IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511123

Tomato Leaf Disease Detection with a ResNet50-Based Deep Neural Network

Jyothi K[1], Amulya N Dummer[2], Anu H S[3], Deeksha U Harige[4], Deekshita G Naik[5]

Department of Information Science and Engineering, JNN College of Engineering, Shivamogga, India[1,2,3,4,5]

Abstract: Tomato crops are affected by several leaf diseases that can reduce plant health and crop yield. Identifying symptoms early can help farmers take appropriate action, but visual inspection may be time-consuming and difficult when symptoms from different diseases look alike. This study presents a deep learning system that classifies tomato leaf images using ResNet50 and transfer learning. The model predicts one of ten categories: Bacterial Spot, Early Blight, Late Blight, Leaf Mold, Septoria Leaf Spot, Spider Mites (Two-Spotted Spider Mite), Target Spot, Tomato Yellow Leaf Curl Virus, Tomato Mosaic Virus, and Healthy. Images are resized to 224 × 224 pixels and processed by an ImageNet-pretrained ResNet50 model. Training is performed in two stages: the classification head is trained first, followed by fine-tuning of selected deeper layers. The dataset contains 9,027 training images, 1,000 validation images, and 1,004 test images. On the test set, the model achieved 99.10% accuracy, 99.10% balanced accuracy, a 99.10% macro F1-score, and a 99.10% weighted F1-score. Cohen's kappa was 0.9900, and the Matthews correlation coefficient was 0.9901. Grad-CAM was used to visualize image regions influencing predictions. A separate colour-based analysis estimated visible discolouration on selected leaves; this estimate is indicative and has not been validated as a scientific measurement of disease severity.

Keywords: Tomato leaf disease, deep learning, ResNet50, transfer learning, CNN, plant disease detection, Grad-CAM, image classification.

I. INTRODUCTION

Tomato is widely cultivated for household consumption and commercial use. However, tomato plants can develop diseases that appear as spots, discolouration, lesions, mould, or other visible changes on their leaves. If these symptoms are not recognized promptly, plant growth and crop yield may be affected.

Farmers and agricultural specialists commonly identify diseases by examining visible symptoms. Although this approach can be effective when carried out by experienced people, it becomes more challenging when many plants need to be inspected or when different diseases produce similar symptoms. These difficulties create a need for tools that can support faster and more consistent screening.

Computer vision and deep learning provide a way to analyse leaf images automatically. Convolutional neural networks (CNNs) learn visual patterns such as colour, texture, and shape directly from image data, reducing the need to design these features manually. However, training a deep network from the beginning can require substantial labelled data and computing resources. Transfer learning addresses this challenge by adapting a model pretrained on a large image collection to a new task.

This work applies transfer learning with ResNet50 to classify tomato leaf images into ten categories. Grad-CAM is also used to visualize regions that influence the model's predictions. The aim is to develop and evaluate an image-classification approach while making its predictions easier to inspect.

II. RELATED WORK

Earlier plant-disease identification methods often relied on visual inspection or image-processing techniques. Traditional computer-vision approaches commonly extracted features such as colour, texture, shape, and lesion patterns before passing them to a classifier. These methods can be useful, but their performance may depend on the quality of the selected features and the conditions under which images are captured.

Recent studies have explored CNNs and other deep learning architectures for plant and tomato leaf disease classification. Research has examined different CNN architectures, hybrid CNN-based methods, feature fusion, and approaches intended to improve deployment efficiency. Reviews of machine learning and deep learning for plant disease identification also describe the growing use of image-based classification in agriculture.

Building on this research direction, the present work uses an ImageNet-pretrained ResNet50 model and fine-

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511123

tunes selected layers for ten tomato leaf categories. It also includes Grad-CAM visualizations and an indicative colour-based estimate of the visibly affected leaf area.

III. PROPOSED METHODOLOGY

A. System Overview

The system follows an image-classification pipeline, shown in Fig. 1. A tomato leaf image is resized and preprocessed before being passed to ResNet50. The network extracts visual features, and the classification head produces scores for the ten categories. The category with the highest predicted probability is returned as the model's output. Grad-CAM can then be generated to visualize image regions that contributed to the prediction.



