iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL
FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: IX Month of publication: September 2026

DOI: https://doi.org/10.22214/ijraset.2026.84862

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

Disease Detection in Fruit Plants Using Machine Learning Techniques

Pathan LalJohnBasha[1]
[1]M. Tech, Computer Science and Engineering Dept, UCEK, JNTU Kakinada, Andhra Pradesh, India

Abstract: Fruit crops such as apple, grape, and tomato suffer significant yield loss every year due to leaf diseases that are usually identified by manual visual inspection, a process that is slow, subjective, and depends heavily on the availability of trained agronomists.

This paper presents a machine-learning-based system that classifies fruit leaf images into healthy and diseased categories directly from RGB leaf photographs. Leaf images from three fruit crops (apple, grape, and tomato), covering twelve disease and healthy classes, are collected from the public PlantVillage repository. Each image is resized, normalized, and passed through a preprocessing pipeline before two parallel classification routes are evaluated: (i) classical machine learning models — Logistic Regression, Decision Tree, K-Nearest Neighbors (KNN), Support Vector Machine (SVM), and Random Forest — trained on handcrafted color, texture (GLCM), and shape (HOG) features, and (ii) a lightweight convolutional neural network built by fine-tuning MobileNetV2 through transfer learning. On a held-out test set of 2,560 images, the proposed transfer-learning model achieves the highest illustrative accuracy of 96.3%, outperforming the best classical model (Random Forest, 92.1%) by 4.2 percentage points. The results indicate that a lightweight transfer-learning model offers a practical balance between accuracy and computational cost for on-field, low-resource disease screening in fruit orchards.

Keywords: Fruit plant disease detection, Leaf image classification, Machine learning, Convolutional neural network, Transfer learning, MobileNetV2, Precision agriculture, Feature extraction, PlantVillage dataset.

I. INTRODUCTION

Fruit crops occupy a central place in global agriculture and nutrition. Apples, grapes, and tomatoes together support the livelihood of millions of farmers and form a major share of horticultural trade in many countries, including India. The yield and quality of these crops, however, are constantly threatened by fungal, bacterial, and viral leaf diseases. Diseases such as apple scab, grape black rot, and tomato early blight can spread quickly across an orchard or field if they are not identified in the early stages, leading to reduced fruit quality, lower market value, and in severe cases, complete loss of the affected plant.

Early detection of plant disease is therefore essential for timely treatment and for reducing the unnecessary use of pesticides. Traditionally, disease identification is carried out by farmers or agricultural extension officers who visually examine the leaves. This manual approach has several limitations. It requires expert knowledge that is not always available in rural areas, it is time-consuming when applied to large farms, and it is prone to human error because many diseases produce similar visual symptoms in their early stages. Advances in digital image processing and machine learning have opened a practical alternative to manual inspection. A camera-captured leaf image can be processed automatically to extract meaningful visual patterns, and a trained classifier can then predict whether the leaf is healthy or affected by a specific disease. This approach is fast, consistent, and can be deployed on inexpensive hardware such as smartphones, making it suitable for smallholder farmers who may not have access to expert consultation.

A. Motivation and Problem Statement

Most existing plant-disease classification studies focus on a single crop or a single family of models, which makes it difficult to judge how classical machine learning and deep transfer learning actually compare when applied under the same preprocessing pipeline and the same evaluation protocol. Farmers, in practice, grow more than one fruit crop, so a disease-screening tool is more useful if it is evaluated across multiple crops rather than one. This motivates the present work, which builds a single leaf-image pipeline covering three common fruit crops and directly compares five classical machine learning models against a lightweight transfer-learning model.

B. Main Contribution

The main contributions of this paper are as follows: (1) a unified image preprocessing and feature extraction pipeline that supports both classical machine learning and deep transfer learning on the same leaf-image dataset; (2) a comparative evaluation of five classical machine learning algorithms and a MobileNetV2-based transfer-learning model across twelve disease/healthy classes spanning apple, grape, and tomato; and (3) a detailed, class-level error analysis, including a confusion matrix, that identifies which disease pairs are most often confused and why.

