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.2026.85725

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

Mango and Apple Leaf Disease Classification Using CNN and Frequency Domain Analysis

Sharmistha Gupta[1], Soumendu Banerjee[2]
[1]*Department of Computer Science, St. Xavier's College, Burdwan, India*
[2]*Department of Computer Application, Academy of Technology, Hooghly, India*

Abstract: ***Plant diseases significantly affect crop productivity and quality, making early and accurate identification essential for effective agricultural management. While manual inspection of leaf symptoms is possible, it is time-consuming and heavily dependent on domain expertise. This study presents a comparative deep learning framework for the binary classification of healthy and diseased mango and apple leaves using Convolutional Neural Networks (CNNs) and frequency-domain analysis. The dataset comprises 5,394 leaf images—2,659 diseased and 2,735 healthy—evaluated across four distinct architectural configurations. Alongside standard spatial RGB representations, frequency-domain features are extracted using Haar Discrete Wavelet Transform (DWT) and Dynamic Mode Decomposition (DMD). Specifically, a baseline RGB CNN is compared against a channel-concatenated frequency-enhanced CNN, a dual-branch CNN utilizing multiplicative-addition feature fusion to process spatial and frequency streams separately, and a ResNet18 architecture. Model performance is assessed using test accuracy, precision, recall, F1-score, confusion matrices, and ROC-AUC scores. Experimental results show that the baseline RGB CNN achieved a test accuracy of 91.11% (ROC-AUC: 0.9509), whereas direct channel concatenation with DWT and DMD features yielded 85.93% (ROC-AUC: 0.9437) and the dual-branch fusion model achieved 90.19% (ROC-AUC: 0.9470). Among all evaluated approaches, ResNet18 achieved the highest performance with a test accuracy of 97.22% and an ROC-AUC score of 0.9934. These findings indicate that while naive integration of frequency-domain features does not surpass spatial RGB baselines in this experimental setting, deep residual learning provides superior feature representation for automated mango and apple leaf disease classification.***

Keywords: ***Deep learning, Convolutional Neural Network, Haar Discrete Wavelet Transform, Frequency Domain Analysis, Image classification.***

I. INTRODUCTION

Agriculture plays an important role in food production and economic development in India. The health of crops directly influences agricultural productivity and quality. Plant diseases are one of the major factors responsible for reduced crop yield and quality. Diseases affecting leaves often produce visible changes in colour, texture, shape, and surface patterns. Therefore, timely identification of healthy and diseased leaves can support early disease management and help reduce potential crop losses. However, conventional disease identification generally depends on visual inspection and expert knowledge. Such approaches can be time-consuming, particularly when many plants need to be examined, and the accuracy of identification may vary depending on the experience of the observer.

Recent developments in computer vision and deep learning have provided effective approaches for automated analysis of plant leaf images. Image-based classification systems can learn visual patterns directly from leaf images and classify them into different categories. Among various deep learning techniques, Convolutional Neural Networks (CNNs) have been widely used for image classification because of their ability to automatically learn spatial features from images. Instead of relying entirely on manually designed features, CNNs learn hierarchical representations through convolutional layers, activation functions, pooling operations, and classification layers. These characteristics make CNNs suitable for identifying visual differences between healthy and diseased plant leaves.

In the context of plant disease classification, most image-based approaches primarily utilise spatial and colour information available in RGB images. Although these features can provide important information about disease symptoms, some characteristics of leaf images may also be represented in the frequency domain. Changes in leaf texture, edges, patterns, and local variations can produce frequency components that may not be fully represented by conventional spatial-domain analysis. Therefore, incorporating frequency-domain information along with spatial image features may provide additional information for distinguishing healthy and diseased leaves.

