iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** XI **Month of publication:** November 2026

DOI: https://doi.org/10.22214/ijraset.2026.85629

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

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 XI Nov 2026- Available at www.ijraset.com

Quantum-Enhanced Deep Learning Framework for Leaf Disease Prediction and Classification

Pandy Kranthi Kiran[1], Dr. M. Humera Khanam[2], Md. A. Khudhus[3]
[1, 2, 3]*Department of Computer Science and Engineering, Sri Venkateswara University College of Engineering, Tirupati, India*

Abstract: ***Recognising plant leaf diseases from images is difficult because symptoms of different diseases often look alike, and expert inspection is not always available to growers. This paper presents a hybrid quantum deep learning framework that predicts both the plant species and the disease of a leaf image in a single pass. An ImageNet-pretrained MobileNetV2 backbone extracts a high-dimensional feature vector, which a small projection reduces to four values that drive a four-qubit variational quantum circuit built from angle embedding and strongly entangling layers. The resulting Pauli-Z expectation values are concatenated with the classical features and passed to separate plant and disease classification heads. The model is compared with a MobileNetV2-only baseline on an identical stratified train–test split covering six crops and twenty disease and healthy classes. The hybrid model scores higher than the baseline on accuracy, precision, recall and F1-score for both tasks, although the margins are modest and come from a single split. The framework is also deployed as a Flask web application that adds blur and leaf validation, leaf segmentation, Grad-CAM explanations and disease-management guidance. The results are therefore presented as preliminary evidence of the value of a compact quantum feature stage rather than a definitive advantage.***

Keywords: Plant Leaf Disease; MobileNetV2; Quantum Machine Learning; Variational Quantum Circuit; Multi-Task Learning; Grad-CAM

I. INTRODUCTION

Plant diseases reduce crop yield and quality, and many of them first show up on the leaves as spots, discoloration, lesions, blight or rust. Finding these symptoms early allows timely treatment, yet identification in practice still depends largely on visual scouting by trained people. This is slow, subjective and prone to error. Image-based machine learning and deep learning have therefore attracted considerable attention as a way to automate the task, and recent surveys cover the field from classical classifiers to modern deep networks [1], [2].

Many published systems, however, stop at a single disease label. They rarely identify the host plant separately, they are closed-set (an unsupported object is still forced into one of the known classes), they give no indication of which image regions drove the decision, and they say nothing about what a grower should do next. Lightweight convolutional networks such as MobileNetV2 [3] are attractive for this setting because they offer a good balance between accuracy and computational cost and are easy to deploy.

In parallel, quantum machine learning has proposed variational quantum circuits as trainable feature transformations [4]–[6]. Rather than replacing a deep network, a small circuit can be inserted into it as an extra feature-processing stage. This work asks a practical question: does a compact four-qubit circuit placed after a MobileNetV2 backbone add useful information to the classical features for plant and disease recognition? The main contributions of this paper are as follows.

1) A multi-task hybrid model that predicts plant species and disease from one image, using a four-qubit variational circuit whose outputs are fused with the classical MobileNetV2 features.
2) A controlled comparison against a MobileNetV2-only baseline that shares the backbone, the output structure and one fixed train–test split, so that the effect of the quantum stage is the main difference under study.
3) A complete inference pipeline in a web application, with blur and leaf validation, leaf segmentation, dual-stream processing, Grad-CAM visualisation and disease-specific guidance.
4) A candid discussion of what the results do and do not show.

II. RELATED WORK

Balafas et al. [2] surveyed machine learning and deep learning methods for plant disease classification and detection and ran experiments on the PlantDoc dataset. They found that MobileNetV2 [3] gave a useful trade-off between accuracy and training time, which supports choosing a compact backbone when computation matters.

©IJRASET: All Rights are Reserved

A broader review by Bhargava et al. [1] covers computer-vision and artificial-intelligence techniques for detecting, classifying and diagnosing leaf diseases, including machine learning, deep learning and few-shot approaches, and points out the limitations that automated systems still face.