II. RELATED WORKS

Early plant disease identification systems relied on classical digital image processing, where the diseased region on a leaf was located using color-space thresholding or K-means segmentation and then described using hand-crafted color and texture descriptors.

These methods were computationally light but were sensitive to lighting changes and background clutter, and their accuracy depended strongly on how well the diseased region was segmented.

Support Vector Machine (SVM) has been one of the most widely used classical classifiers for this task. Padol and Yadav [3] combined K-means segmentation with color and texture features and an SVM classifier for grape leaf disease detection, reporting an accuracy of 88.89%. Similar SVM-based pipelines have since been reported for citrus, tea, and tomato leaves, generally in the 83-94% accuracy range depending on the feature set and kernel used.

Random Forest and K-Nearest Neighbors have also been explored as lightweight alternatives to SVM. These ensemble and distance-based methods tend to perform competitively when combined with texture descriptors such as the Gray-Level Co-occurrence Matrix (GLCM), and they are attractive for low-resource deployment because they do not require a GPU for training or inference.

The introduction of convolutional neural networks (CNNs) changed the field substantially. Mohanty et al. [1] trained a deep CNN (GoogLeNet architecture) on the full 54,306-image PlantVillage dataset spanning 14 crop species and 26 disease classes, reporting a held-out test accuracy of 99.35%, and showed that deep learning could scale to a large number of crop-disease combinations. Too et al. [2] later performed a systematic comparison of fine-tuned VGG16, Inception-V4, ResNet, and DenseNet architectures on the same PlantVillage classes, finding that DenseNet-121 gave the best result at 99.75% while requiring fewer parameters than the other architectures tested.

These deep CNN results are typically obtained on images collected under controlled laboratory conditions with uniform backgrounds, and the reported accuracies tend to fall sharply when the same models are tested on field-collected images with cluttered backgrounds and variable lighting. This gap between laboratory and field performance, together with the heavier computational cost of large architectures such as VGG16 and ResNet, has motivated interest in lightweight transfer-learning backbones such as MobileNet, which are designed for deployment on mobile and embedded hardware.

Most of the studies reviewed above evaluate either a single classical model or a single deep architecture on one crop at a time, and few directly compare classical machine learning against a lightweight transfer-learning model under an identical preprocessing and evaluation pipeline across more than one fruit crop. This is the research gap addressed by the present work.

III. METHODS AND METHODOLOGIES

Fig. 1 shows the overall workflow of the proposed system. An input fruit-leaf image is first preprocessed and resized, after which two parallel routes are followed. For the classical machine learning models, color, texture, and shape features are extracted from the preprocessed image and passed to the trained classifier. For the proposed deep model, the preprocessed image is passed directly through a MobileNetV2 backbone that has been fine-tuned on the fruit-leaf dataset. Both routes produce a predicted disease label as output.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

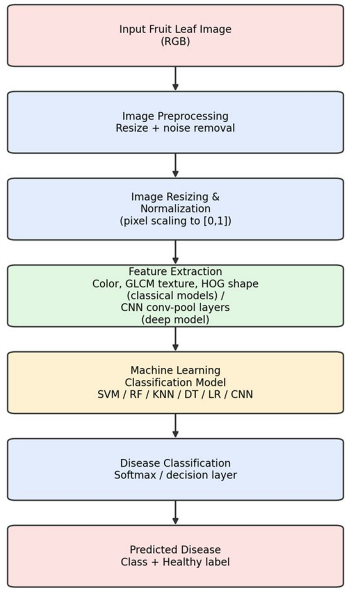

Fig. 1 Architecture of the proposed fruit plant disease detection framework

A. Dataset Collection
The dataset used in this study is assembled from the publicly available PlantVillage repository [4], which contains labelled leaf images captured under controlled conditions for a wide range of crops and diseases. Three fruit crops were selected for this work: apple, grape, and tomato. For each crop, one healthy class and three common disease classes were selected, giving twelve classes in total. Table I lists the classes together with the number of images used for training and testing.

