Early Diagnosis of Tomato Plant Diseases Using Deep Learning and Image Analysis

Poonam Bari[1]*, Melissa Sequeira[1], Prasad Bari[1], Lata Ragha[2]

[1]Fr. C. Rodrigues Institute of Technology, Navi Mumbai, India
[2]Xavier Institute of Engineering, Mumbai, India

Abstract

Tomato is one of the global consumption for food. Tomato plants are very much susceptable to viruses, fungal and bacterial infections leading to various dieases. Traditional methods for detecting the diseases in tomato plants is by manual checking which is very time consuming and inefficient. So, by automating this manual process, it would be easier for farmers to detect the diseases in a shorter span of time. Convolutional neural networks (CNN) are used in this study to automatically classify tomato leaf diseases using deep learning. PlantVillage dataset was used comprising over more than 18,000 images. To adapt with real world conditions, the images were balanced and augmented across all the 10 dieases classes so that the model responses under different lighting and environment conditions. A baseline CNN model and an optimized architecture with additional convolutional layers, dropout regularization, and tuned hyperparameters were implemented and evaluated. The optimized model achieved an accuracy of 99.60% with 100 epochs, significantly outperforming the baseline model accuracy of 98.77% for 60 epochs with early stopping and 94% for 60 epochs without early stopping, with corresponding improvements in recall, precision and F1-score. Performance evaluation included loss–accuracy curves, confusion matrices, and comprehensive classification metrics, confirming reliable multi-class disease identification. The system demonstrates potential applicability for real-world agricultural decision support, although challenges such as class imbalance, varied backgrounds, and overfitting were observed. Proposed improvements include advanced augmentation, class-balancing strategies, and deployment through lightweight keras model. Overall, the findings highlight the effectiveness of CNN-based systems for plant disease detection and provide a foundation for scalable, field-ready agricultural diagnostic tools.

Keywords: Image classification, Data augmentation, Convolution neural network, Leaf disease, Deep learning,

1. Introduction

Tomato (Solanum Lycopersicum) is one of the most economically important crop that plays a vital role in global food supply chains. However, tomato cultivation faces significant challenges from various leaf diseases that can severely impact yield and quality, with more than 25% of total tomato crop production lost annually due to plant diseases (Joshi et al., 2025). Traditional diagnosis through manual inspection is often slow, subjective, and depends heavily on expert availability, which makes large-scale monitoring difficult for farmers (Karthikeyan et al., 2025). The agro-ecosystem is being negatively impacted by the quick rise in the usage of chemicals like fungicides and bactericides to prevent plant diseases, and the high frequency of diseases has an impact on both the amount and quality of produce. This creates an urgent need for automated, accurate, and accessible disease detection systems that can help farmers identify and manage crop diseases efficiently through rapid and effective early disease identification.

To discourse the difficulties, this work recommends creating a unique lightweight CNN architecture trained from scratch, specifically for tomato leaf disease classification, designed to be efficient enough for mobile and edge deployment while maintaining high accuracy. Unlike approaches that rely on heavy pretrained models such as visual geometry group (VGG) network, residual network (ResNet), or Inception network, the method prioritizes computing efficiency while maintaining diagnostic appearance, which makes it suited for real-world farming tasks with limited resources. The model is accomplished on the PlantVillage dataset with systematic preprocessing and augmentation strategies to improve generalization, while implementing hyperparameter tuning to optimize learning behavior. The final model will be prepared for deployment using lightweight formats such as TensorFlow Lite, assisting real inference on smartphones and edge devices commonly used in agricultural settings. This approach aims to provide farmers with an accessible, offline-capable diagnostic tool that can operate without continuous internet connectivity or cloud infrastructure.

Early research established foundational CNN frameworks for tomato leaf disease identification, starting with Brahimi et al. (2017), who presented deep learning for multi-class classification with nine diseases and symptom visualization on limited datasets having 15,000 images only, achieving approximately 95% accuracy using custom CNN architectures with basic convolutional and visualization layers. This foundational work progressed to dataset specific enhancements, that is, Agarwal et al. (2020) introduced tomato leaf disease detection, a fully from-scratch CNN on the PlantVillage dataset having 54,000 images, achieving 97.1% accuracy through optimized 3×3 convolutional layers, Rectified Linear Unit (ReLU) activations, and max-pooling, without any transfer learning dependencies. Trivedi et al., (2021) showed early detection capabilities using high-performance Deep Neural Networks (DNNs), incorporating deeper architectures with 8 convolutional layers, aggressive data augmentation (rotation and flipping), and softmax optimization, achieving 98.2% precision across nine disease classes. Mid-stage developments integrated advanced localization, segmentation, and compact model designs. Aggarwal (2022) developed biological classification frameworks using custom CNNs on pre-segmented leaf images, reaching 96.5% accuracy by emphasizing biologically meaningful features. Nawaz et al. (2022) proposed robust localization classification pipelines combining faster region-based CNN, yielding 98% mean average precision for accurate bounding box detection. Shoaib et al. (2022) combined U-Net-based segmentation with CNN classification pipelines, achieving 97.8% end-to-end accuracy on segmented disease regions. Özbilge et al. (2022) engineered compact CNN architectures with only 1.2 million parameters, achieving highest accuracy of 98.3% through efficient convolutional designs. Chen et al. (2022) modified the AlexNet architecture from scratch, without pretrained weights, for binary healthy versus diseased leaf classification, achieving 99.1% accuracy. Recent advancements emphasized lightweight architectures, attention mechanisms, disease severity estimation, and hybrid learning approaches. Upadhyay & Saxena (2024) enhanced ResNet-50 classifiers with architectural modifications, achieving 99.4% accuracy while still relying on transfer learning. Bedi et al. (2024) introduced the light weight framework, combining convolutional autoencoders with few-shot learning to estimate disease severity, achieving 96.7% accuracy across severity levels. Angulakshmi et al. (2024) applied hybrid deep learning techniques using ensemble CNN models, achieving 98% accuracy across ten disease classes. Karthikeyan et al. (2025) presented a comprehensive review of CNN evolution for tomato disease detection, noting that from-scratch CNN models consistently lag by approximately 2–4% compared to transfer learning approaches. Dhanya et al. (2025) integrated spatial attention mechanisms to improve focused feature extraction, achieving 99.2% accuracy. Dhar et al. (2025) combined molecular validation techniques with deep learning models for early blight severity detection, reporting 98.5% accuracy under controlled conditions. Joshi et al. (2025) utilized hybrid data augmentation strategies to improve generalization, achieving 99.3% accuracy for precision disease diagnosis.

First, shallow network depths that is less convolution layers, fail to capture hierarchical disease progression patterns, such as the transition of early blight from initial spotting to advanced necrosis stages, resulting in performance saturation at approximately 96–98% accuracy. Second, there is a lack of systematic hyperparameter optimization; kernel sizes, strides, dilation rates, and batch normalization placements are largely selected heuristically, leading to suboptimal feature learning for tomato-specific disease symptoms, as observed in several custom CNN-based studies. Third, the dominance of pre-trained backbone models limits true from-scratch deploy ability, while complex hybrid frameworks incorporating attention mechanisms or convolutional auto encoders significantly increase computational overhead and inference latency, often exceeding 200 milliseconds, thereby reducing suitability for real-time and edge-based agricultural applications (Dhanya et al., 2025). To overcome these gaps, this study presents a 10-layer CNN has been developed entirely from scratch, without relying on any pretrained models. The network architecture was carefully designed to progressively learn low-level to high-level disease features, and systematic hyperparameter tuning was performed to optimize kernel sizes, number of filters, dropout rates, batch normalization, and learning rate. This approach ensures improved feature extraction, stable convergence, and robust disease classification performance.