Fig. 1 System overview of the proposed pipeline

Pipeline: Input image → Image preprocessing → ResNet50 feature extraction → Fine-tuning and classification → Predicted class → Grad-CAM visualization.

A tomato leaf image is first given as input. The image is converted to 224 × 224 pixels so that it matches the input requirement of the ResNet50 model. The processed image is then passed through the pretrained network. The ResNet50 model extracts different levels of visual information from the image. The final classification part uses these features to determine the leaf class among the 10 classes. After obtaining the prediction, Grad-CAM can be used to visualize the regions that contributed to the prediction.

B. ResNet50 Model

ResNet50 is a 50-layer deep neural network. One of its important features is that it uses **residual connections**. In a normal deep neural network, increasing the number of layers can make training difficult. ResNet solves this problem by adding shortcut connections between layers. Instead of learning the complete mapping directly, the network learns a residual function. The basic residual operation can be represented as:

$$y = F(x) + x \qquad (1)$$

where x is the input, F(x) is the transformation learned by the network, and y is the output.

In our project, we did not train ResNet50 completely from scratch. We used the ImageNet-pretrained ResNet50 and adapted it for our ten tomato leaf classes. This helped us make use of the features already learned by the model.

C. Dataset

The dataset used in the experiment contains ten categories of tomato leaf images. The images are divided into training, validation, and test sets, as listed in Table I. The classes are: (1) Bacterial Spot, (2) Early Blight, (3) Late Blight, (4)

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511123

Leaf Mold, (5) Septoria Leaf Spot, (6) Spider Mites, (7) Target Spot, (8) Tomato Yellow Leaf Curl Virus, (9) Tomato Mosaic Virus, and (10) Healthy. Sample images of each class are shown in Fig. 2.

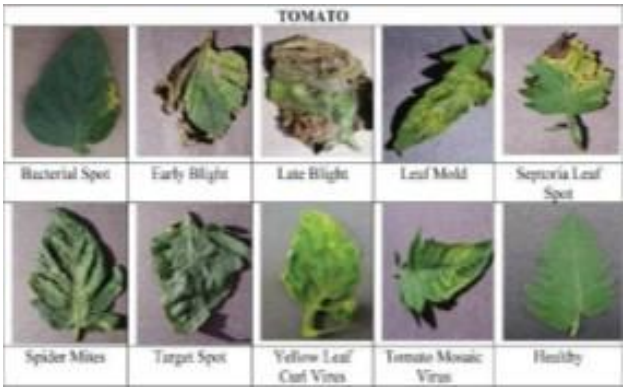


Fig. 2 Sample leaf images from the ten classes TABLE I DATASET SPLIT

Dataset	Number of Images
Training	9,027
Validation	1,000
Testing	1,004
Total	11,031

Each image is resized to 224 × 224 pixels to match the model's expected input dimensions. The same input preparation procedure should be applied consistently during training, validation, and testing.

D. Transfer Learning

We used transfer learning in two stages.

1) Phase 1 – Training the Classification Head:

In the first stage, the newly added classification head is trained for five epochs with a learning rate of 0.001. The purpose of this stage was to allow the newly added classification layer to learn how to separate the ten tomato leaf categories using the features provided by the pretrained ResNet50 model.

2) Phase 2 – Fine-Tuning:

In the second stage, selected deeper layers, mainly layers 3 and 4 as reported in the implementation, are fine-tuned for twelve epochs with a lower learning rate of 0.0001. The lower learning rate is intended to adapt pretrained representations to tomato leaf images without making unnecessarily large changes to the weights learned during pretraining. Therefore, the complete training process consisted of 17 epochs.

E. Disease Classification

The classification head produces a score for each of the ten classes. The Softmax function converts these scores into probabilities:

$$P(y_i) = e^{z}i \,/\, \Sigma_{j=1}^{10}\, e^{z}j \quad (2)$$

where z_i is the score produced for a particular class i. The class with the largest probability is selected as the prediction. For example, if Early Blight receives the highest probability, the system labels the input image as Early Blight.