Explainability has become a recurring theme. Nigar et al. [7] paired a deep classifier for 38 plant disease categories with LIME to explain individual predictions and reported an accuracy of 99.69%, arguing that interpretable output matters as much as the score. Oad et al. [8] combined VGG16, VGG19, ResNet101V2 and InceptionV3 through voting and again used LIME to highlight influential regions. Both works show the value of explanations, but they rely on heavier models or ensembles and use LIME, whereas the present work uses a single compact network with Grad-CAM.

Transfer learning with pretrained networks is well established. Daalah et al. [9] fine-tuned ResNet-50, EfficientNet-B0, MobileNetV2 and AlexNet on tomato leaves from PlantVillage; ResNet-50 achieved the highest accuracy (97.32%), while MobileNetV2 and EfficientNet-B0 offered a better balance between accuracy and efficiency. Hybrid architectures are also being explored: Prommakhot et al. [10] combined convolutional feature extraction with BiLSTM and transformer-based sequential learning on PlantVillage and reported 97.88% accuracy. Our model is hybrid in a different sense. It couples MobileNetV2 features with a variational quantum circuit instead of a sequential or attention-based module.

The quantum component draws on work that surveys quantum machine learning [4], treats quantum circuits as kernel feature maps [5], and demonstrates quantum-enhanced feature spaces for supervised learning [6]. To our knowledge from the studies reviewed here, the plant-disease works above do not combine a CNN backbone with a quantum circuit and a dual-output plant-and-disease design in a deployable application, which is the gap this paper addresses.

III. EXISTING SYSTEM

A conventional plant disease recognition system accepts a leaf image, applies image preprocessing, extracts visual features with a convolutional neural network, and predicts one or more disease classes. Such a system performs well when the input resembles the training domain. However, the classifier itself does not necessarily separate a supported leaf from an unrelated object or an unsupported plant. In this work, a MobileNetV2-only multi-task CNN serves as the baseline existing system. It gives plant and disease predictions on the same dataset split and with the same evaluation procedure as the proposed hybrid model, so the classical and quantum-enhanced approaches can be compared directly..

IV. PROPOSED METHODOLOGY

The proposed system combines image validation, leaf segmentation, MobileNetV2 feature extraction, a four-qubit variational quantum circuit, multi-task classification, Grad-CAM and disease-information retrieval for six plants (apple, tomato, potato, corn, grape and guava) and 20 disease and healthy classes. An uploaded image first passes a blur check and a leaf-versus-non-leaf validation model, and a segmentation module then produces a leaf mask. A masked stream is used for plant identification and the original stream for disease identification, and both go through the same trained hybrid network. The output shows the plant and disease labels with confidence values, disease-specific guidance and a Grad-CAM heat map, with a plant–disease consistency check applied as post-processing.

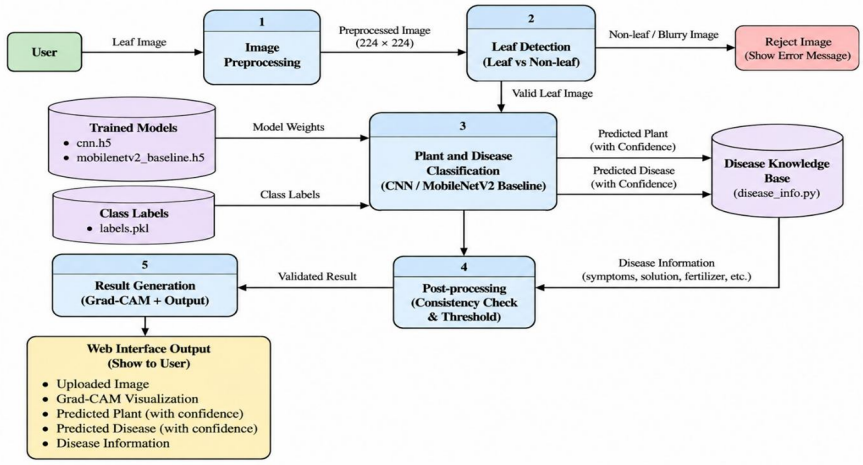


Fig.1 Block Diagram of Proposed System

A. *Dataset and Preprocessing*

The dataset covers six crops and 20 disease and healthy classes. Up to 300 images were sampled per class, and a single stratified 80/20 split (seed 42) was reused for all experiments, giving 1,200 test images. Images are resized to 224 × 224 pixels, converted to RGB and normalised with the MobileNetV2 preprocessing function. Training uses random flips and a mild contrast change.

