ISSN: 2458-8989

Natural and Engineering Sciences

NESciences, 2026, 11 (3): 85-96 doi: 10.28978/nesciences.263008

A Hybrid Mobile Net–Inception CNN for Efficient Tomato Leaf Disease Classification in Precision Agriculture

A. Harish Kumar[1]*, Rajendra Naik Bhukya[2], Aravalli Sainath Chaithanya[3]

[1]Department of Electronics and Communication Engineering, University College of Engineering, Osmania University, Hyderabad, Telangana, India.
ORCID: 0000-0002-0514-0120, Email: allam.hari405@gmail.com,

[2]Department of Electronics and Communication Engineering, University College of Engineering, Osmania University, Hyderabad, Telangana, India.
ORCID: 0000-0002-9872-1433, Email: rajendranaikb@gmail.com,

[3]Department of Electronics and Communication Engineering, University College of Engineering, Osmania University, Hyderabad, Telangana, India.
ORCID: 0000-0002-8317-0827, Email: chaitanya.aravalli@gmail.com

Abstract

This study presents a hybrid MobileNet–Inception convolutional neural network for efficient tomato leaf disease classification in precision agriculture. The model integrates depthwise separable convolutions from MobileNet with multi-scale feature extraction using Inception modules to achieve a balance between accuracy and computational efficiency. A dataset of 5,000 tomato leaf images across 10 classes was compiled from Mendeley Data and Kaggle and split into training, validation, and testing sets in an 80:10:10 ratio. The model was trained for 50 epochs using the Adam optimizer (learning rate = 0.001), with dropout and L2 regularization applied to improve generalization. Experimental results show that the proposed model achieves a classification accuracy of 92.6%, with macro-precision, recall, and F1-score of 0.94, 0.93, and 0.93, respectively, outperforming benchmark models including AlexNet, VGG16, VGG19, InceptionV3, DenseNet121, ResNet50, MobileNet, and EfficientNetB0. In addition, the model demonstrates low inference latency (154 ms; 6.49 FPS), making it suitable for real-time applications. Analysis of confusion matrices indicates minor misclassifications among visually similar disease classes. These results highlight the model's effectiveness and practical applicability for real-time tomato disease detection, supporting timely decision-making in precision agriculture systems.

Keywords

Tomato leaf disease classification, Hybrid MobileNet–Inception CNN, precision agriculture, lightweight models, real-time inference.

Available online: 01/09/2026

**Corresponding Author: A. Harish Kumar, E-mail: allam.hari405@gmail.com*

Introduction

Plant diseases pose a significant threat to global food security, leading to substantial economic losses and environmental degradation (Gai & Wang, 2024). Among major crops, tomato plants are highly susceptible to diseases such as Tomato Mosaic Virus, Bacterial Spot, Target Spot, Late Blight, Early Blight, Leaf Mold, Septoria Leaf Spot, Spider Mites, and Tomato Yellow Leaf Curl Virus. Many of these diseases exhibit visually similar symptoms, making timely and accurate diagnosis challenging.

Traditional laboratory-based diagnostic techniques, including PCR and ELISA, are reliable but often costly, time-consuming, and require expert knowledge (Li et al., 1998). Conventional image processing approaches also struggle with variability in leaf appearance and overlapping disease characteristics (Rahman et al., 2023). In recent years, deep learning—particularly convolutional neural networks (CNNs)—has improved plant disease detection through automatic feature extraction and enhanced classification performance (Vengaiah & Konda, 2023). However, existing CNN-based models often require large annotated datasets, face challenges in distinguishing visually similar classes, and involve high computational complexity, limiting their practical deployment.

A key research gap remains in developing models that achieve high classification performance while maintaining computational efficiency for real-world agricultural applications. Many existing approaches prioritize either accuracy or efficiency, with limited success in achieving both simultaneously.

To address this gap, this paper proposes a Hybrid MobileNet–Inception CNN model for tomato leaf disease classification. The model integrates MobileNet's depthwise separable convolutions with Inception modules for multi-scale feature extraction, enabling efficient and discriminative feature representation. The objective of this study is to develop an accurate and computationally efficient model for real-time tomato leaf disease classification on resource-constrained devices.

The main contributions of this work are as follows:

1. Proposed a Hybrid MobileNet–Inception CNN architecture that combines depthwise separable convolutions with multi-scale feature extraction for improved tomato leaf disease classification.
2. Developed a balanced multi-class tomato leaf dataset comprising 10 disease categories to enhance model robustness and minimize class bias.
3. Established a comprehensive evaluation framework using comparative analysis, standard performance metrics, and confusion matrix analysis to assess the model against existing CNN architectures.
4. Demonstrated that the proposed model achieves 92.6% classification accuracy with low inference latency of 154 ms, highlighting its computational efficiency over conventional deep CNN models.
5. Demonstrated the practical suitability of the proposed model for real-time deployment on resource-constrained mobile and embedded platforms.

The remainder of this paper is organized as follows. Section 2 reviews related work on plant disease detection using deep learning techniques. Section 3 describes the proposed methodology, including the Hybrid MobileNet–Inception architecture. Section 4 presents the implementation details, including dataset preparation and training configuration. Section 5 discusses the experimental results and performance evaluation. Finally, Section 6 concludes the paper and outlines future research directions.

