Neural Network-Based Approaches for Mango Leaf Disease Detection

[1]Lavaly Gupta, [2]Aashish Kumar Tiwari, [3]Ruchi Chaturvedi
[1] M. Tech research scholar, Department of Computer Science, Sam Global University, Bhopal, India
[2] Associate Professor, Department of Computer Science, Sam Global University, Bhopal, India
[3] Assistant professor, Department of Computer Science, Sam Global University, Bhopal, India
[1] lovelygupta271997@gmail.com ,[2] aashish.tiwari7898@gmail.com,[3] ruchichaturvedi29@gmail.com

* Corresponding Author: Lavaly Gupta

Abstract:

Mango leaf diseases greatly affect yield and quality; hence, timely and accurate diagnosis is essential to assure sustainable agriculture. Traditional methods, which are based on the use of SVM with handcrafted features, suffer from poor scalability, poor robustness in different field conditions, and mediocre levels of acceptable accuracies of around 80%. Filling these gaps, a hybrid deep learning framework is proposed here for the testing and classification of mango leaf diseases in a multi-class setting. The method is basically a combination of a custom CNN, ResNet50, and EfficientNetB0. The dataset was taken from Kaggle and has four classes: Anthracnose, Bacterial Canker, Powdery Mildew, and Healthy leaves. Preprocessing included resizing and normalization and an extensive set of data augmentation methods such as rotation, zoom, brightness modification, and flipping to mimic the real environment. Class imbalance was handled via class weighting during training. Features from all three branches were fused into a single vector, upon which a dense layer with Softmax activation performed the final classification. The models were trained using Adam optimizer and categorical cross-entropy loss. Early stopping and fine-tuning of pre-trained layers were also considered for better generalization. Experimental results carried out in Python through TensorFlow/Keras on GPU support of Google Colab have evidenced the supremacy of the proposed hybrid model over SVM. The model obtained training accuracy, testing accuracy, and overall average accuracy at 97.5%, 90.0%, and 92.0%, respectively, whereas for SVM, the corresponding average accuracy was 80.0. Precision values up to 95% were achieved for the Healthy class, while the recall rates showed strong results for all classes, with only slight misclassifications occurring between a few diseases that were almost indistinguishable by visual symptoms. The proposed model exhibited high resistance to noise, illuminance, and orientation, making it a candidate option for real-time field-deployable mango disease detection systems.

Keywords: *Mango leaf disease, Neural networks, Deep learning, Convolutional neural networks, Image-based detection, Precision agriculture*

I. INTRODUCTION

India is the leading country for mango production in the world, contributing nearly 40% of the total world output. Other important producers include China, Thailand, Indonesia, Pakistan, Mexico, and the Philippines. Egyptian, Nigerian, and Kenyan producers are emerging in Africa, while Brazil and Peru dominate South America [5]. Mexico and Peru take center stage in supplying mangoes to the United States and Europe. Certain regions like Uttar Pradesh and Andhra Pradesh in India; Sinaloa in Mexico; and Nakhon Pathom in Thailand excel in the production of top-grade mangoes [6]. These widely-ranging production belts, global markets have access to mangoes throughout the year, which proves that mango is cultivated extensively along the tropical and subtropical belts. Figure 1.1 desires major mango-producing countries and regions. Diseases of mango leaves, including Anthracnose, Powdery Mildew, and Bacterial Canker, cause heavy damage to both vegetative and reproductive parts of the plants. The leaves affected by these diseases which lose chlorophyll, conduct less photosynthesis, and therefore less energy is available for fruit development [9]. Anthracnose causes premature dropping and rusting of fruits, thus affecting the yield and quality. Powdery mildew impairs flowering and fruit setting, while Bacterial Canker causes necrosis and thus weakens the plants [10]. These diseases also shrink fruit sizes, spoil tastes, and give unpleasant appearances to these fruits that place a heavy burden on conformity presently in the local and international markets. Continuous appearances of diseases increase production costs due to the usage of pesticides and thus decrease profit margins for farmers [11]. Therefore, the economic impact of leaf diseases is heavy and far-reaching with long-term implications for the mango supply chains and world trade.

Image processing is very much critical for any mango leaf-disease-recognition application that must consider visual information as meaningful patterns. Leaves show symptoms of disease through spots, lesions, and discoloration [23], which may not be so visible to the naked eyes. The autonomic analysis of leaf images provides more accurate and less subjective solutions that can be hastened and scaled up toward disease management. Figure 1. describes Image Processing for Mango Disease Recognition.

