Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X
Volume 3, Issue 10, 2025

APPLE LEAF DISEASE CLASSIFICATION USING DEEP CNNS

Muhammad Zia[1], Muhammad Uzair Khan[2], Rohail Ali Khan[3], Dawar Awan[4]
Saadia Tabassum*[5], Muhammad Lais[6], Muhammad Sohail Khan[7]

[1,2,3,4,6,7] *Department of Electrical Engineering Technology, Shuhada-e-APS University of Technology, Nowshera*
*[5] *Department of Electronics Engineering Technology, Shuhada-e-APS University of Technology, Nowshera*

*[5] saadia@uotnowshera.edu.pk

DOI: https://doi.org/10.5281/zenodo.17383357

Keywords
Apple diseases, Deep CNN, Plant Diseases, Black Rot, Apple Scab, Cedar Apple Rust

Article History
Received: 26 August 2025
Accepted: 04 October 2025
Published: 18 October 2025

Copyright @Author
Corresponding Author: *
Saadia Tabassum

Abstract

The presence of foliar diseases in apple crops poses a significant threat to both yield and fruit quality, necessitating the development of reliable automated detection mechanisms. This study presents a deep learning–based framework employing a Convolutional Neural Network (CNN) for the accurate classification of apple leaf conditions into four distinct categories: Apple Scab, Black Rot, Cedar Apple Rust, and Healthy. A total of 3,171 annotated images were sourced from the PlantVillage dataset and divided into training (70%), validation (20%), and testing (10%) subsets. The proposed CNN model incorporates four convolutional modules, each consisting of batch normalization and max-pooling layers, followed by a fully connected dense layer with a dropout rate of 0.5 to mitigate overfitting. Model training was conducted for 40 epochs using the Adam optimization algorithm. Experimental evaluation demonstrated that the network attained an overall classification accuracy of 96.85% on the test dataset, with F1-scores ranging from 93.8% to 97.9% across the individual disease classes. The findings confirm that the proposed CNN architecture provides a robust and efficient approach for early and automated detection of apple leaf diseases, thereby supporting precision agriculture and timely crop management interventions.

INTRODUCTION

A Apples (Malus domestica) represent one of the most widely cultivated and economically significant fruit crops globally, contributing approximately 15% to total world fruit production and serving as a key driver of agricultural income. As reported by the Food and Agriculture Organization (FAO) in 2023, global apple output surpassed 97 million tonnes, with China, the United States, and Poland ranking among the leading producers (Figure 1) [1]. The international apple trade exceeded USD 8.3 billion in 2023, underscoring its substantial economic relevance across both fresh and processed product markets, including juice, cider, and vinegar [2]. Beyond economic significance, apple farming sustains millions of small- and large-scale growers, supporting food security, employment generation, and rural development in numerous regions [3]. However, the profitability of apple production is highly dependent on efficient disease surveillance and management, as foliar and fruit diseases can drastically impact both yield and quality.

Among the various biotic stresses affecting apple orchards, fungal and bacterial pathogens pose the most critical threats to sustainable production. Apple Scab (*Venturia inaequalis*) and Black Rot (*Botryosphaeria obtusa*) are recognized as the most destructive diseases due to their rapid spread and high economic losses [4], [5]. In regions with humid climatic conditions, uncontrolled apple scab infections can result in yield reductions exceeding 70% [6]. Likewise, black rot leads to fruit decay and branch cankers, compromising both the current harvest and long-term orchard health [7]. These infections also degrade fruit appearance, reduce storage life, and limit market acceptability, ultimately affecting export competitiveness [8]. Historical outbreaks across Europe and North America prompted transitions from reactive fungicide spraying toward predictive epidemiological modeling and the cultivation of disease-resistant varieties [9]. Nevertheless, the emergence of new pathogen races and the development of fungicide resistance continue to challenge conventional control strategies, necessitating the integration of advanced detection and diagnostic tools [10].

