OPEN ACCESS

Volume: 13

Special Issue: 2

Month: January

Year: 2026

E-ISSN: 2582-0397

P-ISSN: 2321-788X

Citation:
Murkar, Seema, et al. "Lightweight and Explainable Deep Learning Models for Pomegranate Disease Diagnosis: A Case Study with Maharashtra Observations." *Shanlax International Journal of Arts, Science and Humanities,* vol. 13, no. 2, 2026, pp. 43–52.

DOI:
https://doi.org/10.34293/sijash.v13iS2-i1-Jan.10460

Lightweight and Explainable Deep Learning Models for Pomegranate Disease Diagnosis: A Case Study with Maharashtra Observations

Ms. Seema Murkar
Research Scholar, Department of Computer Science and Engineering
Mangalayatan University, Aligarh, India
Assistant Professor, Information Technology and Data Science
Vidyalankar School of Information Technology, Mumbai, India

Dr. Manoj Varshney
Associate Professor, Department of Computer Engineering and Applications
Mangalayatan University, Aligarh, India

Dr. Mandar Sohani
Professor
Vidyalankar Institute of Technology, Mumbai, India

Abstract
Pomegranate is a major fruit crop in Maharashtra with major cultivation marked in the districts Solapur, Satara, Sangali, Ahmednagar and Latur. But farmers in these districts face significant challenges from disease outbreaks. Traditional methods of disease diagnosis are often time-consuming and inefficient in generating accurate results. In the recent year, deep learning has emerged as an effective solution for automation in disease detection. However, deep learning models are often computationally heavy and lack interpretability, restricting their real-world applications. But the lite versions of these models are exceptionally small in size and hence perform well in real time. So, this study tests the effectiveness of MobileNetV3-Lite and EfficientNet-Lite models on the available dataset comprising of 5,099 images of pomegranate fruit from various regions of India; this study focuses on diseases prevalent in Maharashtra orchards to ensure contextual relevance. Further, integration of Grad-CAM visualizes affected regions within image by creating heatmaps and hence providing insight into the reasoning behind the model's predictions. The results show that EfficientNet achieved highest classification accuracy of 99.50%, in contrast MobileNetV3 achieved marginally lower accuracy 98.48% but its lite model exhibited a significant smaller size 6.17 MB along with inference time of 0.019 s/image, making it more suitable for real-time and low resource agricultural applications.

Keywords: Pomegranate Disease Classification, Deep Learning, EfficientNet, MobileNetV3, Grad-CAM, MobileNet-Lite

Introduction

The agricultural sector forms a foundational pillar for global economies and food security, with specific crops often holding particular regional significance [11][12]. India ranks first in pomegranate cultivation [1] with Maharashtra pomegranate cultivation representing a major role in the economy of many farmers [2]. But the sudden outbreak of pomegranate plant diseases, such as anthracnose, bacterial blight, cercospora and fruit rot, leads to significant yield reductions and economic losses. Hence it is not only important to detect disease early but to have accurate diagnosis for safeguarding crop productivity [3][13]. Traditional methods of diagnosis mainly depend on expert's visual inspection and are highly time-consuming, labour intensive and sometimes inaccurate [14][15].

Recent study in deep learning has implemented powerful methods for AI-based image disease diagnosis in agriculture [4][16]. Convolutional Neural Networks (CNNs) have illustrated excellent predictive performance in classifying plant diseases from images [17]. Inspite of these favourable results, real time implementation of complex deep learning models in agricultural applications creates specific hurdles. Most of these models achieve high accuracy but also require substantial computational resources, making them unsuitable for edge devices or environments with limited infrastructure [5][18]. Furthermore, the inherent "black-box" nature of deep learning models often obscures their decision-making processes, hindering trust and adoption by end-users like farmers and agronomists [19].

This research addresses these limitations by testing lightweight and explainable deep learning models tailored for pomegranate disease diagnosis in Maharashtra. To provide enhanced crop management and sustainable agricultural practices in Maharashtra this study emphasizes on lightweight deep learning models MobileNetV3-lite and EfficientNet-Lite due to their remarkable accuracy and low memory requirements making them more suitable for mobile and edge devices. To make this approach readily acceptable by farmers, explainable AI model Grad-CAM has been applied on the images to interpret results by creating heatmaps on the affected regions on the image and hence making this model more suitable for agriculture applications.