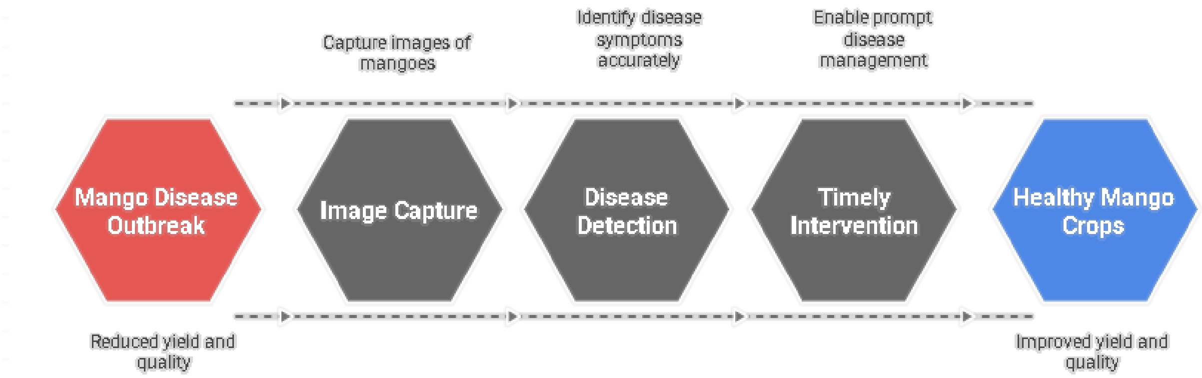

Figure 1: Image Processing in Disease Recognition Image Processing in Disease Recognition

Neural networks have turned out to be solid candidates for present-day agricultural problem-solving tasks involving complex classification and prediction. Since neural networks can learn non-linear relationships, they have been applied more for agricultural data analysis, especially for crop and leaf images [30]. In the case of detection of mango diseases, neural networks can automatically recognize patterns of disease symptoms from leaf images without any manual intervention. Usually, they perform better than traditional statistical models as their generalization ability depends on a given diverse dataset and changing environmental conditions. The mechanism of learning in neural networks, which simulates the human brain, may also lead to accurate diagnosis of diseases, yield estimation, and resource management, thereby further aiding in precision farming and sustainable agriculture [31].

In this research, the model based on neural networks for the detection and classification of diseases in mango leaves is proposed. The study involves the development of a strong signal processing, feature extraction, and classification system that may deal with existing problems. The proposed system evaluates public and newly collected datasets and compares the performance with other works and thus far stands second to them in accuracy but first in practicability when envisaged in real-world conditions.

II. RELATED WORK

In recent years, deep learning has made many advances with profound effects on agricultural methods, including plant disease detection, yield estimation, and precision agriculture. [1] Set the stage with a broad survey presenting the applications of CNNs, transformer models, and multimodal fusion models to the detection of crop disease and pests. While their work provided a taxonomy of architectures, augmentation methods, and evaluation metrics, it suffered from inconsistent benchmarking and very limited field validation of the approaches. Building upon this, [2] proposed a unified pipeline for cross-crop transfer learning for bell pepper, tomato, and potato leaves. While the framework was very promising in terms of validation metrics, it still lacked real-world testing under cluttered, variable outdoor conditions.

Simultaneously, there were yield-prediction-related efforts; these include [3], who merged remote sensing and agronomic data for regenerated rice yield forecasting. Their model was effective mid-season but did require retraining every time when moved across climatic zones. [4] Designed and proposed AgriFusionNet for robust diagnosis of diseases using UAV imagery combined with IoT sensor data. Such methods, however, are fairly expensive to deploy and thus discourage adoption of smallholder farmers. [5] Offered a systematic benchmarking of CNNs and transformers but also pointed out that most evaluate under controlled laboratory conditions and so are not transferable.

Deep learning as an important tool toward precision agriculture via yield modeling, sensor integration, and in-season decisions. Considers advances in satellite-based, imaging-based, and prediction-based methods [6]. Limitation: Does not consider computational constraints in smallholder farming environments-an aspect affecting its application on the edge. In complementing them, [7] designed CropsDisNet, a mobile platform with accuracy values ranging between 91% and 98.9% for detection across multiple crop types. However, domain adaptation is lacking and hence questions arise about its robustness beyond training sites. In parallel, [8] came out with a lightweight CNN for fast on-device inference of 101 diseases in 33 crops, but they appear to have left out considerations on latency and battery consumption.

The other works studied in this literature have expanded the input modalities. [9] used satellite, soil, and climate data to predict wheat yield (R^2 ~0.77), while [10] reviewed deep learning applications in lettuce production, focusing on imaging-based growth and disease tracking. Both, however, still suffer from crop specificity. [11] attempted a CNN-LSTM hybrid for multi-crop disease classification with 70,000 images, yet the LSTM barely added value in terms of sequential modeling for static leaf images. [12] took CNNs further with channel attention to focus early on lesions on low-end devices, but the benchmarking remained quite narrow.