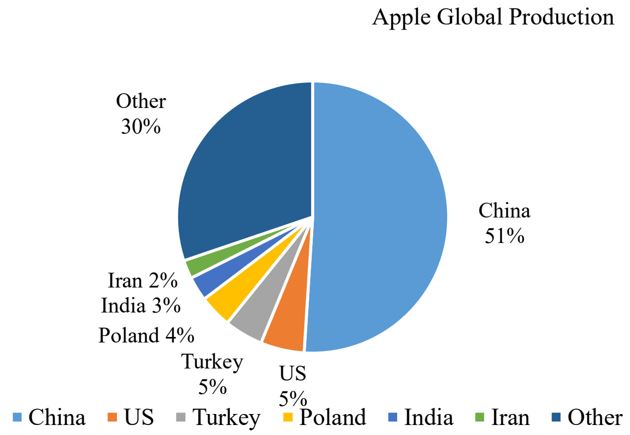

Figure 1. Top Countries by Apple Production

Modern apple production systems face additional challenges stemming from climate change, intensive monoculture, and transboundary pathogen dissemination, which collectively increase the frequency and unpredictability of disease outbreaks [11]. Traditional methods such as field scouting and laboratory-based diagnostics, although reliable, are resource-intensive and impractical for continuous monitoring across large orchards [12]. Consequently, precision agriculture and artificial intelligence (AI)-driven disease surveillance have gained prominence as efficient alternatives. Recent research demonstrates that machine learning (ML) and deep learning (DL) techniques—particularly Convolutional Neural Networks (CNNs)—outperform traditional image processing methods by autonomously extracting hierarchical visual features from diseased leaf images [13], [14]. These AI-based systems facilitate rapid, automated, and field-deployable disease recognition using images captured through mobile devices, unmanned aerial vehicles (UAVs), or Internet of Things (IoT) sensors [15]. The integration of CNN-based models into digital agriculture platforms holds significant potential for early disease detection, real-time crop monitoring, and intelligent decision support. As agriculture continues to embrace digital transformation, CNN architectures are anticipated to serve as the backbone of next-generation smart orchard management systems, enabling data-driven, predictive, and sustainable disease control frameworks [16], [17].

Literature review

The classification of apple leaf diseases using CNNs has seen rapid development in recent years, yet achieving robust field-level performance remains challenging. Zhang et al. proposed a ResNet-34 based CNN model and reported an accuracy of 93.765 % on a 3,642-image dataset, highlighting the inherent difficulties in generalizing from lab-collected images to varied real-world scenarios [18]. The authors in [19] developed MCFFA-Net, a stacked ensemble combining MobileNetV2, DenseNet201 and InceptionResNetV2 with multi-scale dilated residual modules and channel attention; although its accuracy was lower (90.86 %), the work emphasizes robustness in complex backgrounds and model fusion benefits. Object detection approaches have also been investigated to locate diseased regions rather than just classify whole-leaf images. Liu et al. proposed YOLOX-ASSANano, a lightweight real-time detector that achieves mAP = 91.08 % on a multi-scene apple leaf disease dataset, prioritizing detection speed and deployability in orchards over absolute classification accuracy [20]. Similarly, improved Faster R-CNN architectures (e.g. using Res2Net + FPN + RoIAlign + soft-NMS) have been proposed, but report only modest AP figures (e.g. 63.1 % AP) when applied in complex real-field settings, underscoring the difficulty of lesion localization in natural conditions [21].

Another study addresses the challenge of early detection of apple leaf diseases, which significantly affect crop yield and industry revenue [22]. The authors developed an expert-labeled dataset of about 9,000 high-quality apple leaf images representing major foliar diseases. A two-stage deep learning framework was proposed: a lightweight model for classifying leaves as healthy, diseased, or damaged, followed by a detection stage for localizing disease symptoms. The system achieved 88% classification accuracy and a mean average precision (mAP) of 42%, demonstrating strong performance even for small lesions. The approach shows promise as a practical diagnostic tool for farmers and can be extended to other crops in the future. The authors in [23] explores a deep learning–based automated framework for classifying apple leaf diseases using transfer learning with models like VGG16, ResNetV2, InceptionV3, and MobileNetV2. Experimental results show that the ResNetV2 model combined with the Adam optimizer achieved the highest classification accuracy of 94%. Yet another study compares several CNN architectures—ResNet50, AlexNet, Inception, VGG16, and InceptionResNetV2—for classifying apple leaf conditions, including apple scab, cedar rust, black rot, and healthy leaves [24]. Using a dataset of 8,800 images, the results indicate that the VGG16 model achieved the best multiclass classification accuracy of 92.5%.

A recent work proposed an AI-enabled Apple Leaf Disease Classification (AIE-ALDC) framework designed for precision agriculture [25]. The approach integrates orientation-based data augmentation and Gaussian filtering to enhance image quality and reduce noise. It employs a Capsule Network (CapsNet) for feature extraction, optimized through a Water Wave Optimization (WWO) algorithm, and utilizes a Bidirectional Long Short-Term Memory (BiLSTM) network for final classification. Experimental evaluations demonstrate that the AIE-ALDC model outperforms existing state-of-the-art methods, confirming its effectiveness and novelty in accurately identifying apple leaf diseases. Extending these advancements, another study utilized the DenseNet-121 architecture to identify multiple apple leaf diseases using three distinct strategies—regression, multi-label classification, and a focus loss function [26]. Using a dataset of 2,462 images covering six disease types, the proposed methods achieved accuracies of 93.51%, 93.31%, and 93.71%, respectively. These results surpassed the conventional multi-classification approach based on the cross-entropy loss function, which attained 92.29%, demonstrating the improved effectiveness of tailored loss functions and model adaptations for disease identification. The authors in [27] compared multiple optimization techniques, including a customized PowerSign algorithm, to enhance model performance. Among them, the Adagrad optimizer delivered the most stable results, achieving 92% training, 91% validation, and 91% testing accuracy. The optimized model was further integrated into a mobile application and evaluated in real field conditions, where it demonstrated high accuracy, provided the camera was positioned close to the leaf with clear focus.

In summary, current literature indicates that while CNN architectures can reach high accuracy under controlled settings, the real challenge lies in robust field performance and lightweight real-time deployment. Most methods reporting very high accuracy do so on benchmark or cleaned datasets; in contrast, detection-level AP, mobile-scale deployability, and field generalization often yield lower figures. Our work builds on this insight by focusing on apple leaf disease classification using CNNs specifically tailored for robustness under field variability, leveraging detection, segmentation, or augmentation techniques to close the gap between lab and orchard performance.