Literature Review

The Role of Deep Learning in Agricultural Disease Diagnosis

As compared to conventional methods of disease diagnosis like visual inspection, deep learning techniques have proven to be accurate in the agricultural disease diagnosis [4][14]. The critical task of plant pathologists of identifying disease from visual symptoms has been made easy by the advancement of Convolutional Neural Networks (CNNs) in image recognition [17]. Many researchers have successfully tested CNNs' efficiency by achieving high accuracy in classifying disease across different crops such as turmeric, eggplant and corn [20][21][12]. For instance, the CNN attained 99% accuracy for turmeric leaf and 98.85% for corn disease identification [20][12]. This technique has proven its efficiency and timely diagnosis over labour intensive and human bias error prone manual methods [15]. There is substantial reduction of crop losses with the application of deep learning techniques by detecting disease at early stage and hence enabling farmers to take proactive measures [3].

Challenges of Model Efficiency and Interpretability

Despite the high accuracy of deep learning models in agricultural diagnostics, their practical deployment encounters significant hurdles related to efficiency and interpretability [5][6]. Many advanced deep learning architectures are accurate but require lot of memory and time for computation due to their numerous parameters and complex operations. So deploying these models on low resource devices such as mobile phones and small sensor based devices is very challenging [18][21]. These problems limit advanced deep learning models; it is very difficult to implement these models in real-time processing and easy accessibility among farmers [5]. Additionally, the "black-box" nature of deep learning models negatively impacts user trust and awareness because the predictions made by the model are without transparent reasoning [19]. To verify model's result agronomists and farmers require to know the interpretation behind the disease diagnosis than just the name of the particular disease, so that they can validate the model output against their expertise and hence implement precise treatment effectively [10].

Advances in Lightweight Architectures for Field Deployment

To overcome the computational limitations of deep learning models, significant advancements have occurred in the development of lightweight architectures [21]. These models have been optimized by reducing parameter count and processing overhead while still maintaining high accuracy, making them appropriate for implementing on mobile phones and sensor based electronic devices [18]. MobileNetV3 and EfficientNet-Lite architectures have been designed to achieve prominent results in this category [7][8]. For instance, MobileNetV3 offers fine-tuned, class-based features with less computational overhead and hence makes a foundation for compact frameworks for disease classification [7]. Similarly, EfficientNet models have proven their efficiency by effectively managing the tradeoff between accuracy and model size [22][8]. Techniques such as depthwise separable convolutions and neural architecture search are incorporated in these models to achieve optimized yet high performance design making it suitable in real-time field conditions [21].

Explainable AI Techniques in Plant Disease Identification

In critical domains like agricultural diagnostics, enabling trust and practical application of deep learning models is particularly challenging [19]. In recent years, Explainable Artificial Intelligence (XAI) techniques have proven their importance in providing methods to understand and visualize the decision-making process of these opaque models. Gradient-weighted Class Activation Mapping (Grad-CAM) that generates visual heatmaps on an image region which are most relevant to the model's predictions emerges as a leading XAI technique [9][10]. By producing heatmaps that indicate most prominent features, Grad-CAM allows users to see which parts of a plant leaf, for example, the model focused on when identifying a disease [10]. This visual feedback not only increases transparency but also enables validation of the model's reasoning against agronomic expertise. The integration of Grad-CAM with deep learning models, such as ResNet, has successfully demonstrated its utility in explaining classifications of corn leaf diseases, providing human-interpretable insights [10].