©IJRASET: All Rights are Reserved

Wavelet-based analysis is one approach for representing image information at different frequency levels. The Haar Discrete Wavelet Transform (DWT) can decompose an image into different subbands containing low-frequency and high-frequency information. Such decomposition can provide additional representations of image texture and structural characteristics. In this study, Haar DWT is considered as a complementary feature-analysis technique together with CNN-based spatial feature extraction. Dynamic Mode Decomposition (DMD) is also incorporated as an additional feature-analysis technique to provide frequency-related information. The resulting frequency features are combined with the original RGB information to investigate whether the additional representations can improve leaf disease classification. The present study focuses on the classification of mango and apple leaf images into two categories: healthy and diseased. The dataset used in the study contains a total of 5,394 images. Among these, 2,659 images belong to the diseased category, while 2,735 images belong to the healthy category. Thus, the dataset provides a relatively balanced distribution between the two classes. The classification task is formulated as a binary classification problem, where healthy leaves are assigned label 0 and diseased leaves are assigned label 1. A CNN-based model is first developed as a baseline approach using RGB leaf images. This provides a reference against which the effect of additional frequency-domain information can be evaluated. Subsequently, we incorporate Haar DWT and DMD-based information into the input representation. A further architecture separates the RGB and frequency-domain information into different processing branches. The RGB information is processed through a CNN branch, while the frequency information is processed through a separate feature-extraction branch. The outputs of these branches are subsequently combined through feature fusion before classification. This architecture enables spatial and frequency-related information to contribute jointly to the final classification decision. In addition to the custom CNN-based approaches, ResNet18 is investigated as another deep learning model. ResNet18 uses residual connections to facilitate the learning of deeper feature representations. In the implemented system, the conventional CNN architecture is replaced by a ResNet18-style architecture containing residual blocks, batch normalisation, ReLU activation, convolutional layers, adaptive average pooling, and a final two-class classification layer. This provides an additional basis for evaluating the effectiveness of the proposed CNN and frequency-enhanced approaches. The experimental results demonstrate the performance differences among the four implemented deep learning approaches for mango and apple leaf disease classification. The baseline Convolutional Neural Network (CNN), trained using RGB leaf images, achieved a test accuracy of 91.11% with an ROC-AUC score of 0.9509. The classification report indicated precision and recall values of 0.91 and 0.92, respectively, for healthy leaves, and 0.91 and 0.90 for diseased leaves.

To investigate the contribution of frequency-domain information, Haar Discrete Wavelet Transform (DWT) and Dynamic Mode Decomposition (DMD) were incorporated into the CNN input representation. However, this frequency-enhanced CNN achieved a test accuracy of 85.93% and an ROC-AUC score of 0.9437, indicating that the direct integration of frequency-derived channels did not improve the performance of the baseline model in the present experimental setting.

A dual-branch CNN architecture was subsequently developed to process RGB image information and frequency-derived features through separate feature-extraction branches. The two branches were combined using a multiplicative-addition feature-fusion mechanism. This architecture achieved a test accuracy of 90.19% and an ROC-AUC score of 0.9470. Although the dual-branch architecture provided a different feature-learning strategy, its performance remained slightly below that of the baseline CNN.

Furthermore, ResNet18 was implemented to investigate the performance of a residual deep learning architecture. The model achieved the highest test accuracy of 97.22%, with a disease-class F1-score of 0.9706 and recall of 0.9538. These results indicate that ResNet18 provided the best classification performance among the evaluated approaches.

The main contributions of this study can be summarized as follows:

1) A deep learning-based framework is developed for binary classification of healthy and diseased mango and apple leaf images.
2) Haar DWT and DMD are investigated as complementary frequency-domain feature-analysis techniques for extracting additional information from leaf images.
3) A frequency-enhanced CNN architecture is implemented by combining RGB image information with Haar DWT and DMD-derived channels.
4) A dual-branch CNN architecture is developed to separately process spatial RGB features and frequency-derived features, followed by multiplicative-addition feature fusion.
5) A comparative experimental evaluation is conducted using baseline CNN, frequency-enhanced CNN, dual-branch CNN, and ResNet18 architectures.
6) The classification performance is assessed using accuracy, precision, recall, F1-score, confusion matrices, and ROC-AUC wherever available.

Overall, this study presents a comparative investigation of spatial-domain and frequency-domain feature utilisation in deep learning-based mango and apple leaf disease classification.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The experimental findings indicate that the effectiveness of frequency-domain information depends on how the extracted features are incorporated into the learning architecture. Among the evaluated models, ResNet18 achieved the highest test accuracy of 97.22%, demonstrating its effectiveness for the binary classification task considered in this study.

