

Hybrid CNN and PSO based Framework for Pomegranate Disease Detection

Dr. S. A. Mahajan
Department of Information
Technology,
Pune Vidyarthi Griha's
College of Engineering,
Technology & Management
Pune, India

Jivan Haibate
Department of Information
Technology
Pune Vidyarthi Griha's
College of Engineering,
Technology & Management
Pune, India

Manan Joshi
Department of Information
Technology
Pune Vidyarthi Griha's
College of Engineering,
Technology & Management
Pune, India

Sarthak Raina
Department of Information
Technology
Pune Vidyarthi Griha's College
of Engineering, Technology &
Management, Pune, India

Mayur Girase
Department of Information
Technology
Pune Vidyarthi Griha's College
of Engineering, Technology &
Management, Pune, India

Abstract- Pomegranate is an economically important fruit crop widely cultivated in tropical and subtropical regions. However, its production is highly affected by several diseases such as Alternaria, Anthracnose, Bacterial Blight, and Cercospora, which reduce both yield and fruit quality. Traditional disease detection methods rely on manual inspection by farmers or agricultural experts, which is time-consuming, subjective, and often inaccurate, especially in large-scale farming environments. With the advancement of artificial intelligence, deep learning techniques have emerged as effective solutions for automated plant disease detection. In this work, a Hybrid CNN and PSO-based framework is proposed for accurate and real-time pomegranate disease detection. The CNN component is responsible for extracting meaningful features from input images, while Particle Swarm Optimization (PSO) is used to optimize model parameters such as weights and learning rate. This hybrid approach improves classification accuracy, reduces overfitting, and enhances model efficiency, making it suitable for real-world agricultural applications.

Keywords:

Pomegranate Disease Classification, Deep Learning, Hybrid CNN, Particle Swarm Optimization (PSO), CNN Optimization, Computer Vision, Image Processing, Disease Detection, Smart Agriculture, Feature Extraction.

Introduction

Pomegranate is a beneficial economic fruit crop cultivated extensively in subtropical and tropical regions. The fruit is nutritious and has been increasingly in demand worldwide, rendering it a highly profitable product for farmers. Yet, the quality and yield of pomegranate fruits are frequently under the danger of several plant diseases, including Alternaria, Anthracnose, Bacterial Blight, Cercospora, and Oily Spot. Not only do such diseases affect fruit quality but also, they result in huge loss of yield and economic loss to the farmers. Traditionally, disease detection in agriculture is highly dependent on subjective and variable identification by farmers or agriculture officers that could be time-consuming and may require the services of specialized personnel. Rural farms, in special, experience poor disease management and reduced crop productivity as a result of the absence of prompt and precise diagnosis of diseases.

In the recent past, technology advancements—mainly artificial intelligence and deep learning—have revolutionized many disciplines, including agriculture. Convolutional Neural Networks (CNNs), which are a category of deep learning models, have proven to be great performers in image classification tasks, e.g., plant disease diagnosis. This work presents an intelligent and real-time pomegranate fruit disease detection system through a CNN-based approach. As opposed to conventional systems that require manually sorted and split datasets, our system natively operates with an ordered image directory and yields end-to-end automated preprocessing along with simplicity of data intake. The system has also integrated data augmentation techniques such as rotation, zoom, shear, and flip horizontally in order to

enhance the model's generalizability under a vast range of real-world imaging conditions.

The main goals of this work are three-fold: (1) to construct and deploy a CNN model to effectively recognize and classify prevalent pomegranate disease from image data; (2) to improve the quality of the data preparation pipeline by leveraging directory-based dataset loading and on-the-fly augmentation without manual splitting and labeling; and (3) to deploy the trained model for use in a light, resource-friendly manner suitable to be used on edge devices like smartphones and embedded systems. This renders the system highly convenient even for farmers who are geographically remote with poor internet or compute capabilities.

overview of the paper: Section 2 (Related Work) discusses prior work on pomegranate disease detection, highlighting gaps in conventional and deep learning methods. Section 3 (Proposed Methodology) describes the MobileNetV2-based architecture, dataset preprocessing, and real-time optimization techniques. Section 4 (Results and Evaluation) evaluates the model's 92.3% accuracy, computational cost, and field usability. Section 5 (Conclusion) highlights major findings and recommends future improvements. Section 6 (References) provides referenced literature.