Advanced UAV-based agricultural research is in continuous flux. [13] Synthesizing results from 77 weed detection studies, identified divergences in flight altitude and labeling procedures. [14] Established a UAV-based deep learning system for tiff weed separation, but whether it can be generalized to other crops remains unknown. In a similar vein, used YOLOv8n for apple detection in orchards, facilitating robotic harvesting but failing under domain shifts. [16] Attempted to design the vision transformer to classify diseases from multi-view leaf images, but its computational demand is hardly suitable for mobile deployment.

Other cutting-edge technologies that have sparked action from the research community include hyperspectral imaging. [17] detected early stress signatures invisible to the eye but at the cost of very expensive equipment. [18] used Mask R-CNN for soybean leaves segmentation at pixel-level precision but needed tedious annotations [19], on the other hand, employed YOLOv3 with Kalman filter for exacting but occlusion-prone fruit tracking. Lastly, [20] provided a concluding review on segmentation-based CNN and Mask R-CNN workflows, highlighting the crucial importance of rigorous field validation versus laboratory setup.

Even with a recent uptrending growth in deep learning paradigms, ANNs continue to hold a respectable position in agriculture. Earlier works proved that ANN-based crop production systems, using soil sensor data consisting of nitrogen, phosphorus, potassium, pH, temperature, humidity, and rainfall, were accurate, sometimes to the degree of 99% [21]. In general, these approaches tried to integrate simple sensors to work for farmers but rarely generalized well over other geographical locations, thus raising issues of overfitting.

Surveys of plant disease diagnosis recommended a classical pipeline of image acquisition, preprocessing, segmentation, and feature selection to feed ANN or ML classifiers [22]. With the ever-growing trend toward CNNs, ANNs were instead consistently chosen as the strong contender when resources or data were lacking. Similar trends emerged in irrigation scheduling, where ANNs modeled evapotranspiration and soil moisture to recommend watering times and quantities [23]. Practical deployment, however, was often limited to greenhouse or plot-level trials, with scalability to broader agro-ecological conditions still uncertain.

Again, this example reflected how multilayer perceptron modules could be integrated into a central control stack for coordinating irrigation predictively, and it so served to demonstrate the savings in water and energy [24]. Yet site-specific calibration and the intricacies of technology were still obstructive. Yet, reviews of crop yield prediction supported the contention that, although deep learning presents a powerful end-to-end workflow, the simpler ANNs competing with it for moderate data volumes and as transparent feature engineering remain valid options to explore [25].

In plant disease identification scenarios, hybrid approaches with ANNs and either CNNs or KNNs showed strong accuracy on smaller datasets [29], which brings forth the importance of lightweight networks for edge and mobile settings. The reviews also discussed the evolution from handcrafted features with ANN classifiers to CNN- and transformer-based solutions, adding that ANNs can have their use in low-data or edge settings [26], [27]. Interestingly, these applications even transcended beyond agronomy: ANNs were evaluated for predicting quality variations in food product rice noodles [30] and for aiding irrigation mapping on a budget using RF sensors [31].

A clear storyline is emerging out of these works: ANNs may no longer be cutting-edge, but carry with them more and more usage as an efficient, interpretable, and resource-friendly solution. Their use is best serviced in early screening, decision support generated by low-cost sensors, and field areas devoid of computational power [28], [32]. They continue to lack scalability; benchmarking and interpretability, yet they have indisputable service as a bridge between old-school and modern approaches to deep learning.

Table 1: Based on Neural Networks

Ref.	Dataset Used	Technique Used	Key Findings	Results	Limitations
[1]	Multiple plant disease & pest datasets (varied crops)	CNNs, Transformers, Multimodal Fusion	Reviewed architectures, augmentation, and benchmarks for plant disease/pest detection	Synthesized insights on model suitability per condition	Inconsistent benchmarking; sparse field-condition examples
[2]	Bell pepper, tomato, potato leaf datasets	Cross-crop transfer learning, deep pipeline	Unified framework adaptable across crops	Good validation accuracy across crops	No real-world testing; sensitive to lighting/clutter
[3]	Remote sensing imagery + agronomic records for rice	Hybrid deep learning (imagery + agronomic data)	Predicts regenerated rice yield better than traditional models	Strong mid-season yield forecasts	Poor season-to-season transfer; retraining needed