II. LITERATURE REVIEW

Earlier research established public image datasets and demonstrated that convolutional neural networks can learn disease-related patterns directly from leaf photographs. Hughes and Salathé[1] introduced an open repository of plant-health images to support the development of image-based diagnostic tools. Mohanty, Hughes, and Salathé[2] trained deep CNNs on 54,306 leaf images and reported high accuracy under controlled test conditions, while also showing that performance dropped substantially on images collected under different conditions. Sladojevic et al.[3] applied deep neural networks to classify plant diseases from leaf images, demonstrating the value of learned visual representations. Brahimi et al.[4] studied tomato disease classification and symptom visualisation, illustrating how CNN predictions can be used alongside visual interpretation. Together, these studies established deep learning as a promising approach but highlighted the importance of dataset diversity and realistic evaluation.

Subsequent work expanded the methods and crops investigated. Fuentes et al.[5] developed a deep-learning-based detector for real-time recognition of tomato diseases and pests, focusing on detection in complex images. Ramcharan et al.[6] explored image-based cassava disease diagnosis using deep learning, showing the relevance of such methods to crops important in developing regions. Ferentinos[7] evaluated several deep-learning models for plant disease detection and diagnosis across multiple plant–disease combinations. Barbedo[8] examined factors affecting the use of deep learning for plant disease recognition and discussed issues such as image acquisition, symptom variability, and dataset quality. These studies indicate that model performance depends not only on the architecture but also on the crop, disease symptoms, image conditions, and training data.

Prior research also investigated transfer learning, mobile image capture, and comparative evaluation of CNN architectures. Too et al.[9] compared deep-learning models for plant disease identification and reported that fine-tuned pretrained networks can be effective for this task. Picon et al.[10] studied deep CNNs for crop-disease classification using images captured by mobile devices, bringing attention to the gap between laboratory datasets and images acquired in practical settings. Abade, Ferreira, and Vidal[11] systematically reviewed CNN-based plant-disease recognition studies and summarised common datasets, crop types, and methodological trends. Dhaka et al.[12] surveyed deep CNNs for predicting plant leaf diseases, comparing model families, preprocessing methods, and evaluation practices. These works collectively support using CNNs and residual architectures as strong baselines, while emphasising consistent data splits and evaluation protocols for meaningful comparisons.

More recent reviews have synthesised the growing body of work and identified continuing challenges. Tugrul, Elfatimi, and Eryigit[13] and S. Das et al[17] reviewed CNN-based plant-leaf disease detection and discussed the benefits and limitations of deep models. Das et al.[14] provided a systematic analysis of deep learning for plant-disease detection and classification, including datasets, methods, and open challenges. Salka et al.[15] reviewed CNN models and datasets, noting the use of architectures such as VGG, GoogleNet, EfficientNet, and ResNet. Across this literature, most systems learn from spatial RGB appearance, while texture, edges, and fine-scale lesion patterns may also be represented through frequency-oriented analysis. This motivates the present study: Haar Discrete Wavelet Transform (DWT) and Dynamic Mode Decomposition (DMD) are treated as complementary feature-analysis/preprocessing techniques, not as deep-learning models. The learned classifiers are CNN-based networks and ResNet18, with an additional multi-branch design that fuses RGB and frequency-derived features.

III. RESEARCH METHODOLOGY

The proposed study investigates deep learning-based classification of healthy and diseased mango and apple leaf images. We have used four different model configurations, which are implemented and evaluated to examine the effectiveness of conventional spatial features, frequency-domain features, feature fusion, and residual learning.

