

Smart Agriculture System for Pomegranate Disease Detection and Yield Prediction

¹ Hanamant R Jakaraddi, ¹ Tejaswini P B, ² Aditya U Diwan

¹ Department of MCA, Acharya Institute of Technology, Bangalore, Karnataka, India-560107

² Department of Computer Application, Acharya Institute of Graduate Studies, Bangalore, Karnataka, India-560107

DOI: <https://doi.org/10.51583/IJLTEMAS.2025.1408000077>

Abstract— This project is about using deep learning to find and identify different diseases in pomegranate fruits. The main goal is to help farmers catch diseases early and protect their crops to get better harvests. We used a dataset of around 2,000 pictures of pomegranates. These images were divided into five groups: Healthy, Alternaria, Anthracnose, Bacterial Blight, and Cercospora. Before training, we prepared the images by resizing them, normalizing the colors, and adding some changes to make the model learn better (this is called data augmentation). We trained two models: one was a simple CNN (Convolutional Neural Network) that we made ourselves, and the other was a pre-trained DenseNet121 model that we fine-tuned. Both models were trained for 20 rounds (called epochs). The DenseNet121 model did better, getting 94.2% accuracy on the validation set, while the custom CNN got 87.5%. We also made a web application using Flask, where users can upload images of pomegranates and get real-time results. The app tells whether the fruit is healthy or has a disease, gives tips on how to deal with the disease, and shows an estimate of how much crop might be lost.

Keywords — Deep Learning, Convolutional Neural Networks (CNN), DenseNet121, Image Classification, Precision Agriculture.

I. Introduction

Pomegranate (*Punica granatum* L.) is a popular fruit grown in many parts of the world. It is known for its great taste, health benefits, and use in cooking [1]. But growing pomegranates comes with challenges, especially when it comes to diseases. Some common diseases like Alternaria fruit rot, Anthracnose, Bacterial Blight, and Cercospora leaf spot can damage the leaves, stems, and fruits [2]. These diseases can cause a lot of damage, reducing both the quality and quantity of the fruit.

To keep pomegranate farming successful, it's very important to spot these diseases early. This helps farmers take action before things get worse. Usually, experts check the plants by looking at them carefully. But this method takes a lot of time and may not always be right, especially in big farms [2]. Sometimes, if a disease is found too late, it spreads quickly and becomes harder to control. Farmers may then have to use more chemicals, which is not good for health or the environment.

With today's need for safer and smarter farming, better ways to detect diseases are required. That's why this study focuses on building an automatic and accurate system to find diseases in pomegranate fruits custom-built Convolutional Neural Network (CNN), and the other is a pre-trained model called DenseNet121. The goal is to see which model works better in spotting diseases from images of pomegranates [2].

The project doesn't stop there. To make things easier for farmers, the models are added to a simple web application. This app can give instant disease predictions by uploading fruit images. It also shares helpful tips on how to prevent the disease and how much yield might be lost. This tool can help farmers make better decisions, reduce crop loss, and support eco-friendly farming practices

II. Literature Survey

Narayanan et al. (2022): This paper presents a hybrid convolutional neural network (CNN) model for the classification of banana plant diseases. The approach combines multiple CNN architectures to improve accuracy and robustness in identifying disease symptoms from leaf images. The model is evaluated on a large dataset of banana leaf images, achieving promising results in plant disease diagnosis. [1]

Abouelmagd et al. (2024): This paper introduces a reshaped capsule neural network for tomato leaf disease classification. The capsule network's ability to handle spatial hierarchies in image data is leveraged to improve classification accuracy. The authors demonstrate the effectiveness of the model using a dataset of tomato leaf images, showing significant improvements over traditional CNN methods. [2]

Ahmad et al. (2021): This study investigates the use of color and texture features from leaf images for plant disease identification. By combining these features with machine learning algorithms, the authors propose a robust method for diagnosing plant diseases. The approach shows improved accuracy in identifying disease types in various crops, particularly

tomatoes. [3]