F. Training and Validation

The training set is used to update model parameters, while the validation set is used to monitor performance on images that are not used for those parameter updates. The training curves reported for this experiment show accuracy increasing and loss decreasing during training (Fig. 3). The dashed line in the graph marks the point at which fine-tuning

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511123

begins. After fine-tuning, the accuracy increases considerably and reaches close to 99%. The training and validation losses decrease during training, which shows that the model is learning useful features from the tomato leaf images.

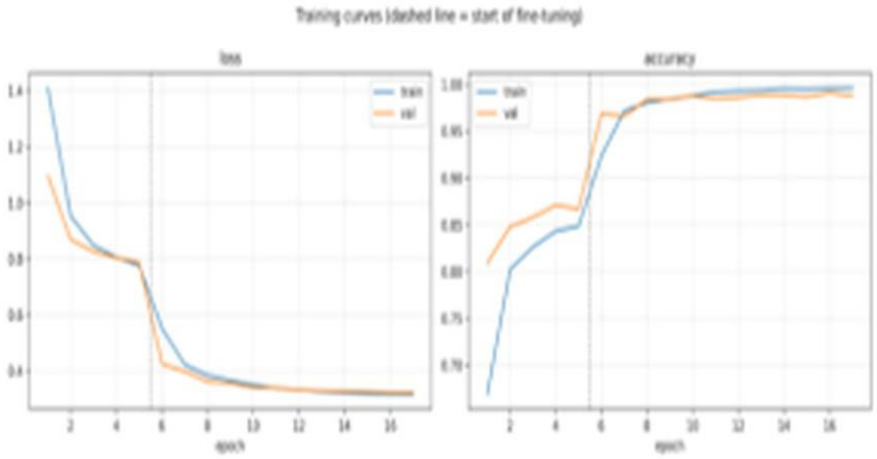


Fig. 3 Training curves (dashed line marks the start of fine-tuning)

G. Grad-CAM Visualization

One limitation of a normal deep learning classifier is that we can see the prediction but cannot directly understand which part of the image influenced that prediction. To address this, we used Grad-CAM. Grad-CAM produces a heatmap over the input image. The highlighted regions indicate areas where the model is giving more attention while making its prediction.

For example, Fig. 4 shows Grad-CAM results for different correctly classified diseases. The model predicted Bacterial Spot with 95.7% confidence, Early Blight with 92.9%, Late Blight with 93.1%, and Leaf Mold with 97.0%. Similar visualizations are provided for the remaining classes.

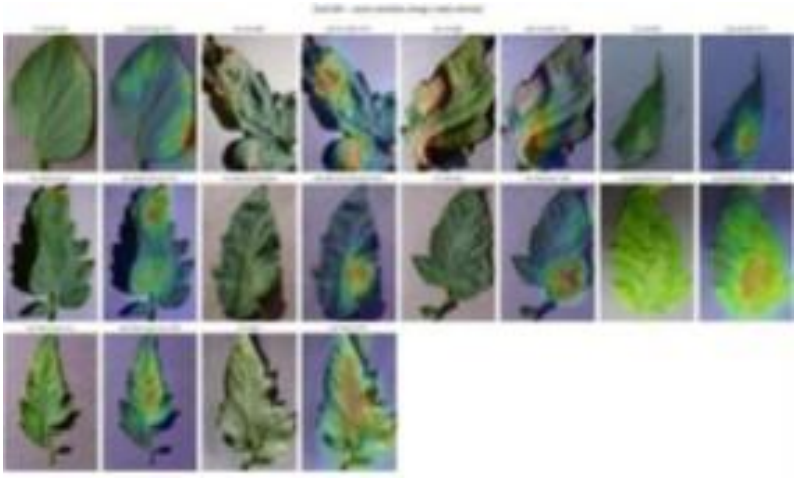


Fig. 4 Grad-CAM visualizations for correctly classified samples

The Grad-CAM results are useful for understanding model behaviour. However, the highlighted region should not be considered an exact disease segmentation because Grad-CAM shows **model attention**, not a medically or agriculturally validated disease boundary.