Methodology

3.1 Data Acquisition and Pre-processing

In this research, the dataset was sourced from the open-access PlantVillage repository, which provides a standardized collection of plant disease images. Specifically, the subset utilized comprises 3,171 labeled images of apple leaves, distributed across four categories: Healthy, Apple Scab, Black Rot, and Cedar Apple Rust. Each image corresponds to an individual apple leaf sample, depicting either a normal leaf or one showing distinct disease symptoms, as illustrated in 2.

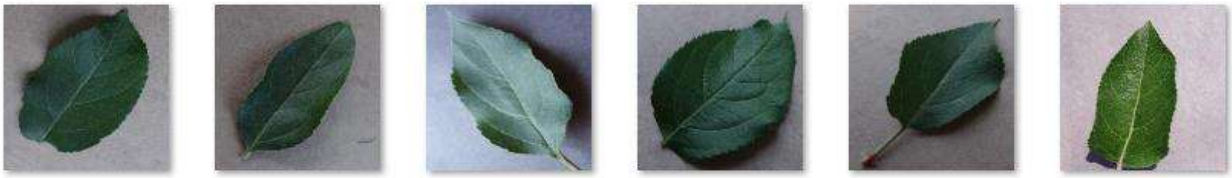

a) Healthy Apple Leaf

b) Cedar Apple Rust Apple leaf

c) Black Rot Apple Leaf

d) Scab Apple Leaf

Figure 2. Apple Leaves Dataset

For effective and unbiased model training, the dataset was partitioned into three subsets: training, validation, and testing. The final allocation consisted of 2,218 images for training, 636 images for validation, and 317 images for testing. Within these splits, the Healthy category contained 1,151, 330, and 164 samples, respectively; Cedar Apple Rust comprised 192, 55, and 28 samples; Black Rot included 434, 125, and 62 samples; while Apple Scab contained 441, 126, and 63 samples.

All images were preprocessed by resizing them to 128 × 128 pixels and normalizing pixel intensity values between 0 and 1 to ensure uniform scaling. This normalization enhanced numerical stability during training and accelerated convergence of the CNN model. The preprocessing step also standardized image dimensions and illumination, allowing the network to more effectively learn consistent and discriminative visual features for accurate classification of apple leaf diseases. The complete dataset preparation workflow is illustrated in Figure *3*.

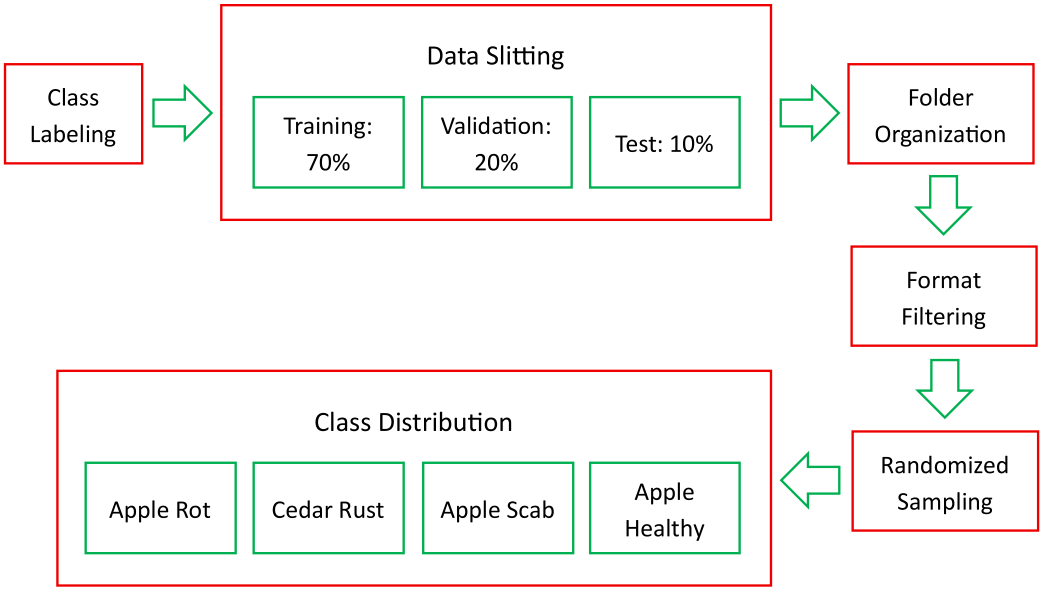

Figure 3. Dataset Division

3.2 Experimental Setup

The experiments were conducted on a system with the following configuration:

- Processor: Intel Core i7
- RAM: 16 GB
- GPU: NVIDIA CUDA-enabled
- Framework: Adam optimizer

3.3 CNN Model Architecture

A deep CNN architecture was specifically designed and implemented using the TensorFlow and Keras frameworks to perform multi-class classification of apple leaf images into four distinct categories: Healthy, Apple Scab, Black Rot, and Cedar Apple Rust, as illustrated in Figure 4. Prior to model training, all images were standardized to an input resolution of 128 × 128 pixels in the RGB color space to ensure consistent feature extraction and computational efficiency.

The proposed network architecture comprises four sequential convolutional modules, each integrating multiple Conv2D layers followed by max-pooling operations. These convolutional blocks progressively extract hierarchical features, beginning with fundamental visual cues such as edges and color gradients and advancing toward more abstract, disease-specific texture patterns. Following the convolutional stages, the output feature maps are flattened and connected to a fully connected dense layer consisting of 256 neurons activated by the Rectified Linear Unit (ReLU) function to introduce non-linearity and enhance model learning capacity. To mitigate overfitting and improve model generalization, a dropout layer with a dropout rate of 0.5 was incorporated.

For the classification stage, a Softmax-activated output layer with four neurons was employed, corresponding to the four target disease classes. The network was trained for 40 epochs using a batch size of 32, and the optimization process was governed by the Adam optimizer with a learning rate of 0.0001. The categorical cross-entropy loss function was utilized to effectively manage the multi-class prediction task. The total parameter count for the model was 5,371,684, out of which 5,369,764 were trainable, reflecting the network's high representational capacity.

Upon evaluation, the proposed CNN achieved a peak validation accuracy of 97.8%, demonstrating its strong capability in distinguishing between healthy and diseased apple leaf images. The results validate the model's robustness, effective feature learning, and high generalization performance in complex agricultural image classification scenarios.

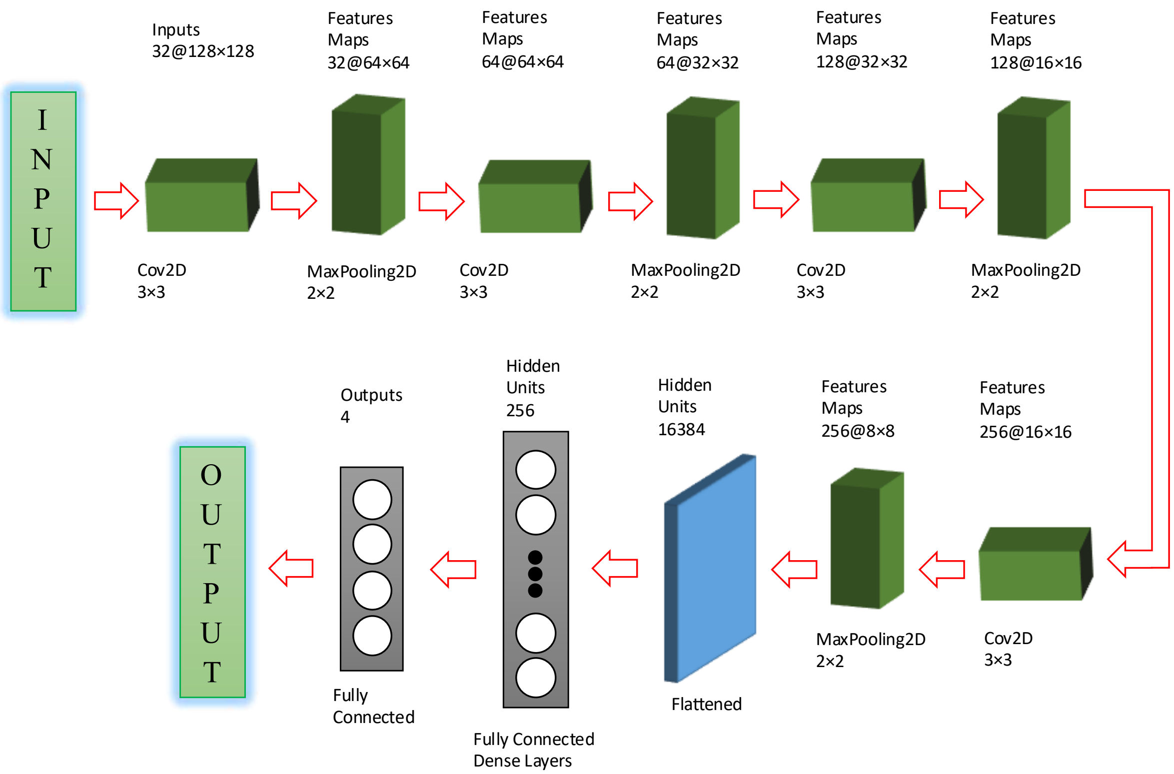

Figure 4. Deep CNN structure for actual data

3.3.1 Training and Validation

Table 1 provides a comprehensive overview of the principal hyperparameters adopted during the training phase of the proposed deep CNN. To improve the model's ability to generalize to unseen data, real-time data augmentation techniques were applied using the ImageDataGenerator utility available in Keras. This augmentation process introduced controlled variations such as horizontal flipping, small rotations, and zoom transformations, thereby enriching the training dataset and reducing sensitivity to overfitting.

During training, both accuracy and loss metrics were continuously tracked across each epoch to evaluate the learning progression and model stability. The network exhibited a smooth and consistent convergence pattern, reflecting the effectiveness of the chosen optimizer and hyperparameter settings.

Upon completion of training, the model achieved a final training accuracy of 99.86% and a validation accuracy of 97.80%, accompanied by a validation loss of 13%. These results demonstrate that the model successfully captured discriminative features of the apple leaf dataset while maintaining a strong balance between fitting and generalization.

The relatively small validation loss, combined with the minimal gap between training and validation performance, indicates that the implemented regularization strategies and data augmentation techniques effectively mitigated overfitting. Consequently, the proposed CNN model exhibits high reliability and robustness in accurately classifying apple leaf diseases under varying image conditions.

Table 1. Hyper parameters of the proposed Deep CNN model.

Parameters	Value
Batch Size	32
Epoch	40
Learning rate	0.0001
Dropout Value	0.5
Training set size	2218
Validation set size	636
Test size	317

3.3.2 Testing and Evaluation

The testing phase was performed to rigorously assess the generalization capability and predictive reliability of the developed CNN on unseen data samples. In this stage, the trained model was reloaded and evaluated using an independent test dataset consisting of 317 apple leaf images, all of which were completely excluded from the training and validation processes to ensure unbiased performance assessment.

To obtain a comprehensive understanding of the model's effectiveness, several quantitative performance metrics were computed, including accuracy, precision, recall, and F1-score. Additionally, a confusion matrix was generated to provide insight into the class-wise distribution of correct and incorrect predictions, enabling the identification of potential misclassification trends among the four target categories—Healthy, Apple Scab, Black Rot, and Cedar Apple Rust.

The proposed CNN achieved an overall test accuracy of 96.85%, signifying a strong ability to correctly classify apple leaf conditions. Furthermore, the F1-scores for individual classes ranged between 93.8% and 97.9%, highlighting the model's balanced performance across both majority and minority classes. The high F1-scores confirm that the model maintained an optimal trade-off between precision and recall, effectively capturing disease-specific visual features even in challenging test samples.

Overall, these results demonstrate that the developed deep CNN exhibits robust generalization and high diagnostic reliability, making it well-suited for practical implementation in automated apple leaf disease detection systems.

Results and discussion

This section presents a comprehensive evaluation of the proposed Deep CNN for apple leaf disease classification. The model's performance is critically analyzed using established evaluation metrics—accuracy, precision, recall, and F1-score—to quantify classification effectiveness. In addition, a confusion matrix is employed to assess class-wise prediction consistency, while comparative analyses with traditional machine learning and transfer learning models provide deeper insight into the network's robustness, feature learning ability, and generalization performance.

4.1 Dataset Statistics

A well-balanced and diverse dataset is crucial for developing a robust deep learning model. The dataset utilized in this study was partitioned into training, validation, and testing subsets using a 70:20:10 ratio to ensure both adequate learning exposure and fair performance evaluation. Table 2 summarizes the image distribution across the four classes—Healthy, Apple Scab, Black Rot, and Cedar Apple Rust.

Table 2: Dataset distribution across training, validation, and testing subsets.

Class	Train	Validation	Test
Apple Healthy	1151	330	164
Apple cadar rust	192	55	28

Apple black rot	434	125	62
Apple scab	441	126	63
Test	2218	636	317

The dataset composition reveals that while the healthy leaf class has a relatively higher sample count, the disease classes are moderately imbalanced. To address this imbalance and enhance generalization, real-time data augmentation was employed during training, including rotation, flipping, and scaling transformations. Such augmentation increased the diversity of training samples, helping the model learn more invariant disease-specific patterns. The use of stratified sampling ensured proportional representation of each class across all subsets, reducing the likelihood of biased learning outcomes.

4.2 Training and Validation Performance

The performance of the proposed CNN was compared with several widely used transfer learning architectures, including AlexNet, ResNet, VGG16, and InceptionV3. Figure 5 illustrates the comparative validation accuracies obtained during the training phase.

Among all evaluated architectures, AlexNet achieved the lowest accuracy of 87.34%, primarily due to its relatively shallow structure and limited feature extraction depth. ResNet and VGG16 demonstrated improved performance, achieving 92.56% accuracy, benefiting from deeper feature hierarchies and residual learning mechanisms. InceptionV3, which utilizes multi-scale convolutional filters, reached 94.32%, showing superior capacity to capture complex texture and lesion patterns.

However, the proposed Deep CNN model outperformed all competing models, attaining a validation accuracy of 98.19%. This substantial improvement can be attributed to its carefully optimized architecture, which integrates multiple convolutional layers with dropout regularization and batch normalization. These design choices enabled efficient feature abstraction while preventing overfitting. The consistent decline in validation loss and stabilization of accuracy across epochs indicate convergent and stable learning behavior.

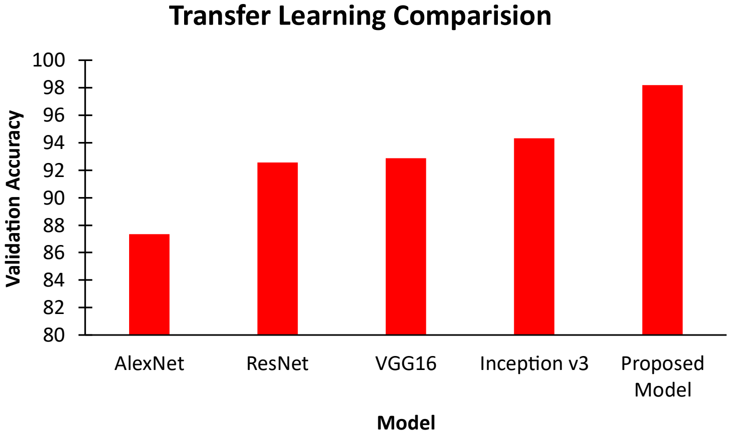

Figure 5. Validation accuracy of the popular transfer learning approaches

4.3 Testing Performance

The final evaluation was performed on an independent test dataset containing 317 images unseen during model training and validation. To establish a performance baseline, conventional classifiers—Logistic Regression, Support Vector Machine (SVM), Decision Tree, and K-Nearest Neighbors (KNN)—were trained and evaluated under identical conditions. The comparative results, presented in Figure 6, show that these classical algorithms achieved moderate accuracies of 80.99%, 50.69%, 72.23%, and 87.86%, respectively.

In contrast, the proposed Deep CNN model achieved a remarkable testing accuracy of 96.85%, highlighting the superiority of automated hierarchical feature extraction over hand-crafted feature-based approaches. The CNN's ability to capture spatial and texture-based variations in diseased leaf regions contributes significantly to this improvement. Furthermore, the inclusion of data augmentation and dropout layers enhanced the model's generalization ability, minimizing overfitting.

To further evaluate model robustness, precision and F1-score analyses were conducted (Figure 7 and Figure 8). The proposed CNN achieved notably higher precision across all classes, confirming its reliability in correctly identifying diseased samples without generating false positives. Similarly, the F1-score, which combines both precision and recall, indicated balanced performance across categories, ensuring consistent detection of both common and rare diseases.

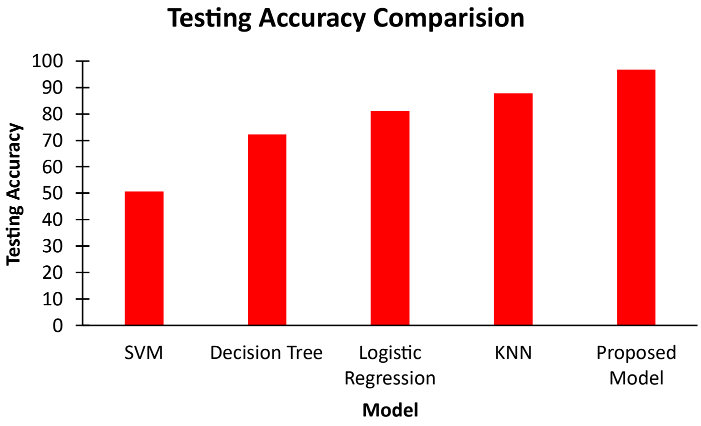

Figure 6. Validation accuracy of different models

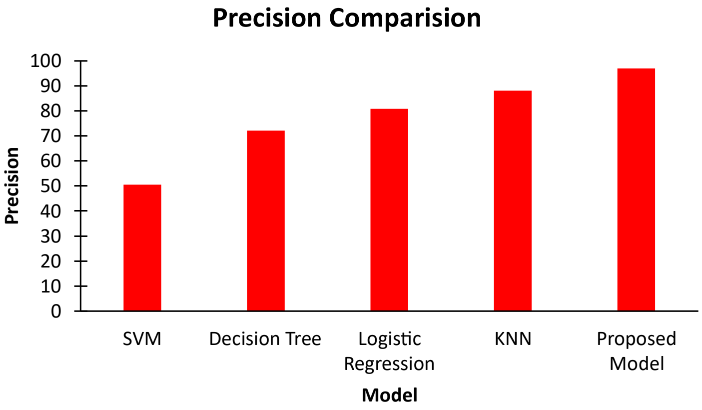

Figure 7. Precision values of the different models

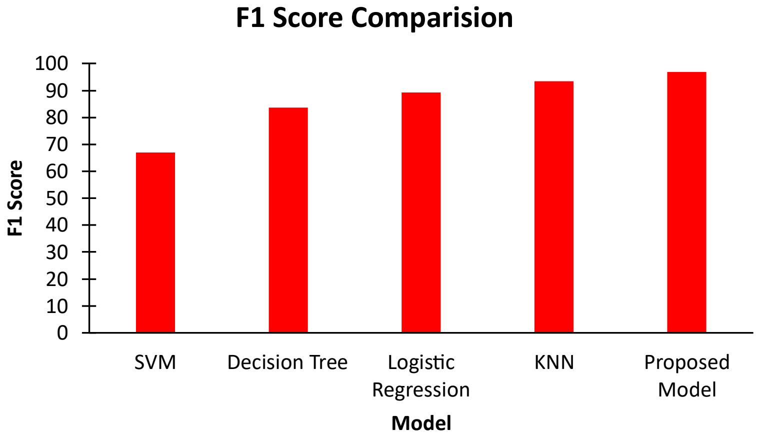

Figure 8. F1 Scores of different models

The test set evaluation yielded an accuracy of 96.85%, precision of 96.97%, recall of 96.85%, and F1-score of 96.86% (Figure 9). These consistent scores reflect the model's strong generalization capacity and effective learning of disease-specific visual cues. The close correspondence between validation and test results further supports the CNN's stability and predictive reliability on unseen data.

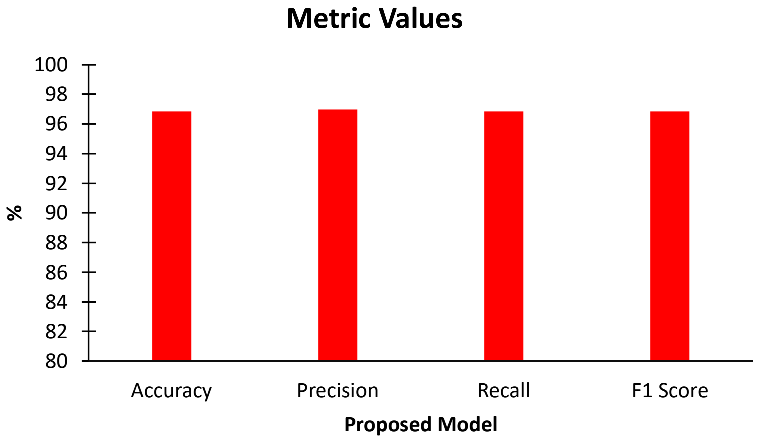

Figure 9. Model Performance

4.4 Confusion Matrix Analysis

A detailed examination of the confusion matrix (Table 3) provides insights into class-wise prediction behavior. The matrix demonstrates that the model correctly classified the majority of samples across all categories, with only a few misclassifications observed between visually similar disease types.

Table 3. Confusion Matrix of the Proposed Deep CNN Model on Test Set

True \ Predicted	Apple_scab	Black_rot	Cedar_apple_rust	Healthy
Apple_scab	61	0	0	2
Black_rot	1	59	0	1
Cedar_apple_rust	1	0	26	1

Healthy	2	0	0	161

Most misclassifications occurred between Apple Scab and Healthy leaves, which share similar texture and coloration in early disease stages, making visual discrimination challenging even for human experts. Despite this, the error rate remained minimal, confirming the CNN's capacity to differentiate subtle disease features such as lesion shape, color intensity, and vein pattern irregularities. This analysis reinforces that the proposed model maintains high precision and reliability even in visually ambiguous scenarios.

4.5 Per-Class F1-Score Ranking

To further assess class-level performance, the F1-score for each category was computed, as summarized in Table 4.

Table 4. F1 Score of the Proposed Deep CNN Model.

Class	F1-Score
Apple_healthy	0.9787
Black_rot	0.9752
Cedar_apple_rust	0.9630
Apple_scab	0.9385

The results reveal that the model achieved the highest F1-score (97.87%) for healthy leaves, reflecting its superior ability to identify normal leaf conditions with minimal false positives. The slightly lower F1-score for Apple Scab (93.85%) indicates that this disease's early-stage symptoms exhibit close visual similarity to healthy leaves and other diseases, occasionally leading to minor misclassification.

Nevertheless, all F1-scores exceeding 93% confirm that the CNN maintained a balanced trade-off between precision and recall across all categories. This uniform performance indicates that the model learned disease-specific spatial and texture representations effectively, ensuring consistent diagnostic reliability. The overall results validate that the proposed CNN is well-suited for automated apple disease detection and can potentially be deployed for real-time agricultural health monitoring applications.

4.6 Discussion

The comparative analysis demonstrates that the proposed Deep CNN model not only surpasses conventional classifiers but also outperforms established transfer learning frameworks. This performance gain can be attributed to its customized architectural design, careful hyperparameter tuning, and effective regularization strategies. The model's ability to generalize across varying disease appearances reflects its adaptability to field conditions where illumination, background, and leaf orientation vary significantly.

Moreover, the consistently high evaluation metrics across multiple datasets suggest that the model successfully captures disease-specific morphological traits, making it an efficient diagnostic tool. Future improvements could include the integration of attention mechanisms or multi-scale feature extraction layers to further enhance sensitivity toward fine-grained lesion patterns.

Overall, the findings substantiate the proposed Deep CNN's capability as a robust, scalable, and reliable solution for automated apple leaf disease detection, contributing to precision agriculture and early crop disease management.

Conclusion

In this study, a Deep CNN was developed and rigorously evaluated for the automated detection and classification of apple leaf diseases, including Apple Scab, Black Rot, Cedar Apple Rust, and Healthy leaves. The entire experimental pipeline—from dataset preprocessing and augmentation to model training and evaluation—was implemented using Python with TensorFlow and Keras, ensuring reproducibility and computational efficiency. The proposed CNN, incorporating multiple convolutional, pooling, and fully connected layers along with Batch Normalization and Dropout regularization, effectively captured disease-specific visual features and demonstrated strong learning stability. The model achieved a remarkable overall accuracy of 96.85% and an F1-score of 96.86% on unseen test data, indicating excellent generalization performance. Confusion matrix analysis and per-class F1-scores confirmed consistent classification across all categories, with only minor misclassifications in visually similar disease types. While the results validate the robustness and reliability of the proposed model, future improvements could involve advanced data augmentation and class balancing to further enhance generalization, as well as integrating transfer learning architectures such as ResNet50, InceptionV4, or EfficientNet to optimize feature extraction and computational efficiency. Additionally, deploying the model on edge or IoT-based devices could enable real-time, field-level disease detection without reliance on continuous internet connectivity. Extending the framework to multiple crop species and a wider range of plant diseases would further enhance scalability, making it a promising tool for intelligent and sustainable agricultural disease management.

REFERENCES

1. Faostat, F. and A.C. Production, *Food and agriculture organization of the united nations, 2010.* Roma, Italy, 2016.
2. Vučkovski, B.G., V. Babić, and V. Paraušić, *Fruit export through the prism of the high-value market and GLOBALG. AP certification: an analysis from Serbia.* BizInfo Blace, 2023. **14**(2): p. 37-46.
3. Narayanasamy, P., *Phyllosphere Microbial Plant Pathogens: Detection and Crop Disease Management: Volume 1 Nature and Biology*. 2024: CRC Press.
4. Gladieux, P., et al., *On the origin and spread of the scab disease of apple: out of central Asia.* PloS one, 2008. **3**(1): p. e1455.
5. Bowen, J.K., et al., *Venturia inaequalis: the causal agent of apple scab.* Molecular plant pathology, 2011. **12**(2): p. 105-122.
6. Shafi, S.M., et al., *An overview of apple scab, its cause and management strategies.* EC Microbiol, 2019. **15**: p. 1-4.
7. Švara, A., et al., *Phenotyping, genetics, and "-omics" approaches to unravel and introgress enhanced resistance against apple scab (Venturia inaequalis) in apple cultivars (Malus× domestica).* Horticulture Research, 2024. **11**(2): p. uhae002.
8. Michalecka, M., et al., *Population structure of Venturia inaequalis, a causal agent of apple scab, in response to heterogeneous apple tree cultivation.* BMC evolutionary biology, 2018. **18**(1): p. 5.
9. Zhou, C., et al., *A Multimodal Parallel Transformer Framework for Apple Disease Detection and Severity Classification with Lightweight Optimization.* Agronomy, 2025. **15**(5): p. 1246.
10. Pecknold, P.C., *Apple scab of flowering crabapples.* 2016.
11. Khan, A.I., S. Quadri, and S. Banday, *Deep learning for apple diseases: classification and identification.* International Journal of Computational Intelligence Studies, 2021. **10**(1): p. 1-12.
12. Wang, G., et al., *An enhanced lightweight model for apple leaf disease detection in complex orchard environments.* Frontiers in Plant Science, 2025. **16**: p. 1545875.
13. Yang, Q., S. Duan, and L. Wang, *Efficient identification of apple leaf diseases in the wild using convolutional neural networks.* Agronomy, 2022. **12**(11): p. 2784.
14. Cheng, H. and H. Li, *Identification of apple leaf disease via novel attention mechanism based convolutional neural network.* Frontiers in Plant Science, 2023. **14**: p. 1274231.
15. Ali, M.U., et al., *AppleLeafNet: a lightweight and efficient deep learning framework for diagnosing apple leaf diseases.* Frontiers in Plant Science, 2024. **15**: p. 1502314.
16. Muhammad, Z., et al., *Automated Classification of Healthy and Infected Grape Leaves Using a Deep Convolutional Neural Network.* Physical Education, Health and Social Sciences, 2025. **3**(3): p. 116-128.
17. Bi, Z., et al., *Apple leaf disease severity grading based on deep learning and the DRL-Watershed algorithm.* Scientific Reports, 2025. **15**(1): p. 30071.

18. Zhang, D., H. Yang, and J. Cao, *Identify apple leaf diseases using deep learning algorithm.* arXiv preprint arXiv:2107.12598, 2021.
19. Ahmed, M.R., et al. *MCFFA-Net: Multi-contextual feature fusion and attention guided network for apple foliar disease classification*. in *2022 25th International Conference on Computer and Information Technology (ICCIT)*. 2022. IEEE.
20. Liu, S., et al., *An improved lightweight network for real-time detection of apple leaf diseases in natural scenes.* Agronomy, 2022. **12**(10): p. 2363.
21. Gong, X. and S. Zhang, *A high-precision detection method of apple leaf diseases using improved faster R-CNN.* Agriculture, 2023. **13**(2): p. 240.
22. Khan, A.I., et al., *Deep diagnosis: A real-time apple leaf disease detection system based on deep learning.* computers and Electronics in Agriculture, 2022. **198**: p. 107093.
23. Alsayed, A., A. Alsabei, and M. Arif, *Classification of apple tree leaves diseases using deep learning methods.* International Journal of Computer Science & Network Security, 2021. **21**(7): p. 324-330.
24. Savla, D., et al. *Apple leaf disease detection and classification using CNN models*. in *International Conference on Computing in Engineering & Technology*. 2022. Springer.
25. Al-Wesabi, F.N., et al., *Artificial Intelligence Enabled Apple Leaf Disease Classification for Precision Agriculture.* Computers, Materials & Continua, 2022. **70**(3).
26. Zhong, Y. and M. Zhao, *Research on deep learning in apple leaf disease recognition.* Computers and electronics in agriculture, 2020. **168**: p. 105146.
27. Özden, C., *Apple leaf disease detection and classification based on transfer learning.* Turkish Journal of Agriculture and Forestry, 2021. **45**(6): p. 775-783.