B. *Feature Extraction*

A MobileNetV2 backbone pretrained on ImageNet extracts a 1280-dimensional feature vector from each image through global average pooling.

C. *Quantum Feature Enhancement*

Dense layers reduce the feature vector to four values, which are loaded into a four-qubit variational circuit by angle embedding. Strongly entangling layers then transform the state, and the Pauli-Z expectation value of each qubit gives four quantum features. These are concatenated with the 1280 classical features to form a 1284-dimensional fused vector.

D. *Multi-Task Classification and Training*

Two heads, each with a 256-unit dense layer, dropout and softmax, predict the plant (6 classes) and the disease (20 classes). The model is trained with the summed cross-entropy loss using Adam (learning rate 0.0001), a batch size of 32 and 20 epochs. A MobileNetV2-only model with the same heads serves as the baseline.

TABLE I
MODEL AND TRAINING CONFIGURATION

Item	Setting
Input Size	224 × 224 × 3
Backbone	MobileNetV2,ImageNet-pretrained
Feature vector (global average pooling)	1280
Quantum projection	Dense 128 → Dense 4 (tanh)
Quantum circuit	4 qubits, angle embedding, strongly entangling layers
Quantum output	4 Pauli-Z expectation values
Fused feature vector	1284
Classification heads	256-unit dense, dropout, softmax (6 plant classes, 20 disease classes)
Loss	Sparse categorical cross-entropy (both outputs)
Optimiser	Adam, learning rate 0.0001
Batch size / epochs	32/20
Data	Up to 300 images per class, stratified 80/20 split (seed 42), 1,200 test images

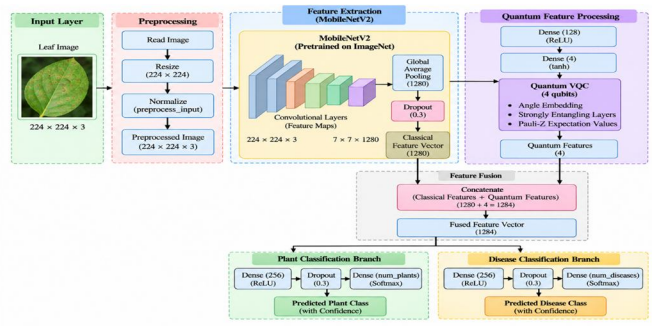


Fig. 2 Multi-Task Classification

E. Deployment Pipeline

The application first computes a blur diagnostic using the variance of the Laplacian, and an image below the configured threshold is rejected with a request for a clearer photo. Images that pass are examined by a separate MobileNetV2-based leaf detector trained with artificial leaves and unrelated objects as negatives and PlantVillage and PlantDoc images as positives. Inputs for which both the plant and disease confidence are very low are treated as unsupported plants. A segmentation module then produces a binary leaf mask, and an image whose leaf area is too small is rejected. Two streams are formed: the masked stream is used for plant identification so that background clutter has less influence, and the original stream is used for disease identification so that spot and texture detail is preserved. Both streams pass through the same trained network, which returns a probability distribution for each task. Because the two heads are separate, a post-processing step checks that the predicted disease is consistent with the predicted plant. Finally, Grad-CAM [11] overlays an attention map on the input, and a knowledge base returns symptoms, management advice and fertiliser guidance for the predicted disease.

V. SYSTEM ARCHITECTURE

The architecture starts with an uploaded leaf image, which goes through preprocessing and validation before classification. MobileNetV2 is used without its original classification top, and global average pooling produces a 1280-dimensional feature representation. In the hybrid model, this representation passes through a 128-unit dense projection and is then reduced to four values that feed the four-qubit variational quantum circuit.