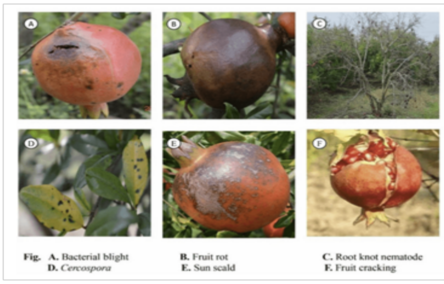

Figure 1 Pomegranate Plant Common Diseases Observed in Solapur District, Maharashtra [23]

Methodology

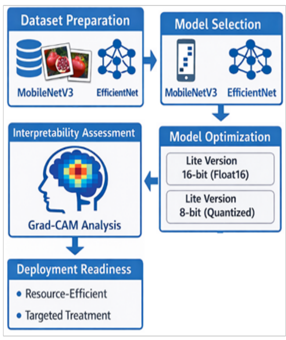

Figure 2 Proposed Model Workflow

Dataset Preparation and Splitting

A dataset comprising 5099 images was used to test the deep learning lightweight models. Each image of this dataset has resolution of 3120 x 3120 pixels. These images are divided into 5 classes, one as healthy class and remaining 4 are pomegranate diseases — bacterial blight, anthracnose, cercospora fruit spot, alternaria fruit spot — which are also found in Maharashtra orchards [23]. This approach ensures that the models are trained on data directly relevant to the local agricultural context, addressing typical pathologies affecting pomegranate crops. The dataset was then organised into three subsets to ensure unbiased model evaluation by splitting it into the ratio of 70:15:15 where 70% for training, 15% for testing and remaining 15% for validation. All images are then resized to 224 x 224 pixels, the requirements of the pre-trained models. Further, to improve generalization and reduce overfitting, data augmentation techniques such as rotation, zooming, horizontal flipping, and brightness variation were applied to the training set. Validation and test images were untouched.

Table 1 Pomegranate Fruit Sample Images from Dataset

Category	Sample Image
Alternaria	
Anthracnose	
Bacterial Blight	
Cercospora	

Table 2 Dataset Pre-splitting

Category	Training	Testing	Validation	Total
Alternaria fruit spot	620	132	134	886
Anthracnose	816	174	176	1166
Bacterial blight	676	144	146	966
Cercospora fruit spot	441	94	96	631
Healthy	1014	217	219	1450
Total				5099

Model Selection

This study is entirely based on pre-trained deep learning models using a transfer learning approach. Two lightweight models were selected: MobileNetV3-Large and EfficientNet. Both models were initialized with ImageNet weights and used as feature extractors by removing their original classification layers.

MobileNetV3

MobileNetV3 uses depthwise separable convolutions which reduces the number of parameters, computational cost and make it more lightweight. To enhance the representation power of the model, Squeeze-and-Excitation (SE) blocks are used to compute channel wise attention weights and then scale the feature maps. To provide better performance, the Hard-Swish activation function is used rather than the traditional ReLU function.

EfficientNet

EfficientNet is based on compound scaling which uses a fixed scaling coefficient at each scale. This allows the model to capture multi-scale and hierarchical features without incurring any high computational cost. The model implements MBConv (mobile inverted bottleneck convolution) blocks, which combine depthwise separable convolutions with expansion layers and squeeze-and-excitation modules. This structure allows EfficientNet to gain rich feature extraction while maintaining accuracy. As a result, EfficientNet shows remarkable improvement during the fine-tuned phase as deep layers acquire more domain-specific disease patterns.

Training and Evaluation

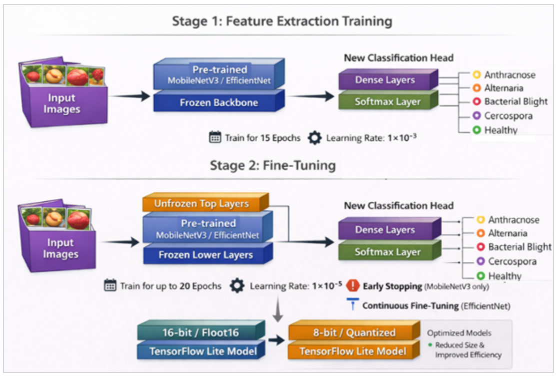

Figure 3 Transfer Learning with MobileNetV3 and EfficientNet

As shown in Figure 3, the proposed model is implemented in two stages: the first stage represents feature extraction and the second stage holds for fine tuning. During the first stage, all convolution layers of the backbone were frozen and the newly added classification head was trained. The classification head was composed of a global average pooling layer, a fully connected layer and a Softmax output layer. This stage was trained for 15 epochs using Adam optimizer with learning rate of 1×10^{-3}, which is very standard for preventing overfitting.