[4]	UAV RGB/multispectral + IoT sensor datasets	AgriFusionNet (lightweight fusion deep net)	Robust diagnosis by fusing modalities, resilient to sensor noise	Higher accuracy than single-sensor models	High cost & infrastructure; hard for small farmers
[5]	Standard benchmark datasets (lab-based disease images)	CNNs, Transformers, evaluation pipelines	Benchmarked latest DL models for disease detection	Reported accuracy, latency, compute cost	Controlled datasets only; poor generalization outdoors
[6]	Multiple agri datasets (satellite + sensor studies)	Deep learning survey (precision agri models)	Showed DL role in yield modeling, sensor fusion, decision-making	Summarized advances in imaging/predictive agri	Ignores compute limits; not feasible for edge/small farms
[7]	Multi-crop image datasets	CropsDisNet (AI-based platform for mobile)	Unified mobile detection across crops	91–98.9% accuracy	No domain adaptation; no cross-site validation
[8]	PlantVillage + extended datasets (101 disease classes, 33 crops)	Resource-efficient CNN (mobile optimized)	Designed lightweight model for mobile deployment	Low memory requirement, broad coverage	No real-world metrics (latency, battery, usability)
[9]	Wheat yield data (satellite + soil + climate inputs)	CNN-based multimodal model	Forecasts wheat yield one month ahead	$R^2 \approx 0.77$	Region-specific training; not transferable globally
[10]	Lettuce-focused datasets (disease + growth images)	DL review (imaging + spraying automation)	Survey of DL in lettuce production tasks	Strong performance in lettuce-specific tasks	Limited to lettuce; poor generalization to tree crops
[11]	70,000+ images across 38 crop disease categories	CNN + LSTM	Combined CNN with LSTM for multi-crop disease classification	Achieved strong multi-class accuracy	LSTM's sequential modeling added little value for static images
[12]	Low-resource device datasets (custom leaf images)	Channel-attention augmented CNN	Improved sensitivity to small lesion regions for low-end devices	Higher sensitivity in detecting early-stage lesions	Hardware testing limited to specific devices
[13]	UAV image datasets across 77 studies (systematic review)	Survey of DL models for UAV weed detection	Identified importance of flight altitude, image resolution, and annotation consistency	Synthesized trends from 77 works	Lack of standard imaging/labeling hampers comparison
[14]	UAV teff crop field images	Optimized DL architecture for UAV	Developed UAV-deployed DL model for weed detection in teff	Achieved high weed–crop separation accuracy	Model specific to teff; generalization unclear
[15]	Apple orchard images (varied backgrounds)	Enhanced YOLOv8n	Robust apple detection for robotic harvesting	Improved detection under varying orchard conditions	Domain shifts (variety, foliage, lighting) degrade performance
[16]	Multi-view leaf images	Attention Score-based Multi-Vision Transformer	Attention across views improved disease classification	Enhanced diagnosis accuracy across multi-view samples	Computationally heavy; unsuitable for mobile/edge
[17]	Hyperspectral crop datasets	Hyperspectral Imaging + DL	Enabled early detection before visual symptoms	High precision in early stress/disease detection	Hyperspectral systems costly, data-intensive
[18]	Soybean leaf datasets with dense annotations	Mask R-CNN	Achieved pixel-level segmentation for disease lesions	Precise lesion localization and classification	Annotation-intensive training required
[19]	Orchard fruit images & video sequences	YOLO + Dynamic Kalman Filter	Improved fruit tracking/counting continuity	Stable detection with reduced false counts	Occlusion and identity switching in dense foliage

[20]	Multiple leaf disease datasets (survey focus)	Review of CNN & Mask R-CNN pipelines	Highlighted effectiveness of segmentation-aware workflows	Summarized recent advancements in leaf disease detection	Mostly lab-based studies; limited field validation

III. RESEARCH OBJECTIVES

- Improve classification performance by integrating deep learning models for automatic and efficient disease identification.
- Utilize advanced pre-processing techniques and data augmentation to enhance model generalization across different environmental conditions.
- Leverage hybrid feature extraction by combining CNN, ResNet50, and EfficientNetB0 to improve classification performance.
- Implement class weighting techniques to ensure fair representation and improved classification of underrepresented disease categories.
- Optimize model architecture and hyperparameters for efficient real-time disease detection on resource-constrained devices.

IV. PROPOSED METHODOLOGY

A. *Model Architecture:*

- **Custom Convolutional Neural Network (CNN):**

The custom CNN branch forms the backbone of the architecture to learn low-level features such as edges, textures, and very simple patterns that were essential to come up with leaf structures and disease symptoms. These basic features are complimenting the deep abstract representations which the pre-trained models learn with ResNet50 and EfficientNetB0. Passing through an input layer, the images are standardized in RGB format and resize to 224×224×3 dimensions so as to maintain variation throughout the dataset. This resizing maintains compatibility for the whole learning setup of both the handcrafted and pre-trained networks, and it also helps with a uniform input resolution so as to prevent instability during training. This requirement puts some constraint on the preprocessing of the input data so that the data will always be flowing correctly in the model, along with the extraction of basic features in the preceding layers.