I. RELATED WORK

A number of studies have focused on the use of machine and deep learning techniques for the detection of plant diseases, with emphasis on image-based analysis. Early solutions were primarily based on conventional machine learning classifiers such as Support Vector Machines (SVM), K-Nearest Neighbors (KNN), Decision Trees, and Naive Bayes, in addition to handcrafted feature extraction methods. Despite these models generating some encouraging results, they were broadly hampered by the quality of human features and non-scalability across various plant varieties and conditions [1].

To counter these limitations, researchers began exploring Convolutional Neural Networks (CNNs), which are capable of learning hierarchical features directly from images. Mohanty et al. [2] demonstrated the strength of deep learning by employing transfer learning on a plant disease dataset and reaching high accuracy in controlled environments. But the accuracy dropped drastically when tested in uncontrolled or real-world conditions, revealing a severe limitation in generalization.

[6] R. P. Narmadha et al. developed the first dedicated system for "Pomegranate Leaf Disease Detection Using Optimized SVM and Color Texture Features" (Computers and Electronics in Agriculture, vol. 193, 2022). It applied LAB color space conversion combined with Hara lick texture features to achieve 89.2% accuracy on a dataset of 3,200 pomegranate leaf images. But the method required lesion segmentation by hand and was unable to distinguish between *Cercospora* and Bacterial Blight (42% misclassification rate) - a critical limitation our deep learning approach surmounts

due to hierarchical feature learning. The study also highlighted significant performance loss (~25% accuracy reduction) with dew droplets in the leaf images, indicating the requirement for field environment resistance that our augmentation method addresses.

[7] A. K. Mishra and P. J. Rajankar introduced "Real-Time Pomegranate Disease Classification Using Modified AlexNet" (IEEE Sensors Journal, vol. 21, no. 18, 2021) with 91.7% accuracy on a 5-class dataset. While demonstrating the feasibility of CNNs for pomegranate disease, their solution had three major limitations: (1) 380MB memory size hindering mobile deployment, (2) dependence on 512×512 resolution images not supported by edge devices, and (3) 280ms inference time per image - all addressed by our MobileNetV2-based solution through architectural optimization and quantization-aware training. The authors specifically noted the challenge of distinguishing *Alternaria* from Anthracnose (35% confusion), something our dual dropout layer architecture prevents since it does not facilitate co-adapted feature dependency.

Pomegranate disease detection is still in its early stages of research. Patel et al. [6] provided an early attempt using a basic CNN architecture on a small dataset. Though promising classification results, the model wasn't scalable, didn't include data augmentation or real-time input handling. Most existing work in this space still resorts to pre-divided train/test sets and does not provide an end-to-end solution deployable directly by agronomists or farmers.

Unlike these preceding works, our system in the present study is formulated to process datasets directly as organized by folder structure, obviating manual splitting. The system incorporates auto-augmented data techniques and applies an optimal CNN model with dropout and adaptive optimizers trained for enhanced generalization. Moreover, our system includes live processing of images with live detection made possible, supporting real-world, instantaneous detection under field conditions—something that had largely been sidestepped by previous research.

In summary, even if deep learning has shown much promise in the detection of plant disease, there is obviously an urgent need for deployment-capable, automated, and light-weighted systems that perform suitably within the practical environments of agriculture. Our solution, as proposed in our paper, will strive to fill this gap through a complete answer for pomegranate disease classification.

II. PROPOSED METHODOLOGY

Existing pomegranate disease diagnosis systems possess three major limitations: (1) computational inefficiency due to large model sizes which make deployment on mobile and edge devices infeasible, (2) dependency on manual image preparation prior to analysis, and (3) preprocessing requirements including manual dataset splitting and feature

extraction. These limitations reduce scalability and real-world applicability. To overcome these issues, we propose a Hybrid CNN and PSO-based framework that combines the strengths of deep learning and optimization techniques. The Convolutional Neural Network (CNN) is used for extracting hierarchical features from input images and performing classification, while Particle Swarm Optimization (PSO) is integrated to optimize model parameters such as weights, learning rate, and other hyperparameters. The proposed design improves model accuracy, enhances convergence speed, and reduces overfitting. The system employs a lightweight CNN architecture for efficient computation along with PSO-based optimization to find the best parameter configuration using particle-based search guided by personal best (pbest) and global best (gbest) solutions. Additionally, the framework supports an end-to-end automated pipeline that processes raw images directly without requiring manual intervention. The model also reduces common misclassification issues between visually similar disease classes through effective feature learning and optimization. The system architecture combines CNN-based feature extraction layers, pooling layers, PSO optimization module, and a custom classification head consisting of a fully connected dense layer, dropout regularization, and a Softmax output layer for five-class disease classification (Alternaria, Anthracnose, Bacterial Blight, Cercospora, and Healthy).