Ahmed et al. (2023): This review covers the current progress, applications, and challenges of deep learning modelling techniques. The paper examines a variety of deep learning models, their advantages, and their real-world applications, particularly in fields like plant disease detection. Challenges such as model interpretability and data limitations are also discussed. [4]

Alzubaidi et al. (2021): This extensive review explores deep learning concepts and CNN architectures, focusing on their challenges and applications. The authors provide an overview of state-of-the-art techniques, including their role in image recognition and disease detection. The paper concludes with insights on future directions in deep learning for various domains, including agriculture. [5]

Angel et al. (2021): This paper provides a deep learning-based approach for disease detection in tomatoes. Using convolutional neural networks, the authors propose a model for automatic identification of tomato plant diseases from leaf images. The approach offers high accuracy and can be applied in real-time farming environments. [6]

Chakrabarty et al. (2024): The authors propose interpretable fusion model combining lightweight CNN & transformer architectures for rice leaf disease identification. This hybrid model aims to improve diagnostic accuracy while maintaining computational efficiency, making it suitable for resource-constrained environments. [7]

Sharath et al. (2020): This paper discusses disease detection in plants with convolutional neural networks (CNNs). The authors explore how CNN-based models can effectively classify plant diseases based on visual symptoms and provide results from experiments conducted on various plant species, including detection accuracy and limitations. [8]

Diana et al. (2023): This paper presents a convolutional capsule network for image-based disease differentiation in grape leaves. The proposed model enhances the ability to capture fine-grained spatial features of grape leaf diseases, resulting in improved classification performance compared to conventional CNNs. [9].

System Architecture




Fig-1: Pomegranate Disease Detection System-Simplified Architecture

This diagram explains how a disease prediction web app works. The app is built using Flask, which is a tool for creating websites, and it uses machine learning and data tools to help detect diseases from images. At the heart of the app is the Web App part, made up of Flask Core and Route Handlers. These parts take care of what the user does on the website and decide how to handle their requests.

When someone uses the app and uploads a picture, their request goes to the next layer called ML & Data Services. This layer processes the image, makes predictions about the disease, and then shows information about that disease to the user. The image first goes through the Image Preprocessor, which uses helpful tools called Image Utils to prepare the picture for analysis. Then, a machine learning model, like a Custom CNN or DenseNet121, studies the image to find signs of disease. This part is called the Prediction Engine.

After the app figures out what disease might be in the image, the Disease Info Handler takes over. It looks up details about the disease and sends that information back to the user so they understand what it means. All the data, including the uploaded images, predictions, and disease details, are saved and managed using MongoDB, which is a database system.

The app also has several helper tools grouped under Utilities. These include the File Handler, which helps save files; Image Utils, which fixes or adjusts images if needed; and Validation Utils, which checks if the user's input is correct and makes sure everything works smoothly. Additionally, there is a Storage section that holds all uploaded images, static files (like web page designs), and model files. Keeping these files ready helps the app run faster and more efficiently.

All these parts work together in a smart and organized way to create a complete system. This system can predict diseases accurately and is built to be easy to update or improve in the future.

III. Methodology




Fig-2: Workflow Diagram

The pomegranate disease detection system employs a sophisticated deep learning pipeline integrating mathematical optimization and computer vision techniques. The methodology encompasses systematic data preprocessing, model architecture design, and probabilistic classification frameworks.

Data Preprocessing and Augmentation:

The preprocessing pipeline standardizes input images through pixel normalization and geometric transformations. The normalization formula ensures consistent input scaling:

$$X_{norm} = \frac{X_{pixel} - \mu}{\sigma}$$

where X_{pixel} represents raw pixel intensities, μ is the mean pixel value across the dataset, and σ denotes standard deviation. This z-score normalization centers data around zero with unit variance, accelerating convergence during gradient descent optimization.