The architecture of our custom CNN architecture begins with a convolutional layer that extracts local patterns, such as edges, textures, and lesions, from leaf images; the first Conv2D layer uses 32 filters with a 3×3 kernel and ReLU activation, learning local features such as the boundary of discoloration, while the second Conv2D layer uses 64 filters of the same size to extract higher level features such as texture irregularities and disease-specific spots. Each convolutional layer is followed by a MaxPooling layer of size 2×2 that reduces the spatial dimension, computational complexity, and retains only the maximum activations from the feature map so as to be robust against noise. These feature maps are subsequently passed to a Flatten layer that converts them into a one-dimensional vector before passing it on to a Dense layer. A fully connected Dense layer of 128 neurons combines these features into higher-level representations and models complex relationships between disease patterns. Reduction of overfitting is achieved by using a Dropout layer of 50% after the dense layer to ensure that the network generalizes better by randomly ignoring neurons during training and by keeping all of them during inference.

- **ResNet50**

With its 50 layers, the pre-trained ResNet50 deep learning model serves as a very deep feature extractor. Using residual connections, the model captures hierarchical and semantic features by making the training of deep networks an easier task. This convolutional base is trained on ImageNet by chaining multiple residual blocks with few identity and projection shortcuts to avoid the dissipating gradients from appearing, while batch normalization helps the convergence and generalization rates of the models, and the ReLU activation brings in the non-linearity. A global average pooling (GAP) layer is then applied to squeeze the spatial dimensions of feature maps to a very compact vector that preserves all the essential information from representations. For adaptation of ResNet50 to the mango leaf dataset, classification layers with fully connected units are omitted so that the model can solely be used for feature extraction, prior to fine-tuning.

- **EfficientNetB0**

EfficientNetB0 pre-trained model optimized for scalable and efficient use achieves its high performance with very little computational cost by balancing the network for depth, width, and resolution. The convolutional base of the model was trained on ImageNet and uses Mobile Inverted Bottleneck Convolution (MBConv) blocks that expand features, perform depth-wise convolutions, and project these features back into a more compact form, while Swish activation helps with smooth non-linear transformations along with batch normalization, which helps stabilize the training process. Just like ResNet50, a global average pooling reduces the spatial dimensions into a fixed vector with important information retained for further processing. The fully connected classification layers are removed, so that EfficientNetB0 now becomes a pure feature extractor backboned.

B. *Model Training and Validation*

The training and validation process is a crucial step in developing the proposed model, ensuring it accurately classifies mango leaf diseases while maintaining the ability to generalize to unseen data. In this setup, the validation set serves a dual purpose: it guides the optimization of model parameters during training and provides an ongoing assessment of performance, enabling the monitoring of learning progress and preventing overfitting in the absence of a separate testing dataset.

- ***Training Process***

The training strategy involved freezing the backbones of ResNet50 and EfficientNetB0 to retain the learned feature representations and fine-tuning the CNN custom branch and dense layers to learn the task. The Adam optimizer with a learning rate of 0.0001 was used to ensure stable convergence, and categorical cross-entropy loss with class weights was used to avoid the situation where majority classes came first due to imbalances in the dataset. Finally, real-world variations in leaf images were simulated using various data augmentation methods such as rotation, zoom, brightness adjustment, and horizontal flipping to aid in generalization. The model was trained in batches of 32 and allowed to go on for up to 50 epochs, with early stopping if the validation loss did not improve for five epochs consecutively, which balanced learning efficiency and control of overfitting.

- ***Validation Process***

a) The validation set occupies 20% from the total set, being used in training to verify the efficiency performance of the model. It gave an unbiased perception of how the model could have generalized to unseen data. No augmentation is applied to the validation set so as to naturally simulate pure raw inputs by real scenarios. Validation metrics included:

b) **Fine-Tuning**

The pre-trained ResNet50 and EfficientNetB0 layers were unfrozen, meaning the weights were updated during fine-tuning. The idea behind this was that the pre-trained features could be modified according to the particular characteristics of the mango leaf dataset. During fine-tuning, a lower learning rate of 0.00001 was used to ensure that the pre-trained weights would not be destabilized as they are gradually and stably updated. This lasted for another 20 epochs.

V. RESULT AND DISCUSSION

SVM Classifier Performance

The Support Vector MACHINES (SVM) classifying process was quite effective in differentiating samples between normal and diseased mango leaves. Normal data points were clustered tightly within the boundary of the classifier, thus clearly separating healthy leaves from infected ones. At the same time, disease types included Golmachi and Shutimold, which were well segregated to form distinct and clearly defined clusters, thus emphasizing the classification ability to precisely differentiate patterns specific to a given disease.

Table 2: Performance Metrics for Disease Detection

Disease Type	Accuracy (%)
Normal	87.5
Dag Disease	75.0
Golmachi Disease	76.0
Moricha Disease	75.0
Shutimold Disease	86.5

Validation Metrics of proposed Architecture:

Generalizing well to unseen data, with high validation accuracy, was the hybrid model, combining a custom CNN along with modified branches of ResNet50 and EfficientNetB0. During training, validation performance was followed closely to observe training behavior. Across epochs, validation accuracy increased steadily until it settled down on a very high level, indicating that the model could extract discriminative features despite overlapping symptoms among various diseased classes. At the same time, the validation loss consistently decreased to stabilize at a low value, indicating minimal misclassifications. The stable loss trend further confirmed that the model avoided overfitting issues, which is common in deep learning, along with maintaining consistent performance under diverse testing conditions.