Literature Review

Recent studies have extensively explored deep learning approaches for tomato leaf disease classification, demonstrating significant improvements in automated disease detection. However, existing methods vary in terms of model complexity, dataset dependency, and real-world applicability. Several works have focused on high-accuracy deep learning models**.** For instance, (Javidan et al., 2025) compared multiple architectures and reported strong performance using MobileNet and ensemble methods, achieving above 96% accuracies. Similarly, (Zhou et al., 2021) proposed a restructured residual dense network (RRDN) that combines residual and dense connections, achieving 95% accuracy with improved training stability. In (Kokate et al., 2023), a CNN trained on the PlantVillage dataset achieved 95.53% accuracy using data augmentation techniques. While these approaches demonstrate high classification performance, they largely rely on controlled datasets and may not generalize well to real-world agricultural conditions.

Other studies have emphasized ensemble and deep architectures to maximize performance. In (Sharma et al., 2025), a MobileNet and ResNet50-based ensemble achieved 99.91% accuracy. Likewise, (Hossain et al., 2023) demonstrated high accuracy (99.53%) using pre-processing techniques combined with deep CNN models. Similarly, (Sainath Chaithanya & Rachana, 2023) employed a fine-tuned ResNet50 model for papaya leaf disease classification, achieving 87.95% accuracy. Although these methods yield excellent results, their high computational complexity and reliance on preprocessing pipelines limit their suitability for deployment on resource-constrained devices.

In contrast, some research has explored lightweight models for efficient deployment. The study in (Agarwal et al., 2020) developed a compact CNN model with reduced storage requirements (1.5 MB), achieving an average accuracy of 91.2%. While such models are computationally efficient and suitable for mobile platforms, they often struggle to distinguish between visually similar disease classes due to limited feature extraction capability.

In addition to model-specific approaches, comparative studies such as (Sharma et al., 2022) highlight that while conventional CNNs can outperform deeper architectures like ResNet50 in certain cases, performance improvements often depend on dataset characteristics and model tuning, indicating a lack of consistency across approaches.

Despite these advancements, several key limitations persist. Many models rely on benchmark datasets such as PlantVillage, which do not adequately represent real-world variability, including lighting conditions, occlusions, and background noise. Furthermore, there exists a trade-off between accuracy and computational efficiency: deep and ensemble models achieve high accuracy but are computationally expensive, whereas lightweight models are efficient but may lack discriminative power for similar disease classes.

To address these limitations, this study proposes a Hybrid MobileNet–Inception CNN model that integrates the computational efficiency of depthwise separable convolutions with the multi-scale feature extraction capability of Inception modules. This combination enables improved discrimination of visually similar disease patterns while maintaining low computational cost, making the model suitable for real-time deployment in practical agricultural environments.

Methodology

Dataset Preparation

The dataset used in this study was obtained from the publicly available Tomato Leaf Disease dataset on Mendeley Data (Solapure et al., 2024). As the original dataset contains an unequal number of images per class, a balanced subset was constructed by selecting 500 images for each disease category to ensure uniform class distribution.

The choice of a fixed sample size per class was motivated by the need to enable unbiased model training and consistent performance evaluation across classes. Such balanced sampling is commonly adopted in classification tasks to avoid dominance of majority classes and to ensure that evaluation metrics accurately reflect model performance. Additionally, selecting a moderate and uniform sample size helps maintain sufficient intra-class variability while keeping computational requirements manageable.

For most categories, the 500 images were randomly sampled from Mendeley (Table 1) to preserve diversity within each class. In the case of Tomato Mosaic Virus, only 299 images were available; therefore, an additional 201 images were sourced from publicly available Kaggle repositories ("Tomato Leaf Disease Detection Dataset," 2025) with comparable quality, resolution, and disease characteristics to maintain dataset consistency.

Table 1 presents the class-wise distribution of original and selected images used in the experiment. Figure 1 shows one representative image from each class, illustrating variations in leaf texture, shape, and disease symptoms.

Table 1: Class-wise distribution of original and selected tomato leaf images used in the experiment.

Class Name	Total Images in Dataset	Images Used for Experiment
Bacterial Spot	1702	500
Early Blight	800	500
Late Blight	1528	500
Leaf Mold	762	500
Septoria Leaf Spot	1417	500
Spider Mites (Two-Spotted)	1341	500
Target Spot	1124	500
Tomato Yellow Leaf Curl Virus	4286	500
Tomato Mosaic Virus	299 (+201 from Kaggle)	500
Healthy	1272	500