In the second stage, fine-tuning was performed to make the model more appropriate for pomegranate disease domain classification. The last 40 layers were used from the first stage trained model keeping all early layers frozen, so that the model learns correctly and does not suffer with overfitting. Fine tuning was compiled with 20 epochs and a reduced learning rate of 1×10^{-5} to ensure stable convergence.

For MobileNetV3, early stopping at the sixth epoch was observed during the fine-tuning stage because there was no considerable improvement in the accuracy of the model in the last three epochs. Also, to prevent overfitting, early stopping was employed in the model. This behaviour aligns with MobileNetV3's lightweight architecture, as it typically learns very fast. In contrast, EfficientNet was trained for all 20 epochs in the second stage and could achieve remarkable accuracy in the last epoch.

Model Optimization and Quantization

To enable real time implementation of these trained models in resource-constrained agriculture environments, the models are further optimized using post-training quantization. Each model was converted to 16-bit (Float16) and 8-bit (Int8) versions against original 32-bit versions. This significantly reduced model size and inference time while preserving classification performance.

Interpretability Analysis

To address the black-box nature of deep learning models, Grad-CAM was implemented for generating visual understanding of decisions made by the trained models. Grad-CAM does not require any model modification or retraining. It extends the concept of Class Activation Mapping (CAM) that produces image captions which are then utilized in fully connected layers of the trained model to provide visual explanation of the model's prediction. The Grad-CAM generates heatmaps on the affected regions of the image, allowing local farmers to validate it with expert opinion.

Results and Discussions

The performance of the proposed models MobileNetV3-Large and EfficientNet was assessed using accuracy for both Stage 1 and Stage 2. Additionally, precision, recall, F1-Score and confusion matrix were generated to analyze class-wise prediction performance.

Table 3 Model Performance Comparison

Model	Accuracy (%)		Precision	Recall	F1-Score
	Stage 1	Stage 2			
MobileNetV3	98.54	98.46	0.9871	0.9869	0.9869
EfficientNet	98.71	99.50			

Table 4 Confusion Matrix

<table>
<tr><th>Model</th><th>Confusion matrix</th></tr>
<tr><td>MobileNetV3</td><td>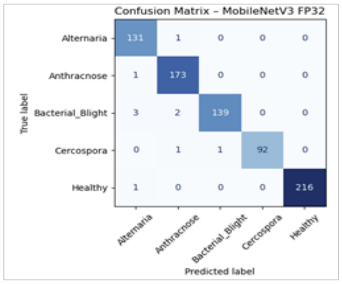
</td></tr>
<tr><td>EfficientNet</td><td>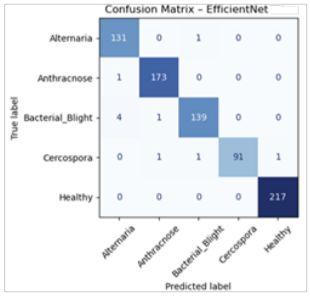
</td></tr>
</table>

In Stage 1, where the backbone networks were kept frozen, both models achieved reasonable classification performance. EfficientNet slightly outperformed MobileNetV3 across all evaluation metrics, indicating stronger feature extraction capabilities even without fine-tuning.

Table 5 Model Accuracy Comparison

Model	Stage 2 accuracy (%)	FP16 accuracy (%)	INT8 accuracy (%)
MobileNetV3	98.46	98.68	95.17
EfficientNet	99.50	98.68	95.13

Table 6 Model Size Comparison

Model	Original size (MB)	FP16 size (MB)	INT8 size (MB)
MobileNetV3	14.84	6.17	3.56
EfficientNet	31.45	8.33	4.99

Model quantization significantly reduced storage requirements. INT8 quantization resulted in the highest compression, making both models suitable for deployment on edge and mobile devices without compromising structural integrity.

Table 7 Inference Time Comparison

Model	Regular (ms)	FP16 (ms)	INT8 (ms)
MobileNetV3	243.56	15.45	8.36
EfficientNet	296.29	38.60	29.88

Quantized models demonstrated faster inference times compared to their regular counterparts. MobileNetV3 consistently achieved lower latency than EfficientNet, highlighting its suitability for real-time applications such as mobile-based crop disease detection systems.