The complete methodology consists of the following stages:

1) Dataset collection and organisation.
2) Image preprocessing and class labelling.
3) Development of a baseline CNN model.
4) Extraction of frequency-domain information using Haar DWT and DMD.
5) Development of a frequency-enhanced CNN model.
6) Development of a dual-branch CNN with feature fusion.
7) Implementation of ResNet18.
8) Model training, testing, and performance evaluation.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The four approaches are independently trained and evaluated to compare their classification capabilities.
The dataset[16] used in this study consists of mango and apple leaf images collected from the publicly available Kaggle dataset. The images are organised into two main categories: healthy and diseased. The dataset contains a total of 5394 images, distributed as follows:

Category	Number of images
Diseased leaves	2659
Healthy leaves	2735
Total	5394

The classification problem is formulated as a binary classification. Each image is assigned a corresponding numerical label. The relatively balanced distribution of the two categories provides a suitable basis for evaluating binary classification performance.
Before training the deep learning models, the leaf images are prepared for computational processing.

The preprocessing stage includes:
- Loading images from the respective healthy and diseased folders.
- Assigning numerical labels according to their categories.
- Resizing images to a uniform resolution of 224 × 224 pixels.
- Converting the images into tensor representations.
- Normalising RGB image values before model training.

The RGB images contain three channels corresponding to red, green, and blue colour information.
For the frequency-enhanced approaches, additional channels are generated using Haar DWT and DMD. These channels are combined with the original RGB representation to provide complementary image information.
A conventional Convolutional Neural Network (CNN) is initially developed as the baseline classification model. The primary objective of this model is to identify the spatial and visual characteristics that distinguish healthy and diseased mango and apple leaves.

The baseline CNN consists of the following processing stages:
1) Convolutional layers: Extract visual features such as edges, textures, and patterns from the input images.
2) ReLU activation: Introduces non-linearity into the network.
3) Max-pooling layers: Reduce the spatial dimensions of the extracted feature maps.
4) Flattening: Converts the feature maps into a one-dimensional representation.
5) Fully connected layers: Learn the relationships between extracted features and classification categories.
6) Dropout: Helps reduce overfitting during training.
7) Output layer: Produces predictions for the two classes, healthy and diseased.

The baseline CNN receives the original RGB images as input and serves as the reference model for subsequent experiments. To investigate the contribution of frequency-related information, we incorporate two additional feature-analysis techniques into the experimental framework: Haar Discrete Wavelet Transform (DWT) and Dynamic Mode Decomposition (DMD).

IV. COMPARATIVE TABLE AND OBSERVATION

The classification report, confusion matrix and ROC-AUC curve images are shown in the Annexure. Here we have shown the comparative accuracy and ROC-AUC values using a table of four different models:

SI NO.	Model / Approach	Test accuracy (%)
1	Baseline CNN	91.11
2	CNN with Haar DWT and DMD	85.93
3	Dual-Branch CNN with Feature Fusion	90.19
4	ResNet18	97.22

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

From the above comparison table, we can observe that ResNet18 achieved the highest test accuracy of 97.22%, outperforming the other three approaches. The Baseline CNN achieved 91.11% accuracy, while the frequency-enhanced CNN achieved 85.93%. The Dual-Branch CNN with Feature Fusion achieved 90.19% accuracy. The results indicate that incorporating frequency-domain features did not improve classification accuracy in these experiments.

V. FUTURE WORK

Future work may focus on multi-class disease classification, improved frequency-domain analysis, and explainable AI (XAI), where we can incorporate techniques such as Grad-CAM to highlight disease-affected regions, make classification results more interpretable, and enable real-time disease detection.

