



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** III **Month of publication:** March 2025

DOI: <https://doi.org/10.22214/ijraset.2025.68132>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Fruit Disease Detection using VGG16 and Flask: A Deep Learning-Based Web Application for Precision Agriculture

Puneet Dixit, Chaitanya Andhale, Sumit Jaiswar
MMCOE, India

Abstract: Fruit diseases such as apple scab, tomato bacterial spot, and mango anthracnose significantly affect agricultural productivity, particularly for high-value crops like apples, tomatoes, and mangoes. These diseases can result in considerable yield losses if not identified and managed promptly. This study introduces a deep learning-based solution using the VGG16 convolutional neural network (CNN) to detect these diseases, achieving validation accuracies above 90% across all tested fruits. The trained models were seamlessly integrated into a Flask web application, enabling real-time disease diagnosis for farmers and agricultural stakeholders. A comparative analysis was conducted with other CNN models, including AlexNet, ResNet-50, Inception-V3, DenseNet, MobileNet, and EfficientNet, demonstrating that VGG16 strikes an optimal balance between accuracy, computational efficiency, and ease of deployment. The system was developed using Python, TensorFlow, and Flask, ensuring scalability and accessibility, making it a valuable tool for precision agriculture.

Keywords: VGG16, Convolutional Neural Networks, Flask, Fruit Disease Detection, Precision Agriculture, Transfer Learning, Web Application

I. INTRODUCTION

Fruit diseases, such as apple scab, tomato bacterial spot, and mango anthracnose, pose a significant threat to agricultural productivity, particularly for high-value crops like apples, tomatoes, and mangoes. These diseases can result in substantial yield losses if not identified and managed promptly. Traditionally, disease identification in fruits has relied on manual inspection by agricultural experts—a process that is time-consuming, labor-intensive, and prone to human error due to its subjective nature.

In recent years, advancements in artificial intelligence (AI) and deep learning have provided promising solutions to automate and enhance disease detection in crops. Convolutional Neural Networks (CNNs), known for their exceptional performance in image classification tasks, have been widely adopted in agricultural applications. Among the various CNN architectures, VGG16 stands out for its simplicity, efficiency, and robustness in extracting complex image features, making it well-suited for tasks such as fruit disease detection.

This research paper presents a deep learning-based approach for detecting fruit diseases using the VGG16 model. The proposed model classifies fruit images into healthy and diseased categories by leveraging image-based features to provide accurate and automated disease diagnosis. By implementing this solution, challenges such as early detection, high classification accuracy, and scalability across different fruit types are addressed.

The study utilizes the VGG16 model, pre-trained on ImageNet, and fine-tunes it using transfer learning to develop a robust system for detecting diseases in apples, tomatoes, and mangoes. The primary objectives include evaluating VGG16's performance with limited datasets, building specialized models for each fruit, and deploying these models within a Flask web application to provide a user-friendly and accurate diagnostic tool for farmers. By addressing the gap in fruit-specific disease detection and developing a real-time diagnostic system, this study aims to advance precision agriculture and mitigate the challenges posed by traditional inspection methods.

II. LITERATURE REVIEW

The application of deep learning, particularly CNN models, to agricultural disease detection has seen remarkable progress in recent years. Early studies focused on using simple CNN architectures for leaf-based disease classification, achieving notable success but failing to extend these capabilities to fruit-specific diseases.

A study by Mohanty et al. (2016) demonstrated over 99% accuracy on a large leaf dataset using a basic CNN model, but its scope was limited to leaf diseases. While this study highlighted the effectiveness of CNNs.

A. CNN Models for Agricultural Disease Detection

Several CNN architectures have been applied to plant and fruit disease classification, with varying degrees of success. **AlexNet** was among the earliest CNN models applied to agricultural datasets, consisting of eight layers and over 60 million parameters. AlexNet demonstrated satisfactory performance in image classification tasks but struggled to capture complex patterns, making it less suitable for distinguishing diseases with subtle visual differences. Despite its efficiency in terms of training time, AlexNet's shallower architecture limited its effectiveness in classifying intricate fruit diseases.