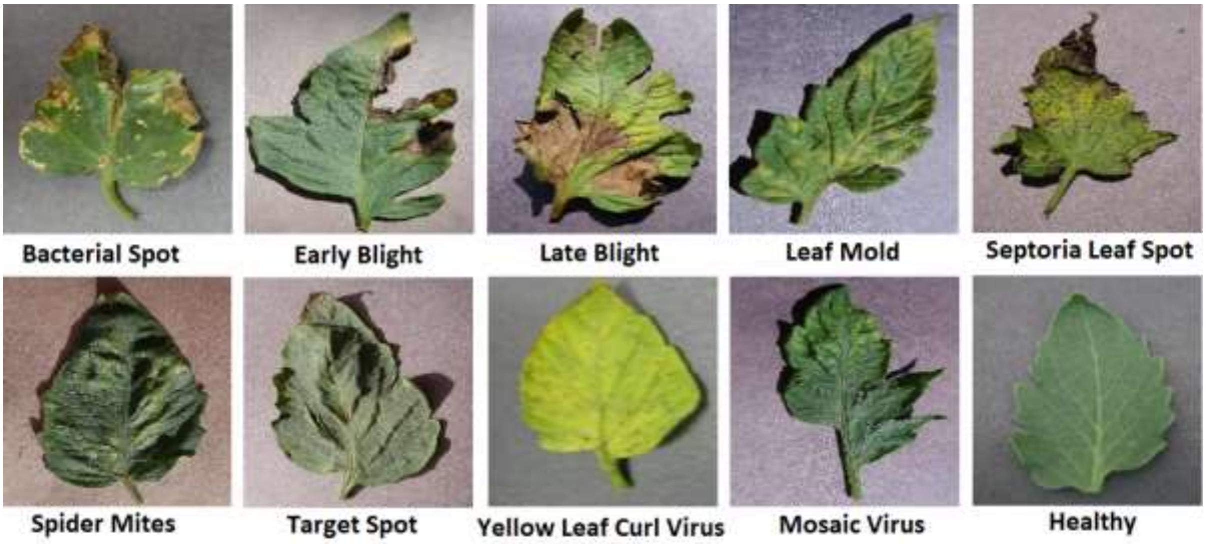

Fig 1: Sample images from each tomato leaf class

Proposed Hybrid MobileNet-Inception CNN

The proposed Hybrid MobileNet–Inception CNN integrates the lightweight depthwise separable convolution strategy of (Howard et al., 2017) MobileNet with the multi-scale feature extraction capability of the (Szegedy et al., 2015) Inception architecture.

The architecture (Figure 2) is organized into three major stages:

1. MobileNet-based shallow feature extraction
2. Multi-scale feature learning using Inception modules
3. Deep feature refinement and classification

The model processes RGB tomato leaf images of size 256×256×3 and contains approximately 508,338 total parameters, making it suitable for lightweight and real-time agricultural disease diagnosis applications. Among these, 505,394 are trainable parameters and 2,944 are non-trainable parameters.

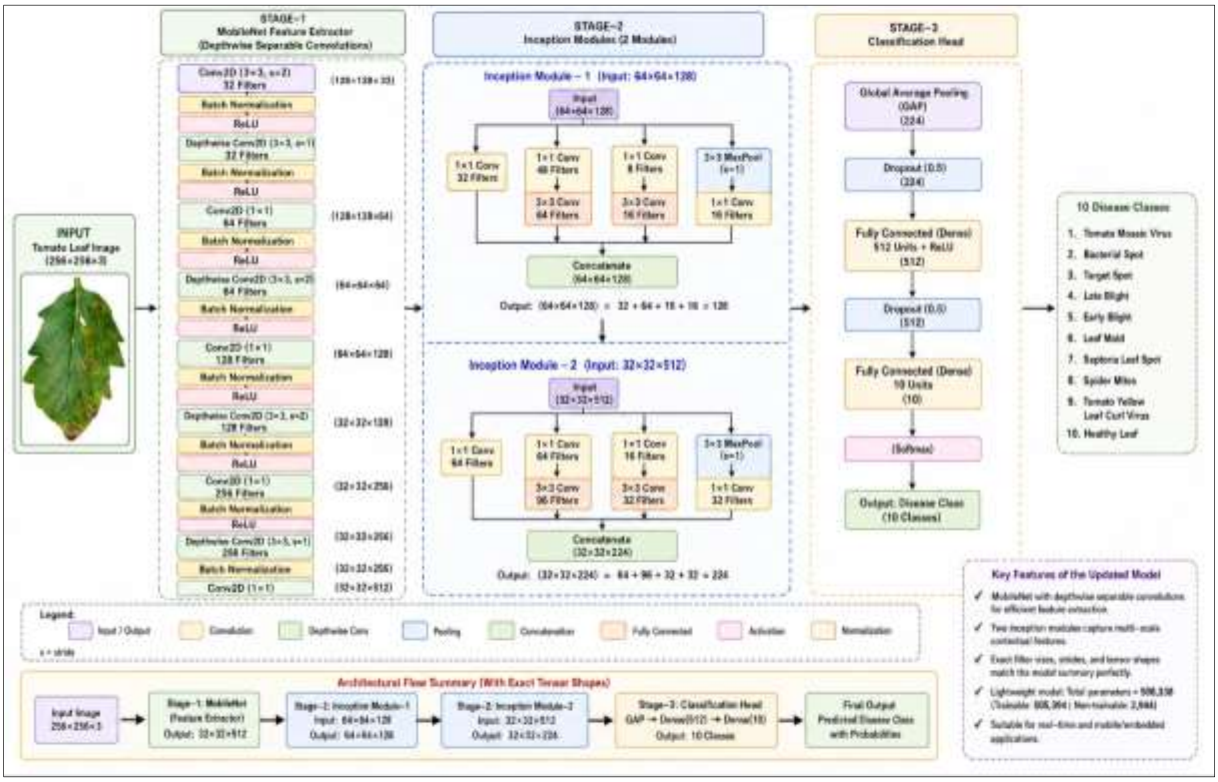

Fig 2: Hybrid MobileNet–Inception CNN Architecture

Stage-1: MobileNet-Shallow Feature Extraction:

The network begins with a standard convolution layer using a 3×3 kernel with 32 filters and stride 2. This layer performs initial feature extraction while reducing the spatial resolution from 256×256 to

128×128. Batch Normalization and ReLU activation are applied to stabilize training and introduce non-linearity.
Subsequently, multiple depthwise separable convolution (DSC) blocks are employed. Each DSC block consists of:

- Depthwise convolution (3×3) for channel-wise spatial filtering
- Pointwise convolution (1×1) for inter-channel feature fusion
- Batch Normalization and ReLU activation after each convolution stage

The feature extraction pipeline progressively increases feature depth while reducing spatial dimensions:

- 32 → 64 feature channels
- 64 → 128 feature channels
- 128 → 256 feature channels
- 256 → 512 feature channels

Spatial down-sampling is performed using stride-2 depthwise convolutions, progressively reducing feature-map dimensions while increasing channel depth, as illustrated in Figure 2.
Depthwise separable convolutions substantially reduce parameter count and computational complexity compared with conventional convolutions while preserving essential texture and lesion features required for disease recognition.

Stage-2: Inception-Multi-Scale Feature Learning:
To capture disease characteristics at multiple spatial scales, two Inception modules are incorporated within the architecture.
Inception Module-1
The first Inception module receives feature maps of size 64×64×128 and extracts complementary spatial features using parallel convolutional branches:

- 1×1 convolution branch (32 filters)
- 1×1 → 3×3 convolution branch (48 → 64 filters)
- 1×1 → 3×3 convolution branch (8 → 16 filters)
- 3×3 max pooling → 1×1 convolution branch (16 filters)

The outputs from all branches are concatenated along the channel dimension to generate a multi-scale feature representation of size 64×64×128.
This module enables simultaneous detection of:

- small localized lesions,
- texture irregularities,
- edge discontinuities,
- and broader disease-affected regions.

Following the first Inception module, additional DSC blocks further refine the extracted features and increase feature depth to 256 and 512 channels.
Inception Module-2
A second Inception module is applied after the deeper DSC layers to refine high-level semantic representations and improve discrimination between visually similar disease categories.
The second Inception module operates on feature maps of size 32×32×512 and contains four parallel branches:

- 1×1 convolution branch (64 filters)
- 1×1 → 3×3 convolution branch (64 → 96 filters)
- 1×1 → 3×3 convolution branch (16 → 32 filters)
- 3×3 max pooling → 1×1 convolution branch (32 filters)
- The concatenated output produces a final feature representation of size 32×32×224.

Stage 3- Classification Head:
The resulting feature maps are then processed through:

- Global Average Pooling (GAP) to reduce spatial dimensions and minimize overfitting
- Dropout layer (rate = 0.5) for regularization
- Fully connected dense layer with 512 neurons and ReLU activation
- Additional dropout layer (rate = 0.5)
- Fully connected output layer with 10 neurons
- Softmax activation for multiclass classification
- The final output corresponds to 10 tomato leaf disease categories.

Design Rationale

The proposed Hybrid MobileNet–Inception CNN achieves an effective balance between lightweight computation and robust feature representation.

The use of depthwise separable convolutions significantly reduces:

- parameter count,
- memory usage,
- and computational complexity,

making the model suitable for real-time and resource-constrained agricultural applications.

Meanwhile, the Inception modules improve representational power through multi-scale feature extraction, allowing the network to capture both fine-grained lesion patterns and broader discoloration regions associated with tomato leaf diseases.

The incorporation of Batch Normalization, Dropout regularization, and Global Average Pooling further improves:

- training stability,
- generalization capability,
- and resistance to overfitting.

Overall, the hybrid architecture provides a computationally efficient yet highly discriminative framework for automated tomato leaf disease classification (Zhang et al., 2023).

The architecture follows a sequential yet interleaved feature-learning pipeline. This interconnection strategy enables progressive refinement of features from shallow texture representations to high-level semantic disease patterns.

Implementation details

Dataset Description

This study utilizes a dataset of 5,000 tomato leaf images, comprising 10 classes with 500 images per class. The dataset is split into training, validation, and testing sets in an 80:10:10 ratio (4,000 / 500 / 500). All images are resized to 256×256×3 to ensure uniform input dimensions across all models.

Model Implementation

The proposed Hybrid MobileNet–Inception CNN combines a lightweight MobileNet backbone with Inception modules for multi-scale feature extraction. The classification head consists of fully connected layers followed by a Softmax layer with 10 output classes

Training Configuration

The model is trained using the Adam optimizer with a learning rate of 0.001 and categorical cross-entropy loss (Keras, 2025). A batch size of 16 is used, and training is conducted for 50 epochs. Dropout (rate = 0.4) and L2 regularization are applied to reduce overfitting.

Early stopping and learning rate scheduling were not explicitly employed, as empirical observations showed stable convergence within 50 epochs with minimal validation loss fluctuations. Additionally, the balanced dataset and moderate size allowed effective training without requiring cross-validation, which would significantly increase computational cost without substantial performance gain.