The circuit applies angle embedding followed by strongly entangling layers and returns four Pauli-Z expectation values. These four quantum features are concatenated with the original global features, and the fused representation goes to separate dense classification heads. The plant head predicts six plant categories, and the disease head predicts 20 disease and healthy classes. A Grad-CAM visualisation of the regions that influenced the prediction is also provided.

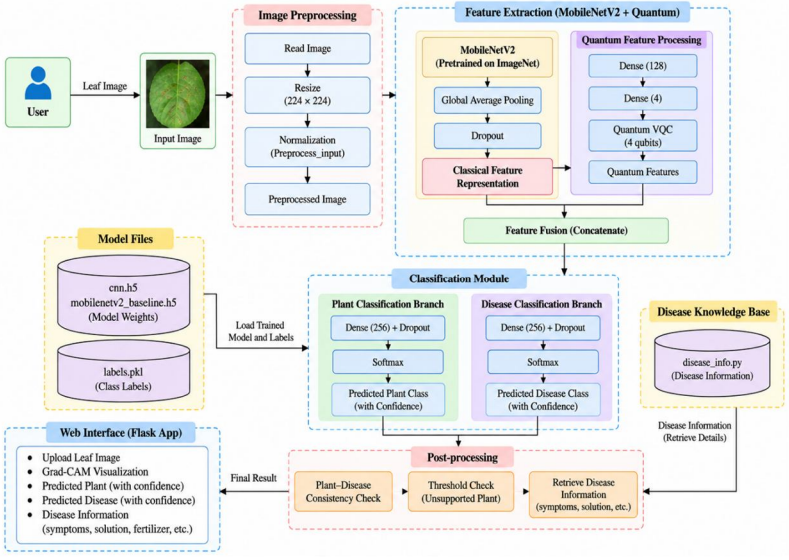


Fig. 3 System Architecture

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue XI Nov 2026- Available at www.ijraset.com

VI. MODELS

A. *MobileNetV2 Feature Extraction*

The backbone is built from inverted residual blocks. Each block expands the channels with a 1 × 1 convolution, filters them with a 3 × 3 depthwise convolution, and projects them back with a linear 1 × 1 convolution. When the stride is 1 and the input and output channel counts match, a skip connection adds the block input x:

$$y = x + W_p(DW(ReLU6(W_e x))),$$
$$ReLU6(a) = \min(\max(0, a), 6)$$

Here W_e is the expansion convolution, $DW(\cdot)$ is the depthwise convolution with its own ReLU6, and W_p is the linear projection. Global average pooling over the final feature maps $A^k \in \mathbb{R}^{HxW}$ gives the 1280-dimensional feature vector f:

$$f_k = (1/HW) \sum_{i=1}^{H} \sum_{j=1}^{W} A^k_{ij}, \quad k = 1, \ldots, 1280$$

B. *Quantum Feature Extraction*

The pooled MobileNetV2 features $f \in \mathbb{R}^{1280}$ are reduced to four angles u, which are embedded into a four-qubit circuit $U(\vartheta)$ of strongly entangling layers. The quantum features are the Pauli-Z expectation values:

$$q_i = \langle \psi(u)| U\dagger(\vartheta)\, Z_i\, U(\vartheta) |\psi(u)\rangle, \quad i = 1, \ldots, 4$$

C. *Feature Fusion*

The quantum and classical features are concatenated before the two classification heads:

$$z = [f; q] \in \mathbb{R}^{1284}$$

D. *Grad-CAM Explanation*

For the predicted disease class c with score y^c, the heat map is:

$$\alpha^c_k = (1/HW) \sum_{i,j} (\partial y^c / \partial A^k_{ij}),$$
$$L^c\text{Grad-CAM} = ReLU(\sum_k \alpha^c_k A^k)$$