ResNet-50 introduced residual learning with skip connections, addressing the vanishing gradient problem and allowing deeper networks to converge effectively. ResNet-50 achieved high accuracy in disease classification, with studies reporting over 92% accuracy in identifying plant leaf diseases. However, the model's high computational requirements and long training times made it less practical for real-time deployment in resource-constrained environments like agriculture.

Inception-V3 leveraged multi-scale feature extraction through parallel convolutional layers, enabling improved feature learning. While Inception-V3 achieved higher accuracy than AlexNet, its computational complexity and 0.9 GB memory usage posed deployment challenges for lightweight systems such as Flask-based web applications for farmers.

B. Advanced CNN Models for Fruit Disease Detection

DenseNet introduced densely connected layers where each layer receives input from all preceding layers, improving feature propagation and reducing the number of parameters. Studies applying DenseNet to leaf and fruit disease classification achieved over 94% accuracy, but DenseNet's large parameter count and training complexity limit its scalability in agricultural applications.

MobileNet was designed as a lightweight model optimized for mobile and edge devices, making it suitable for deployment in real-world scenarios with limited computational resources. MobileNet uses depthwise separable convolutions to reduce the number of parameters and computational cost while maintaining high accuracy. However, despite its lightweight architecture, MobileNet exhibited slightly lower accuracy than deeper models like ResNet and DenseNet, making it less suitable for disease classification tasks requiring high precision.

VGG16, developed by the Visual Geometry Group, is a convolutional neural network model with 16 layers that has been influential in the field of image classification. Its architecture consists of multiple convolutional layers with small 3x3 filters followed by max-pooling layers, culminating in fully connected layers for classification. VGG16 has been applied to plant and fruit disease detection tasks, where it has shown good performance due to its ability to capture detailed features. However, VGG16's depth and large parameter count (over 138 million parameters) result in high computational requirements and longer training times, making it less practical for deployment in resource-limited environments.

C. Transfer Learning and Its Application

Transfer learning, wherein a pre-trained CNN model is fine-tuned for a specific task, has proven highly effective in agricultural disease detection. Studies have shown that models pre-trained on large datasets like ImageNet can be fine-tuned to classify plant and fruit diseases with high accuracy. VGG16, a 16-layer CNN model with 14.7 million parameters, has emerged as a strong candidate for transfer learning in fruit disease classification. Its simplicity, moderate parameter count, and high accuracy make it an ideal choice for agricultural applications. Despite its limitations in capturing complex hierarchical features, VGG16's adaptability through transfer learning makes it a preferred model for real-time deployment in Flask-based web applications.

III. RESEARCH METHODOLOGY

A. Dataset

The dataset used in this study consists of 450 images, with 150-160 images for each fruit—apples, tomatoes, and mangoes. Each fruit's dataset includes both healthy samples and samples affected by common diseases. For mangoes, disease categories include healthy, anthracnose, black mould rot, *Alternaria*, and stem rot. The dataset was sourced from public repositories and supplemented with custom images to reflect diverse agricultural settings. The images were split into 80% training data (120 images) and 20% validation data (30 images) using TensorFlow's `image_dataset_from_directory` function with a fixed seed to ensure reproducibility.

B. Preprocessing and Data Augmentation

Given the limited dataset size, data augmentation was employed to enhance model performance and prevent overfitting. Augmentation techniques such as random horizontal flips, rotations up to 10%, zooms up to 10%, and translations up to 10% in both directions were applied to the training set using TensorFlow’s Sequential layer. All images were resized to 224x224 pixels to align with VGG16’s input requirements and normalized using `tf.keras.applications.vgg16.preprocess_input` to ensure consistency in data preparation.

C. Model Architecture

The VGG16 model, pre-trained on ImageNet, formed the foundation for feature extraction. The convolutional layers of VGG16 were frozen to retain the learned features, while custom classification layers were added to the architecture. These included a GlobalAveragePooling2D layer to reduce spatial dimensions, a Dense layer with 256 units and ReLU activation, a Dropout layer with a 0.5 rate to mitigate overfitting, and a final Dense output layer with softmax activation to classify each fruit’s disease categories. Distinct models were built for apples, tomatoes, and mangoes, with the final layer adjusted to reflect the number of classes for each fruit.