Experimental Setup and Reproducibility

Experiments were conducted on Google Colab using an NVIDIA Tesla T4 GPU. The implementation was carried out using Python with deep learning libraries such as TensorFlow/Keras. All models were trained under identical conditions, including input size, optimizer settings, batch size, and number of epochs, to ensure fair comparison.

Results and Discussions

Training and validation analysis

Figures 3 and 4 illustrate the training and validation performance of baseline models and the proposed model. The proposed model demonstrates faster convergence and lower validation loss, indicating stable learning and improved generalization.

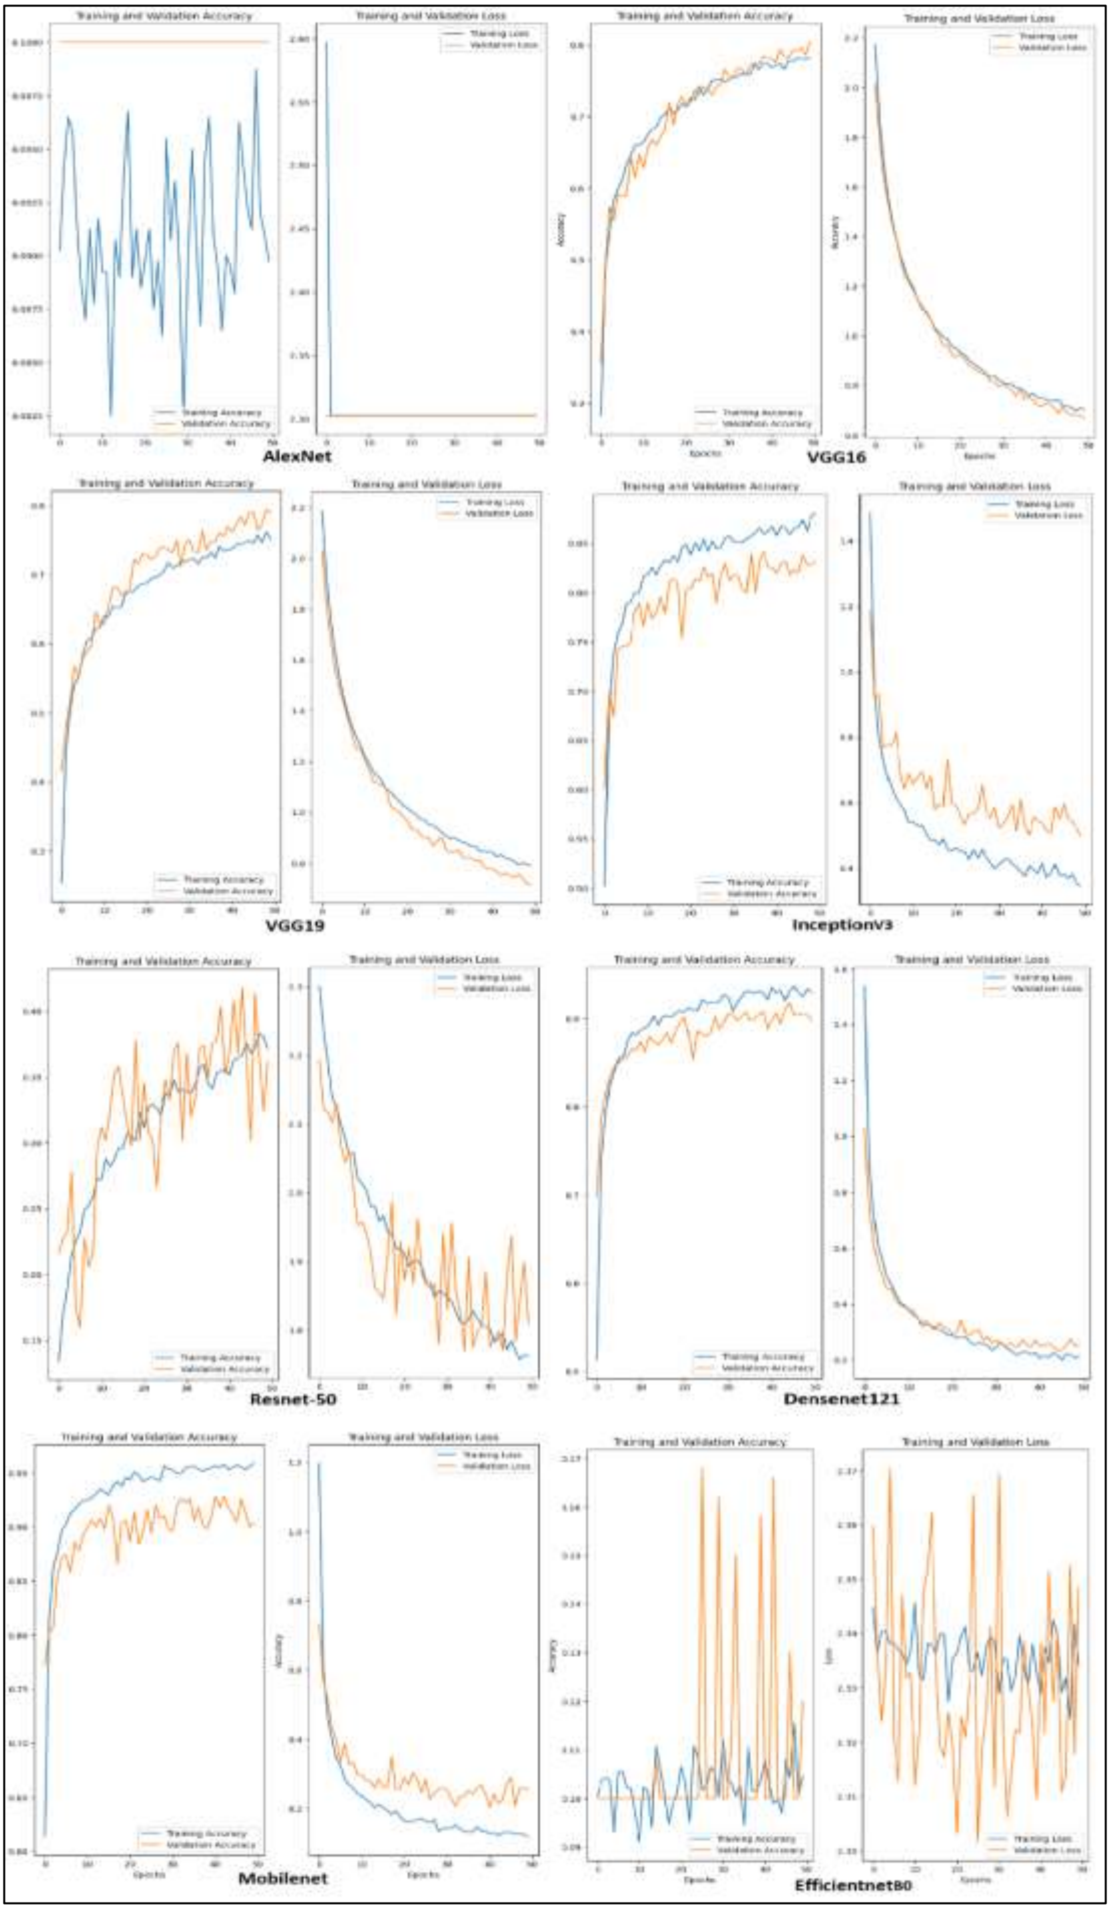

Fig 3: Training and validation accuracy and loss curves for baseline CNN models.

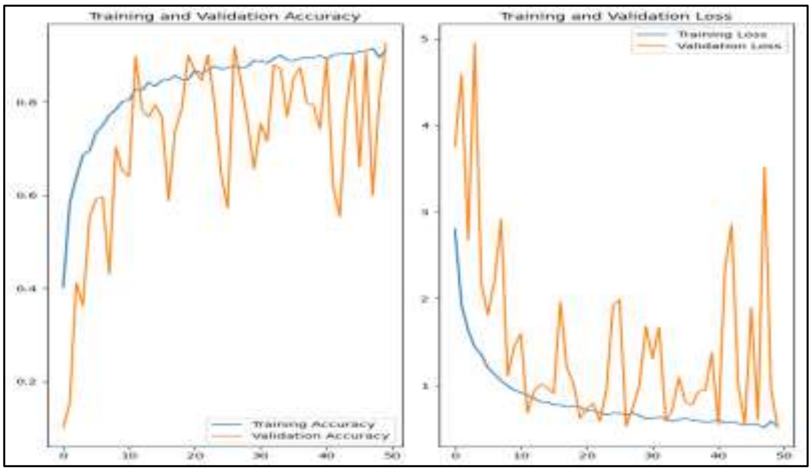

Fig 4: Training and Validation accuracies and loss curves of Proposed model

Comparative Performance Analysis

Table 2 compares the proposed model with standard CNN architectures. The proposed model achieves the highest accuracy (92.6%) with significantly fewer parameters (0.51M) and lower inference time.

Table 2: Comparative evaluation of CNN architectures with proposed model