H. Affected Region Analysis

As an additional part of our project, we performed a colour-based analysis to estimate the portion of the leaf showing visible discolouration. Some example results are given in Table II and Fig. 5. These values give an approximate idea of how much of the leaf appears discoloured in the selected images.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511123

TABLE II EXAMPLE COLOUR-BASED AFFECTED-AREA ESTIMATES

Disease	Confidence	Estimated Affected Area
Bacterial Spot	95.67%	9.19%
Early Blight	92.92%	87.54%
Late Blight	93.10%	64.33%
Leaf Mold	96.98%	7.78%

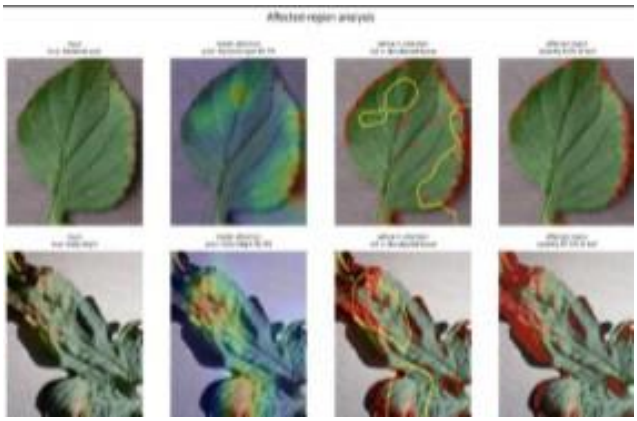


Fig. 5 Affected region analysis

A class-wise colour-based severity analysis was also produced. However, this should be clearly described as an **indicative colour-based estimate**, and not as an exact measurement of disease severity.

IV. EXPERIMENTAL RESULTS

A. Overall Results

The final model was evaluated on 1,004 test images. The obtained results are listed in Table III. These results show that the model performed very well on the test dataset.

TABLE III OVERALL TEST RESULTS

Metric	Result
Accuracy	99.10%
Balanced Accuracy	99.10%
Macro F1-Score	99.10%
Weighted F1-Score	99.10%
Cohen's Kappa	0.9900
Matthews Correlation Coefficient	0.9901

B. Class-Wise Results

The class-wise results obtained from the test set are given in Table IV.

TABLE IV CLASS-WISE TEST RESULTS

Class	Support	Precision	Recall	F1-Score
Bacterial Spot	102	0.9902	0.9902	0.9902
Early Blight	102	1.0000	0.9902	0.9951
Late Blight	100	1.0000	0.9800	0.9899
Leaf Mold	100	1.0000	1.0000	1.0000



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511123

Septoria Leaf Spot	100	0.9900	0.9900	0.9900
Spider Mites	100	0.9524	1.0000	0.9756
Target Spot	100	0.9897	0.9600	0.9746
Tomato Yellow Leaf Curl Virus	100	1.0000	1.0000	1.0000
Tomato Mosaic Virus	100	1.0000	1.0000	1.0000
Healthy	100	0.9901	1.0000	0.9950

The most notable error is the classification of four Target Spot images as Spider Mites. A small number of errors were also reported among Bacterial Spot, Early Blight, Late Blight, and Septoria Leaf Spot. These cases suggest that visual similarities between disease symptoms remain a challenge.

C. Confusion Matrix

The confusion matrix (Fig. 6) gives a clearer idea about the mistakes made by the model. Out of the test images, the correctly classified counts were: Bacterial Spot 101/102, Early Blight 101/102, Late Blight 98/100, Leaf Mold 100/100, Septoria Leaf Spot 99/100, Spider Mites 100/100, Target Spot 96/100, Tomato Yellow Leaf Curl Virus 100/100, Tomato Mosaic Virus 100/100, and Healthy 100/100.