REFERENCES

[1] Hughes, D. P., & Salathé, M. (2015). An open access repository of images on plant health to enable the development of mobile disease diagnostics. arXiv:1511.08060.
[2] Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419. https://doi.org/10.3389/fpls.2016.01419
[3] Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. Computational Intelligence and Neuroscience, 2016, 3289801. https://doi.org/10.1155/2016/3289801
[4] Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for tomato diseases: Classification and symptoms visualization. Applied Artificial Intelligence, 31(4), 299–315. https://doi.org/10.1080/08839514.2017.1315516
[5] Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. Sensors, 17(9), 2022. https://doi.org/10.3390/s17092022
[6] Ramcharan, A., et al. (2017). Deep learning for image-based cassava disease detection. Frontiers in Plant Science, 8, 1852. https://doi.org/10.3389/fpls.2017.01852
[7] Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318. https://doi.org/10.1016/j.compag.2018.01.009
[8] Barbedo, J. G. A. (2018). Factors influencing the use of deep learning for plant disease recognition. Biosystems Engineering, 172, 84–91. https://doi.org/10.1016/j.biosystemseng.2018.05.013
[9] Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. Computers and Electronics in Agriculture, 161, 272–279. https://doi.org/10.1016/j.compag.2018.03.032
[10] Picon, A., et al. (2019). Deep convolutional neural networks for mobile capture device-based crop disease classification. Frontiers in Plant Science, 10, 939. https://doi.org/10.3389/fpls.2019.00939
[11] Abade, A., Ferreira, P. A., & Vidal, F. de B. (2021). Plant diseases recognition on images using convolutional neural networks: A systematic review. Computers and Electronics in Agriculture, 185, 106125. https://doi.org/10.1016/j.compag.2021.106125
[12] Dhaka, V. S., et al. (2021). A survey of deep convolutional neural networks applied for prediction of plant leaf diseases. Sensors, 21(14), 4749. https://doi.org/10.3390/s21144749
[13] Tugrul, B., Elfatimi, E., & Eryigit, R. (2022). Convolutional neural networks in detection of plant leaf diseases: A review. Agriculture, 12(8), 1192. https://doi.org/10.3390/agriculture12081192
[14] Das, P. K., Rupa, S. S., Pumrin, S., Das, U. C., & Hossen, M. K. (2024). Deep learning for plant disease detection and classification: A systematic analysis and review. Current Applied Science and Technology, 24(4), e0259016. https://doi.org/10.55003/cast.2024.259016
[15] Salka, T. D., Hanafi, M., Abdul Rahman, S. M. S. A., Mohamed Zulperi, D., & Omar, Z. (2025). Plant leaf disease detection and classification using convolution neural networks model: A review. Artificial Intelligence Review, 58, 322. https://doi.org/10.1007/s10462-025-11234-6
[16] https://www.kaggle.com/datasets/prasanshasatpathy/leaves-healthy-or-diseased
[17] Das, S., Banerjee, S., Banerjee, P., & Das, N. (2024, September). Securing Comparative Study of Leaf Disease Identification with Different Deep Learning Models Using Hyperledger Fabric. In 2024 4th International Conference on Computer, Communication, Control & Information Technology (C3IT) (pp. 1-6). IEEE.

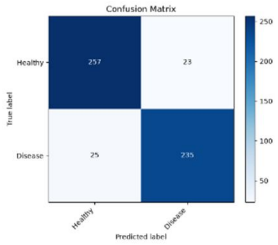
©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Annexure

Table 1: Baseline CNN: a) Classification Report b) Confusion Matrix c) ROC-AUC curve

Test Accuracy: 91.11%

Class	Precision	Recall	F1-score	Support
Healthy	0.91	0.92	0.91	280
Disease	0.91	0.90	0.91	260



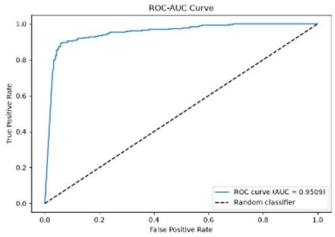
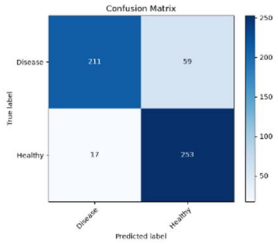


Table 2: CNN with DWT & DMD: a) Classification Report b) Confusion Matrix c) ROC-AUC curve

Test Accuracy: 85.93%

Classification Report:

Class	Precision	Recall	F1-score	Support
Disease	0.93	0.78	0.85	270
Healthy	0.81	0.94	0.87	270



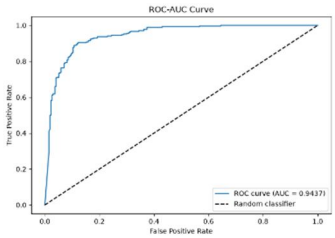
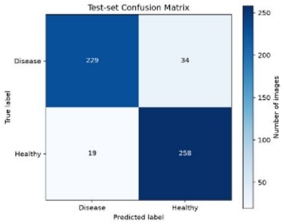