2. Materials and Methods

Tomato leaf infections are a major concern to the agricultural business, causing significant productivity and quality losses. Early diagnosis of these disorders is critical for effective treatment and preventative methods. Manual identification is often slow, subjective, and dependent on expert knowledge (Karthikeyan et al., 2025). Deep learning based image classification offers an efficient alternative by automatically detecting diseases from leaf images. The overall workflow of the proposed system includes data collection, dataset preparation, CNN model construction, model training, model optimization, performance evaluation, and final deployment.

2.1 Data Collection

The image dataset employed in this study was collected from the PlantVillage dataset available on Kaggle (Kaggle, 2016). The dataset contains approximately 18,000 images of tomato leaves belonging to ten classes, including Early Blight, Late Blight, Leaf Mold, Septoria Leaf Spot, Mosaic Virus, Yellow Leaf Curl Virus, Target Spot, Two-Spotted Spider Mite, Bacterial Spot, and Healthy Leaves as shown in Figure 1. All images are in Red-Green-Blue (RGB) format and were organized in class-wise directories to enable automatic labeling.

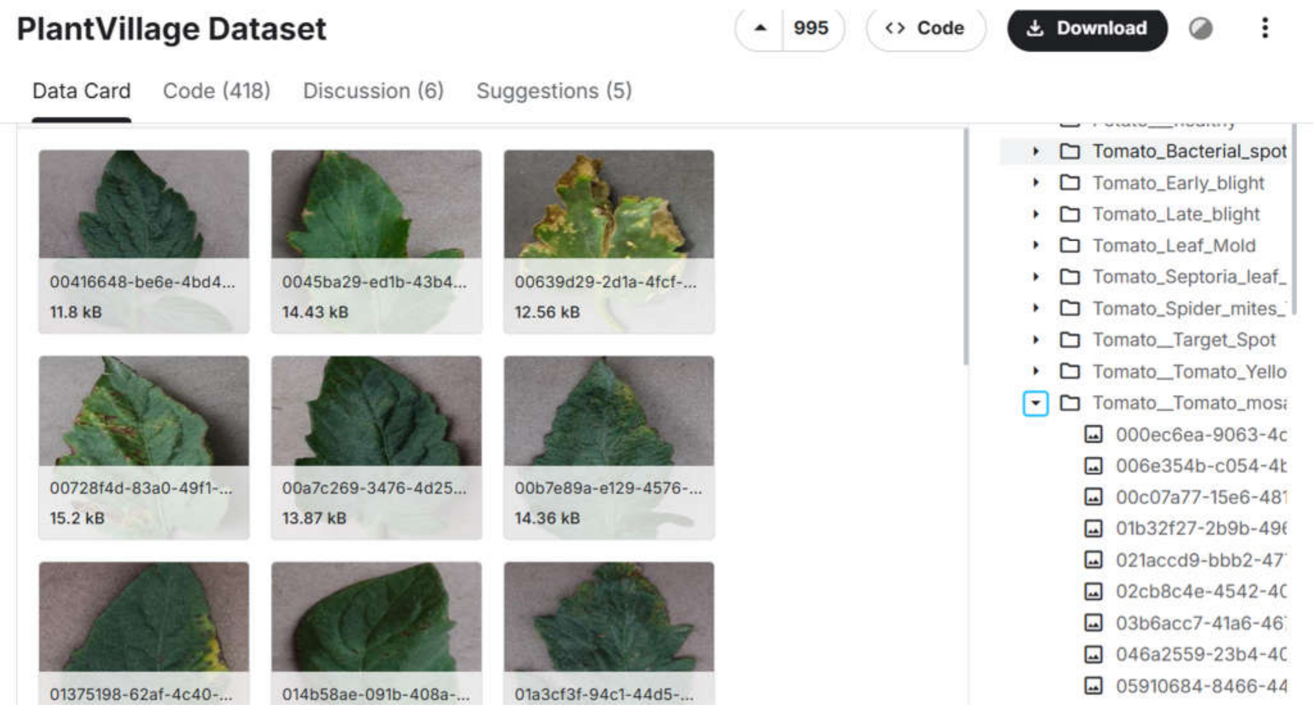

Fig. 1. Sample tomato leaf images from the PlantVillage dataset representing different disease classes.

2.2 Dataset Preparation

Dataset preparation involved preprocessing operations to balance the images across all classes to avoid class imbalance problem, to standardize and enhance the input samples. Each image was resized to 128 × 128 pixels to meet the input dimensions required for the CNN model. For faster model convergence, the value of pixels were standardized to the range of [0, 1].. To increase dataset variability and reduce overfitting, several augmentation techniques were used like random rotation, zoom, horizontal flips, and slight brightness variations. For classes exceeding the target size, excess images were randomly removed. To make sure the reproducibility of dataset splits, a fixed random seed value of 42 was used during train, validation, test splitting test (70%, 20%, 10%). Data augmentation was implemented only to the training set during model training. Table 1 shows how dataset is balanced using strong augmentation. The determined number of photos in a single class was 3,200, thus this quantity was chosen for dataset balance.

Table 1. Dataset balancing using augmentation

Class Name	Original images before balancing	Images after balancing
Tomato Bacterial spot	2127	3200
Tomato Early blight	1000	3200
Tomato Late blight	1909	3200
Tomato Leaf Mold	952	3200
Tomato Septoria leaf spot	1771	3200
Tomato Spider mites	1676	3200
Tomato Target Spot	1404	3200
Tomato YellowLeaf Curl Virus	3200	3200
Tomato Tomato mosaic virus	373	3200
Tomato healthy	1591	3200

2.3 Model Design and Training

A CNN is a special type of feed forward neural network in deep learning widely used for image analysis and pattern recognition. CNNs were inspired by the layered architecture of the human visual cortex. CNNs has the ability to automatically learn hierarchical features including edges, textures, shapes, and disease-specific patterns directly from raw images. CNNs carry out end-to-end learning, which enables reliable and scalable disease classification even under changing lighting and background conditions, in contrast to conventional machine learning techniques hat rely on manually created features.

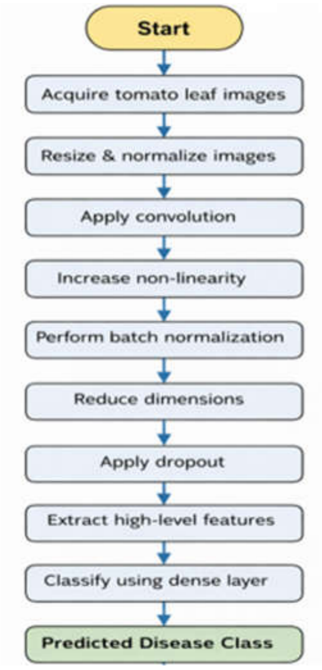

Fig. 2. Flowchart of the CNN for tomato disease detection

The flowchart in Fig. 2 shows the steps of a CNN-based system for tomato leaf disease detection. Images are collected, resized, and normalized, then passed through convolution, activation, batch normalization, pooling, and dropout layers to extract important features. In the end, a dense layer predicts the disease class and categorizes the image. A collection of measures, including training, validation, testing accuracy with the variety of parameters for training in relation to the pre-trained model, have been used to analyze the model's success. CNN have become popular for plant disease identification over the last few years. A custom CNN was developed without utilizing any pretrained networks such as VGG16 or ResNet. The proposed model includes an input layer followed by multiple convolutional blocks. Each block contains a convolution 2D (Conv2D) layer with ReLU activation and a MaxPooling2D layer to extract hierarchical visual features. The output is flattened after feature extraction and fed to fully connected dense layers for classification.