E. *Multi-Task Loss*

Each head uses a softmax output $\hat{p}_t$, and the total loss is the sum of the plant and disease cross-entropies over a batch of B images:

$$\mathscr{L} = \mathscr{L}\text{plant} + \mathscr{L}\text{disease},$$
$$\mathscr{L}_t = -(1/B) \sum_{\beta=1}^{B} \log \hat{p}_{t,\gamma t}^{(b)}$$

F. *Evaluation Metrics*

With true positives TP▪ , false positives FP▪ , false negatives FN▪ , N test images and C classes:

$$Acc = (1/N) \sum_▪ TP▪ ,$$
$$F1▪ = 2P▪ R▪ /(P▪ + R▪)$$
$$Macro\text{-}F1 = (1/C) \sum_{▪=1}^{c} F1▪$$

Here P▪ = TP▪ /(TP▪ + FP▪) and R▪ = TP▪ /(TP▪ + FN▪). Macro-precision and macro-recall are averaged over classes in the same way.

The proposed framework predicts the plant species and the disease of a leaf image in one pass. It has four stages: preprocessing, feature extraction, quantum feature enhancement and multi-task classification

VII. RESULTS AND DISCUSSIONS

A. *Classification Performance*

Table II reports the MobileNetV2 baseline and Table III the hybrid model, each for the plant and disease tasks on the same 1,200-image test set.

TABLE II
MOBILENETV2 BASELINE (TEST SET, %)

Task	Accuracy	Precision	Recall	F1-Score
Plant	95.42	96.10	95.23	95.64
Disease	86.00	86.45	86.00	85.95

©IJRASET: All Rights are Reserved

TABLE III
MOBILENETV2 + QUANTUM MODEL (TEST SET, %)

Task	Accuracy	Precision	Recall	F1-Score
Plant	96.08	96.46	96.62	96.52
Disease	87.58	88.00	87.58	87.50

B. Application Results

The web interface lets a user upload a leaf image and submit it for prediction (Fig. 4). After successful processing, the application displays the predicted plant and disease with their confidence percentages (Fig. 5), together with the Grad-CAM attention map and disease information (Fig. 6). Blurry images, non-leaf images and unsupported plants produce explanatory messages instead of a forced prediction (Fig. 7).