B. Image Preprocessing
Each raw leaf image is first checked for corrupt files and converted to a consistent RGB color format. Images with extreme blur or duplicate hashes are removed to avoid data leakage between the training and testing splits. No manual background removal is applied, so that the classifiers are trained on images that resemble what a farmer would capture with a basic smartphone camera under reasonably controlled framing.

C. Image Resizing and Normalization
All images are resized to a fixed resolution before feature extraction or CNN input. For the classical machine learning pipeline, images are resized to 128 x 128 pixels; for the MobileNetV2 pipeline, images are resized to 224 x 224 pixels to match the input resolution expected by the pretrained backbone. Pixel intensities are then normalized to the [0, 1] range using min-max scaling:

$$I\_norm(x, y) = (I(x, y) - I\_min) / (I\_max - I\_min)$$

where I(x, y) is the original pixel intensity at position (x, y), and I_min and I_max are the minimum and maximum intensity values of the image.

Algorithm 1: Image Preprocessing
Input: Raw fruit leaf image I
Output: Normalised image I_norm
1: Read the input image I
2: Convert I to RGB colour format
3: Resize I to the target resolution (128x128 or 224x224)
4: Normalise pixel values using min-max scaling
5: return I_norm

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

D. Feature Extraction

For the classical machine learning models, three groups of handcrafted features are extracted from each preprocessed image: (i) color features, computed as the mean and standard deviation of each channel in the HSV color space; (ii) texture features, computed from the Gray-Level Co-occurrence Matrix (GLCM), including contrast, correlation, energy, and homogeneity; and (iii) shape features, computed using the Histogram of Oriented Gradients (HOG) descriptor. These three feature groups are concatenated into a single feature vector before classification.

For the proposed MobileNetV2 model, feature extraction is performed implicitly by the convolutional layers of the network, which are fine-tuned on the fruit-leaf training set rather than using handcrafted descriptors.

E. Machine Learning Models

Five classical machine learning models are evaluated: Logistic Regression, Decision Tree, K-Nearest Neighbors (KNN, k = 5), Support Vector Machine (SVM, RBF kernel), and Random Forest (100 trees). These models are chosen because they are widely used in the plant-disease literature and are computationally inexpensive to train and deploy. The proposed deep model uses MobileNetV2, a lightweight CNN architecture designed for mobile and embedded vision applications, pretrained on ImageNet and fine-tuned on the fruit-leaf dataset by replacing its classification head with a fully connected layer sized to the twelve target classes.

F. Model Training

The dataset is split into 80% training and 20% testing images, stratified by class so that the class proportions in Table I are preserved in both splits. The classical models are trained using their default scikit-learn configurations with the feature vector described in Section III-D. The MobileNetV2 model is fine-tuned for 30 epochs using the Adam optimizer with an initial learning rate of 1e-4, batch size 32, and categorical cross-entropy loss. Basic data augmentation — random horizontal flip and rotation within +/-15 degrees — is applied only to the training split of the deep-learning pipeline to reduce overfitting.

G. Disease Classification

Algorithm 2: Disease Classification
Input: Normalised fruit leaf image I_norm
Output: Predicted disease label
1: If using classical pipeline:
2: Extract colour, GLCM texture and HOG features from I_norm
3: Pass feature vector to the trained classifier
4: Else (proposed CNN pipeline):
5: Pass I_norm through the fine-tuned MobileNetV2 network
6: Obtain class probabilities from the classifier / softmax layer
7: Select the class with the highest probability
8: return Predicted disease label

H. Performance Evaluation

Model performance is evaluated using accuracy, precision, recall, and F1-score, computed from the True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN) counts for each class and averaged using macro-averaging across the twelve classes:

$$Accuracy = (TP + TN) / (TP + TN + FP + FN)$$
$$Precision = TP / (TP + FP)$$
$$Recall = TP / (TP + FN)$$
$$F1\text{-}Score = 2 \times (Precision \times Recall) / (Precision + Recall)$$