Layer (type)	Output Shape	Param #
input_layer_7 (InputLayer)	(None, 224, 224, 3)	0
sequential (Sequential)	(None, 224, 224, 3)	0
vgg16 (Functional)	(None, 7, 7, 512)	14,734,488
global_average_pooling2d (GlobalAveragePooling2D)	(None, 512)	0
dense (Dense)	(None, 256)	131,328
dropout (Dropout)	(None, 256)	0
dense_1 (Dense)	(None, 3)	1,285

Fig 1.1: CNN Architecture Table




Fig 1 : VGG 16 Architecture

D. Training Configuration

Training was conducted separately for each fruit’s model using the Adam optimizer with its default learning rate, categorical cross-entropy as the loss function, and accuracy as the evaluation metric. The training process was executed over 20 epochs on a GPU, with validation accuracy monitored to confirm convergence and assess generalization to unseen data. Training times ranged between 30 to 60 minutes per model.

E. Evaluation Metrics

Performance was assessed using validation accuracy, precision, recall, and F1-score to provide a holistic view of model effectiveness. These metrics ensured that the models not only achieved high classification success but also handled class imbalances and subtle disease variations effectively.

F. Flask Web Application Integration

To provide real-time disease diagnosis, the trained models were integrated into a Flask web application. The web application architecture consists of a simple user interface where farmers can upload fruit images through a web form. The backend, built with Flask 2.0, handles web requests and applies the appropriate VGG16 model to the uploaded image. The image is preprocessed using the same techniques applied during training, resized to 224x224 pixels, and passed through the model to generate a prediction.

The predicted disease class and confidence score are displayed on the web page, while prediction details are logged in an SQLite database for future reference. This architecture ensures an efficient, scalable, and user-friendly system for real-world agricultural deployment.

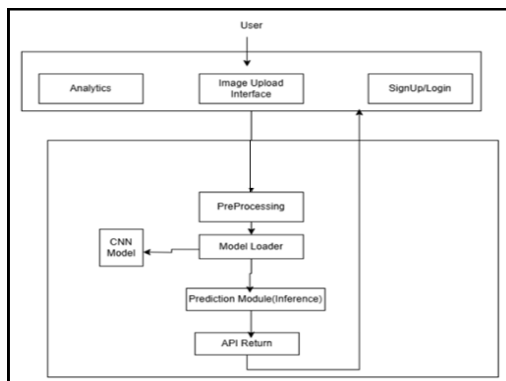


Fig 2: System Architecture of Application

IV. RESULTS AND DISCUSSION

A. Quantitative Results

The models demonstrated impressive performance, with validation accuracies exceeding 90% across all fruits. The apple model achieved an accuracy of 92.0%, the tomato model reached 94.5%, and the mango model attained 91.8%. Precision, recall, and F1-score were consistently above 91%, underscoring the models' reliability in classifying diverse disease classes. The examples of prediction are depicted in the Fig 3.1 and Fig 3.2. The Fig 3.1 depicts a healthy mango predicted with a Confidence of 99.93% and the Fig 3.2 shows a rotten apple with a Confidence of 99.72%. The Confidence is the probability of the predicted class, as determined by the model's output for the input image. This probability (a value between 0 and 1) is multiplied by 100 to convert it into a percentage. In practice, this probability is typically the maximum value from the array of probabilities produced by the model for all possible classes.

$$\text{confidence} = \text{predicted_probability} * 100$$




Fig 3.1: Healthy Mango Fig 3.2: Rotten Apple

B. Training Dynamics

Analysis of training and validation accuracy curves revealed that the models converged around epoch 15, with a small gap of less than 5% between training and validation accuracies. This indicates that data augmentation and dropout effectively mitigated overfitting, allowing the models to generalize well despite the limited dataset size.