Three different neural network models were developed and evaluated to determine the optimal architecture. Model 1, consisting of 3 layers and trained for 60 epochs, resulted in a higher loss, and early stopping was applied, terminating training at epoch 21 to prevent overfitting. Model 2 used a moderately deeper architecture with 5 layers and was trained for 60 epochs; hyperparameter tuning was performed, leading to a lower loss compared to Model 1. Model 3 employed a deeper network with 10 layers and was trained for 100 epochs, along with hyperparameter tuning, and achieved the lowest loss among all models, indicating superior learning performance. Based on this comparative analysis, Model 3 was selected as the most effective model.

Fig. 3 shows the building of the planned 10-layer CNN designed for tomato leaf disease classification. The model accepts an RGB input image of size 128 × 128 × 3, followed by a sequence of five convolutional blocks. Each block consists of Conv2D layers with ReLU activation, batch normalization, and max-pooling, enabling hierarchical feature extraction while progressively reducing spatial dimensions as shown in Table 2. The number of filters increases from 32 to 512, allowing the network to learn both low-level texture features and high-level disease-specific patterns. Global Average Pooling is used instead of flattening to reduce parameter count and overfitting. The retrieved features are passing via fully connected layers and a softmax to forecast one of the ten tomato leaf disease classes.

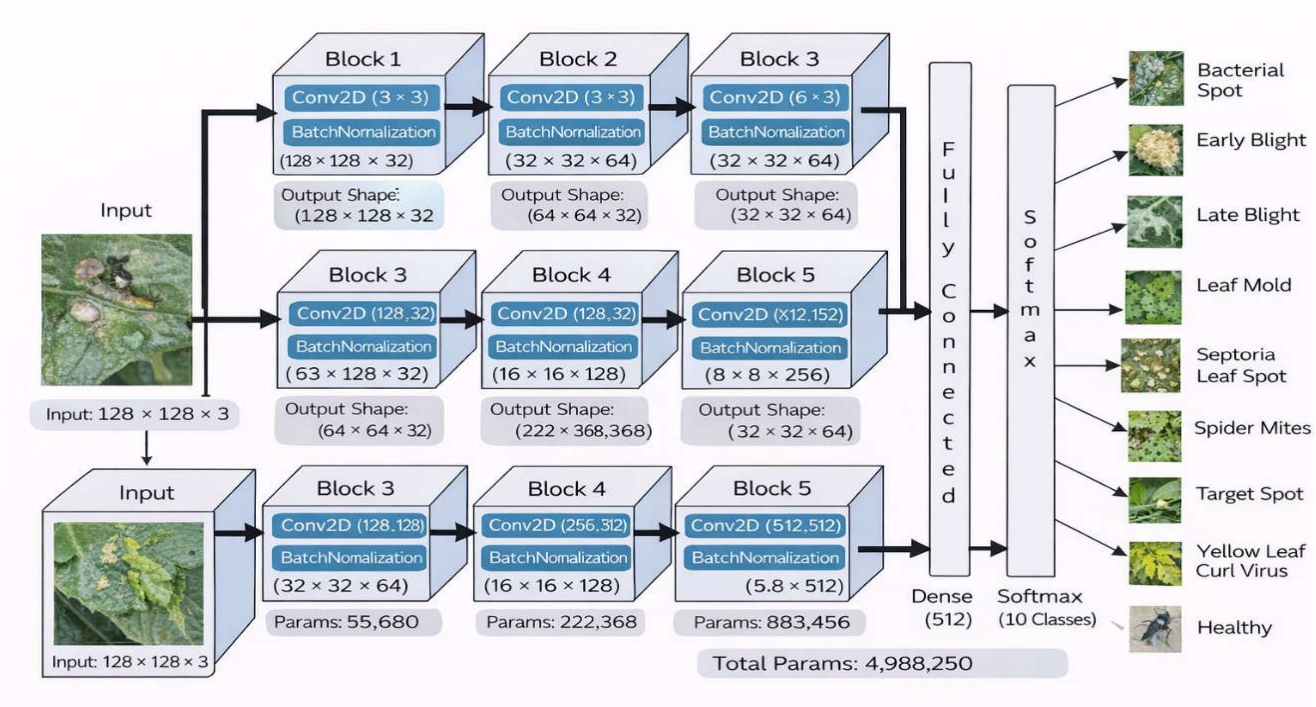

Fig. 3.10 layer CNN model structure

Table 2. Optimized CNN model layers

Stage	Layer Type	Filters / Units	Output Size
Input	RGB Image	–	128 × 128 × 3
Block 1	Conv2D + ReLU	32	128 × 128 × 32
	MaxPooling2D	–	64 × 64 × 32
Block 2	Conv2D + ReLU	64	64 × 64 × 64
	MaxPooling2D	–	32 × 32 × 64
Block 3	Conv2D + ReLU	128	32 × 32 × 128
	MaxPooling2D	–	16 × 16 × 128
Block 4	Conv2D + ReLU	256	16 × 16 × 256
	MaxPooling2D	–	8 × 8 × 256
Block 5	Conv2D + ReLU	256	8 × 8 × 256
Feature Aggregation	Global Average Pooling	–	256
Fully Connected	Dense + ReLU	256	256
Regularization	Dropout	–	256
Output	Dense + Softmax	10 classes	10

The model was built with the Adam optimizer and categorical cross-entropy loss. Accuracy was chosen as the major evaluation indicator. The CNN was initially trained with three layers during 20-60 epochs. Early stopping was used with a patience value of 10 to end training when no change in validation loss was detected after the 21st epoch. But later, after increasing the number of convolution layers from 3 to 5 and then to 10 (without early stopping), the validation loss was improving significantly after initial overfitting till 10th epoch.

3. Results and Discussions

Evaluation metrics are essential to assess model performance and verify the ability of the system to generalize effectively to unseen images. This study uses accuracy, precision, recall, and F1-score as key performance indicators.

3.1 Accuracy

According to equation (1), accuracy defines the overall correctness of the model by calculating the proportion of correctly classified instances out of the total number of instances.

$$\boldsymbol{Accuracy} = \frac{\textbf{True Positive} + \textbf{True Negative}}{\textbf{True Positive} + \textbf{True Negative} + \textbf{False positive} + \textbf{False Negative}} \tag{1}$$

3.2 Precision

The equation (2) demonstrates precision, which measures how many of the instances predicted as positive by the model are actually correct. It reflects the model's ability to avoid false positives.

$$\boldsymbol{Precision} = \frac{\boldsymbol{True\ Positive}}{\boldsymbol{True\ Positive} + \boldsymbol{False\ Positive}} \tag{2}$$

3.3 Recall

Recall (Sensitivity) indicates the model's ability to correctly identify all actual positive instances. It reflects how well the model avoids false negatives which is represented in equation (3).

$$\boldsymbol{Recall} = \frac{\boldsymbol{True\ Positive}}{\boldsymbol{True\ Positive} + \boldsymbol{False\ Negative}} \tag{3}$$

3.4 F1-Score

According to equation (4), F1-score is the harmonic mean of precision and recall, providing a balanced measure when class distribution is uneven.

$$\boldsymbol{F1\text{-}Score} = \mathbf{2} \times \frac{\boldsymbol{Precision} \times \boldsymbol{Recall}}{\boldsymbol{Precision} + \boldsymbol{Recall}} \tag{4}$$

A confusion matrix and classification report were generated to analyze class-wise prediction performance.
The baseline CNN model, that is, model 1 achieved an accuracy of 86.29%, while the optimized CNN model, that is model 3 improved performance significantly, achieving a best accuracy of 99.6%. Model 3 demonstrated stable convergence with lower validation loss and higher generalization capability. Training and testing accuracies, as well as loss values, were compared and are summarized in Table 3.