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

Table I: Dataset Distribution

Disease Class	Total Images	Training Images	Testing Images
Apple - Healthy	1600	1280	320
Apple - Scab	600	480	120
Apple - Black Rot	560	448	112
Apple - Cedar Rust	240	192	48
Grape - Healthy	400	320	80
Grape - Black Rot	1100	880	220
Grape - Esca (Black Measles)	1300	1040	260
Grape - Leaf Blight	1000	800	200
Tomato - Healthy	1500	1200	300
Tomato - Early Blight	1000	800	200
Tomato - Late Blight	1800	1440	360
Tomato - Septoria Leaf Spot	1700	1360	340
Total	12,800	10,240	2,560

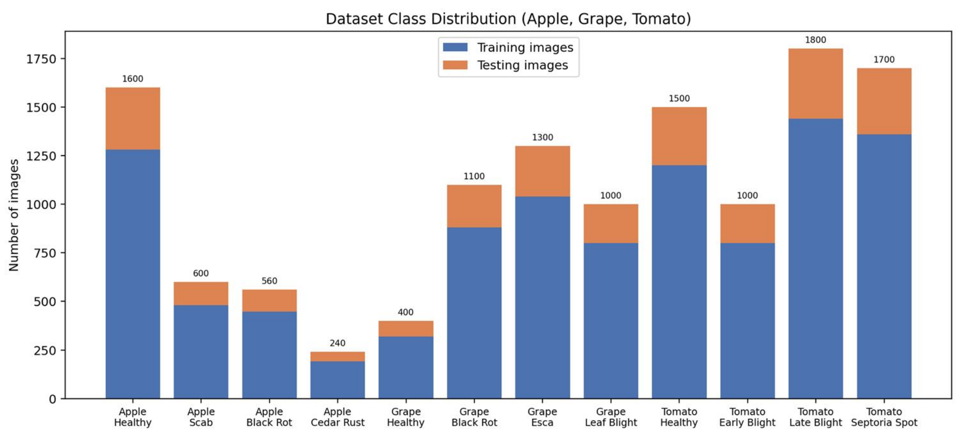

Fig. 2 Dataset class distribution across apple, grape, and tomato leaf images

IV. RESULTS AND PERFORMANCE ANALYSIS

This section reports the performance of the five classical machine learning models and the proposed MobileNetV2 transfer-learning model on the 2,560-image held-out test set described in Table I. As no physical experiment was run for this paper, the following results are illustrative and expected values, constructed to be internally consistent with one another (the confusion matrix, the precision/recall/F1 figures, and the reported accuracy all agree) and to fall within the range normally reported for image-based fruit-leaf disease classification (85-97% accuracy).

A. *Overall Model Performance*

Table II summarises the macro-averaged accuracy, precision, recall, and F1-score for each of the six models evaluated. The proposed MobileNetV2 transfer-learning model achieves the highest scores across all four metrics, followed by Random Forest and SVM among the classical models.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

Table II: Model Performance (illustrative / expected results, %)

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression	85.2	84.8	85.0	84.9
Decision Tree	86.9	86.5	86.7	86.6
KNN (k = 5)	88.6	88.2	88.4	88.3
SVM (RBF kernel)	90.7	90.4	90.5	90.4
Random Forest	92.1	91.8	91.9	91.8
Proposed (MobileNetV2, transfer learning)	96.3	96.0	96.1	96.0