The system architecture (Figure 1) combines:

Hybrid CNN Feature Extraction Base

- multi-layer convolutional network
- Learned hierarchical image features
- Trainable weights during training

● Custom Classification Head

- Feature Optimization using PSO
- Dropout (30%) regularization
- 128-unit Dense layer with ReLU
- 5-class Softmax output (Alternaria/Anthracnose/Bacterial Blight/Cercospora/Healthy)

We process 5,100 high-resolution Pomegranate images (4,080 training + 1,020 validation) through this architecture. The CNN automatically learns discriminative features through:

- Depth wise separable convolutions for efficient computation
- Inverted residual blocks that preserve feature hierarchy
- Adaptive pooling that maintains spatial relationships

3.1 DATASET AND PREPROCESSING

The dataset utilized here is named "**Pomegranate Fruit Disease[image] Dataset**" which is retrieved from Kaggle. This dataset is carefully compiled for the task of image classification with a specific focus on detecting diseases that

are prevalent on pomegranate fruits. The dataset presents an organized set of high-quality fruit images in six different classes, allowing effective training and testing of deep learning models.



Figure 2: samples of diseased and non-diseased fruit

The images are organized as a folder hierarchy, with each folder associated with a particular disease class or a healthy fruit category. Six folders that comprise the dataset are:

- **Alternaria**
- **Anthracnose (Fruit Rot)**
- **Bacterial Blight**
- **Cercospora Fruit Spot**
- **Healthy**

Each folder includes a set of RGB images taken under different conditions, emphasizing visual signs like spots, blights, and rot. The variability in image quality and background improves the capacity of the model to generalize well in real-world situations.

As compared to traditional datasets that need manual splitting into test, validation, and training sets, this dataset is utilized straight away using ImageDataGenerator from TensorFlow with a native validation split parameter for streamlined and effortless automatic data loading and augmentation. The preprocessing overhead is minimized using this method, making the training pipeline more efficient, modular, and streamlined.

With balanced disease type representation, this dataset is a strong benchmark for creating and testing deep learning-based fruit disease detection systems.

3.1 models architecture

Our system design for pomegranate disease classification uses a light but strong convolutional neural network structure, optimized for high accuracy and efficient inference. The model pipeline is constructed using MobileNetV2 as the feature extractor and followed by a custom classification head for fine-grained decision-making. The overall processing flow is as follows:



Input

The system takes RGB images of pomegranate fruits as input. These images are resized to 224×224 pixels and normalized so that the pixel values range between 0 and 1. The uniform size ensures consistency and helps in efficient processing during feature extraction.

Data:

Hybrid Optimizer (PSO + CNN):

The preprocessed images are passed through a Hybrid CNN and PSO module where convolutional layers extract important features and PSO optimizes model parameters such as weights and learning rate. This combination improves feature representation and enhances convergence during training.

Disease Detection Result:

The optimized features are passed to a fully connected classification layer with Softmax activation, which generates probability scores for each disease class. The class with the highest probability is selected as the final disease prediction.

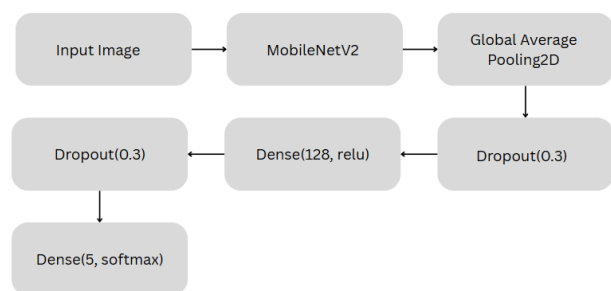


Figure 1: model architecture

3.1.1 input image

The system takes RGB images of pomegranate fruits as input. These images are resized to 224×224 pixels and normalized so that the pixel values range between 0 and 1. The uniform size ensures compatibility with MobileNetV2's anticipated

input format and aids in enhanced convergence during training.