Table 3. Comparison between all trained CNN Models

Criteria	Model 1	Model 2	Model 3(best)
Train Accuracy	86%	98.77%	99.6%
Precision	87.45%	98.77%	99.77
Recall	85.85%	98.77%	99.68%
Fscore	85.66%	98.77%	99.721%

Fig. 4, Fig. 5 and Fig. 6 presents the confusion matrices of all trained CNN models. The model 1 shows noticeable misclassifications across visually similar disease classes. The intermediate model that is, model 2 reduces class confusion, while model 3 demonstrates a strong diagonal dominance, indicating highly accurate class-wise predictions. This confirms the effectiveness of architectural depth and hyperparameter optimization in improving classification reliability.

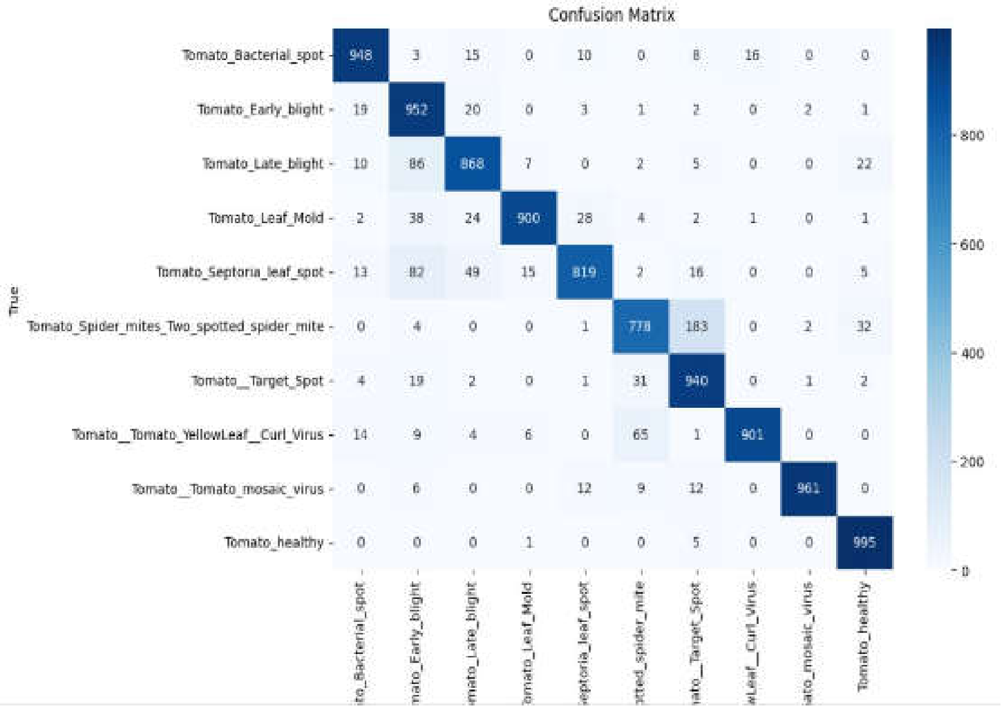

Fig. 4. Confusion matrix of model 1

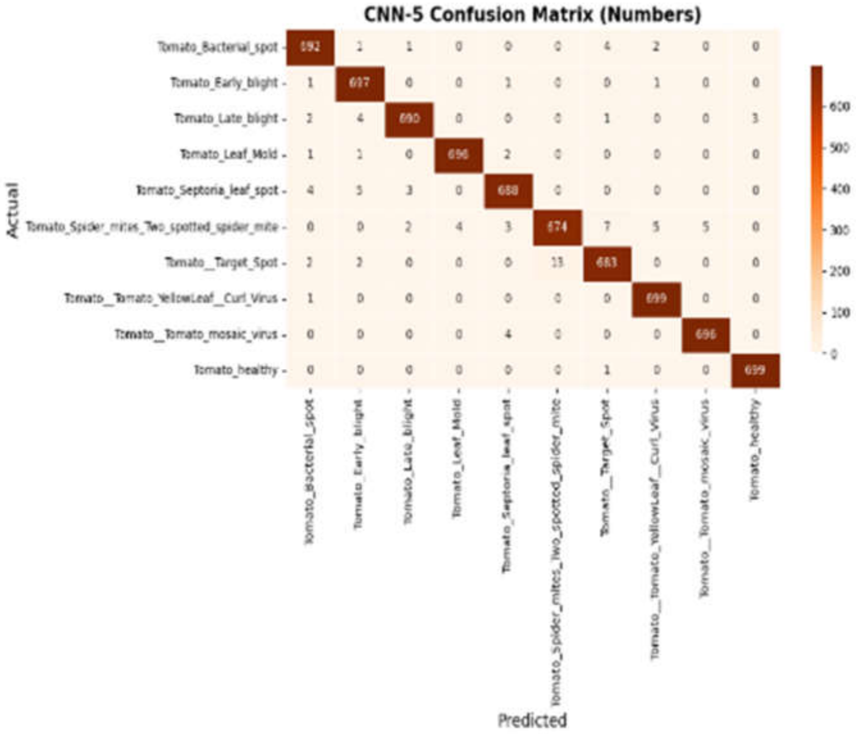

Fig. 5. Confusion matrix of model 2

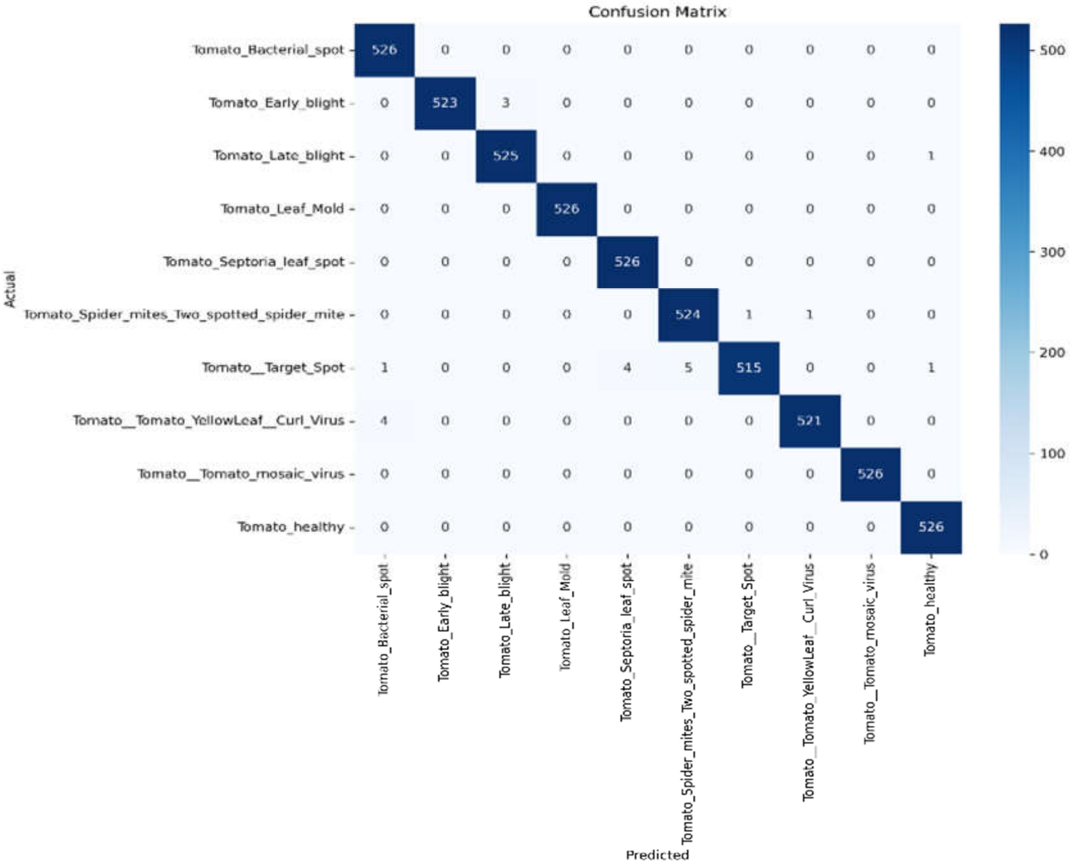

	Tomato_Bacterial_spot	Tomato_Early_blight	Tomato_Late_blight	Tomato_Leaf_Mold	Tomato_Septoria_leaf_spot	Tomato_Spider_mites_Two_spotted_spider_mite	Tomato__Target_Spot	Tomato__Tomato_YellowLeaf__Curl_Virus	Tomato__Tomato_mosaic_virus	Tomato_healthy
Tomato_Bacterial_spot	526	0	0	0	0	0	0	0	0	0
Tomato_Early_blight	0	523	3	0	0	0	0	0	0	0
Tomato_Late_blight	0	0	525	0	0	0	0	0	0	1
Tomato_Leaf_Mold	0	0	0	526	0	0	0	0	0	0
Tomato_Septoria_leaf_spot	0	0	0	0	526	0	0	0	0	0
Tomato_Spider_mites_Two_spotted_spider_mite	0	0	0	0	0	524	1	1	0	0
Tomato__Target_Spot	1	0	0	0	4	5	515	0	0	1
Tomato__Tomato_YellowLeaf__Curl_Virus	4	0	0	0	0	0	0	521	0	0
Tomato__Tomato_mosaic_virus	0	0	0	0	0	0	0	0	526	0
Tomato_healthy	0	0	0	0	0	0	0	0	0	526

Fig. 6. Confusion matrix of model 3 (optimized and best)