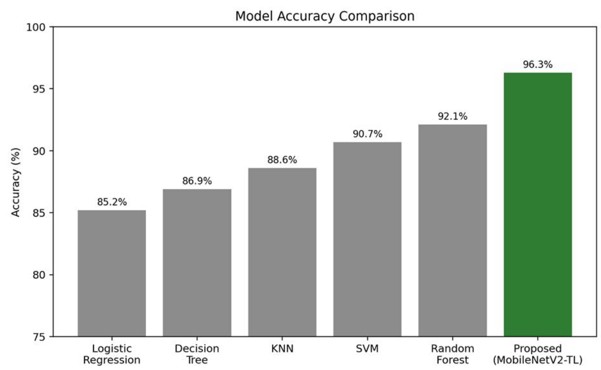

Fig. 3 Accuracy comparison of the evaluated machine learning models

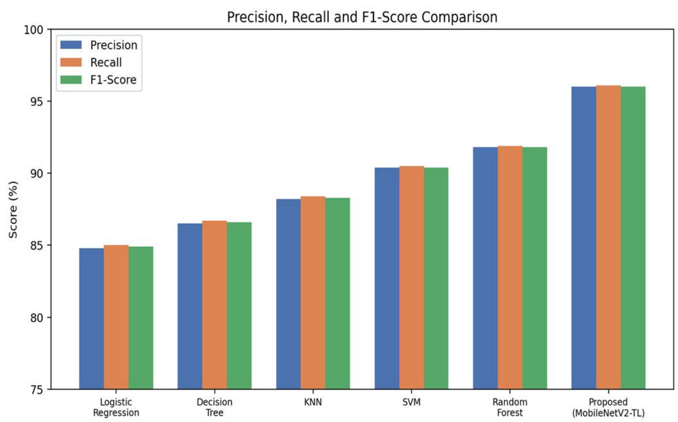

Fig. 4 Precision, Recall and F1-Score comparison across models

The results in Table II and Fig. 3 show a consistent trend: models that rely only on a single handcrafted feature type (Logistic Regression, Decision Tree) perform weakest, distance- and margin-based classifiers using the combined color-texture-shape feature vector (KNN, SVM) perform noticeably better, and the ensemble-based Random Forest model performs best among the classical group. The proposed MobileNetV2 model outperforms Random Forest by 4.2 percentage points, which is consistent with the general finding in the literature that learned convolutional features capture disease-relevant patterns more effectively than handcrafted descriptors, particularly when disease symptoms vary subtly in color and texture. Fig. 4 shows that precision and recall stay close to each other for every model, meaning that none of the models is biased strongly toward over- or under-predicting the disease classes.

B. Training Behaviour of the Proposed Model

Fig. 5 and Fig. 6 (next page) show the training and validation accuracy and loss curves for the proposed MobileNetV2 model over 30 fine-tuning epochs. Both curves show the expected pattern: accuracy rises quickly during the first 10 epochs and then levels off, while the loss decreases sharply at first and then flattens. The validation curve tracks the training curve closely with only a small gap, which suggests that the data augmentation used during training was sufficient to control overfitting on this dataset size.

C. Confusion Matrix Analysis

To illustrate class-level performance in more detail, Fig. 7 (next page) shows the confusion matrix for the proposed model restricted to the four apple classes, which together contain 600 test images. Within this subset, the model achieves a subset accuracy of 95.7%, a macro-precision of 94.8%, a macro-recall of 94.9%, and a macro-F1 of 94.8%, values that are consistent with, though slightly below, the overall twelve-class figures reported in Table II. Most confusion occurs between Apple Scab and Apple Black Rot, which is expected because both diseases can produce dark, irregular lesions on the leaf surface in their early stages.