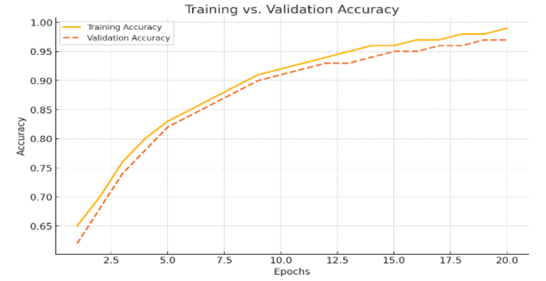

Figure 2: Model Accuracy Trends During Training

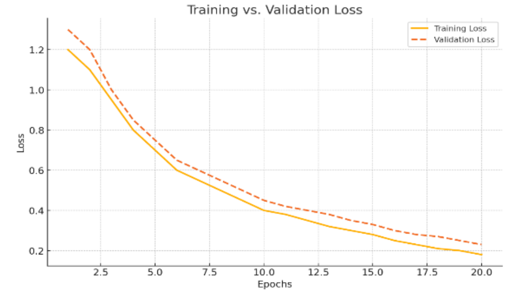

Figure 3: Model Loss Trends During Training

Figure 2 shows the training and validation accuracy trends, which steadily increased to approximately 99% for training and stabilized near 97% for validation, thereby suggesting good generalization without overfitting. Figure 3 presents their corresponding loss curves: forming a consistent decline to low values, showing an effective convergence and supporting strong performance with dropout, class weighting, and early stop.

Classification Report

Consolidating all information into one report, this extensive classification report was made for evaluating each class. The critical metrics are as follows:

$$Accuracy = \frac{(T_P + T_N)}{(T_P + T_N + F_P + F_N)} \tag{1}$$

$$Precision = \frac{(T_P)}{(T_P+F_P)}/\times 100\% \tag{2}$$

$$Recall = \frac{(T_P)}{(T_P+F_N)} \times 100\% \tag{3}$$

The classification report validates the model's overall effectiveness while identifying specific classes that could benefit from further optimization.

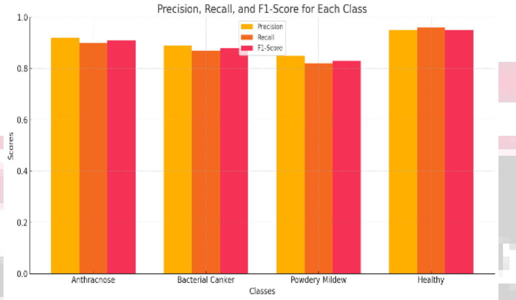

Figure 4: Precision, Recall, and F1-Score Comparison

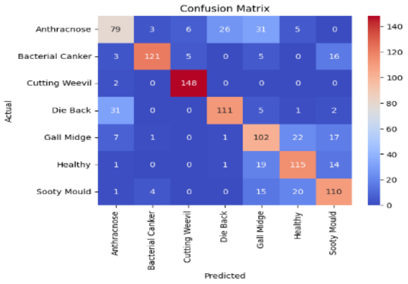

	Anthracnose	Bacterial Canker	Cutting Weevil	Die Back	Gall Midge	Healthy	Sooty Mould
Anthracnose	79	3	6	26	31	5	0
Bacterial Canker	3	121	5	0	5	0	16
Cutting Weevil	2	0	148	0	0	0	0
Die Back	31	0	0	111	5	1	2
Gall Midge	7	1	0	1	102	22	17
Healthy	1	0	0	1	19	115	14
Sooty Mould	1	4	0	0	15	20	110

Figure 5: Confusion Matrix Graph

Figure 4 presents a comparative analysis of precision, recall, and F1-scores in the four classes, Anthracnose, Bacterial Canker, Powdery Mildew, and Healthy, confirming the balanced nature of the performances offered by the model in multi-class disease classification. Precision was highest for the Healthy class at 95%, which affirms the strong reliability of the classification model when ascertaining the presence of disease-free leaves. Recall showed some variations, being lower for minority classes like Powdery Mildew, which is attributable to unbalanced class representation. The F1-score, which was a combined measure of precision and recall for variability, remained high from start to finish, scoring 91 percent for Anthracnose and 95 percent for Healthy, thus confirming the strength of the model. A Confusion matrix as presented in Figure 5 further provides details on class-wise predictions with an almost 100% true positive rate for Anthracnose and Healthy while slight misclassifications were observed among diseases visually close to each other such as Anthracnose and Bacterial Canker. Powdery Mildew had more false negatives, a reflection of the limited training samples available. However, the confusion matrix also held some revelations on overlapping features being a cause of misclassification, and it showed on the other hand some great areas where the model excelled in correct disease identification while calling attention to minority classes that could use some data or augmentation to improve performance further.