Fig. 7 illustrates the class-wise precision, recall, and F1-score obtained by model 3. High and consistent metric values across all ten tomato leaf disease classes indicate balanced learning and strong generalization. Minor variations observed in certain disease categories can be attributed to visual similarity and background complexity in the dataset.

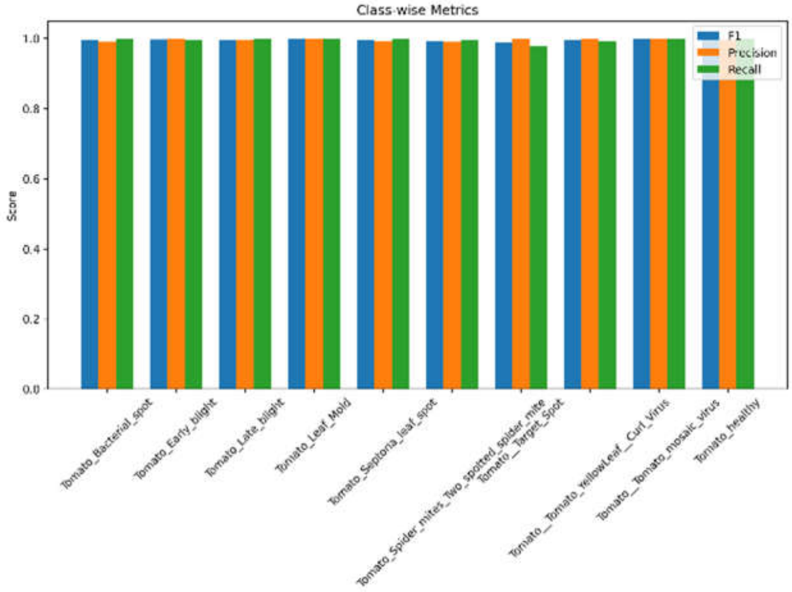

Fig. 7. Evaluation metrics of each classes

This Fig. 8 illustrates the precision-recall (PR) and Receiver-operating characteristic (ROC) curves for all tomato leaf disease classes. The curves indicate high precision and recall values with ROC curves closely approaching the ideal diagonal, demonstrating strong class discrimination and reliable classification performance across all disease categories. The colors of the other disease classes are not visible because all classes show almost the same precision and recall values. This makes their precision-recall curves overlap and appear as a single line on the graph.

Hyperparameter tuning was conducted by adjusting learning rates from 0.01 to 0.0001, batch sizes, number of filters, dense layer units, and dropout rates. Increasing convolutional depth to 10 layers and carefully selecting kernel sizes improved feature extraction capability. Reducing the learning rate from 0.01 to 0.0001 improved convergence stability. These optimizations collectively contributed to the improved model accuracy.

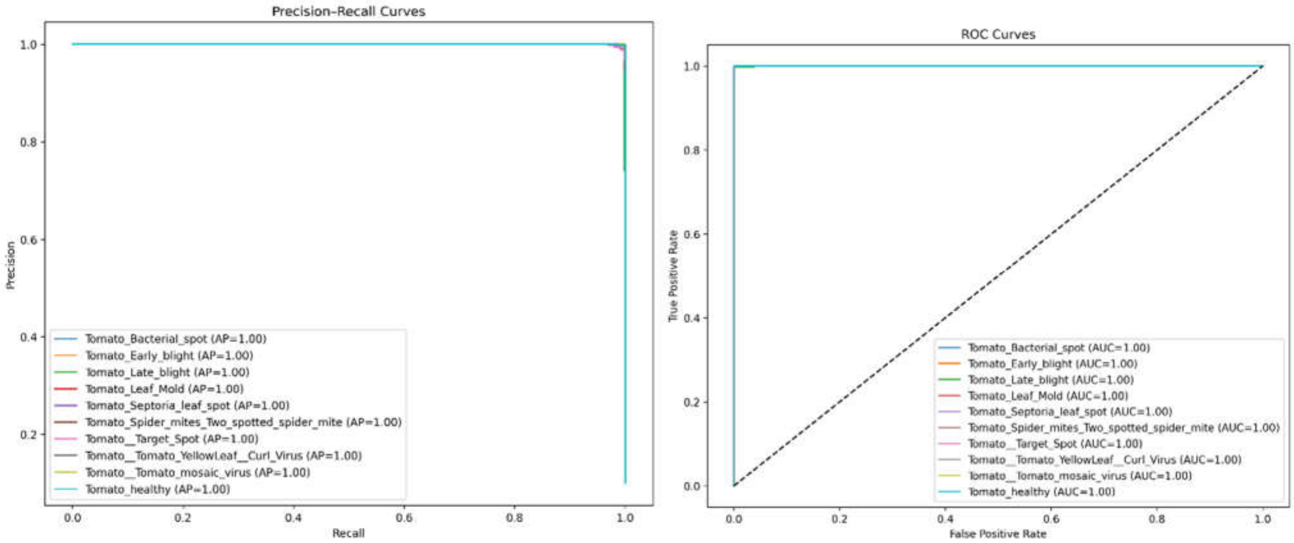

Fig. 8. PR and ROC curves

Fig. 9 shows the training and validation accuracy and loss curves for model 1. The fluctuations in validation performance indicate limited feature extraction capability with higher overfitting and early convergence, highlighting the need for deeper architectures and further optimization.