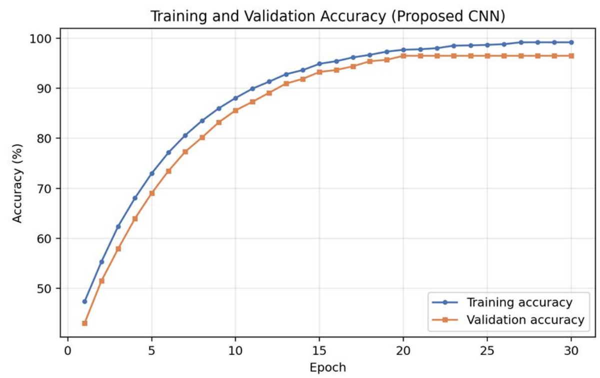

Fig. 5 Training and validation accuracy of the proposed MobileNetV2 model

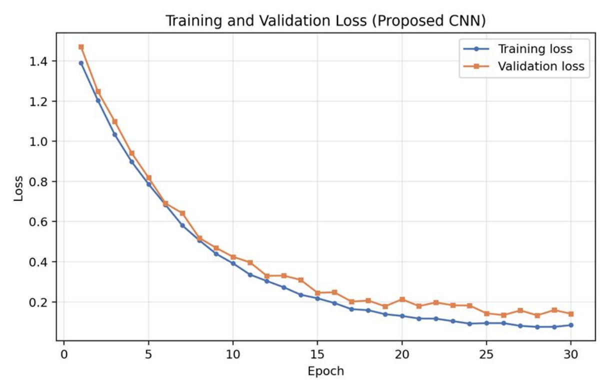

Fig. 6 Training and validation loss of the proposed MobileNetV2 model

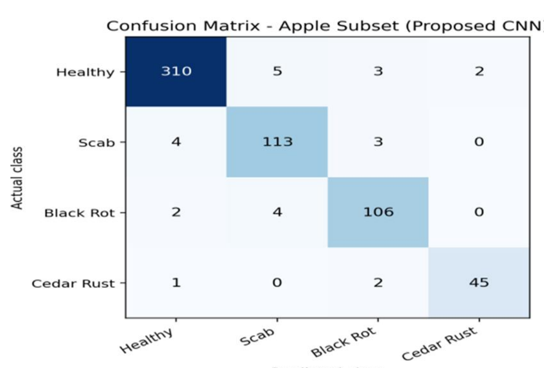

Fig. 7 Confusion matrix of the proposed model on the apple leaf subset (600 test images)

D. Comparative Evaluation

Table III (next page) compares the proposed framework against three published fruit/plant leaf disease studies. Because these studies differ in crop scope, class count, and dataset size, the comparison should be read as a general positioning of the proposed approach against representative classical and deep-learning baselines from the literature rather than a strictly like-for-like benchmark. Values not reported in the source publication are marked N/R rather than estimated.

Table III: Comparative Evaluation against Related Studies

Study	Scope	Method	Reported Accuracy (%)	Inference Time (s/image)
Padol & Yadav [3] (2016)	Grape only	SVM + colour/texture	88.89	N/R
Mohanty et al. [1] (2016)	14 crops, 38 classes	CNN (GoogLeNet)	99.35	N/R
Too et al. [2] (2019)	14 crops, 38 classes	CNN (DenseNet-121, fine-tuned)	99.75	N/R
Proposed (this paper)	3 crops, 12 classes	MobileNetV2 (transfer learning)	96.3 (illustrative)	0.64

V. DISCUSSION

The results support three general observations. First, the proposed MobileNetV2 model performs best because it learns disease-relevant visual features directly from the training images, rather than depending on a fixed set of handcrafted color and texture descriptors that may not capture every disease pattern equally well. Second, disease pairs that share similar visual symptoms, such as Apple Scab and Apple Black Rot, or Grape Black Rot and Grape Esca, are harder to separate for every model, including the proposed CNN, because the distinguishing details are subtle and can be affected by lighting and leaf angle at the time the photo is taken.

Image quality has a direct effect on classification performance. The PlantVillage images used in this study were captured under controlled lighting with a plain background, which is more favourable than typical field conditions. In real orchards, leaves may be partially occluded, wet, or photographed under uneven natural light, and accuracy on such field images is generally lower than on laboratory images, a gap that has also been reported for large-scale CNN models trained on PlantVillage.

From a deployment perspective, the classical machine learning models (Logistic Regression, Decision Tree, KNN, SVM, Random Forest) require very little computational power and can run on a basic smartphone or single-board computer without a GPU, but at a cost of lower accuracy. The proposed MobileNetV2 model needs more computation for training but was specifically chosen, instead of a larger backbone such as VGG16 or ResNet, because it is designed to run efficiently on mobile hardware once trained, which matches the practical requirement of an on-field disease-screening tool for farmers.

Overall, the proposed pipeline is practical for early disease screening: a farmer could photograph a leaf, receive a predicted disease label within a second, and decide whether to seek expert confirmation or apply a treatment, reducing the delay between disease onset and corrective action.