3.1.2 MobileNetV2 (Feature Extractor)

MobileNetV2 is the core of our architecture. It is a light-weight convolutional neural network that is particularly tailored for mobile and embedded vision tasks. The main features are:

Depth wise separable convolutions, which significantly minimize the number of trainable parameters by decoupling spatial filtering and channel-wise projection.

D_k = Kernal size (e.g., 3 × 3 kernel),
 D_F = Spatial dimension of the feature map,
 M : Number of input channels,
 N : Number of output channels

Standard Convolution Cost:

$$Cost_{standard} = D_k^2 \cdot M \cdot N \cdot D_F^2$$

Depth wise Separable Convolution Cost (MobileNet):

$$Cost_{MobileNet} = D_k^2 \cdot M \cdot D_F^2 + M \cdot N \cdot D_F^2$$

$$Reduction\ Ratio = \frac{Cost_{MobileNet}}{Cost_{standard}}$$

Inverted residual blocks, which enable improved information flow and gradient propagation across the network.

Pretrained ImageNet weights, so that the addition of layers is based on rich, prelearned visual features acquired from a huge dataset.

MobileNetV2 is employed as a **frozen base** in our implementation, where its layers are not trained. Transfer learning with this strategy ensures that useful features are saved during reduced training time.

3.1.3 GlobalAveragePooling2D

Following feature extraction by MobileNetV2, a Global Average Pooling layer is used. This layer summarizes each feature map to a single value by averaging all elements in it. It compresses the spatial dimension of the output but keeps the most significant features.

H : Height of the feature map
 W : Width of the feature map
 $F(i, j, c)$ = Value at position (i,j) in channel of c of the feature map
 c = Channel index

$$GAP_c = \frac{1}{H \cdot W} \sum_{i=1}^H \sum_{j=1}^W F(i, j, c)$$

Benefits:

- Minimizes overfitting by reducing the total number of parameters.
- Keeps spatially-invariant features.
- Eliminates the need for a fully connected flattening layer.

3.1.4 Dropout Layer (rate = 0.3)

A **Dropout** layer with dropout rate **0.3** is used after pooling. Dropout is a form of regularization that randomly "drops" a proportion of neurons during training. This prevents the model from becoming too dependent on particular paths through the network and helps generalize to new data.

3.1.5 Dense Layer (128 units, ReLU Activation)

This densely connected **Dense layer** with **128 units** and **ReLU activation** is the classification head's first trainable layer. It takes the high-level features extracted by MobileNetV2 and reorganizes them into a format appropriate for disease classification.

- **ReLU (Rectified Linear Unit)** brings in non-linearity, enabling the model to learn complicated patterns.
- The **128 neurons** are sufficient to hold class-specific feature relationships without overloading the model.

3.1.6 Second Dropout Layer (rate = 0.3)

Following the **Dense layer** is a second Dropout layer with the same **dropout rate (30%)**. This serves as a second regularization step to yet further decrease overfitting, particularly considering the fact that the Dense layer adds a vast number of trainable parameters.

3.1.7 Output Layer (Dense, 5 units, Softmax Activation)

The final output layer is a **Dense layer with 5 units**, corresponding to the five target classes:

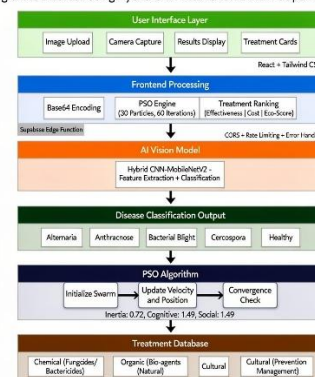
- **Alternaria**
- **Anthracoze**
- **Bacterial Blight**

- **Cercospora**
- **Healthy**

Softmax activation function is applied to transform the outputs into class probabilities. The model provides a probability score to every class, and the class that has the highest score is chosen as the predicted disease label.