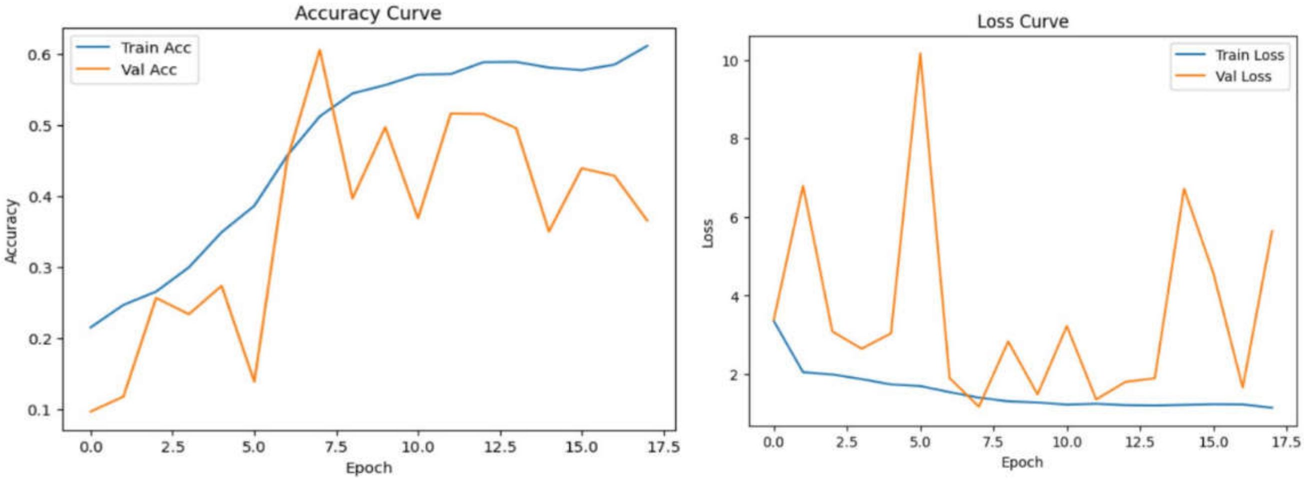

Fig. 9. Accuracy and Loss of Model 1

Fig. 10 presents the learning curves of model 2. Compared to the baseline model, the training and validation curves show improved stability and higher accuracy, indicating better feature learning and reduced underfitting

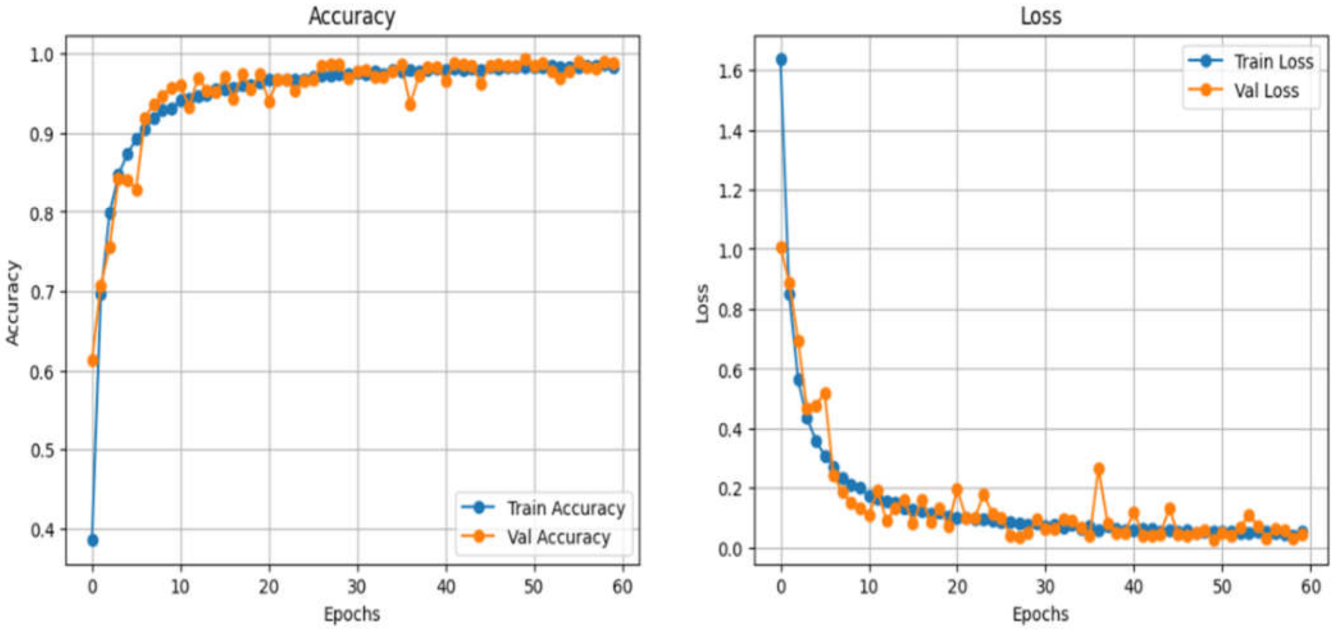

Fig 10. Accuracy and Loss of Model 2

Fig. 11 depicts the performance of model 3. The near arrangement between training and validation curves and the consistently low loss values demonstrate stable convergence, strong generalization, and minimal overfitting, resulting in the best overall classification performance.

.