Table 8 Grad-CAM Visualization Results

Original image	Image after applying Grad-CAM
Category — Anthracnose	
	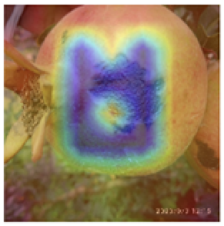
Category — Cercospora	
	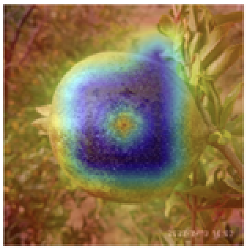
Category — Bacterial blight	
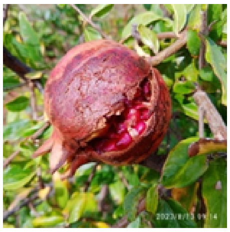	
Category — Alternaria	
	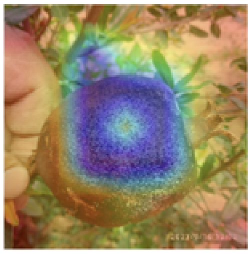
Category — Healthy	
	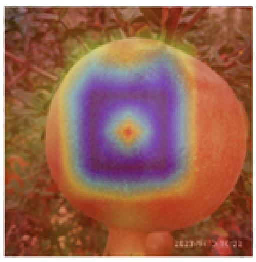

Grad-CAM visualizations were used to interpret model predictions and identify discriminative regions influencing classification decisions. The highlighted regions correspond closely to disease-affected areas such as lesions, discoloration, and necrotic spots. Fine-tuned models exhibited more focused and accurate attention maps, indicating improved feature localization.

Conclusion

This study presented a two-stage transfer learning approach for plant disease classification using MobileNetV3 and EfficientNet architectures. Experimental results showed that fine-tuning significantly improved classification performance across all metrics. EfficientNet achieved superior accuracy, while MobileNetV3 demonstrated faster inference and lower memory requirements. Model optimization using TensorFlow Lite further reduced model size and inference latency, making the system suitable for real-time and edge deployment. Grad-CAM analysis validated the interpretability and reliability of model predictions. Future work will focus on expanding the dataset, incorporating additional crop diseases, and deploying the system in real-world agricultural environments.

References

1. Horticultural Science at Glance. 2018.
2. D. P. Rajput and Jahanara, Knowledge and attitude of the pomegranate growers about recommended cultivation technology in Pune District of Maharashtra, India. Asian Journal of Agricultural Extension, Economics & Sociology, 42(6), 190–194. 2024. doi: 10.9734/ajaees/2024/v42i62478.
3. A. Khakimov, I. Salakhutdinov, A. Omolikov, and S. Utaganov, Traditional and current-prospective methods of agricultural plant diseases detection: A review. IOP Conference Series: Earth and Environmental Science, 951(1), 012002. 2022. doi: 10.1088/1755-1315/951/1/012002.
4. C. Subbarayudu and M. Kubendiran, A comprehensive survey on machine learning and deep learning techniques for crop disease prediction in smart agriculture. Nature Environment and Pollution Technology, 23(2), 619–632. 2024. doi: 10.46488/nept.2024.v23i02.003.
5. Q. Tian et al., Enhancing practicality of deep learning for crop disease identification under field conditions: insights from model evaluation and crop-specific approaches. Pest Management Science, 80(11), 5864–5875. 2024. doi: 10.1002/ps.8317.
6. Z. Chen, Y. Cao, Y. Liu, H. Wang, T. Xie, and X. Liu, A comprehensive study on challenges in deploying deep learning based software. Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, 750–762. 2020. doi: 10.1145/3368089.3409759.
7. T. Samson Adekunle, M. Oladayo Lawrence, O. Omotayo Alabi, A. A. Afolorunso, G. Nse Ebong, and M. Abiola Oladipupo, Deep learning technique for plant disease detection. Computer Science and Information Technologies, 5(1), 49–56. 2024. doi: 10.11591/csit.v5i1.p49-56.
8. S. C. K., J. C. D., and N. Patil, Cardamom plant disease detection approach using EfficientNetV2. IEEE Access, 10, 789–804. 2022. doi: 10.1109/access.2021.3138920.
9. V. Dhore, A. Bhat, V. Nerlekar, K. Chavhan, and A. Umare, Enhancing explainable AI: a hybrid approach combining GradCAM and LRP for CNN interpretability. 2024. doi: 10.48550/arXiv.2405.12175.
10. K. Gopalan, S. Srinivasan, P., M. Singh, S. K. Mathivanan, and U. Moorthy, Corn leaf disease diagnosis: enhancing accuracy with ResNet152 and Grad-CAM for explainable AI. BMC Plant Biology, 25(1). 2025. doi: 10.1186/s12870-025-06386-0.
11. C. Messina et al., Toward a general framework for AI-enabled prediction in crop improvement. Theoretical and Applied Genetics, 138(7). 2025. doi: 10.1007/s00122-025-04928-6.
12. P. Pawaskar, H. K. Yogish, B. Pakruddin, and Y. Deepa, A novel and efficient deep learning models for assessing AI's impact on disease diagnosis in agriculture. Indian Journal of Agricultural Research. 2025. doi: 10.18805/ijare.a-6388.