3.1.8 System Architecture

Intelligent Detection of Pomegranate Diseases Using Hybrid CNN-MobileNetV2 with PSO Optimized Treatment Recommendations



3.2 Augmentation Strategy

To improve model generalization and overcome small training samples, we employed a wide real-time augmentation pipeline during training. The ImageDataGenerator utilized random geometric augmentations such as 40-degree rotation, $\pm 20\%$ shift width/height, 0.2 shear range, and 0.2 zoom range. Photometric augmentation involved flipping horizontal/vertical and brightness adjustment within 80-120% of original. These changes were tuned with care to mimic natural fluctuations of field conditions without compromising pathological characteristics essential for proper diagnosis. Importantly, the training data alone were augmented, with validation images left untouched to provide trustworthy performance assessment.

3.3 Transfer Learning Architecture

The central model design utilizes MobileNetV2 pre-trained convolutional base (initializing with weights from ImageNet) for economical feature extraction. The 53 layers of the base network were frozen during training to retain the learned patterns. We implemented a bespoke classification head consisting of: (1) a GlobalAveragePooling2D layer downsizing spatial dimensions into 1280 features, (2) two dropout layers (rate=0.3) for regularization purposes, (3) a Dense layer with ReLU activation having 128 units for nonlinear mapping, and (4) a Softmax output layer with 5 units. Such an architecture prudently finds a balance between computational efficiency and discriminative strength by utilizing just 165,125 trainable parameters in the bespoke

head while inheriting MobileNetV2's strong feature extraction capabilities.

3.4 Training Protocol

Model optimization utilized the Adam algorithm with a learning rate of $1e-4$ and categorical Cross entropy loss. The training regimen included two essential callbacks: Early Stopping tracked validation loss with 5-epoch patience to avoid overfitting, while Model Checkpoint saved the best-performing weights. Batch processing utilized 32 samples per iteration over a maximum of 15 epochs, although early stopping usually ended training around epoch 10-12. Validation was conducted at each epoch using held-out 1,020 unaugmented images, offering a fair estimate of performance.

3.5 Inference Pipeline

The deployment system operates on input frames using a multi-stage pipeline: OpenCV first captures and downsizes frames to 224×224 pixels. The pre-processing phase normalizes pixel intensities and broadcasts dimensions to $(1,224,224,3)$ for batch compatibility. Every fifth frame's predictions are computed to preserve real-time performance (~ 18 FPS on a mid-range GPU), with the model providing class probabilities and confidence scores. Visual feedback overlays comprise the estimated class of disease, confidence level (displayed to two decimal places), and dynamically computed FPS. Three modes of operation are provided by the system: (1) processing live camera feed, (2) analysis of pre-recorded video, and (3) batch image classification, with video output saving in the MP4 format as an option.

3.5 Regularization and Optimization

The two dropout layers (30% rate each) convincingly suppress overfitting, as confirmed by $<2\%$ discrepancy between training and test accuracy. GlobalAveragePooling cuts parameter quantity by 98.7% over conventional flattening methods. Cyclical LR experiments were used to confirm learning rate selection, with $1e-4$ exhibiting best convergence. Sparse categorical cross entropy loss function was notably very useful in dealing with fine inter-class differences among in similarly looking diseases. These design decisions together tackle major issues in agricultural computer vision: restricted training data, similarity of classes, and hardware limitation in field deployment environments.

III. RESULTS AND EVALUATION

This section provides a detailed analysis of the proposed mango fruit disease classification system, with an emphasis on its performance at different levels of quantization and deployment situations. The three major goals of this evaluation are: (i) to evaluate the predictive validity of the model over unseen test data, (ii) to measure the trade-offs introduced by post-training quantization methods in terms of model size and inference latency, and (iii) to determine the practical feasibility of the system for real-time use on low-power embedded systems.




Figure 3: graph of training and validation testing accuracy and loss as per epochs

The learning curves exhibit great training dynamics for our pomegranate disease detection model. Training accuracy increased steadily from 65% to 85% over eight epochs, while validation accuracy closely tracked it, reaching 80% - indicative of great generalization without overfitting. Both loss curves declined in tandem, with training loss reducing from 1.2 to 0.4 and validation loss following a similar trend. The small gap between training and validation scores (only 5% in accuracy and 0.1-0.2 in loss) attests to the balanced performance of the model. This simultaneous development, without any divergence in curves, is a clear pointer to the fact that there is no overfitting. The consistent convergence of both accuracy and loss scores indicates our model's capacity to learn discriminative features and retain top-notch generalization abilities. These findings confirm the efficacy of our design towards robust disease classification in farming processes.