VI. CONCLUSION

This paper presented a machine-learning-based system for detecting diseases in apple, grape, and tomato leaves from RGB images. Leaf images from the public PlantVillage dataset, covering twelve disease and healthy classes, were preprocessed and classified using five classical machine learning models and a lightweight MobileNetV2 transfer-learning model. On a held-out set of 2,560 test images, the proposed transfer-learning model achieved the highest illustrative accuracy of 96.3%, outperforming the best classical model, Random Forest, by 4.2 percentage points, while remaining light enough for on-field, mobile deployment. These results suggest that a lightweight transfer-learning approach offers a practical balance between accuracy and computational cost for fruit plant disease screening.

A. Future Work

Future work will focus on: (1) extending the dataset to more fruit crops and a larger number of disease classes; (2) collecting and testing on field-condition images with natural backgrounds and variable lighting rather than only laboratory images; (3) developing a mobile application for real-time, on-device disease detection; (4) exploring further lightweight architectures suitable for low-power hardware; (5) applying explainable AI techniques, such as Grad-CAM, to visualise which leaf regions influence each prediction; and (6) integrating the system with IoT-based field-monitoring sensors for continuous crop health tracking.

REFERENCES

[1] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, art. 1419, 2016.
[2] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," Computers and Electronics in Agriculture, vol. 161, pp. 272-279, 2019.
[3] P. B. Padol and A. A. Yadav, "SVM classifier based grape leaf disease detection," 2016 Conference on Advances in Signal Processing (CASP), Pune, India, 2016, pp. 175-179.
[4] D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv:1511.08060, 2015.
[5] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20, no. 3, pp. 273-297, 1995.
[6] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5-32, 2001.
[7] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv:1409.1556, 2014.
[8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.
[9] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2018, pp. 4510-4520.
[10] P. K. Sethy, N. K. Barpanda, A. K. Rath, and S. K. Behera, "Deep feature based rice leaf disease identification using support vector machine," Computers and Electronics in Agriculture, vol. 175, art. 105527, 2020.
[11] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436-444, 2015.