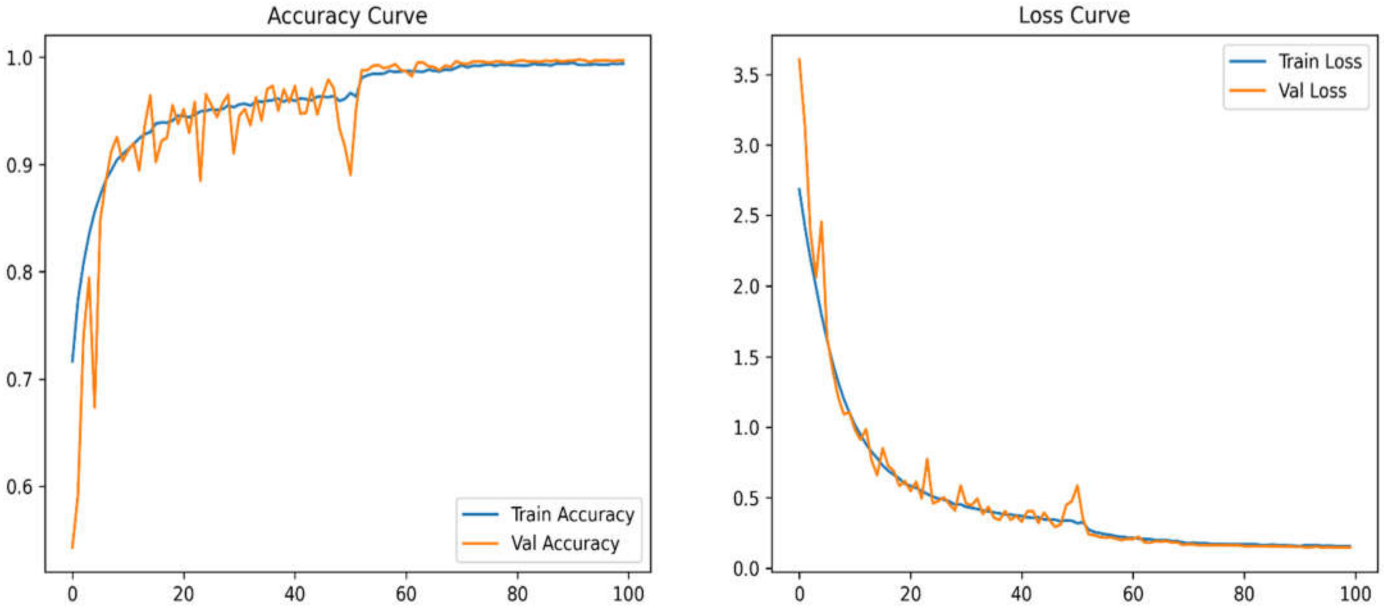

Fig. 11. Accuracy and Loss of Model 3

4. Model Deployment

Model 3 was saved in .keras file format which was later converted to .h5 to enable easy reloading for inference along with weights and architecture together. The model can be integrated into a web interface or a mobile application using TensorFlow Lite, enabling farmers and agricultural technicians to perform real-time leaf disease analysis in the field.

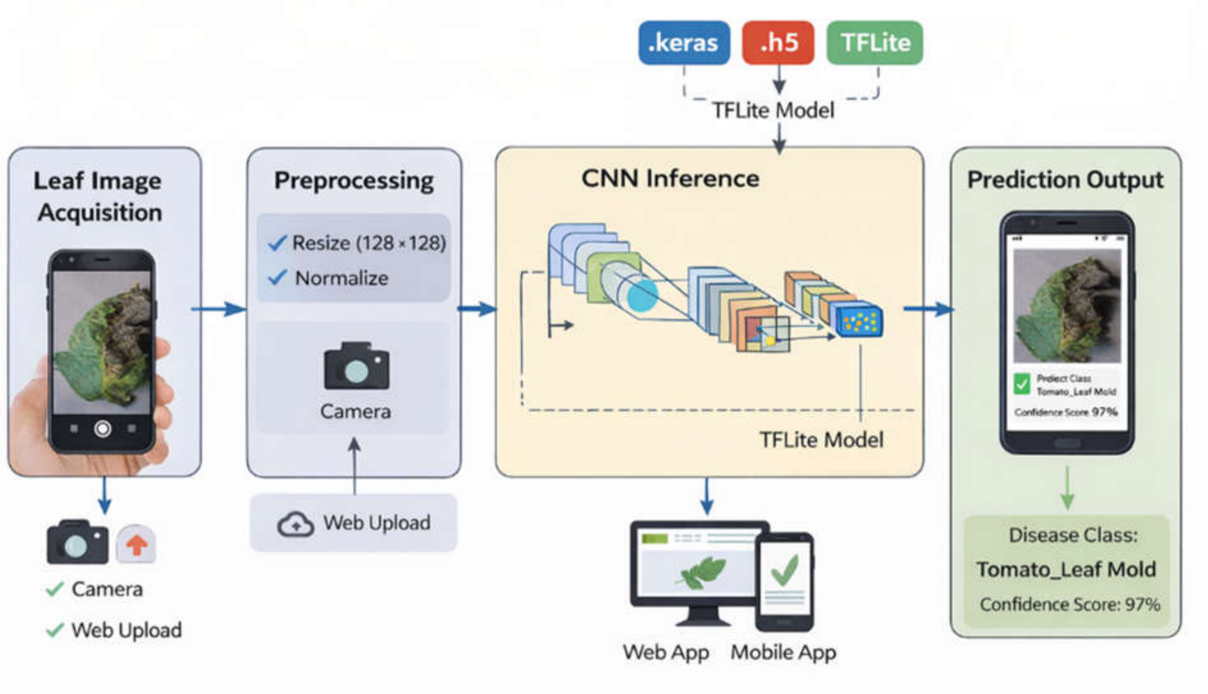

Fig. 12. Deployement architecture of proposed system

The deployment architecture in Fig. 12 illustrates the end-to-end workflow of the tomato leaf disease detection system. Leaf images are captured using a mobile camera or uploaded through a web interface and undergo preprocessing steps such as resizing to 128 × 128 pixels and normalization. The preprocessed images are then passed to the trained CNN model, converted into TensorFlow Lite format for efficient inference. The model performs real-time disease classification, and the predicted disease class along with the confidence score is displayed through a web or mobile application.

The performance of the optimized model, that is, model 3, was tested by doing practical demonstration on the set of test images. Some of the images from the test set were wrongly classified, defining the challenging cases due to disease patterns with similar visual features. The Fig.13 illustrates the limitations of the model under difficult conditions.

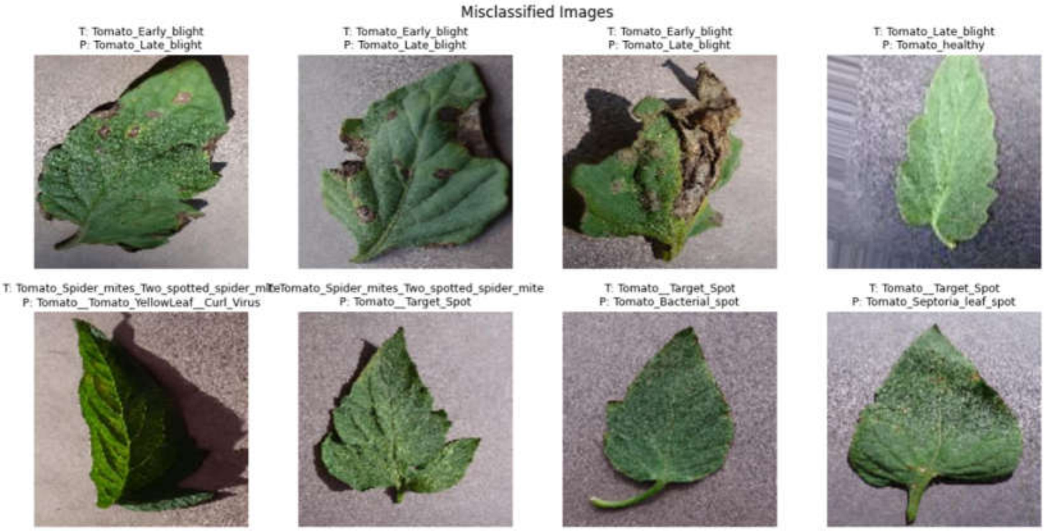

Fig. 13 . Test images that were misclassified by the model

After strong augmentation, adding more layers to build the model and doing hyperparameter tuning, model 3 with the highest accuracy of 99.8% was deployed on streamlit. Screenshots of the deployed interface in Fig. 14 shows the model successfully predicting disease classes in real-time, with high confidence scores, even under varied input conditions. The deployment results confirm that model 3 is not only accurate in controlled testing but also robust and practical for real-world usage, enabling farmers and agricultural technicians to efficiently detect tomato leaf diseases in the field.