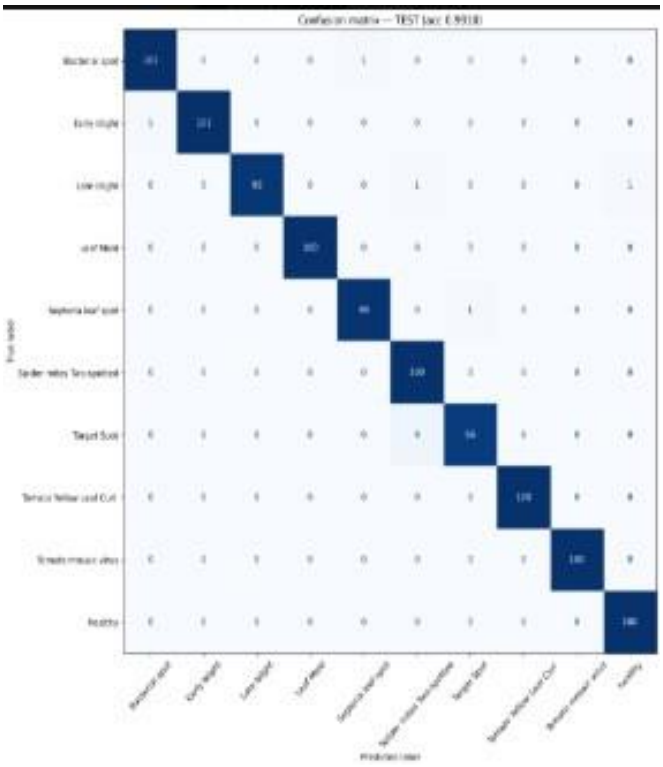


Fig. 6 Confusion matrix on the test set

The model classifies **995 out of 1,004 test images correctly**, giving the reported test accuracy of 99.10%. One noticeable confusion is between Target Spot and Spider Mites, where four Target Spot images are classified as Spider Mites. There are also a few individual mistakes involving Bacterial Spot, Early Blight, Late Blight, and Septoria Leaf Spot.

D. Analysis of Training Performance

The training graph shows that the model's performance improves throughout the training process. Initially, the accuracy is lower because only the classification head is being trained. After the fine-tuning stage begins, there is a



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511123

clear improvement in both training and validation accuracy. At the same time, the loss continues to decrease. The training and validation curves eventually become very close to each other, with accuracy approaching 99%. This indicates that the model learned useful features for distinguishing between the different tomato leaf disease classes.

V. DISCUSSION

The obtained results show that ResNet50 can perform very well for tomato leaf disease classification. The final test accuracy of 99.10% indicates that the model was able to identify most of the images correctly in the test set.

One important reason for the good performance is the use of transfer learning. Instead of training the entire ResNet50 model from the beginning, we started with ImageNet-pretrained weights. This gave the model useful image features, which were then adapted to the tomato leaf dataset through fine-tuning.

The class-wise results show that the model does not perform equally on every category. Some classes achieved perfect scores, while Target Spot and Late Blight had slightly lower recall. The confusion matrix shows that most of the errors are limited to a small number of visually similar cases.

The Grad-CAM results provide another way of looking at the predictions. In many correctly classified examples, the highlighted areas are located around visible changes on the leaf. However, Grad-CAM should only be used for understanding model attention and should not be treated as a confirmed disease segmentation method.

VI. STRENGTHS AND LIMITATIONS

A. Strengths

The main strengths of this system are:

- It can classify ten different tomato leaf categories.
- It uses ResNet50 transfer learning instead of training a deep model completely from scratch.
- It achieves 99.10% test accuracy.
- It provides class-wise precision, recall, and F1-score.
- The confusion matrix helps us understand the classification errors.
- Grad-CAM provides a visual explanation of the model's predictions.
- An additional colour-based analysis is provided for estimating affected leaf area.

B. Limitations

There are some limitations that need to be considered. The first limitation is that the reported results depend on the dataset and test split used in our experiment. Therefore, the same accuracy cannot automatically be expected for images captured in real agricultural fields.

The second limitation is the severity estimation. The affected-area percentage is obtained using a colour-based method and is only an approximate indication. It has not been validated as an exact disease-severity measurement. Another limitation is that Grad-CAM shows where the model is paying attention, but it does not provide a confirmed boundary of the diseased region.

Finally, the model can make mistakes when different diseases have similar visual symptoms. The misclassified examples demonstrate this issue.

VII. FUTURE WORK

There are several ways in which this project can be improved in the future. First, the model can be tested on a larger dataset containing images collected from real farms. Such images would contain different lighting conditions, backgrounds, leaf orientations, and levels of disease.

Second, additional data augmentation methods can be studied to make the model more robust to changes in image quality and environment. The severity estimation can also be improved by training a separate model using properly annotated diseased regions.

Finally, the model can be optimized and converted into a mobile or web application so that tomato leaf disease detection can be performed using images captured directly from a smartphone.