Table 3: Dual branch CNN: : a) Classification Report b) Confusion Matrix c) ROC-AUC curve

Test-set Accuracy: 90.19%

Test-set Classification Report:

Class	Precision	Recall	F1-score	Support
Disease (0)	0.9234	0.8707	0.8963	263
Healthy (1)	0.8836	0.9314	0.9069	277



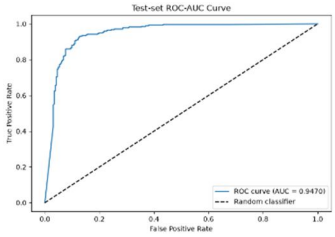
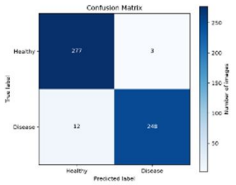


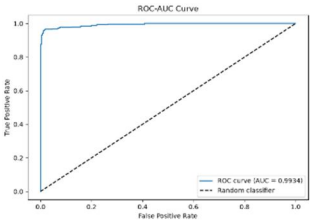
Table 4: ResNet18: : a) Classification Report b) Confusion Matrix c) ROC-AUC curve

Test Accuracy: 97.22%

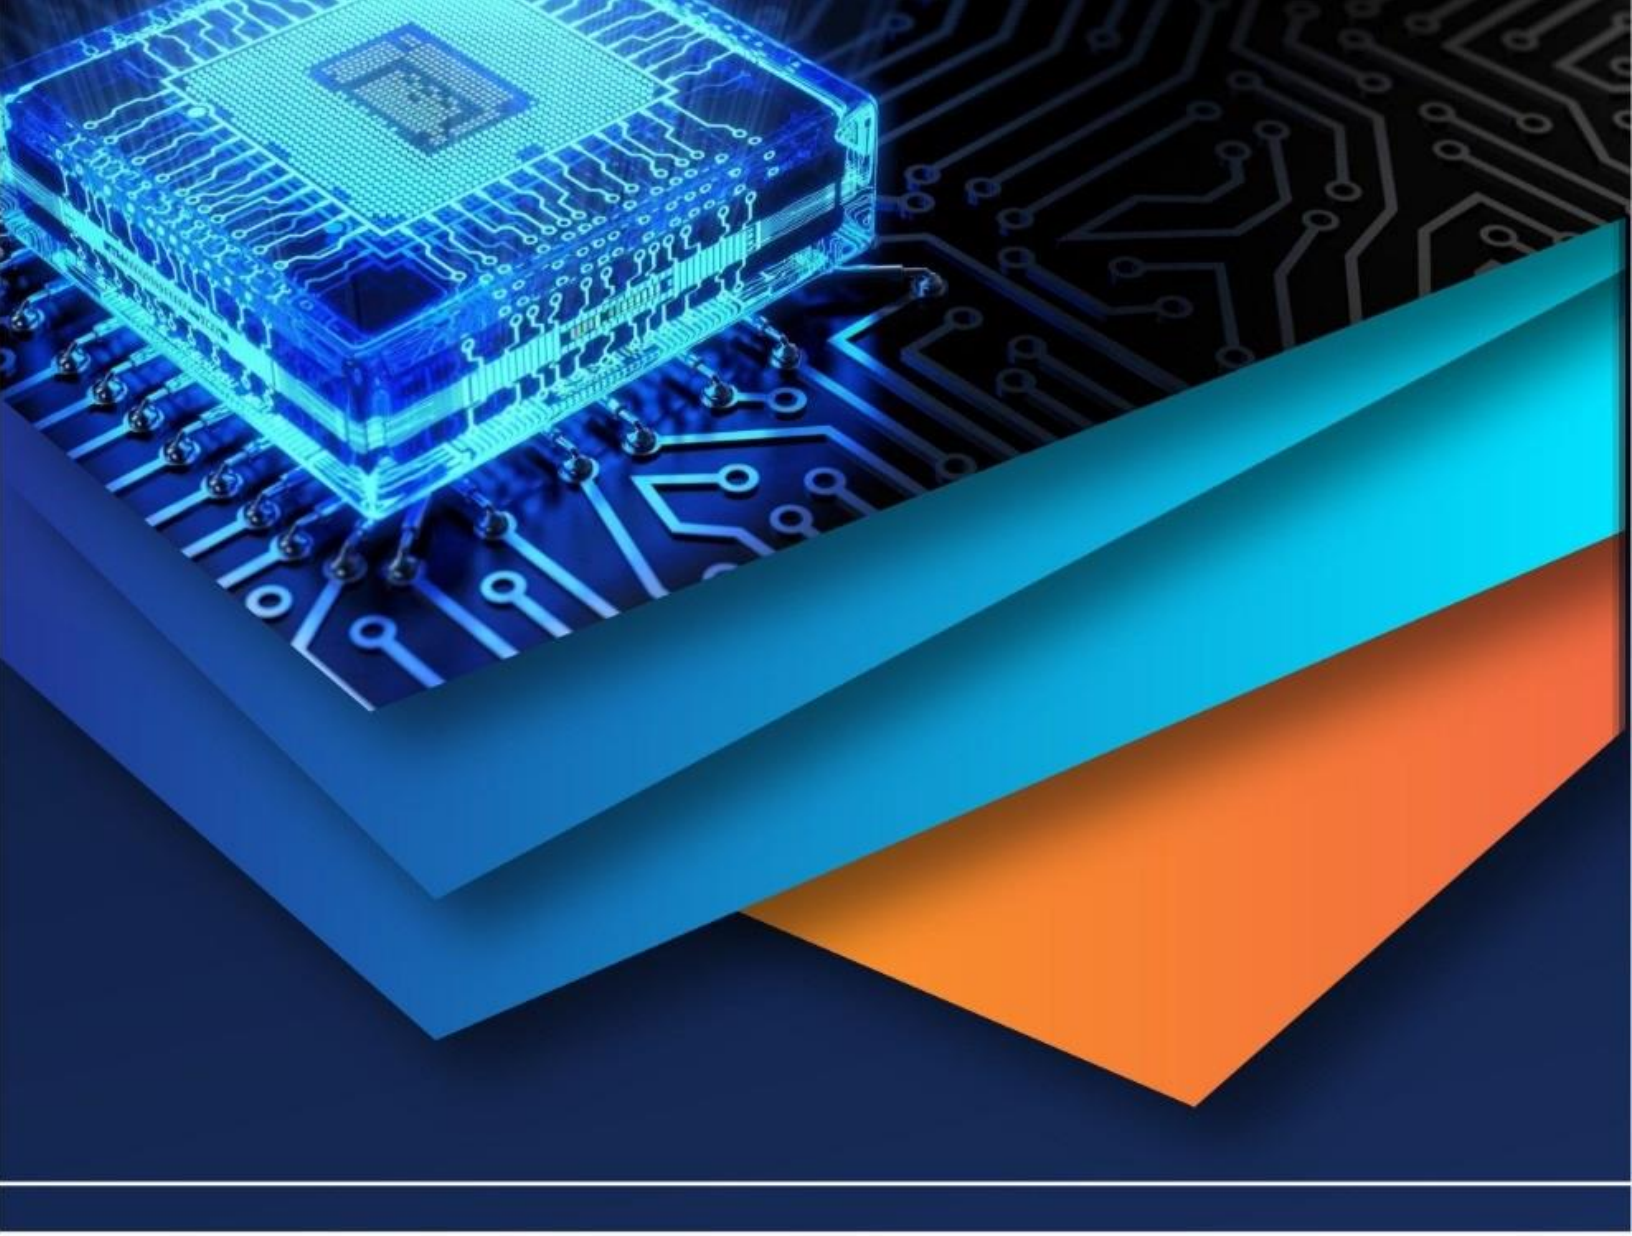
ResNet18 Classification Report:

Class	Precision	Recall	F1-score	Support
Healthy	0.9585	0.9893	0.9736	280
Disease	0.9880	0.9538	0.9706	260





©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



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)