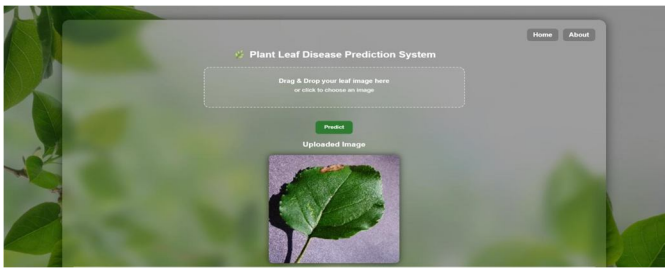


Fig. 4 Leaf-image upload interface

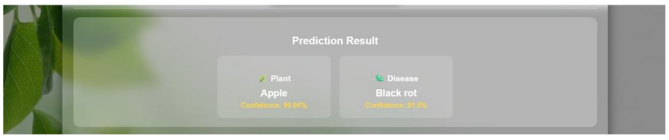


Fig. 5 Predicted plant and disease with confidence percentages

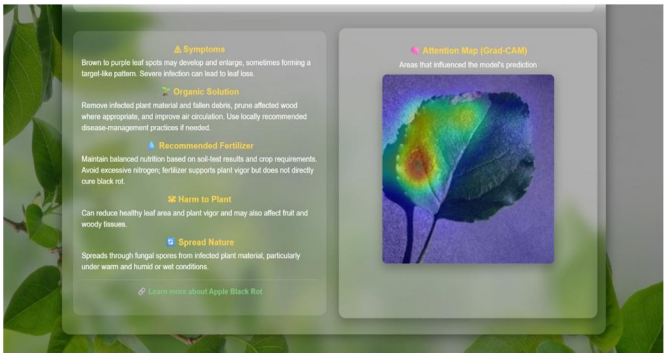


Fig. 6 Disease-management guidance and Grad-CAM attention map

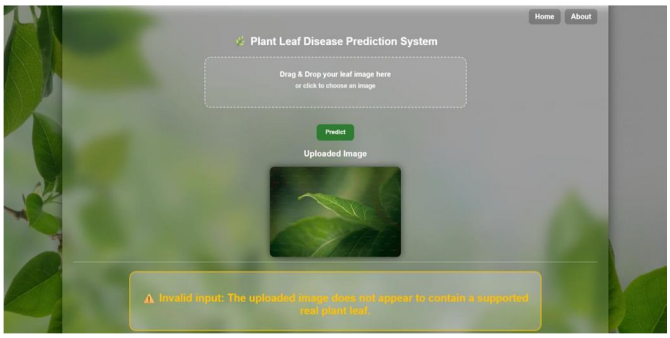


Fig. 7 Application response to an invalid or unsupported input

C. Discussion

The hybrid model's improvement is consistent across all metrics, which suggests that the four quantum features carry some information that complements the 1280 classical ones. The size of the effect needs careful reading, though.

On 1,200 test images, 0.66 percentage points of plant accuracy corresponds to about 8 images, and 1.58 points of disease accuracy to about 19 images. The results also come from one training run on one split, without repeated seeds, confidence intervals or a significance test, so part of the difference could be random variation. In addition, the hybrid model has extra trainable dense layers (the 128-unit projection and the four-unit reduction) that the baseline lacks, so the present design cannot separate the contribution of the quantum circuit from that of the added classical capacity. Replacing the circuit with an equally sized classical layer would answer this directly.

Plant recognition is the easier task for both models, with accuracy above 95%, while disease recognition over 20 visually similar classes is harder, at roughly 86 to 88%. The accuracies are also lower than the 97 to 99% reported in some of the studies reviewed in Section II [7], [9], [10]. Those studies use different datasets, class counts and protocols, so their figures are not directly comparable with a six-crop, 20-class problem with at most 300 images per class.

D. Limitations

The data come from a fixed domain, so the results do not establish performance on images taken under varied field conditions.

The disease-management text gives general guidance and has not been verified by agronomists.

The validation steps of the application (blur check, leaf detection and segmentation) were demonstrated through example inputs but not evaluated quantitatively, and the results come from a single split of one dataset.

VIII. CONCLUSION AND FUTURE SCOPE

This paper presented a MobileNetV2-based multi-task framework that adds the output of a four-qubit variational quantum circuit to the classical features, so that the plant species and the disease are predicted from one leaf image, together with a web application offering input validation, Grad-CAM explanations and disease guidance. Against a MobileNetV2-only baseline on the same split, the hybrid model improved plant accuracy from 95.42% to 96.08% and disease accuracy from 86.00% to 87.58%. These gains are small, come from a single split, and are not separated from the effect of the added dense layers; they indicate promise rather than a proven advantage for the quantum stage. Future work will repeat the experiments over several seeds with statistical testing, add a classical ablation of the quantum branch, explore other circuit designs, extend the dataset with more crops and field-captured images, validate the management advice with agricultural experts, and optimise the model for mobile devices.

REFERENCES

[1] A. Bhargava et al., "Plant leaf disease detection, classification, and diagnosis using computer vision and artificial intelligence: A review," *I*EEE Access, vol. 12, pp. 37443–37469, 2024, doi: 10.1109/ACCESS.2024.3373001.

[2] V. Balafas et al., "Machine learning and deep learning for plant disease classification and detection," IEEE Access, vol. 11, pp. 114352–114377, 2023, doi: 10.1109/ACCESS.2023.3324722.

[3] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 4510–4520, doi: 10.1109/CVPR.2018.00474.

[4] J. Biamonte et al., "Quantum machine learning," Nature, vol. 549, pp. 195–202, 2017, doi: 10.1038/nature23474.

[5] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," Phys. Rev. Lett., vol. 122, no. 4, Art. no. 040504, 2019, doi: 10.1103/PhysRevLett.122.040504.

[6] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," Nature, vol. 567, pp. 209–212, 2019, doi: 10.1038/s41586-019-0980-2.