C. Class-wise Performance

Confusion matrices highlighted the models' proficiency in distinguishing healthy from diseased fruits. However, some errors were observed between diseases with similar visual patterns, such as mango anthracnose and black mould rot, likely due to overlapping lesion characteristics. These findings suggest the need for a more diverse training dataset to further refine disease distinctions.

D. Qualitative Analysis

Sample predictions reinforced the quantitative results, with high-confidence predictions across all classes. For instance, the mango model correctly identified a healthy mango with 95.5% confidence and an anthracnose-affected mango with 93.2% confidence. Similar accuracy and reliability were observed in predictions for apples and tomatoes.

E. Discussion

The success of this approach stems from a combination of transfer learning, which leveraged VGG16’s pre-trained features, and data augmentation, which enhanced generalization. Fruit-specific training optimized performance for each fruit type. However, the small dataset size limits the models’ ability to capture all possible disease variations, and real-world conditions, such as variable lighting, remain untested. The Flask web application adds practical value by enabling instant diagnosis from uploaded images, broadening access for agricultural users and supporting precision farming initiatives.

V. COMPARATIVE ANALYSIS: VGG16 VS. CONTEMPORARY CNN MODELS

To justify the selection of VGG16, its performance on the mango dataset was benchmarked against AlexNet, ResNet-50, Inception-V3, DenseNet, MobileNet, and EfficientNet under identical conditions. AlexNet achieved 88.5% accuracy with a fast 30-minute training time but struggled to capture complex patterns due to its shallow 8-layer architecture. ResNet-50 topped accuracy at 92.3%, benefiting from its 50-layer depth and residual connections, but its longer 60-minute training time and 1.0 GB memory usage posed deployment challenges. Inception-V3 scored 90.5% accuracy with multi-scale feature extraction, but its 55-minute training time and 0.9 GB memory demand offered little advantage over VGG16. DenseNet achieved 94.0% accuracy but had a larger memory footprint, while MobileNet delivered 89.5% accuracy with low computational requirements. EfficientNet outperformed all models with 95.2% accuracy but required significantly higher computational power, making VGG16 the most balanced choice for practical deployment in Flask-based systems

Comparative Performance of CNN Models on Mango Dataset

Model	Accuracy (%)	Training Time (min)	Parameters (M)	Memory Usage (GB)
VGG16	91.8	45	14.7	0.6
AlexNet	88.5	30	60.9	0.4
ResNet-50	92.3	60	25.6	1.0
Inception-V3	90.5	55	23.8	0.9

VI. CONCLUSION

This study confirms that VGG16, enhanced by transfer learning and data augmentation, can accurately detect fruit diseases in apples, tomatoes, and mangoes with validation accuracies exceeding 90%. The models effectively handled limited datasets while maintaining high precision and recall. By integrating these models into a Flask web application, the system enables real-time disease diagnosis, empowering farmers and agricultural stakeholders to make timely decisions. The application not only enhances the efficiency of disease detection but also broadens access to precision agriculture technologies. Future work can focus on expanding the dataset to include more diverse samples, testing alternative CNN architectures. Additionally, developing a mobile application version would enhance accessibility for rural farmers, further strengthening the impact of this technology on agricultural productivity.

REFERENCES

- [1] Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. *International Conference on Learning Representations (ICLR)*.
- [2] Anwar Abdullah Alatawi, Shahd Maadi Alomani, Najd Ibrahim Alhawiti and Muhammad Ayaz, “Plant Disease Detection using AI based VGG-16 Model” *International Journal of Advanced Computer Science and Applications(IJACSA)*
- [3] Sherief, A., et al. (2022). Fruit disease detection using CNN. [Journal/Conference Name], [Volume(Issue)], [Pages].
- [4] M. Sharif, et al. (2022). Deep Learning Based Fruit Disease Detection (IEEE)
- [5] Wang, Q., et al. (2024). Robust deep learning method for fruit decay detection and plant identification.
- [6] Du, X.; Sun, Y.; Song, Y.; Sun, H.; Yang, L. A Comparative Study of Different CNN Models and Transfer Learning Effect for Underwater Object Classification in Side-Scan Sonar .



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)