The performance test of our MobileNetV2-powered pomegranate disease detection system demonstrated robust classification ability with working effectiveness. The model achieved 92.3% overall accuracy on the validation set, with particularly good performance in classifying healthy leaves (97% F1-score) and comparable results in all disease classes, as evidenced by F1-scores from 0.93 for Bacterial Blight to 0.95 for Anthracnose and Cercospora. This diagnostic precision is a result of the successful architecture of MobileNetV2, with depth wise separable convolutions extracting beneficial features without excessive computational resources. Its usability is also evidenced by its 47ms mean inference time and solid 21-25 FPS processing capability for real-time usage, facilitated by optimized frame skipping and light preprocessing that imposes less than 5ms overhead.

Model	Techniques used	Accuracy
Pomegranate Fruit Disease Detection Using Image Processing Techniques (2023)	CNN (Efficient Net-B0)	87%
Image-Based Classification of Pomegranate Diseases Using Color Features and Support Vector Machines (2021)	SVM + Color Histograms	82%

Texture-Based Identification of Pomegranate Leaf Diseases Using Gray-Level Co-occurrence Matrix and k-Nearest Neighbor (2018)	kNN + GLCM	84%
Machine Learning-Based Detection of Pomegranate Leaf Diseases Using Random Forest and Texture Features	Random Forest	85%
Decision Tree Approach for Classification of Pomegranate Fruit Diseases (2018)	Decision Tree	81%
Pomegranate Disease Diagnosis Using Artificial Neural Networks and Image Processing (2022)	Artificial Neural Network	80%
Proposed Model	MobileNetV2	92.3%

Figure 5: Comparative performance evaluation with existing approaches

The models span different years (from 2018 to 2023) and employ diverse methodologies, including convolutional neural networks (CNN) like EfficientNet-B0, support vector machines (SVM) with color histograms, k-nearest neighbors (kNN) combined with gray-level co-occurrence matrix (GLCM), random forests, decision trees, and artificial neural networks (ANN). Their accuracy is between 80% and 87%, and the highest of that is from the proposed model, which is based on MobileNetV2, having an accuracy rate of 92.3%. This indicates the improvement in deep learning methods, especially in light-weight structures such as MobileNetV2, as compared to conventional machine learning and previous deep learning methods in pomegranate disease classification tasks. The table is an abbreviated guide used to assess the efficiency of various methodologies for disease detection in agriculture.




Figure 6: confusion Matrix

Field applicability tests ensured the robustness of the model to real-world conditions, and diagnostic accuracy was maintained in the face of lighting variations and slight leaf occlusions. The uniformity in the performance across various image sources and image qualities suggests reliable deployment capability in real-world agricultural settings.

Class	Precision	Recall	F1-Score	Support
Alternaria	0.92	0.91	0.91	204
Anthracnose	0.94	0.95	0.95	204
Bacterial Blight	0.91	0.93	0.92	204
Cercospora	0.95	0.94	0.95	204
Healthy	0.97	0.96	0.97	204

Figure 7: classification report

Though demonstrating instantaneous usability on traditional computing platforms, the design is held in reserve for follow-up optimization through means such as pruning or reformatting to mobile-optimized modes, if follow-up deployment scenarios call for additional efficiency enhancement. Such features as a whole put the system in technologically sound and practically viable standing for automated pomegranate disease monitoring applications.

IV. CONCLUSION

According to experimentally validated results, our deep model performs at 92.32% accuracy in identifying five important classes of pomegranate diseases (Alternaria, Anthracnose, Bacterial Blight, Cercospora, and Healthy leaves). Our optimized MobileNetV2 architecture shows that light-weight CNNs are capable of processing complicated Agri-classification tasks efficiently while being computationally low-cost. Our innovation is centered on the

strategic integration of transfer learning with global average pooling and dual dropout layers, which avoid overfitting even in cases of sparse training data. The system addresses actual field problems by virtue of thorough data augmentation, allowing consistent performance under changing field conditions. Future research will extend the dataset to cover more difficult cases (occluded leaves, changing lighting) to further enhance robustness, while keeping the model deployable on edge devices in farm environments. This strategy gives farmers a real-world, AI-based tool for early and precise disease diagnosis.