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511123

VIII. CONCLUSION

In this work, we developed a deep learning-based system for detecting tomato leaf diseases using ResNet50 transfer learning. The system classifies tomato leaf images into ten categories, including nine disease classes and one healthy class.

The model uses ImageNet-pretrained ResNet50 and is trained in two stages. The first stage trains the classification head for five epochs, followed by twelve epochs of fine-tuning of selected deeper layers. The input images are resized to 224 × 224 pixels before being passed to the model. The model achieved 99.10% accuracy on the test dataset, along with 99.10% balanced accuracy, 99.10% macro F1-score, and 99.10% weighted F1-score. The Cohen's kappa and Matthews correlation coefficient were 0.9900 and 0.9901, respectively.

The confusion matrix shows that most of the test images are correctly classified, although a small number of errors occurred between visually similar diseases. Grad-CAM was used to understand the regions receiving attention from the model, while colour-based analysis was used to provide an approximate indication of the affected leaf area.

Overall, the results show that ResNet50 with transfer learning is a promising approach for tomato leaf disease classification. However, further testing with real-world field images and better disease-region annotation is needed before the system can be considered for practical agricultural deployment.

REFERENCES

[1] S. G. E. Brucal, L. C. M. de Jesus, and L. A. Samaniego Jr., "Development of a Localized Tomato Leaf Disease Detection using YoloV9 Model via RoboFlow 3.0," Asia Pacific College, School of Engineering, Makati, Philippines, 2024.

[2] H. D. Gadade and D. K. Kirange, "Tomato Leaf Disease Diagnosis and Severity Measurement," *IEEE*, 2020.

[3] K. M. Hosny, W. M. El-Hady, F. M. Samy, E. Vrochidou, and G. A. Papakostas, "Multi-Class Classification of Plant Leaf Diseases Using Feature Fusion of Deep Convolutional Neural Network and Local Binary Pattern," *IEEE Access*, vol. 11, pp. 62307–62325, 2023.

[4] M. Gehlot and M. L. Saini, "Analysis of Different CNN Architectures For Tomato Leaf Disease Classification," in *2020 5th IEEE International Conference on Recent Advances and Innovations in Engineering (ICRAIE)*, 2020.

[5] A. Kaur, V. Kukreja, M. Kumar, A. Choudhary, and R. Sharma, "Multiclass Classification of Tomato Leaf Diseases using a Hybrid of LSTM and CNN Model," in *2024 IEEE 9th International Conference for Convergence in Technology (I2CT)*, Pune, India, 2024.

[6] X. Wang and J. Liu, "An Efficient Deep Learning Model for Tomato Disease Detection," *Plant Methods*, vol. 20, no. 61, 2024.

[7] K. Roy, S. S. Chaudhuri, J. Frnda, S. Bandopadhyay, I. J. Ray, S. Banerjee, and J. Nedoma, "Detection of Tomato Leaf Diseases for Agro-Based Industries Using Novel PCA DeepNet," *IEEE Access*, vol. 11, pp. 14983–15005, 2023.

[8] S. J. Basha, H. Bommala, V. P. Kumar M., N. Siva Kumar, T. S. R. Sai, and D. S. Kumar, "A Comprehensive Review on Deep Learning and Machine Learning Approaches for Tomato Leaf Disease Identification," in *Proc. 8th International Conference on Communication and Electronics Systems (ICCES 2023)*, IEEE, 2023.

[9] V. Balafas, E. Karantoumanis, M. Louta, and N. Ploskas, "Machine Learning and Deep Learning for Plant Disease Classification and Detection," *IEEE Access*, vol. 11, pp. 114352–114389, 2023.

[10] M. J. Hasan, S. Mazumdar, and S. Momen, "Deployable Deep Learning for Cross-Domain Plant Leaf Disease Detection via Ensemble Learning, Knowledge Distillation, and Quantization," *IEEE Access*, vol. 13, pp. 140313–140336, 2025.

[11] M. I. Hossain, S. Jahan, M. R. Al Asif, M. Samsuddoha, and K. Ahmed, "Detecting Tomato Leaf Diseases by Image Processing through Deep Convolutional Neural Networks," *Smart Agricultural Technology*, vol. 5, 100301, 2023.