Model	**Accuracy (%)**	**Macro Avg.**			**Parameters (M)**	**Inference Time**	**FPS**
		Precision	**Recall**	**F1-score**			
AlexNet	10.0	0.01	0.10	0.02	956.56	956.56	1.05
VGG16	79.8	0.81	0.80	0.80	968.07	968.07	1.03
VGG19	79.0	0.81	0.79	0.79	1320.36	1320.36	0.76
Inceptionv3	83.2	0.84	0.83	0.83	346.80	346.80	2.88
ResNet-50	36.2	0.48	0.36	0.33	393.42	393.42	2.54
DenseNet-121	89.8	0.90	0.90	0.90	376.08	376.08	2.66
MobileNet	90.4	0.91	0.90	0.90	191.50	191.50	5.22
EfficientNet-B0	12.0	0.03	0.12	0.04	215.28	215.28	4.65
Proposed Model	**92.6**	**0.94**	**0.93**	**0.93**	**154.05**	**154.05**	**6.49**

The unusually low performance of certain baseline models (e.g., AlexNet and EfficientNet-B0) can be attributed to training from scratch on a relatively small dataset without transfer learning. These architectures typically require large-scale datasets and pre-trained weights to perform effectively. In contrast, the proposed model is specifically designed for efficient learning on limited data, contributing to its superior performance.

Inference Efficiency

The proposed model achieves an inference time of 154.05 ms (6.49 FPS), outperforming deeper architectures while maintaining high accuracy. This demonstrates its suitability for real-time deployment on resource-constrained devices.