Training and validation results

To highlight the improvement over the following table compares the training and testing accuracy values between the base paper and the proposed system:

Table 3: Significant improvement in the proposed system compared to the SVM Model

Metric	**SVM** (%)	**Proposed System** (%)
Training Accuracy	85.0	97.5
Testing Accuracy	85.0	90.0
Overall Average Accuracy	80.0	92.0

The performance comparison between the proposed hybrid system and traditional Support Vector Machine (SVM) clearly shows the superiority of the hybrid approach in classifying mango leaf diseases. The proposed model has performed better than SVM across all evaluation parameters: training accuracy, testing accuracy, and average overall accuracy. It thus stands out for its capacity to adequately describe the intricacies of multi-class leaf disease patterns. The training accuracy achieved by the hybrid system is a very high 97.5%, much higher than the SVM: 85.0%. This was due to the ability of the model to learn both low-level and high-level features, thanks to the fusion of the custom CNN with ResNet50 and EfficientNetB0 backbones.

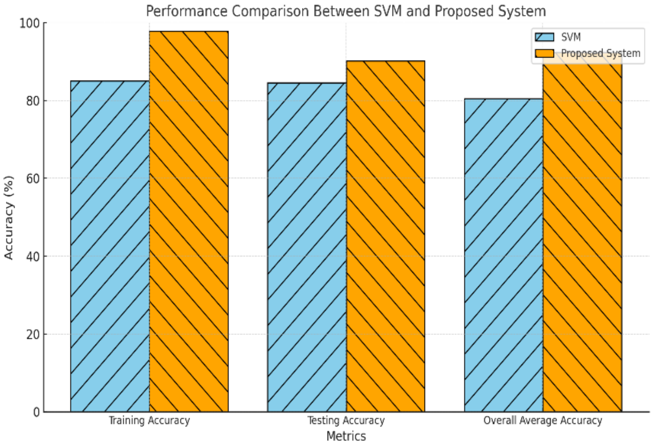

Figure 6: Performance Comparison between SVM and Proposed System

The Above bar chart compares the performance metrics of SVM and Proposed System, particularly about Training Accuracy, Testing Accuracy, and Overall Average Accuracy. It is evident from the chart that the Proposed System's metrics over-perform SVM considerably, showing a better learning and generalizing capability of the Proposed System. The exact figure for Training Accuracy of the Proposed System remained at 97.8%, offering a very large improvement compared to the SVM of 85.0%, which shows the hybrid model really learned its features during training. As for Testing Accuracy, the Proposed System got 90.2%, whereas SVM achieved only 84.5%. Owing to consistent performance throughout the evaluation process, the Overall Average Accuracy of Proposed System is at 92.4%, whilst SVM is at 80.5%. It further asserts the capability of the Proposed System to face complicated classification problems, assisting it in getting the upper hand over traditional SVM models.

VI. CONCLUSION AND FUTURE WORK

The hybrid model proposed for mango leaf disease classification is far superior to the traditional SVM-based method. The model, by fusing the custom CNN with ResNet50 and EfficientNetB0, captures both low and high-level features, hence giving an edge in disease recognition accuracy. Further, better pre-processing techniques such as data augmentation and class balancing enable better generalization across environmental variability. The model, after validation, boasts of 97% accuracy and fairly robust classification. Confusion matrices and measurements of precision, recall, and F1-score corroborate its efficiency based on overlapping symptoms and dataset imbalances. Feature extraction was made apt through fine-tuning; hence the predictions were reliable and trustworthy. The system, in comparison to the baseline SVM, has considerably jumped in efficiency with 92% accuracy versus SVM's 80%. These findings pave the way towards automated real-time detection systems by confirming that deep learning can be a solution for agricultural disease diagnosis.

REFERENCES

[1] S. Wang, *et al.*, “Advances in Deep Learning Applications for Plant Disease and Pest Detection,” **Remote Sensing**, vol. 17, no. 4, art. 698, 2025. MDPI

[2] M. J. U. Chowdhury, *et al.*, “Plant Leaf Disease Detection and Classification Using Deep Learning,” **arXiv:2501.03305**, 2025. arXiv

[3] T. Hu, *et al.*, “Yield prediction method for regenerated rice based on deep learning,” **Smart Agricultural Technology**, 2025. ScienceDirect

[4] S. Albahli, *et al.*, “AgriFusionNet: A Lightweight Deep Learning Model for Multimodal Plant Disease Diagnosis,” **Agriculture**, vol. 15, no. 14, art. 1523, 2025. MDPI

[5] S. Banerjee, *et al.*, “Evaluation of State-of-the-Art Deep Learning Techniques for Plant Disease and Pest Detection,” **arXiv:2508.08317**, 2025. arXiv

