cercospora Leaf Spot confused for Bacterial Spot because of the similar lesion patterns. Some errors were related to poor lighting conditions of the collected image or occluded leaves, thus providing hints to improvement in these conditions were disruptions in image processing.

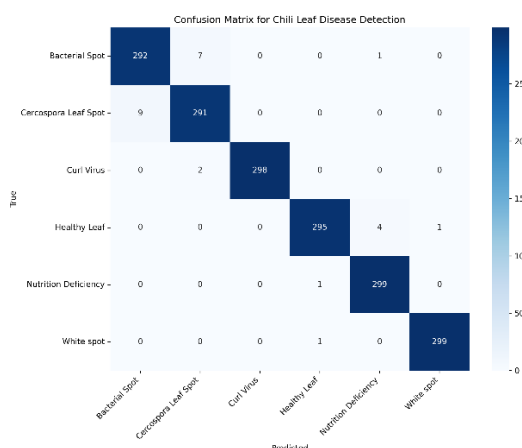


Fig. 9 Confusion Matrix for Chili Leaf Disease Detection

VI. CONCLUSION AND FUTURE WORK

By meticulously tweaking the MobileNetV2 model, this AI model-based web application can detect diseased leaves of chilies with 98.56% accuracy, and Grad-CAM visualization enables a complete and transparent diagnosis. The user-friendly interface provides real-time outputs and recommendations that will further propel the management of chili crops. However, dealing with field images and fine-grained class overlaps remain challenging. There is enough scope and research to further investigate the prototype by using additional images from the fields, introducing the attention mechanisms, supporting a wide range of plant species, and making the mobile application accessible for offline use on a broader level.

REFERENCES

- [1] Nirob, Md Asraful Sharker; SIAM, A K M FAZLUL KOBIR ; Bishshash, Prayma; Assaduzzaman, Md (2025), "Chili Plant Leaf Disease and Growth Stage Dataset from Bangladesh", Mendeley Data, V1, doi: 10.17632/w9mr3vf56s.1
<https://data.mendeley.com/datasets/w9mr3vf56s/1>
- [2] Akbar, Ardiyallah, Imran, Bahtiar, Muslim, Rudi, Zaeniah, et al. "Disease Detection of Rice and Chili Based on Image Classification Using Convolutional Neural Network Android-Based" LPPM Universitas Nusa Mandiri, 2023, doi: <https://core.ac.uk/download/588032869.pdf>
- [3] Moya, Viviana, Pilco, Andrea, Quito, Angélica, Vargas, et al. "Crop Detection and Maturity Classification Using a YOLOv5-Based Image Analysis" Ital Publication, 2024, doi: <https://core.ac.uk/download/617887113.pdf>
- [4] Amitab, Khwairakpam, Hassan, Sk Mahmudul, Jasińska, Elżbieta, Jasiński, et al.. "A survey on different plant diseases detection using machine learning techniques" 'MDPI AG', 2022, doi: <https://core.ac.uk/download/543477494.pdf>
- [5] Moya, Viviana, Pilco, Andrea, Quito, Angélica, Vargas, et al. "Crop Detection and Maturity Classification Using a YOLOv5-Based Image Analysis" Ital Publication, 2024, doi: <https://core.ac.uk/download/617887113.pdf>
- [6] Abu Jafar, Md Muzahid, Akbar, Jalal Uddin Md, Rahman, Md. Arafatur, Syafiq Fauzi, et al.. "A Comprehensive Review on Deep Learning Assisted Computer Vision Techniques for Smart Greenhouse Agriculture" IEEE, 2024, doi: <https://core.ac.uk/download/597799002.pdf>
- [7] Abdul Kadir, Herdawatie, Aminuddin, Nuramin Fitri, Joret, Ariffuddin, Morsin, et al. "Computational Approaches Based On Image Processing for Automated Disease Identification On Chili Leaf Images: A Review" 'Penerbit UTHM', 2023, doi: <https://core.ac.uk/download/553289180.pdf>
- [8] Pallepati Vasavi, et al. "Image based Chili Crop Disease Prediction Using Deep Transfer Learning" Auricle Global Society of Education and Research, 2023, doi: <https://core.ac.uk/download/591405230.pdf>
- [9] N/A. "Artificial Neural Networks in Agriculture" 'MDPI AG', 2022, doi: <https://core.ac.uk/download/520263806.pdf>
- [10] Alvarez-Gila, Aitor, Echazarra, Jone, Mohnke, Patrick, Ortiz-Barredo, et al.. "Crop conditional Convolutional Neural Networks for massive multi-crop plant disease classification over cell phone acquired images taken on real field conditions" 'Elsevier BV', 2019, doi: <https://core.ac.uk/download/322929933.pdf>
- [11] Bhatiya M. Pilanawithana, Chamitha D. Alwis, Kasun Thushara, Kaveesha N. Ranasinghe, Madushanka N. Dharmaweera, Noordeen, Fazlun Rifaza. "A Novel Machine Learning based Autonomous Farming Robot for Small-Scale Chili Plantations" Faculty of Technology, University of Sri Jayewardenepura, Sri Lanka, 2022, doi: <https://core.ac.uk/download/577886396.pdf>