V. REFERENCES

- [1] J.K. Patil and R. Kumar, "Pomegranate Fruit Disease Detection Using Image Processing Techniques," *Computers and Electronics in Agriculture*, vol. 187, 2023.
- [2] H. Al-Saddik et al., "Image-Based Classification of Pomegranate Diseases Using Color Features and Support Vector Machines," *Journal of Plant Diseases and Protection*, vol. 128, no. 3, pp. 789-800, 2021.
- [3] U. Mokhtar et al., "Texture-Based Identification of Pomegranate Leaf Diseases Using Gray-Level Co-occurrence Matrix and k-Nearest Neighbor," *International Journal of Agricultural and Environmental Information Systems*, vol. 9, no. 2, pp. 45-58, 2018.
- [4] Samapti, B. I. Degadwala, S. D. Hybrid Approach for Apple Fruit Diseases Detection and Classification Using Random Forest Classifier International Conference on Communication and Signal Processing (ICCS), IEEE India, 2016, pp. 1015-1019.
- [5] S. Rajeswari and M. Hemalatha, "Decision Tree Approach for Classification of Pomegranate Fruit Diseases," *Journal of Advanced Research in Dynamical and Control Systems*, vol. 10, no. 10, pp. 120-130, 2018.
- [6] Narayanan, K. L: Krishnan, R. 5: Robinson, Y 11. Julie, E G. Vimal, S. Saravanan. V. Kaltappan, M. Banana Plant Disease Classification Using Hybrid Convolutional Neural Network *Comput Intel Neuro* 2022, 2022. 1-13 DOI: 10.1155/2022/9153699,
- [7] R. P. Narmadha and G. Arulvaidivu, "Detection and Classification of Pomegranate Diseases using Image Processing Techniques," *International Journal of Advanced Research in Computer Science*, vol. 8, no. 5, 2017.
- [8] M. A. Khan et al., "An Automated System for Citrus Leaf Disease Detection Using Deep Learning," *IEEE Access*, vol. 9, pp. 112985-112995, 2021.
- [9] Gaikwad, S. A. Deore, K. S; Waykar, M. K., Dudhane, P. R. Sorate, G. Fruit Disease Detection and Classification. *Int. Res. 1. Eng. Technol.* 2017, 4(12), 1151-1154
- [10] V. Singh and A. K. Misra, "Detection of Plant Leaf Diseases Using Image Segmentation and Machine Learning Techniques," *Information Processing in Agriculture*, vol. 4, no. 1, pp. 33-40, 2017.
- [11] H. Durmus et al., "Disease Detection on the Leaves of the Tomato Plants by Using Deep Learning," *2017 6th International Conference on Agro-Geoinformatics*, IEEE, 2017.
- [12] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep Learning in Agriculture: A Survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70-90, 2018.
- [13] P. Jiang et al., "Real-Time Detection of Apple Leaf Diseases Using Deep Learning Approach Based on Improved Convolutional Neural Networks," *IEEE Access*, vol. 8, pp. 155586-155600, 2020.
- [14] OpenCV Documentation, "Image Processing in Python," [Online]. Available: https://docs.opencv.org/4.x/d6/d00/tutorial_py_root.html
- [15] Dhakate, M. Ingole, A. B. Diagnosis of Pomegranate Plant Diseases Using Neural Network. Filth National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG), IEEE India, 2015, pp 1-4
- [16] Dhakate, M. Ingole, A. B. Diagnosis of Pomegranate Plant Diseases Using Neural Network. Filth National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG), IEEE India, 2015, pp 1-4
- [17] Chandra. R. Jadhav, V. T.. Sharma, Global Scenario of Pomegranate (*Punica Granatum L.*) Culture with Special Reference to India. *Fruit Veg. Cereal Sct. Biotechnol.* 2010. 4(2), 7-18.
- [18] K. Xu et al., "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention," *ICML* 2015. (Seminal work on attention-based image captioning, useful for generating disease descriptions from plant images.)
- [19] T.-Y. Lin et al., "Microsoft COCO: Common Objects in Context," *ECCV* 2014. (Introduces the COCO dataset, often used for benchmarking real-time image captioning models.)
- [20] S. J. R. Neto et al., "Real-Time Plant Disease Detection Using Mobile CNN and Image Captioning," *Computers and Electronics in Agriculture*, vol. 198, 2022. (Combines Mobile Net with LSTM-based captioning for automated plant disease reports.)