Convolutional Neural Network Architecture:

The custom CNN implements hierarchical feature extraction through convolutional operations. Each convolution layer applies the mathematical transformation:

$$F_{output}(i,j) = \sigma \left(\sum_{n=1}^m n \cdot K(m,n) + b \right)$$

$$0M - 1 \sum_{n=1}^n n = 0N - 1I(i + m, j + n)$$

where F_{output} represents the feature map at position (i,j) , I is the input tensor, K denotes the learnable kernel weights, b is bias term, & σ represents the ReLU activation function. This operation extracts spatial patterns across six progressive convolutional blocks, each followed by max-pooling for dimensional reduction.

Transfer Learning and Classification:

The DenseNet121 model leverages pre-trained ImageNet weights, implementing dense connectivity patterns for enhanced gradient flow. The final classification employs softmax probability distribution

$$P(y_i | x) = \frac{e^{z_i}}{\sum_{j=1}^C e^{z_j}}$$

where $P(y_i | x)$ represents the probability of class i given input x , z_i is the logit score class i , and $C=5$ represents the total disease categories.

Training Optimization:

The Adam optimizer minimizes sparse categorical cross entropy loss using adaptive learning rates. Data augmentation techniques (horizontal/vertical flips, rotations up to 0.2 radians) artificially expand the training corpus, improving model robustness. The 80-10-10 dataset partitioning strategy ensures unbiased evaluation while dropout layers (rates 0.2-0.3) prevent overfitting. Both models train for 20 epochs with batch size 16, balancing computational efficiency with convergence stability.

IV. Results and Discussion

The implementation and evaluation for the deep learning models for pomegranate disease classification yielded significant results, demonstrating the effectiveness of the proposed system. The performance of both the custom CNN and the DenseNet121 models is presented and discussed in detail below.

The training and validation processes for both models were monitored over 20 epochs. The DenseNet121 model demonstrated superior performance compared to the custom CNN architecture. The final validation accuracy for the custom CNN model was 87.5%, while the DenseNet121 model achievement was significantly higher validation accuracy of 94.2%. This performance gap ($P = .001$) highlights the advantage of leveraging shifts learning with a pre-trained model on large dataset such as ImageNet, which provides a robust foundation for feature extraction.

A detailed analysis using a confusion matrix and classification report revealed the strengths & weaknesses of the DenseNet121 model, which was selected as the primary model for the web application due to its higher accuracy. The overall precision, recall, and F1-score for the DenseNet121 model were 0.94, 0.94, and 0.94, respectively, indicating a balanced performance across all classes.

Table 1 presents the class-specific performance metrics for the DenseNet121 model on the test dataset. The model performed exceptionally well in identifying the "Healthy" class (F1-score = 0.96) and "Anthracnose" (F1-score = 0.95). The "Alternaria" and "Cercospora" classes also showed strong performance (F1-scores of 0.93 and 0.94, respectively). The "Bacterial Blight" class, while still accurately classified (F1-score = 0.92), presented the most confusion, with a small number of instances misclassified as "Cercospora" and vice versa. This is likely due to the visual similarity in the lesions caused by these two fungal diseases.

Disease Class	Precision	Recall	F1-Score
Healthy	0.95	0.97	0.96
Alternaria	0.94	0.92	0.93
Anthracnose	0.96	0.94	0.95
Bacterial Blight	0.91	0.93	0.92


Cercospora	0.94	0.93	0.94
------------	------	------	------

Table-1: Performance metrics of the DenseNet121 model on the test dataset (n=200)

The high performance of the DenseNet121 model aligns with findings from similar studies in plant disease detection, where transfer learning models consistently outperform custom architectures, especially when the available dataset is not extremely large [1, 2]. The success of this approach underscores the value of pre-trained models in agricultural AI applications.

Furthermore, the system's utility extends beyond mere classification. For the predicted disease, it retrieves and displays specific prevention methods and estimated yield impacts. For instance, a prediction of "Bacterial Blight" triggers information on using disease-free planting material and copper-based sprays, along with an estimated yield reduction of 40-80%. This integration of prediction with actionable agricultural advice transforms the system from a diagnostic tool into a comprehensive decision-support system, empowering farmers with knowledge for timely and effective intervention.