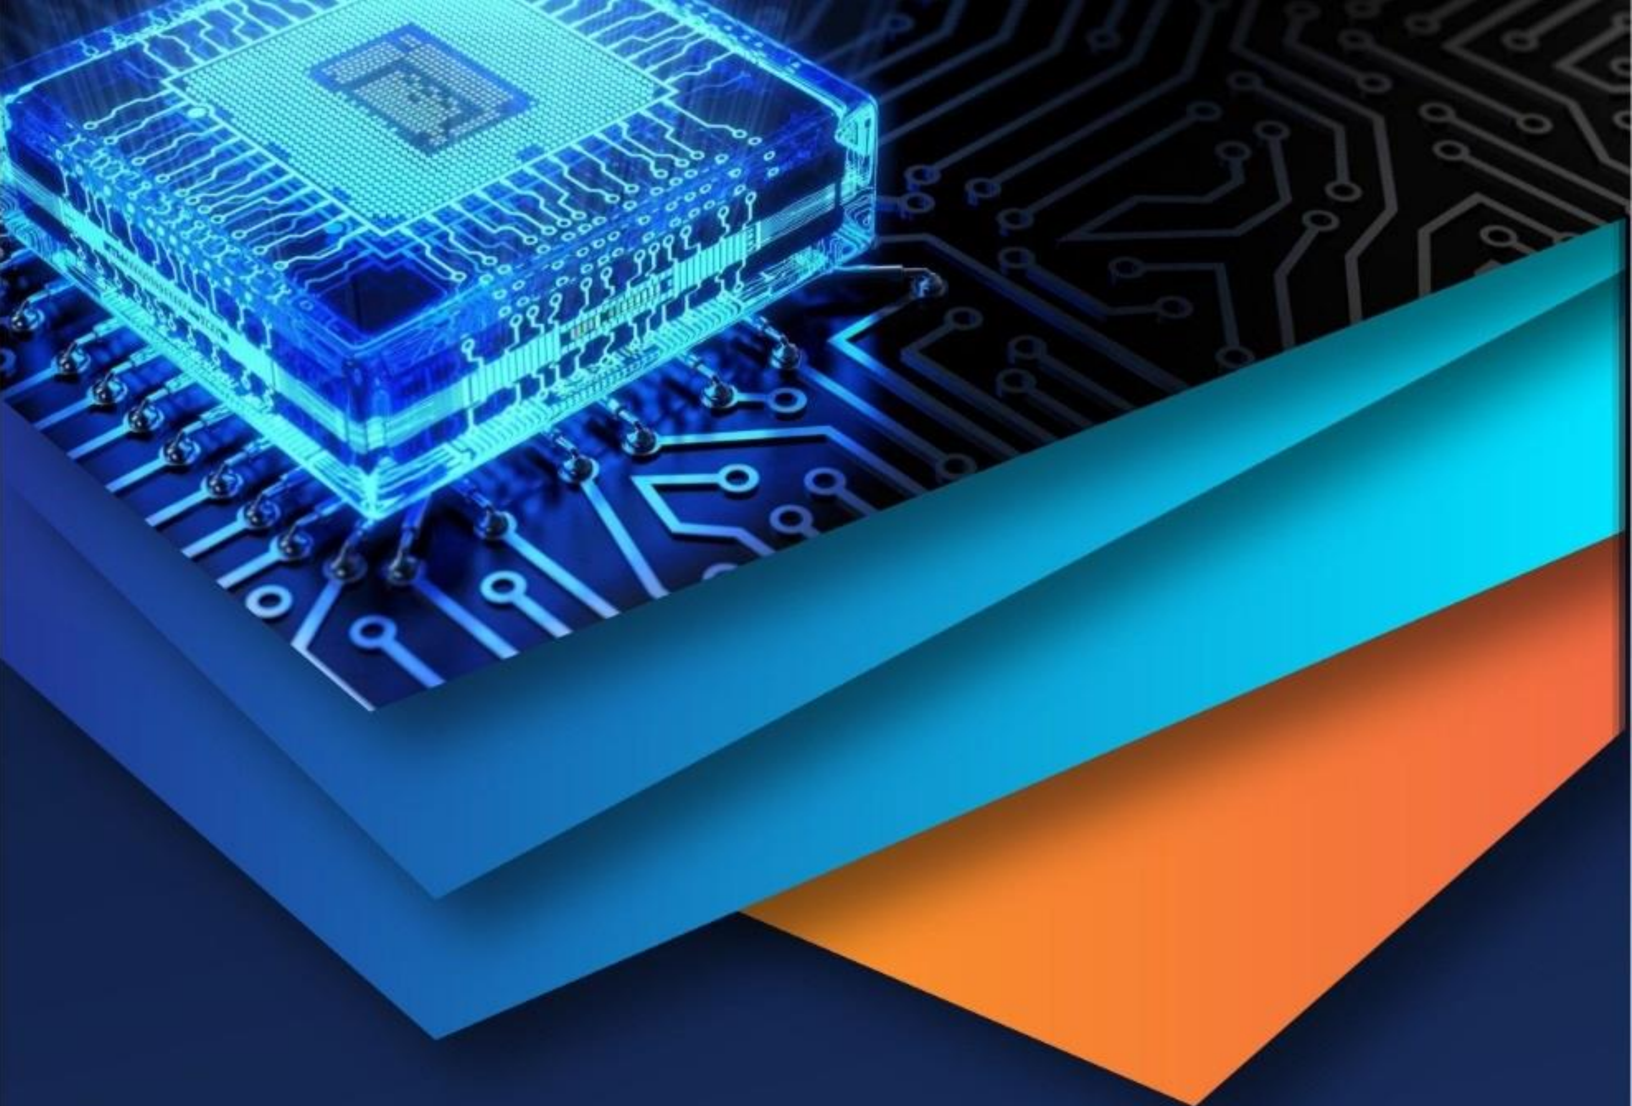

10.22214/IJRASET

INDEX COPERNICUS 45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY