RESEARCH ARTICLE Indian Journal of Agricultural Research, Volume 60 Issue 11: 1775-1783 (November 2026)

AFS-PLDCNet: An Advanced Computational Tool for the Classification of Apple Leaf Diseases

Neha Sawant[1], K.L. Bansal[1]

10.18805/IJARe.A-6605

ABSTRACT

Background: Accurate and timely diagnosis of foliar diseases is the most crucial factor in efforts to maximize crop yield and ensure sustainability. Existing deep learning models, especially single-backbone CNNs, have achieved promising results; however, they often fail to generalize well in different orchard conditions.

Methods: In this study, the AFS-PLDCNet framework has been proposed for robust leaf disease classification. This framework uses an attention-based feature-fusion approach to combine deep representations extracted from EfficientNetV2S, MobileNetV2 and ResNet18. A learnable attention mechanism assigns adaptive weights dynamically to each feature and a lightweight meta-learner is used for classification. A new dataset of 9000 apple leaf images was captured in Himachal Pradesh's orchards, encompassing Alternaria leaf blotch, Marssonina blotch and healthy leaves.

Result: The experimental results demonstrate that AFS-PLDCNet achieved superior classification accuracy compared to existing single-backbone CNNs. The proposed model is well-suited for real-time, field-level leaf disease classification and precision agriculture systems.

Key words: Apple leaf disease, Attention mechanism, Deep learning, Ensemble learning, Feature fusion.

INTRODUCTION

Malus domestica, commonly known as the Apple, belongs to the Rosaceae family and is grown in almost 93 nations globally. Presently, China is at the forefront of apple cultivation, followed by the USA, Turkey, Poland and India in fifth place (Kumar *et al*., 2022). The Himalayan region harbors the majority of apple orchards, with Jammu and Kashmir, Himachal Pradesh and Uttarakhand being the primary apple-producing states (Ripin *et al*., 2024). The production of apples generally depends upon the soil, climate, pests and disease management. The presence of diseases significantly impacts the quantity and quality of fruits, resulting in low yields and inferior quality (Zhong and Zhao, 2020; Mehta *et al*., 2025).

Timely and accurate disease detection is vital as misidentification can lead to the wrong selection of pesticides, resulting in substantial losses (Zhong and Zhao, 2020). Therefore, automated methods are increasingly used to detect and classify diseases accurately, as visual inspection requires field expertise and is time-consuming and costly (Sai Reddy and Neeraja, 2022). Deep learning, particularly convolutional neural networks (CNNs), has become a dominant tool due to its ability to automatically extract discriminative features without complex pre-processing (Cho, 2024), its efficient utilization of GPUs for faster computation (Mustafa *et al*., 2022) and its capability to process extremely large datasets with high accuracy (Barbedo, 2018; Attallah, 2023; Lee and Kim, 2024).

Even though several CNN-based and transfer learning models have shown outstanding results in apple leaf disease classification, there are still several unaddressed issues. The current approaches are based on single backbone designs that do not consider the complementary multi-scale characteristics and others use an ensemble of large-scale models that cannot be deployed in real-time in agricultural settings. The third major limitation is the unavailability of diversified apple leaf disease datasets. The limited available datasets are mostly gathered under controlled settings, which are not equivalent to the varied natural orchards in terms of lighting, orientation and background. All these contribute to poor model generalization and restrict their application in field settings.

The recent attention-based fusion schemes for apple leaf disease, including BLSENet (SE attention with bilinear dual-subnetwork fusion) (Fang *et al*., 2024) and EMA-enhanced lightweight detector (Wang *et al*., 2025) are tested primarily on data from their source and their generalization ability to unseen orchard conditions remains unproven. However, due to the distribution mismatch, domain adaptation approaches have been explored, such as cross-dataset discriminant subspace learning

[1]Department of Computer Science, Himachal Pradesh University, Shimla-171 005, Himachal Pradesh, India.

Corresponding Author: Neha Sawant, Department of Computer Science, Himachal Pradesh University, Shimla-171 005, Himachal Pradesh, India. Email: sawant89neha@gmail.com

ORCID: 0000-0002-0234-8145

How to cite this article: Sawant, N. and Bansal, K.L. (2026). AFS-PLDCNet: An Advanced Computational Tool for the Classification of Apple Leaf Diseases. *Indian Journal of Agricultural Research.* **60(11):** 1775-1783. doi: 10.18805/IJARe.A-6605.

Submitted: 04-06-2026 **Accepted:** 17-07-2026 **Online:** 13-08-2026

(Yan and Wang, 2024). However, none of these studies combine domain adaptation with attention-guided multi-backbone fusion for apple leaf disease classification.

To address these limitations, this paper proposes an attention-based feature-stacked plant leaf disease classification network (AFS-PLDCNet). The proposed framework combines EfficientNetV2S, MobileNetV2 and ResNet18 and an adaptive attention-guided feature fusion scheme that learns dynamically to prioritize the most discriminative features. By utilizing the complementarity of several CNNs, AFS-PLDCNet is more robust, less redundant and computationally efficient for real-time apple leaf disease classification in orchards.

The primary contributions of this paper are as follows:

- A novel, self-curated leaf disease dataset with 9000 images of Alternaria, healthy and Marssonina leaves.
- A novel, attention-based feature-stacked plant leaf disease classification network (AFS-PLDCNet) is proposed.
- The suggested framework's generalization is assessed on an external validation dataset and enhanced by a lightweight domain-specific adaptation of the meta-learner.

Related work

Thakur *et al*. (2022) introduced a lightweight VGG-ICNN, combining VGG16 and InceptionV7, reaching 99.16% accuracy on the PlantVillage dataset. Ali *et al*. (2024) introduced AppleLeafNet, a lightweight 37-layer, 1.33 M parameters, to diagnose apple-leaf diseases in real-time. It reached 98.25% to 98.60% accuracy on the Plant Pathology 2021-FGVC8 dataset. Singh *et al*. (2021) proposed a 19-layer CNN outperforming various state-of-the-art models with 99.2% accuracy on the self-curated dataset. Fu *et al*. (2022) added multiscale layers and a squeeze-and-excitation attention mechanism to AlexNet. The model scored 97.36% accuracy on the AI Studio dataset, outperforming ResNet-50 and GoogLeNet.

The custom CNN structures have shown encouraging results, but their performance is often constrained by limited data and generalization issues, leading researchers to implement transfer learning and fine-tuning with pretrained CNNs. Liu *et al*. (2023) proposed ECA-DCMobileNet, a modified MobileNet-V2, reaching 95% accuracy, outperforming AlexNet, VGG16, ResNet50, ShuffleNet and MobileNetV2. Srinidhi *et al*. (2021) fine-tuned EfficientNet-B7 and DenseNet-121 on the Plant Pathology 2020(FGVC7) dataset to classify apple leaf disease and achieved an accuracy of 99.8% and 99.75%, respectively. The authors recommended considering stacking and ensemble techniques for robustness. Ali *et al*. (2025) fine-tuned an EfficientNet-B0 with global max-pooling, batch normalization and regularization and achieved 99% accuracy. Iftikhar *et al*. (2024) compared AlexNet, MobileNet and ResNet against the proposed fine-tuned E-CNN with SVM and LR. E-CNN achieved 98.17% accuracy.

Although transfer learning has significantly improved apple leaf disease classification, the single-backbone model might not be effective in using complementary feature representations. This weakness has led to the creation of ensemble learning and feature fusion-based convolutional neural network methods, which are the cornerstone of the most recent developments. Gupta *et al*. (2025) presented E-YOLOv8, a lightweight detector for apple leaf disease, achieving 93.9% mAP across three datasets. Zhang *et al*. (2023) enhanced the ResNet-50 with Coordinate Attention (CA) and Weight-Adaptive Multi-scale Feature Fusion (WAMSFF), reaching 98.32% accuracy on the AppleLeaf9 and ATLDSD datasets. Apleni *et al*. (2025) introduced a concatenation- based VGG16, ResNet50 and InceptionV3 ensemble for plant disease classification, reaching 97% accuracy. Sood *et al*. (2025) built a DenseNet121-ResNet50 fusion model for crop diseases and pest detection, achieving 99.13% and 98.67% accuracy on PlantVillage and IP102.

MATERIALS AND METHODS

The proposed methodology is shown in Fig 1. It begins with dataset collection, followed by preprocessing, model training and evaluation on the holdout and external validation sets.

Dataset

In this research, the experiment was performed on two datasets: (a) A self-curated, real-world dataset of leaf disease, created to specifically support this study and (b) an external dataset, AppleLeaf9 (Yang *et al*., 2022), from which only the classes relevant to this study were extracted for external validation.

Self-curated leaf dataset

The dataset for this research consisted of 9,000 images, with 3,000 images per class, *i.e*., Alternaria, Marssonina and healthy apple leaves. All images were captured using a mobile phone at three locations (Kotkhai, Jubbal and Mashobra) in the Shimla district of Himachal Pradesh at varying elevations. Pictures were captured from different angles and under various lighting conditions in a natural setting with unaltered backgrounds. Mild augmentation techniques, including random resized cropping (scale 0.9-1.0), horizontal flipping, rotations up to 15 degrees and

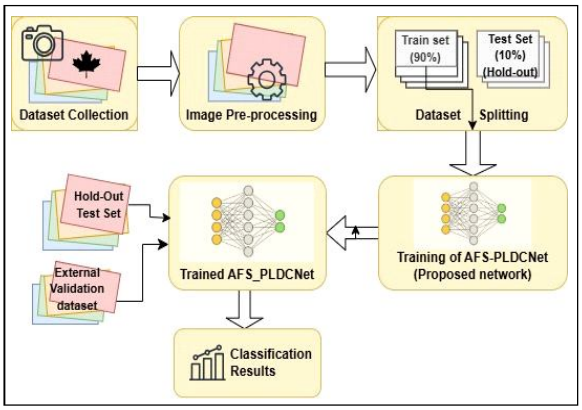


Fig 1: Flow diagram of the proposed methodology.

color jittering, were applied during training. All the images were resized to 224*224 pixels. The dataset was split into 90% training and 10% testing using stratified sampling. Fig 2 shows some sample images from the dataset. The leaves were examined and verified by an expert at Government Progeny cum Demonstration Orchard (PCDO), Annu (Jubbal), under Himachal Pradesh Department of Horticulture. The study was conducted at H.P University during 2022-2025 with field data collected over three growing seasons.

External validation dataset (AppleLeaf9)

The AppleLeaf9 (Yang *et al*., 2022) dataset is used to validate the proposed ensemble model on an unseen domain. This dataset is a fusion of four publicly available datasets. From this, 1388 images of healthy, Alternaria and Marssonina classes were selected to match the categories in the self-curated dataset.

Experimental setup

All the experiments were run on the Kaggle platform using dual Nvidia Tesla T4 GPU acceleration (2 × 16 GB) with PyTorch, TorchVision, NumPy and Scikit-learn libraries.

Ensemble architecture (AFS-PLDCNet)

The proposed AFS-PLDCNet (attention-based feature-stacked plant leaf disease classification network) is a multi-stage framework that combines lightweight base models, cross-validation-based feature extraction, attention-guided feature fusion and a meta-learner for final classification.

Fig 3 is a schematic diagram of the general structure of the AFS-PLDCNet model and its workflow.

Lightweight base model

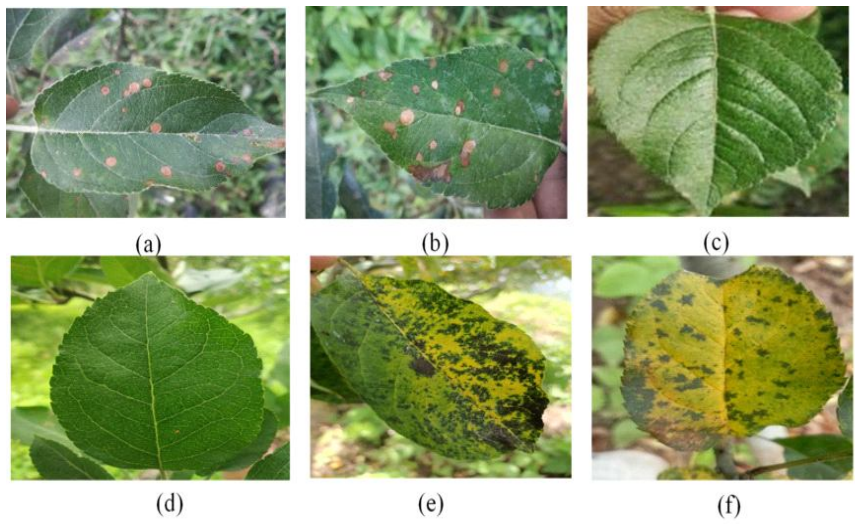
Three widely used lightweight architectures, namely MobileNetV2, EfficientNetV2S and ResNet18, were used as the base models to experiment. MobileNetV2 (Sandler *et al*., 2018) is a small-sized model for mobile/embedded vision applications. It has inverted residual blocks and linear bottlenecks for speed and accuracy. EfficientNetV2S (Tan and Le, 2021) is an improved version of the EfficientNet series, which is optimized to enable faster training and higher predictive accuracy. ResNet18 (He *et al*., 2016) is part of the ResNet family, which uses residual connections to mitigate the vanishing gradient problem, allowing deeper training without performance loss.

K-fold training and feature extraction

For robust learning and maximum generalization, all three lightweight convolutional backbones, MobileNetV2, ResNet18 and EfficientNetV2S, were trained with a 5-fold cross-validation technique. All three backbone networks were fine-tuned for up to 30 epochs using the Adam optimizer (learning rate= 0.0001, batch size= 32, loss= categorical cross-entropy) and with early stopping (patience= 7 epochs) to avoid catastrophic forgetting of pre-trained features and ensure convergence. For the attention module, the backbone features were extracted as the fold-concatenated features (EfficientNetV2S-6400d, MobileNet V2-6400d, ResNet18-2560d, *i.e*., Total: 15360d) and for weighted fusion, the backbone features were extracted as fold-averaged features (1280d, 1280d, 512d).

Attention-based feature fusion

A lightweight attention network is used for estimating the relative importance of each backbone by analyzing the combined feature representation. The attention module comprises three linear layers: Linear (15360→512)→ReLU →dropout, linear (512→128)→ReLU→dropout, linear (128→3) LayerNorm→Softmax, producing three attention weights [α1, α2, α3] (shown in Fig 3). A two-stage training



(a), (b) Alternaria-infected leaves; (c), (d) Healthy leaves; (e), (f) Marssonina-infected.

Fig 2: Sample images from self-curated dataset.

strategy is adopted to ensure meaningful attention-weight learning. During stage 1, the weights of the attention neural network are directly learned from the ground truth classes via a simple classification head, which allows the network to learn from base models which features are most reliable. During stage 2, the attention module is fixed and the updated feature vectors (which are the weighted sum of fold-averaged features) are utilized to train the meta-learner. The decoupled approach has the advantage of avoiding gradient ambiguity, while resulting in per-sample attention weights that exhibit much more variation than joint training approaches:

$$f^{fused} = \alpha_{eff} \cdot \widetilde{f_{eff}} + \alpha_{mob} \cdot \widetilde{f_{mob}} + \alpha_{res} \cdot \widetilde{f_{res}} \quad ...(1)$$

In equation 1, α_{eff}, α_{mob}, α_{res} are the attention weights and, f_eff, f_mob and f_res are the feature vectors derived from EfficientNetV2S, MobileNetV2 and ResNet18, respectively. Zero-padding (denoted by the tilde sign ~ in equation 1) was applied to the features before fusion.

Meta-learner for final classification

After the attention-based feature fusion, the fused features were used to train a meta-learner for final disease classification. The meta-learner comprised two fully connected hidden layers of decreasing width, LayerNorm for stable training and ReLU activations to capture non-linear relationships. At the end, the meta-learner produced an output that included a class derived from the attention-fused feature vectors.

The meta-learner was trained separately by using stratified splits for training and validation datasets. The best-performing model was kept based on the validation performance. This decoupled design enables efficient decision-level learning and improves generalization, which was further verified on an independent dataset excluded from training.

RESULTS AND DISCUSSION

Performance on hold-out test dataset

The performance of the proposed AFS-PLDCNet architecture was tested on a hold-out test set that consists of 10% of the initial dataset. This subset was not used for training the model and has the same data distribution as the training data. AFS-PLDCNet had a total classification accuracy of 99.67% on the hold-out test set with high levels of precision, recall and F1-scores in each of the disease categories.

As shown in the confusion matrix in Fig 4(a), the predictive performance in each of the classes is quite strong, although a few instances exist in which some of the errors are observed. The ROC curves in Fig 4(b) affirm excellent threshold-independent behavior, with almost ideal area under the curve (AUC) values on each of the classes and a micro-averaged area under the curve (AUC) of 0.9984, establishing a strong in-distribution benchmark.

Cross-dataset generalization on external validation dataset (AppleLeaf9)

To assess the cross-domain generalization of the AFS-PLDCNet, the trained model was evaluated on the AppleLeaf9 external validation dataset without any fine-tuning. On the AppleLeaf9 dataset, the model had a total accuracy of 91.55%. This accuracy was significantly lower than simple averaging-based ensembles (94.5%, McNemar's test, $p<0.001$), motivating the adaptation step below.

The confusion matrix in Fig 5(a) shows that most samples are well clustered, but some confusion can be noticed in the Alternaria and Marssonina classes. The ROC curves in Fig 5(b) confirm high threshold-independent classification under a domain shift, with a micro-averaged AUC of 0.9872, indicating strong feature separability on unseen data.

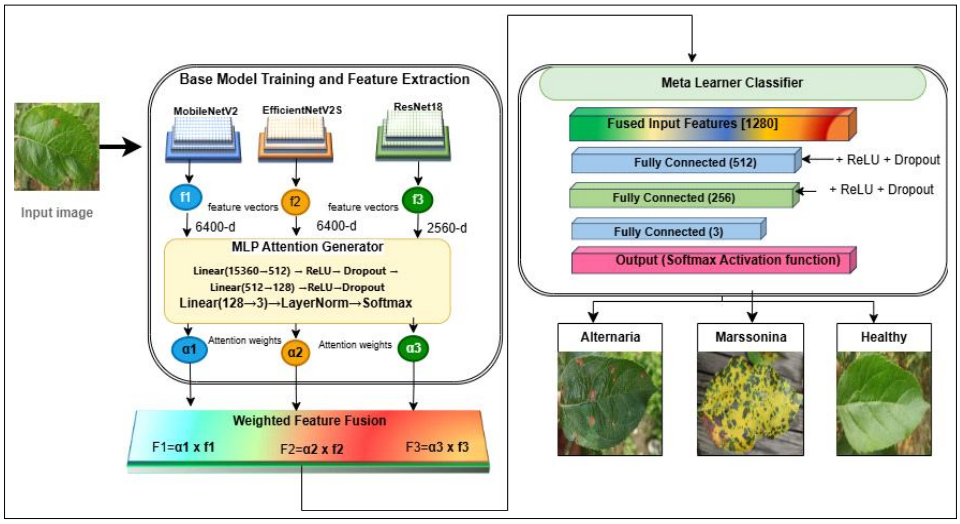


Fig 3: Proposed architecture of AFS-PLDCNet.

Domain adaptation through meta-learner fine-tuning

The cross-dataset analysis shows that, despite AFS-PLDCNet's strong generalization to unseen data, domain shift has adversely affected the model's performance. To address this, the meta-learner was fine-tuned on a stratified 80% subset of AppleLeaf9 while keeping feature extractors fixed and accuracy was calculated on the remaining 20% held-out subset, which was not used during fine-tuning. As shown in Fig 6(a) and (b), accuracy on this held-out subset improved from 92.16% (before adaptation) to 98.51% (after adaptation), compared to 91.55% obtained before adaptation on the complete unseen AppleLeaf9 dataset. This improvement was statistically significant (McNemar's test, $p<0.001$), confirming the gain is not attributable to chance.

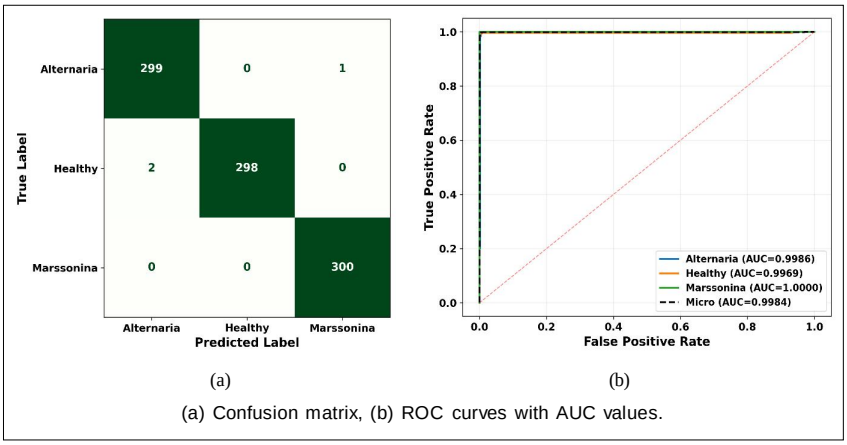
Computational complexity analysis

The practical deployment cost of AFS-PLDCNet was quantified by measuring the floating-point operations (GFLOPs), number of parameters, inference latency for a single image, throughput and on-disk model size, averaged over 50 timed runs with a 10-run warm-up (Table 1).

Comparative analysis of performance

Table 2 presents the classification performance of the three base models and the proposed AFS-PLDCNet under three evaluation settings: hold-out test, external validation and domain adaptation. The number of reported parameters and the size of AFS-PLDCNet are for the lightweight attention and meta-learner fusion modules only, as the three backbone feature extractors are the same and are frozen. Inference time and feature extraction were measured 5 times on each of the three backbones. Although fold-wise feature extraction increases the computational complexity, AFS-PLDCNet is slightly lighter (in terms of trainable parameters, storage size and computation) than each backbone model, as it merely uses a small multilayer perceptron for the fusion stage over features computed by the backbones.

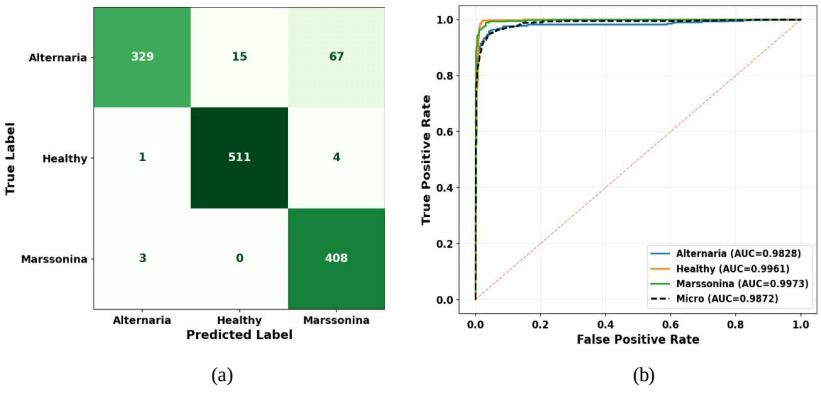
AFS-PLDCNet performed better than the individual backbone networks on the hold-out test set, supporting the benefit of attention-based feature fusion. Performance declined on external datasets due to domain shift in acquisition, imaging devices and composition of the datasets. This result highlights the limitation of direct



(a) (b)

(a) Confusion matrix, (b) ROC curves with AUC values.

Fig 4: AFS-PLDCNet performance on the hold-out test set.



(a) (b)

(a) Confusion matrix; (b) ROC-AUC curves for threshold-independent classification.

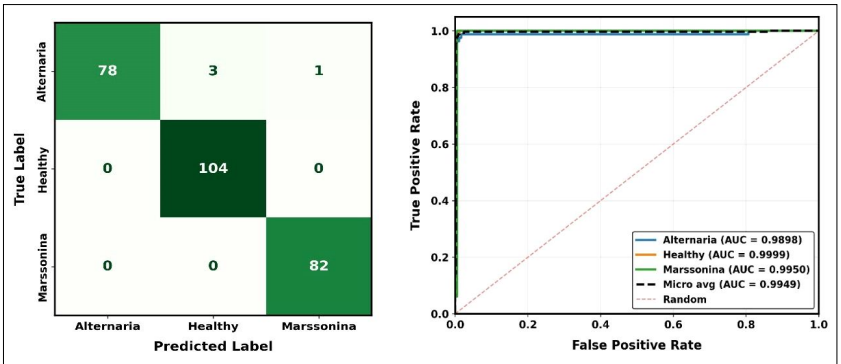
Fig 5: AFS-PLDCNet performance on external validation dataset (AppleLeaf9).

cross-data generalization, which was substantially recovered through lightweight domain adaptation *via* meta-learner fine-tuning, without retraining the feature extractors.

Comparison with ensembling strategies

To isolate the contribution of attention-guided fusion, AFS-PLDCNet was further compared against four simple ensembling strategies applied to the same three backbones (majority voting, soft voting, weighted averaging and max-confidence selection) on the hold-out test set (Table 3). AFS-PLDCNet attained the highest numerical accuracy (99.67%), though McNemar's test showed no statistically significant difference from simple ensembling at this near-ceiling performance level.

Table 4 shows that the proposed AFS-PLDCNet performs better than the existing ensemble and fusion-based frameworks. From the practical point of view, the lightweight modular structure is ideal for mobile and IoT systems for disease diagnosis, allowing farmers to perform on-site detection without the need for high-level computing resources.



(a) (b)

(a) Confusion matrix; (b) ROC-AUC curves illustrating improved threshold-independent classification.

Fig 6: Performance of fine-tuned meta-learner on held-out 20% subset.

Table 1: Computational complexity of the base models and AFS-PLDCNet.

Model	GFLOPs	Params (M)	Inference (ms/img)	Throughput	Size (MB)
EfficientNetV2S	2.89	20.18	16.08	59.5	77.89
MobileNetV2	0.32	2.23	5.02	199.2	8.74
ResNet18	1.82	11.18	2.82	354.2	42.72
AFS-PLDCNet	25.19	8.81	24.61	40.6	33.28

Table 2: Comparative analysis of base model and proposed AFS-PLDCNet across evaluation settings.

Model	Accuracy (%)	Precision	Recall	F1-score	AUC (micro-averaged)
Base Model-EfficientNetV2S (Hold-out test set)	99.56	0.9956	0.9956	0.9956	0.9996
Base Model-MobileNetV2 (Hold-out test set)	99.44	0.9945	0.9944	0.9944	1.0000
Base Model-ResNet18 (Hold-out test set)	99.22	0.9922	0.9922	0.9922	0.9999
Base Model-EfficientNetV2S (External validation dataset-AppleLeaf9)	89.84	0.9149	0.8984	0.8938	0.9833
Base Model-MobileNetV2 (External validation dataset-AppleLeaf9)	92.30	0.9227	0.9230	0.9228	0.9846
Base Model-ResNet18 (External validation dataset-AppleLeaf9)	93.20	0.9333	0.9320	0.9317	0.9874
AFS-PLDCNet (Hold-out test Set)	99.67	0.9967	0.9967	0.9967	0.9984
AFS-PLDCNet (External validation dataset-AppleLeaf9)	91.55	0.9261	0.9155	0.9131	0.9872
AFS-PLDCNet+domain adaptation	98.51	0.9854	0.9851	0.9850	0.9949

Table 3: Comparison of ensembling strategies and AFS-PLDCNet on the hold-out test set.

Method	Accuracy (%)	Precision	Recall	F1 score	AUC
Majority voting	99.56	0.9956	0.9956	0.9956	0.9830
Soft vote	99.56	0.9956	0.9956	0.9956	1.0000
Weighted average	99.56	0.9956	0.9956	0.9956	1.0000
Max confidence	99.56	0.9956	0.9956	0.9956	0.9997
AFS-PLDCNet	99.67	0.9967	0.9967	0.9967	0.9984

Table 4: Comparison of the proposed framework- AFS-PLDCNet with state-of-the-art models.

References	Year	Base model used	Technique used	Data source	Accuracy
(Apleni *et al.*, 2025)	2025	VGG16, ResNet50, InceptionV3	Used concatenation-based feature fusion	Plant village dataset	97%
(Gupta *et al.*, 2025)	2025	E-YoloV8	A light weight E-YOLOV8 framework with CBAM attention module	Public dataset+ custom dataset	93.9%
(Sood *et al.*, 2025)	2025	DenseNet-121 and ResNet-50	Used a hybrid fusion (feature concatenation)	Plant village	99.13%
Proposed framework	2026	MobileNetV2, EfficientNetV2S, ResNet18	An Attention-based feature stacked ensembling technique is used	Custom dataset	99.67%

CONCLUSION AND FUTURE WORK

This paper presents an attention-based feature-stacked plant leaf disease classification network (AFS-PLDCNet), which is used to classify apple leaf disease, namely Alternaria and Marssonina blotch, using a novel self-curated dataset obtained from orchards in Shimla, Himachal Pradesh. The framework combines three lightweight convolutional backbones, *i.e*., EfficientNetV2S, MobileNetV2 and ResNet18, through an adaptive attention-based feature weighting mechanism. The proposed ensemble achieved top-tier performance, with a test accuracy of 99.67% on the hold-out dataset and 91.55% on the external validation dataset (AppleLeaf9). On a stratified 20% held-out subset of this dataset, accuracy was increased from 92.16% to 98.51% after fine-tuning the proposed model. This indicates domain adaptation can effectively recover cross-domain performance. The small size of the meta-learner makes it viable to implement in real-time on resource-limited agricultural monitoring systems, such as mobile and IoT devices.

Despite these strengths, the dataset is relatively small and focuses on only two diseases of apple leaves. In addition, the experiment assesses leaves with symptoms of a single disease, but in real life, there are often several or concomitant infections. In addition, the hold-out to external accuracy gap hints at some overfitting to acquisition conditions, even after domain adaptation. End-to-end inference still requires running all three backbones, so the cost on GPU remains non-trivial. Most prior work also does not report efficiency, which limits Table 4's comparison to accuracy alone. Future research will focus on increasing the dataset and adding more disease varieties, incorporating multi-disease samples and cross-regional data and examining sophisticated attention and transformer-based fusion modalities to enable on-site disease detection using mobile and IoT solutions. Overall, AFS-PLDCNet can be considered a significant step towards the practical, lightweight and interpretable deep learning solutions for precision horticulture and automated disease management in the apple production system.

ACKNOWLEDGEMENT

The authors wish to convey their sincere appreciation to all the researchers whose contributions and findings have been referenced in this study.

Disclaimers

The views and conclusions expressed in this article are solely those of the authors and do not necessarily represent the views of their affiliated institutions. The authors are responsible for the accuracy and completeness of the information provided, but do not accept any liability for any direct or indirect losses resulting from the use of this content.

Data availability

The self-curated dataset used in this study is available from the corresponding author on reasonable request.

Informed consent

Not applicable.

Conflict of interest

The authors declare that there are no conflicts of interest regarding the publication of this article. No funding or sponsorship influenced the design of the study, data collection, analysis, decision to publish, or preparation of the manuscript.

REFERENCES

Ali, H., Shifa, N., Benlamri, R., Farooque, A.A. and Yaqub, R. (2025). A fine-tuned EfficientNet-B0 convolutional neural network for accurate and efficient classification of apple leaf diseases. *Scientific Reports*. **15:** 25732. https://doi.org/10.1038/s41598-025-04479-2.

Ali, M.U., Khalid, M., Farrash, M., Lahza, H.F.M., Zafar, A. and Kim, S.H. (2024). AppleLeafNet: A lightweight and efficient deep learning framework for diagnosing apple leaf diseases. *Frontiers in Plant Science*. **15:** 1502314. https://doi.org/10.3389/fpls.2024.1502314.

Apleni, T., Isinkaye, F.O. and Olusanya, M.O. (2025). Ensemble-based feature fusion for accurate plant disease classification using pre-trained models. *Scientific Reports*. **15:** 41925. https://doi.org/10.1038/s41598-025-25927-z.

Attallah, O. (2023). Tomato leaf disease classification *via* compact convolutional neural networks with transfer learning and feature selection. *Horticulturae*. **9(2):** 149. https://doi.org/10.3390/horticulturae9020149.

Barbedo, J.G.A. (2018). Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification. *Computers and Electronics in Agriculture*. **153:** 46-53. https://doi.org/10.1016/j.compag.2018.08.013.

Cho, O.H. (2024). An evaluation of various machine learning approaches for detecting leaf diseases in agriculture. *Legume Research*. **47(4):** 619-627. doi: 10.18805/LRF-787.

Fang, T., Zhang, J., Qi, D. and Gao, M. (2024). BLSENet: A novel lightweight bilinear convolutional neural network based on attention mechanism and feature fusion strategy for apple leaf disease classification. *Journal of Food Quality*. **2024:** 1-11. https://doi.org/10.1155/2024/5561625.

Fu, L., Li, S., Sun, Y., Mu, Y., Hu, T. and Gong, H. (2022). Lightweight convolutional neural network for apple leaf disease identification. *Frontiers in Plant Science*. **13:** 831219. https://doi.org/10.3389/fpls.2022.831219.

Gupta, C., Gill, N.S., Gulia, P., Duhan, S., Karamti, H., Kumar, A., Alamneh, D.A. and Safra, I. (2025). An enhanced deep learning-based framework for diagnosing apple leaf diseases. *Scientific Reports*. **15:** 39699. https://doi.org/10.1038/s41598-025-23272-9.

He, K., Zhang, X., Ren, S. and Sun, J. (2016). Deep Residual Learning for Image Recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Las Vegas, NV, USA, 2016, pp.770-778. https://doi.org/10.1109/cvpr.2016.90.

Iftikhar, M., Kandhro, I.A., Kausar, N., Kehar, A., Uddin, M. and Dandoush, A. (2024). Plant disease management: A fine-tuned enhanced CNN approach with mobile app integration for early detection and classification. *Artificial Intelligence Review*. **57:** 167. https://doi.org/10.1007/s10462-024-10809-z.

Kumar, M., Barbhai, M.D., Esatbeyoglu, T., Zhang, B., Sheri, V., Dhumal, S., Rais, N., Radha, Al Masry, E.M.S., *et al.* (2022). Apple (*Malus domestica* Borkh.) seed: A review on health promoting bioactivities and its application as functional food ingredient. *Food Bioscience*. **50(B):** 102155. https://doi.org/10.1016/j.fbio.2022.102155.

Lee, W.B. and Kim, B.H. (2024). Machine learning innovations for precise plant disease detection: A review. *Legume Research*. **47(10):** 1633-1638. doi: 10.18805/LRF-799.

Liu, S., Bai, H., Li, F., Wang, D., Zheng, Y., Jiang, Q. and Sun, F. (2023). An apple leaf disease identification model for safeguarding apple food safety. *Food Science and Technology*. **43:** e104322. https://doi.org/10.1590/fst.104322.

Mehta, A.R., Kumar, P., Prem, G., Aggarwal, S. and Kumar, R. (2025). AI-powered innovations in agriculture: A systematic review on plant disease detection and classification. *Indian Journal of Agricultural Research*. **59(9):** 1321-1330. doi: 10.18805/IJARe.A-6371.

Mustafa, H., Umer, M., Hafeez, U., Hameed, A., Sohaib, A., Ullah, S. and Madni, H.A. (2022). Pepper bell leaf disease detection and classification using optimized convolutional neural network. *Multimedia Tools and Applications*. **82:** 12065-12080. https://doi.org/10.1007/s11042-022-13737-8.

Ripin, Khandelwal, K., Gupta, M., Mishra, S., Ekka, S.K., Kujur, R. and Lakra, J. (2024). Development prospects of apple farming in India. *Asian Journal of Advances in Agricultural Research*. **24(1):** 11-18. https://doi.org/10.9734/ajaar/2024/v24i1483.

Sai Reddy, B. and Neeraja, S. (2022). Plant leaf disease classification and damage detection system using deep learning models. *Multimedia Tools and Applications*. **81:** 24021-24040. https://doi.org/10.1007/s11042-022-12147-0.

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. and Chen, L.C. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 2018, pp. 4510-4520. https://doi.org/10.1109/CVPR.2018.00474.

Singh, S., Gupta, I., Gupta, S., Koundal, D., Aljahdali, S., Mahajan, S. and Pandit, A.K. (2021). Deep learning based automated detection of diseases from apple leaf images. *Computers, Materials and Continua*. **71(1):** 1849-1866. https://doi.org/10.32604/cmc.2022.021875.

Sood, V., Shukla, S.S.P., Yadav, A.K., Tiwari, S. and Srivastava, Y. (2025). Towards automated crop protection: Fusion of DenseNet121-ResNet50 model for disease detection and pest recognition. *International Journal of System Assurance Engineering and Management*. **17:** 256-272. https://doi.org/10.1007/s13198-025-02975-2.

Srinidhi, V.V., Sahay, A. and Deeba, K. (2021). Plant Pathology Disease Detection in Apple Leaves Using Deep Convolutional Neural Networks: Apple Leaves Disease Detection Using EfficientNet and DenseNet. 2021 5th International Conference on Computing Methodologies and Communication (ICCMC), Erode, India, 2021, pp. 1119-1127. https://doi.org/10.1109/ICCMC51019.2021.9418268.

Tan, M. and Le, Q.V. (2021). EfficientNetV2: Smaller Models and Faster Training. Proceedings of the 38th International Conference on Machine Learning (ICML). 10096-10106. https://doi.org/10.48550/arXiv.2104.00298.

Thakur, P.S., Sheorey, T. and Ojha, A. (2022). VGG-ICNN: A lightweight CNN model for crop disease identification. *Multimedia Tools and Applications*. **82:** 497-520. https://doi.org/10.1007/s11042-022-13144-z.

Wang, G., Sang, W., Xu, F., Gao, Y., Han, Y. and Liu, Q. (2025). An enhanced lightweight model for apple leaf disease detection in complex orchard environments. *Frontiers in Plant Science*. **16:** 1545875. https://doi.org/10.3389/fpls.2025.1545875.

Yan, H. and Wang, X. (2024). Cross-dataset discriminant subspace learning algorithm for apple leaf diseases identification. *Crop Protection*. **181:** 106690. https://doi.org/10.1016/j.cropro.2024.106690.

Yang, Q., Duan, S. and Wang, L. (2022). Efficient identification of apple leaf diseases in the wild using convolutional neural networks. *Agronomy*. **12(11):** 2784. https://doi.org/10.3390/agronomy12112784.

Zhang, X., Li, H., Sun, S., Zhang, W., Shi, F., Zhang, R. and Liu, Q. (2023). Classification and identification of apple leaf diseases and insect pests based on improved ResNet-50 model. *Horticulturae*. **9(9):** 1046. https://doi.org/10.3390/horticulturae9091046.

Zhong, Y. and Zhao, M. (2020). Research on deep learning in apple leaf disease recognition. *Computers and Electronics in Agriculture*. **168:** 105146. https://doi.org/10.1016/j.compag.2019.105146.