Fig-3: visualizations of training accuracy and loss



The discussion of results is supported by the visualizations of training accuracy and loss (Fig. 2) and the confusion matrix (Fig. 3), which illustrate the model's learning trajectory and classification performance.




Fig-4: Confusion matrix for the DenseNet121 model on the test dataset

V. Conclusion

This study successfully developed an intelligent pomegranate disease detection system using deep learning techniques. The DenseNet121 transfer learning model achieved superior performance with 94.2% validation accuracy, significantly outperforming the custom CNN (87.5%). The Flask-based web application effectively integrates disease classification with actionable agricultural insights, providing farmers with accurate predictions, confidence scores, and comprehensive prevention strategies. The system demonstrates practical utility in early disease detection, potentially reducing crop losses and supporting sustainable agriculture practices. Future enhancements will focus on dataset expansion, mobile integration for real-time field

detection, and incorporating farmer feedback mechanisms to continuously improve diagnostic accuracy and user experience.

References

1. Narayanan, K. L.; Krishnan, R. S.; Robinson, Y. H.; Julie, E. G.; Vimal, S.; Saravanan, V.; Kaliappan, M. Banana Plant Disease Classification Using Hybrid Convolutional Neural Network. *Comput. Intell. Neurosci.* 2022, 2022, 1–13. DOI: 10.1155/2022/9153699.
2. L.M. Abouelmagd, M.Y. Shams, H.S. Marie, A.E. Hassanien An optimized capsule neural networks for tomato leaf disease classification *Eurasip J. Image Video Process.*, 2024 (1) (2024), p. 2
3. N. Ahmad, H.M. Shahzad Asif, G. Saleem, M. Younus Usman, S. Anwar, M. Anjum Riszwan Leaf image-based plant disease identification using colour and texture features *Wirel. Pers. Commun.*, 121 (2) (2021), pp. 1139-1168, [10.1007/s11277-021-09054-2](https://doi.org/10.1007/s11277-021-09054-2)
4. S. Ahmed, M.S.B. Alam, M. Hassan, M.R. Rozbu, T. Ishtiaq, N. Rafa, *et al.* Deep learning modelling techniques: current progress, applications, advantages, and challenges *Artif. Intell. Rev.*, 56 (11) (2023), pp. 13521-13617
5. L. Alzubaidi, J. Zhang, A.J. Humaidi, A. Al-Dujaili, Y. Duan, O. Al-Shamma, *et.al.* Review of deep learning: concepts, CNN architectures, challenges, applications, future directions *J. Big Data*, 8 (2021), pp. 1-74
6. S. Angel, E. Mary, M. Dikshna, B. Athisaya, A.D. Andrushia, M. Neebha Deep learning based disease detection in tomatoes In 2021 3rd International Conference on Signal Processing and Communication (ICPSC), IEEE (2021), pp. 388-392, [10.1109/ICSPC51351.2021.9451731](https://doi.org/10.1109/ICSPC51351.2021.9451731)
7. Chakrabarty, S.T. Ahmed, M.F.U. Islam, S.M. Aziz, S.S. Maidin An interpretable fusion model integrating lightweight CNN and transformer architectures for rice leaf disease identification *Eco. Inform.*, 82 (2024), Article 102718
8. Sharath, D. M.; Kumar, S. A.; Rohan, M. G., Suresh, K.V.; Prathap, C. Disease Detection in Plants Using Convolutional Neural Network. Third International Conference on Smart Systems and Inventive Technology (ICSSIT), IEEE: India, 2020; pp 389–394.
9. A.Diana, Andrushia, T. Mary Neebha, A. Trephena Patricia, S. Umadevi, N. Anand, A. Varshney Image-based disease classification in grape leaves using convolutional capsule network *Soft. Comput.*, 27 (3) (2023), pp. 1457-1470.
10. Z. Guo, Y. Ni, H. Gao, G. Ding, Y. Zeng A dual-branch selective attention capsule network for classifying kiwifruit soft rot with hyper spectral images *Sci. Rep.*, 14 (1) (2024), p. 10664