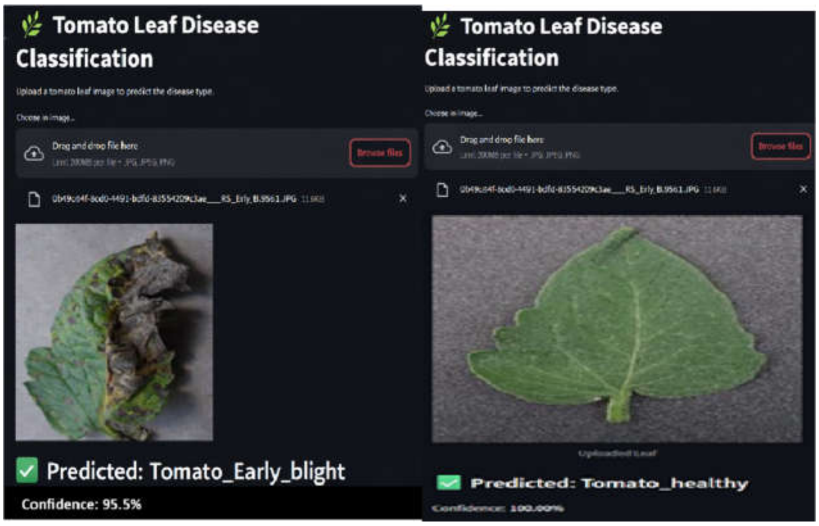
Tomato Leaf Disease Classification

Upload a tomato leaf image to predict the disease type.

Predicted: Tomato_Early_blight

Confidence: 95.5%

Tomato Leaf Disease Classification

Upload a tomato leaf image to predict the disease type.

Predicted: Tomato_healthy

Confidence: 100.00%

Fig 14. Screenshots of deployed web app showing successful predictions

5. Managerial Implication and Conclusions

This study developed a custom CNN for the automated classification of tomato leaf diseases, using the PlantVillage tomato subset as the primary dataset. The complete pipeline from data preprocessing and augmentation to model design, training, optimization, and evaluation, demonstrated that a lightweight CNN built from scratch can achieve strong performance without relying on large pretrained architectures. The proposed 10-layer Conv2D CNN (Model 3) achieved 99.6% test accuracy, 99.77% precision, 99.68% recall, and 99.72% F1-score through systematic progression from 3-layer model1 with accuracy of 86.29% to model 3 through key optimizations, including hyperparameter tuning, i.e. learning rate of 0.0001, 100 epochs (no early stopping), and filter progression from 32 to 256. With < 2M parameters and < 50ms inference, this model 3 built from scratch enables cost-effective mobile deployment for tomato farmers, real-time 10-disease detection on devices. This work establishes optimal hyperparameters for pure CNN architectures, bridging literature gaps in from-scratch models while ensuring sustainable precision agriculture. The lightweight CNN can be expanded to detect diseases in multiple crops and validated on real-world field images to improve generalization beyond the PlantVillage dataset. The explainable artificial intelligence methods such as gradient-weighted class activation mapping can be incorporated to enhance transparency and user trust by highlighting symptomatic regions. Additionally, integrating the model with IoT-based agricultural systems and early-stage disease datasets can support proactive monitoring and more sustainable precision-farming practices.

References

Agarwal, M., Singh, A., Arjaria, S., Sinha, A., & Gupta, S. (2020). ToLeD: Tomato Leaf Disease Detection using Convolution Neural Network. *Procedia Computer Science*, *167*, 293–301. https://doi.org/10.1016/j.procs.2020.03.225

Aggarwal, A. K. (2022). Biological Tomato Leaf Disease Classification using Deep Learning Framework. *International Journal of Biology and Biomedical Engineering*, *16*, 241–244. https://doi.org/10.46300/91011.2022.16.30

Angulakshmi, M., Bc, U., Vv, B. M., Siddharth, V., Venkata, A., Sai, R., & Sr, S. (2024). Tomato Leaf Disease Classification Using Deep Learning Techniques. *International Research Journal of Multidisciplinary Scope*, *5*(3), 359–375. https://doi.org/10.47857/irjms.2024.v05i03.0744

Bedi, P., Gole, P., & Marwaha, S. (2024). PDSE-Lite : lightweight framework for plant disease severity estimation based on Convolutional Autoencoder and Few-Shot Learning. *Frontiers in Plant Science*, *14*, 1–20. https://doi.org/10.3389/fpls.2023.1319894

Brahimi, M., Boukhalfa, K., Moussaoui, A., & Brahimi, M. (2017). Deep Learning for Tomato Diseases : Classification and Symptoms Visualization Deep Learning for Tomato Diseases : Classification and

Symptoms Visualization. *Applied Artificial Intelligence*, *31*(4), 299–315. https://doi.org/10.1080/08839514.2017.1315516

Chen, H., Widodo, A. M., Wisnujati, A., Rahaman, M., Lin, J. C., Chen, L., & Weng, C. (2022). AlexNet Convolutional Neural Network for Disease Detection and Classification of Tomato Leaf. *Electronics*, *11*, 1–17.

Dhanya, R., Mythili, S., Rexeena, X., & Devadas, D. (2025). Detection of Tomato Leaf Diseases Using Deep Learning and Spatial Attention Mechanisms. *International Journal of Research and Scientific Innovation*, *12*(7), 1144–1153. https://doi.org/10.51244/IJRSI

Dhar, T., Ahmad, R., Maya, B., Nigam, S., Sharma, S., Prajwal, R., Dhanger, P., Kumar, A., Kumar, T., & Yeasin, M. (2025). Results in Engineering Molecular validation integrated reliable deep learning approach for early blight severity detection in tomato crops under protected cultivation environments. *Results in Engineering*, *28*. https://doi.org/10.1016/j.rineng.2025.107739

Joshi, K., Hooda, S., Sharma, A., Sonah, H., Deshmukh, R., Tuteja, N., Gill, S. S., & Gill, R. (2025). Current Plant Biology Precision diagnosis of tomato diseases for sustainable agriculture through deep learning approach with hybrid data augmentation. *Current Plant Biology*, *41*, 1–11. https://doi.org/10.1016/j.cpb.2025.100437

Kaggle. (2016). *PlantVillage Dataset*. Kaggle. https://www.kaggle.com/datasets/keerthanak97/plantvillage-dataset/data

Karthikeyan, S., Vamsi, T., Akhil, D., Anwar, S., & Pardhasaradhi, K. (2025). A REVIEW OF TOMATO LEAF DISEASE DETECTION USING CONVOLUTIONAL. *Global Journal For Research Analysis*, *14*(2), 84–86.

Nawaz, M., Nazir, T., Javed, A., Masood, M., Rashid, J., Kim, J., & Hussain, A. (2022). A robust deep learning approach for tomato plant leaf disease localization and classification. *Scientific Reports*, *12*, 1–18. https://doi.org/10.1038/s41598-022-21498-5

Özbilge, E., Ulukök, M., Toygar, Ö., & Özbilge, E. (2022). Tomato Disease Recognition Using a Compact Convolutional Neural Network. *IEEE Access*, *10*, 77213–77224.

Shoaib, M., Hussain, T., Shah, B., Ullah, I., Shah, S. M., Ali, F., & Park, S. H. (2022). Deep learning-based segmentation and classifi cation of leaf images for detection of tomato plant disease. *Frontiers in Plant Science*, *13*, 1–18. https://doi.org/10.3389/fpls.2022.1031748

Trivedi, N. K., Gautam, V., Anand, A., Aljahdali, H. M., Villar, S. G., Anand, D., Goyal, N., & Kadry, S. (2021). Early Detection and Classification of Tomato Leaf Disease Using High-Performance Deep Neural Network. *Sensors*, *21*, 1–15.

Upadhyay, L., & Saxena, A. (2024). Evaluation of Enhanced Resnet-50 Based Deep Learning Classifier for Tomato Leaf Disease Detection and Classification. *Journal of Electrical Systems*, *20*(3), 2270–2282.