Confusion Matrix and error analysis

The confusion matrix (Figure 5) of proposed model indicates that most classes are correctly classified, with minor misclassifications occurring between visually similar diseases such as leaf spot-related classes.

Fig 5: Confusion matrix comparing the proposed model against baseline CNN models on the validation dataset.

Class-wise analysis shows that diseases with overlapping visual features, such as Early Blight and Late Blight, exhibit higher confusion rates. This is primarily due to similarity in lesion patterns and color variations. These errors suggest that incorporating attention mechanisms or segmentation-based feature localization techniques (Chaithanya et al., 2025) could further improve discrimination.

Classified outputs

Figure 6 presents representative classification outputs of the proposed model across all 10 tomato leaf disease categories. The model demonstrates strong capability in capturing subtle variations in lesion patterns, texture, and color, enabling accurate differentiation between visually similar classes. Compared to conventional deep and ensemble models, which often achieve high accuracy at the cost of computational complexity, and lightweight models, which may lack discriminative power, the proposed approach provides an effective balance between performance and efficiency. This is achieved through the hybrid integration of MobileNet and Inception modules, which enables efficient feature extraction along with multi-scale representation learning. These characteristics support its practical applicability for real-time tomato leaf disease detection.

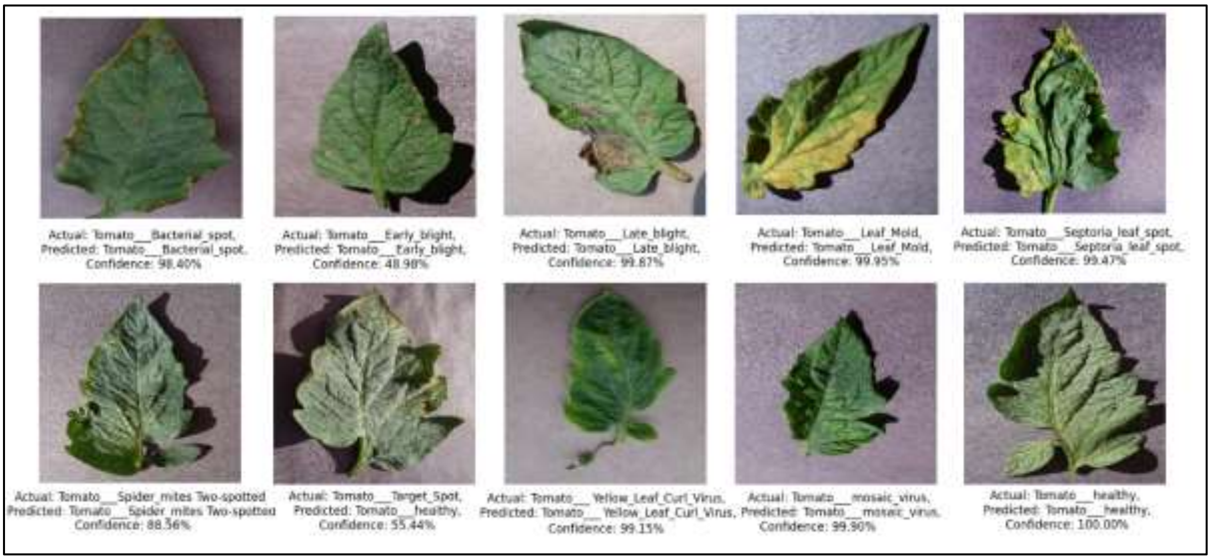

Fig 6: Representative classification outputs of the proposed hybrid model across the 10 tomato leaf disease categories

Conclusion and Future Scope

This study presented a Hybrid MobileNet–Inception CNN architecture for tomato leaf disease classification, combining depthwise separable convolutions with multi-scale feature extraction. The proposed model achieved a classification accuracy of 92.6% with an inference time of 154 ms per image (6.49 FPS), demonstrating an effective balance between performance and computational efficiency. Comparative evaluation with established CNN architectures, including AlexNet, VGG16, VGG19, ResNet50, InceptionV3, DenseNet121, MobileNet, and EfficientNetB0, confirmed the effectiveness of the proposed approach. Furthermore, confusion matrix analysis indicated minimal misclassification, even among visually similar disease classes, highlighting the model's discriminative capability. These results support the potential applicability of the model for real-time disease detection in precision agriculture on resource-constrained devices.

Despite these promising results, several limitations remain. The model was trained and evaluated on a curated and balanced dataset, which may not fully represent real-world variability such as varying illumination, occlusions, and complex backgrounds. Additionally, the use of a fixed dataset size and absence of field-level validation limit the generalizability of the model in practical agricultural environments. The reliance on controlled data sources may also affect performance when deployed under diverse environmental conditions.

Future work will focus on addressing these limitations by incorporating more diverse and unstructured field datasets, applying advanced data augmentation and class-balancing techniques, and integrating attention mechanisms to enhance feature localization. Further improvements can include extending the framework to multi-crop and multi-disease classification and deploying the model in mobile and IoT-based systems for real-time, on-site disease diagnosis. Such advancements can contribute to more robust and scalable solutions for sustainable precision agriculture.