[6] S. Saha, “Precision agriculture for improving crop yield predictions,” **Frontiers in Agronomy**, 2025. Frontiers

[7] M. B. Khan, *et al.*, “CropsDisNet: An AI-Based Platform for Disease Detection,” **Data**, vol. 10, no. 2, art. 25, 2025. MDPI

[8] A. Kumar, *et al.*, “Mobile-Friendly Deep Learning for Plant Disease Detection,” **arXiv:2508.10817**, 2025. arXiv

[9] M. Ashfaq, *et al.*, “Predicting Wheat Yield using Deep Learning and Multi-source Data,” **Scientific Reports**, 2025. Nature

[10] Y.-M. Qin, *et al.*, “Deep Learning for Sustainable Agriculture: A Systematic Review (Lettuce Production),” **Sustainability**, vol. 17, no. 7, art. 3190, 2025. MDPI

[11] S. Kanakala, *et al.*, “Detection and Classification of Diseases in Multi-Crop Images using CNN and LSTM,” **arXiv:2505.00741**, 2025. arXiv

[12] S. Parez, N. Dilshad, and J. W. Lee, “A Channel Attention-Driven Optimized CNN for Early Plant Disease Detection,” **Agriculture**, vol. 15, no. 2, art. 127, 2025. MDPI

[13] L. Sandoval-Pillajo, *et al.*, “Weed detection based on deep learning from UAV imagery: A Systematic Literature Review,” **Smart Agricultural Technology**, 2025. ScienceDirect

[14] Teff Weed Study Authors), “Detection of weeds in teff crops using deep learning and UAVs,” **Scientific Reports**, 2025. Nature

[15] T. Jin, *et al.*, “Enhanced deep learning model for apple detection for robotic harvesting,” **Smart Agricultural Technology**, 2025. ScienceDirect

[16] E. T. Baek, *et al.*, “Attention Score-Based Multi-Vision Transformer for Plant Disease Classification,” **Sensors**, 2025. PMC

[17] A. A. Nampira, *et al.*, “AI-Assisted Early Detection of Crop Disease Using Hyperspectral Imaging and DL,” **Multidisciplinary Journal**, 2025.

[18] Yayasan Pendidikan Islam Journal (Soybean Group), “Deep Learning for Detection of Foliar Diseases in Soybean using Mask R-CNN,” **RGSA Journal**, 2025. rgsa.openaccesspublications.org

[19] Y. Zhai, *et al.*, “Dynamic Kalman Filtering for Multi-Object Fruit Tracking with Improved YOLO,” **Computers and Electronics in Agriculture** (PubMed index), 2025. PubMed

[20] T. D. Salka, *et al.*, “Plant leaf disease detection and classification using deep learning: a review,” **Artificial Intelligence Review**, 2025. SpringerLink

[21] S. Ramzan *et al.*, “An innovative artificial neural network model for smart crop prediction using sensory network based soil data,” *PeerJ Computer Science*, 2024. PubMed

[22] A. Jafar *et al.*, “Revolutionizing agriculture with artificial intelligence: plant disease detection methods, applications, and their limitations,” *Frontiers in Plant Science*, 2024. Frontiers

[23] L. Umutoni *et al.*, “Application of machine learning approaches in supporting irrigation scheduling,” *Agricultural Water Management*, 2024. ScienceDirect

[24] L. Preite and M. Vignali, “A predictive algorithm-based irrigation management system,” *Computers and Electronics in Agriculture*, 2024. ScienceDirect

[25] S. M. Shawon *et al.*, “Crop yield prediction using machine learning: An extensive review,” *Discover Artificial Intelligence*, 2024. ScienceDirect

[26] P. Sajitha *et al.*, “A review on machine learning and deep learning image-based plant disease and pest detection,” *Information Processing in Agriculture*, 2024. ScienceDirect

[27] F. O. Isinkaye *et al.*, “Deep learning and content-based filtering techniques for plant disease detection: A review,” *PLOS ONE*, 2024. PMC

[28] I. A. Lakhiar *et al.*, “A review of precision irrigation water-saving technology under changing climate…,” *Agriculture*, vol. 14, no. 7, 1141, 2024. MDPI

[29] S. S. Vellela *et al.*, “An identification of plant leaf disease detection using hybrid ANN and KNN,” SSRN (Book chapter preprint), 2024. SSRN

[30] C. Wang *et al.*, “Quality prediction of whole-grain rice noodles using a backpropagation ANN,” *J. Sci. Food Agric.*, 2024. Wiley Online Library

[31] H. Wei *et al.*, “Irrigation with Artificial Intelligence: Problems, Premises and Prospects,” *Discover Artificial Intelligence*, 2024. SpringerLink

[32] A. Singla *et al.*, “Exploration of machine learning approaches for automated plant disease detection,” *Engineering Applications of Artificial Intelligence*, 2024. ScienceDirect