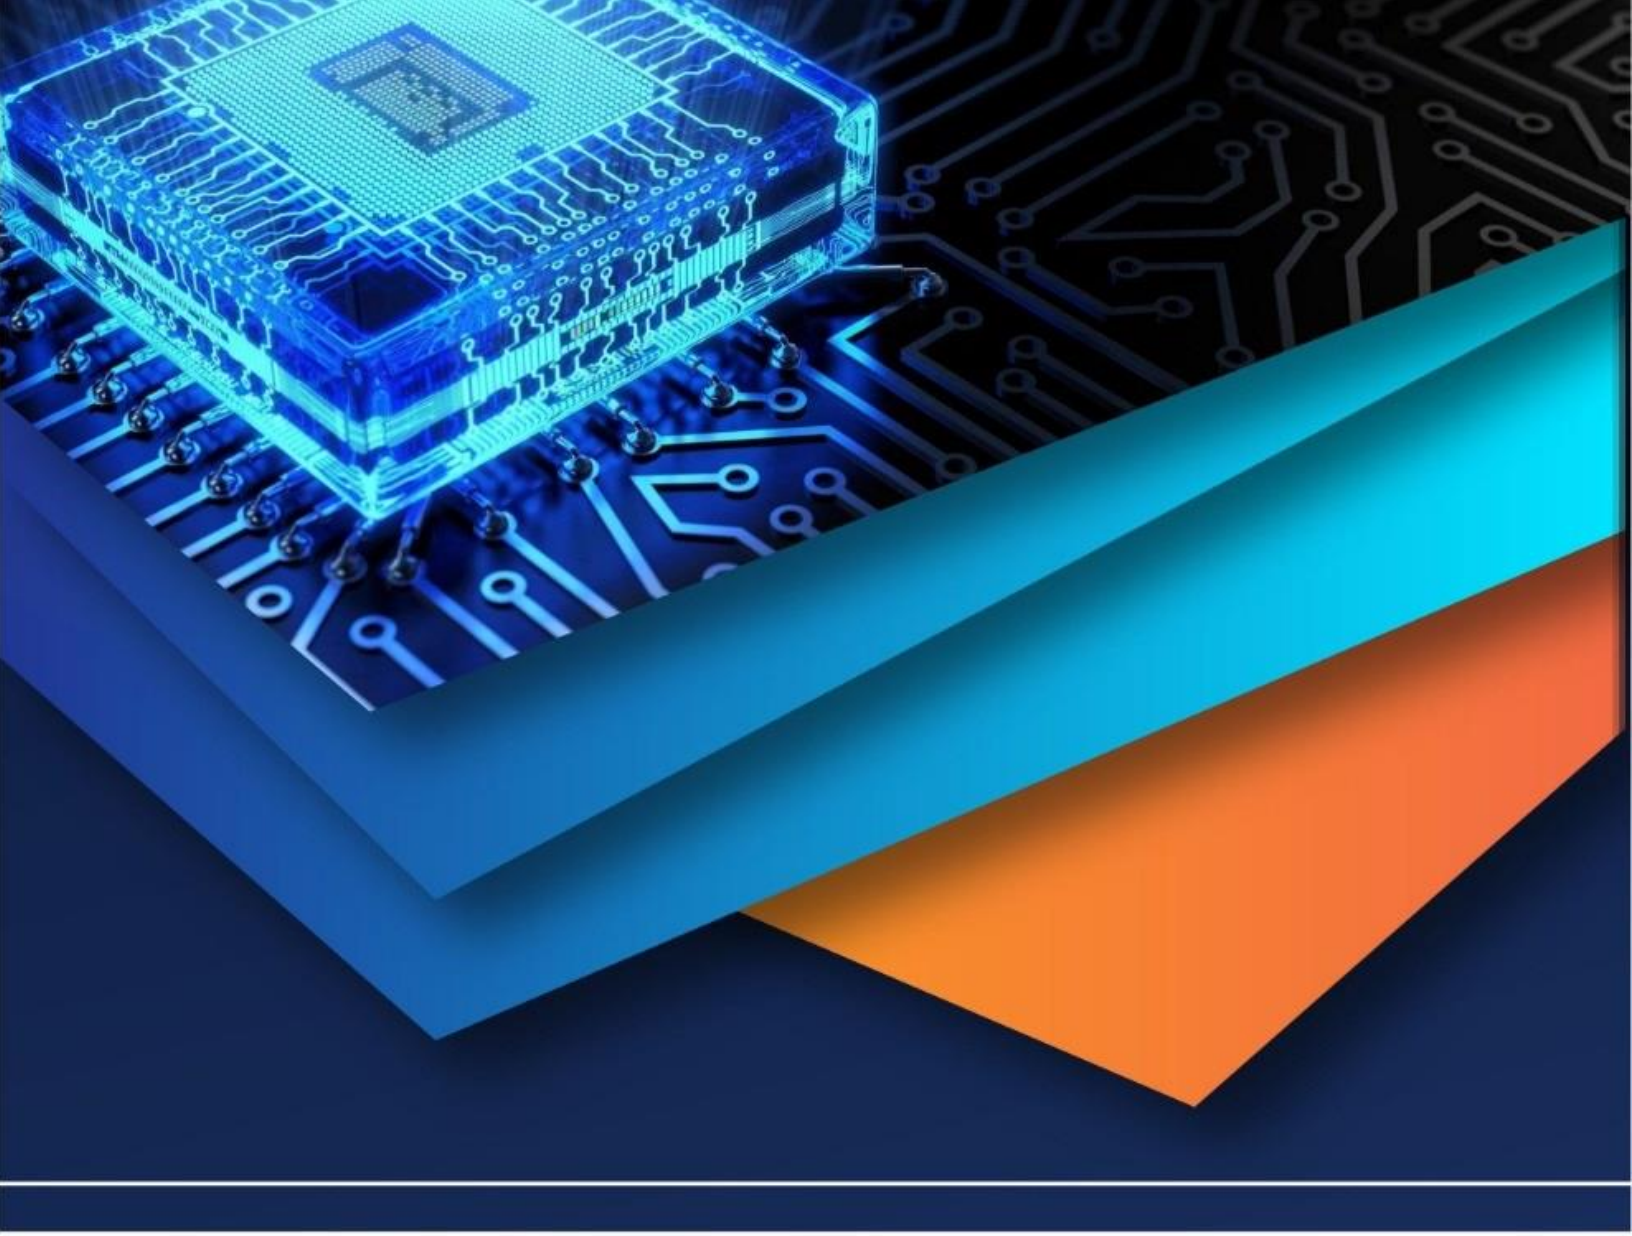
[7] N. Nigar, H. M. Faisal, M. Umer, O. A. Oki, and J. M. Lukose, "Improving plant disease classification with deep-learning-based prediction model using explainable artificial intelligence," IEEE Access, vol. 12, pp. 100005–100014, 2024, doi: 10.1109/ACCESS.2024.3428553.

[8] A. Oad et al., "Plant leaf disease detection using ensemble learning and explainable AI," IEEE Access, vol. 12, pp. 156038–156049, 2024, doi: 10.1109/ACCESS.2024.3484574.

[9] A. Daalah, S. Aldousari, S. Alrashed, H. Alenezi, A. Alhajri, and A. Eleyan, "Plant leaf disease classification using fine-tuned pre-trained CNN models," in Proc. 6th Int. Conf. Bio-engineering for Smart Technologies (BioSMART), 2025, doi: 10.1109/BioSMART66413.2025.11046055.

[10] A. Prommakhot, J. Onshaunjit, W. Ooppakaew, G. Samseemoung, and J. Srinonchat, "Hybrid CNN and transformer-based sequential learning techniques for plant disease classification," IEEE Access, vol. 13, pp. 122876–122887, 2025, doi: 10.1109/ACCESS.2025.3586285.

[11] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. *Vis. (ICCV)*, 2017, pp. 618–626, doi: 10.1109/ICCV.2017.74.



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)