13. Y. Dandawate and R. Kokare, An automated approach for classification of plant diseases towards development of futuristic decision support system in Indian perspective. 2015 International Conference on Advances in Computing, Communications and Informatics (ICACCI), 794–799. 2015. doi: 10.1109/icacci.2015.7275707.
14. P. K. Das, S. S. Rupa, S. Pumrin, U. C. Das, and Md. K. Hossen, Deep learning for plant disease detection and classification: a systematic analysis and review. Current Applied Science and Technology, e0259016. 2024. doi: 10.55003/cast.2024.259016.
15. E. al. Ashok Kumar, Deep learning for automated diagnosis of plant diseases: a technological approach. Journal of Electrical Systems, 20(1), 01–08. 2024. doi: 10.52783/jes.655.
16. S. Senthil Pandi, A. K. Reshmy, S. Vinodh Kumar, and P. Kumar, Towards precision agriculture: harnessing deep learning for accurate plant disease diagnosis. 2024 International Conference on Communication, Computing and Internet of Things (IC3IoT), 1–6. 2024. doi: 10.1109/ic3iot60841.2024.10550238.
17. J. J. R. Angelina, S. J. Subhashini, M. V. Krishna, V. V. Naveen, K. V. Krishna, and C. Yu-Chun, Plant disease detection using deep learning. 2024 10th International Conference on Communication and Signal Processing (ICCSP), 1494–1498. 2024. doi: 10.1109/iccsp60870.2024.10543534.
18. Q. Liu, S. Zhou, and J. Lai, EdgeMedNet: lightweight and accurate U-Net for implementing efficient medical image segmentation on edge devices. IEEE Transactions on Circuits and Systems II: Express Briefs, 70(12), 4329–4333. 2023. doi: 10.1109/tcsii.2023.3291168.
19. I. Laktionov, G. Diachenko, D. Rutkowska, and M. Kisiel-Dorohinicki, An explainable AI approach to agrotechnical monitoring and crop diseases prediction in Dnipro Region of Ukraine. Journal of Artificial Intelligence and Soft Computing Research, 13(4), 247–272. 2023. doi: 10.2478/jaiscr-2023-0018.
20. F. Rajeena P. P., A. S. U., M. A. Moustafa, and M. A. S. Ali, Detecting plant disease in corn leaf using EfficientNet architecture—an analytical approach. Electronics, 12(8), 1938. 2023. doi: 10.3390/electronics12081938.
21. Z. Xie, C. Li, Z. Yang, Z. Zhang, J. Jiang, and H. Guo, YOLOv5s-BiPCNeXt, a lightweight model for detecting disease in eggplant leaves. Plants, 13(16), 2303. 2024. doi: 10.3390/plants13162303.
22. S. Mandava et al., Deep learning architecture, EfficientNetB1, excels in accurate diagnosis and prognosis of acute lymphoblastic leukemia through intelligent computer vision. Blood, 146(Supplement 1), 552–552. 2025. doi: 10.1182/blood-2025-552.
23. S. Bagal, M. S. Patil, S. Thite, A. A. Chandanshive, and N. M. Pise, Survey of diseases and pest on pomegranate from Solapur district. 2022.