Author contributions

Author 1 was responsible for the conceptualization and development of the proposed Hybrid MobileNet–Inception model, as well as software implementation, experimentation, data curation, validation, formal analysis, and preparation of the original draft of the manuscript. Author 2 contributed to formal analysis, resource support, and supervision of the study, and was actively involved in reviewing and refining the manuscript. Author 3 provided technical inputs and contributed to validation, formal analysis, and manuscript review and editing. All authors have read and approved the final version of the manuscript.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgements

1. The authors gratefully acknowledge the support of the University Grants Commission (UGC), India, for providing the research fellowship that enabled this work. The authors also extend their gratitude to Dept. of ECE, University college of Engineering, Osmania University, Hyderabad and Keshav Memorial Institute of Technology, Hyderabad for providing the necessary facilities and academic support to carry out this research.
2. Furthermore, the authors would like to thank the dataset sources "Mendeley Data" for providing the "Tomato Leaf Disease Dataset" available at https://data.mendeley.com/datasets/zfv4jj7855/1, and "Kaggle" for providing the "Tomato Leaf Disease Detection Dataset" available at https://www.kaggle.com/datasets/kaustubhb999/tomatoleaf

References

1. 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
2. Chaithanya, A. S., Devi, L. N., & Srividya, P. (2025). Instance segmentation for PCB defect detection with Detectron2. International Journal of Electrical and Computer Engineering (IJECE), 15(4), 4172–4180. https://doi.org/10.11591/ijece.v15i4.pp4172-4180
3. Gai, Y., & Wang, H. (2024). Plant disease: A growing threat to global food security. Agronomy, 14(8), 1615. https://doi.org/10.3390/agronomy14081615
4. Hossain, M. I., Jahan, S., Al Asif, M. R., Samsuddoha, M., & Ahmed, K. (2023). Detecting tomato leaf diseases by image processing through deep convolutional neural networks. Smart Agricultural Technology, 5, 100301. https://doi.org/10.1016/j.atech.2023.100301
5. Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv. https://doi.org/10.48550/arXiv.1704.04861
6. Javidan, S. M., Ampatzidis, Y., Banakar, A., Asefpour Vakilian, K., & Rahnama, K. (2025). An intelligent group learning framework for detecting common tomato diseases using simple and weighted majority voting with deep learning models. AgriEngineering, 7(2), 31. https://doi.org/10.3390/agriengineering7020031
7. Keras. (2025). Losses. Keras API Documentation. Keras API Documentation
8. Kaustubh B. (2025). Tomato leaf disease detection dataset [Data set]. Kaggle. https://www.kaggle.com/datasets/kaustubhb999/tomatoleaf
9. Kokate, J. K., Kumar, S., & Kulkarni, A. G. (2023). Classification of tomato leaf disease using a custom convolutional neural network. Current Agriculture Research Journal, 11(1). http://dx.doi.org/10.12944/CARJ.11.1.28
10. Li, R. H., Wisler, G. C., Liu, H.-Y., & Duffus, J. E. (1998). Comparison of diagnostic techniques for detecting tomato infectious chlorosis virus. Plant Disease, 82(1), 84–88. https://doi.org/10.1094/PDIS.1998.82.1.84
11. Rahman, S. U., Alam, F., Ahmad, N., Hossain, M. S., & Muhammad, G. (2023). Image processing based system for the detection, identification and treatment of tomato leaf diseases. Multimedia Tools and Applications, 82, 9431–9445. https://doi.org/10.1007/s11042-022-13715-0
12. Sainath Chaithanya, A., & Rachana, M. (2023). Identification of diseased papaya leaf through transfer learning. Indian Journal of Science and Technology, 16(48), 4676–4687. https://doi.org/10.17485/IJST/v16i48.2690
13. Sharma, J., Al-Huqail, A. A., Almogren, A., Doshi, H., Jayaprakash, B., Bharathi, B., Ur Rehman, A., & Hussen, S. (2025). Deep learning based ensemble model for accurate tomato leaf disease classification by leveraging ResNet50 and MobileNetV2 architectures. Scientific Reports, 15(1), 13904. https://doi.org/10.1038/s41598-025-98015-x
14. Sharma, R., Panigrahi, A., Garanayak, M., Chakravarty, S., Paikaray, B. K., & Pattanaik, L. (2022). Tomato leaf disease detection using machine learning. In Proceedings of Advances in Computational Intelligence, Its Concepts & Applications (ACI 2022) (Vol. 3283, pp. 115–123). CEUR Workshop Proceedings. CEUR Workshop Proceedings Paper
15. Solapure, V., SmartAgroTech DY, & Jawale, A. (2024). Tomato leaf disease dataset [Data set]. Mendeley Data. https://doi.org/10.17632/zfv4jj7855.1

16. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 1–9). https://doi.org/10.1109/CVPR.2015.7298594
17. Vengaiah, C., & Konda, S. R. (2023). A review on tomato leaf disease detection using deep learning approaches. International Journal on Recent and Innovation Trends in Computing and Communication, 11(9s), 647–664. https://doi.org/10.17762/ijritcc.v11i9s.7479
18. Zhang, A., Lipton, Z. C., Li, M., & Smola, A. J. (2023). Dive into deep learning. Cambridge University Press. Dive into Deep Learning
19. Zhou, C., Zhou, S., Xing, J., & Song, J. (2021). Tomato leaf disease identification by restructured deep residual dense network. IEEE Access, 9, 28822–28831. https://doi.org/10.1109/ACCESS.2021.3058947