Afghan Agricultural Sciences Journal
2026, vol. 1(1):1-14
DOI:10 xxxxxxx/xxxx

Comparative analysis of optimization algorithms for maxvit-based tomato leaf disease detection in precision agriculture

Rafiullah Rafi [1]*, Sarajudin Saeid [2], Alireza Soleimanipour [3], Abbas Rezaei Asl [3], Sediqullah Hammas [4]

[1]Department of Agricultural Engineering, Faculty of Plants Sciences, Afghanistan National Agricultural Sciences and Technology University, Kandahar, Afghanistan.

[2]Department of Plant Protection, Faculty of Plants Sciences, Afghanistan National Agricultural Sciences and Technology University, Kandahar, Afghanistan.

[3]Department of Biosystem Engineering, Gorgan University of Agricultural Sciences and Natural Resources, Gorgan, Iran.

[4]Department of Agricultural Economics and Extension, Faculty of Agriculture, Logar University, Logar, Afghanistan.

A B S T R A C T

The detection methods used to identify crop diseases require both accuracy and speed because they play a vital role in achieving better crop yields and sustainable agricultural practices. This study conducts an extensive assessment of six optimization algorithms, which include AdamW, Adam, Stochastic Gradient Descent (SGD), Adagrad, RMSprop, and SophiaG, to evaluate their performance within the Multi-Axis Vision Transformer framework for classifying tomato leaf diseases. In this study, a subset of the PlantVillage dataset was employed for model training and evaluation. The experiments were conducted in PyTorch on Google Colaboratory with NVIDIA Tesla T4 GPU acceleration. The MaxViT-Tiny model used pretrained ImageNet-1k weights to establish its initial state, which developers used to train the model through different optimizers with fixed hyperparameters. Model performance was evaluated using F1-score, precision, recall, and accuracy metrics calculated with scikit-learn, along with confusion matrix analysis and convergence curve visualization. AdamW achieved the best performance results among all tested optimizers because it reached an F1 Score and Recall combined with an overall accuracy of 99.89%. SophiaG and Adam achieved similar results, while Adagrad showed the least classification success with 98.32% accuracy, which demonstrated how standard adaptive optimizers fail to manage high-dimensional visual tasks in controlled environments. The results demonstrate that advanced optimization techniques enable technological progress in agricultural smart systems and deep learning capabilities used for plant disease research, which depends on transformer models.

Keywords: AdamW, multi-axis vision transformer, optimization algorithms, SophiaG, tomato leaf diseases

Article Type: Research
Received Date: 24 January 2026
Accepted Date: 05 July 2026
Corresponding Author Email: rafiullah3.ku@gmail.com
ISSN: xxxx-xxxx

1. INTRODUCTION

Global food security is increasingly threatened by the combined impacts of pollinator decline, changing climatic conditions, and the widespread occurrence of plant diseases (Amin et al., 2022). Plant diseases represent one of the major constraints to agricultural productivity, causing substantial yield losses, economic damage, and disruptions in food supply chains worldwide (Barman et al., 2024; Javidan et al., 2025; Krishna et al., 2025; Lee & Lee, 2025; W. Liu & Zhang, 2025; Upadhyay et al., 2025). Consequently, the development of rapid, accurate, and automated disease detection systems has become a critical component of sustainable agricultural production and precision agriculture.

Tomato (*Solanum lycopersicum* L.), a member of the Solanaceae family, is among the most economically important vegetable crops cultivated worldwide (Collins et al., 2022; Panno et al., 2021). The world recognizes tomatoes as one of the most popular crops because their diverse culinary applications and their high nutritional value make them valuable worldwide as a food source (Agarwal et al., 2020; Sakkarvarthi et al., 2022). However, tomato production is severely affected by numerous fungal, bacterial, viral, and pest-related diseases, resulting in considerable reductions in yield and quality (J. Liu & Wang, 2020; Shanthi et al., 2024). Conventional disease diagnosis methods rely heavily on visual inspection and laboratory-based analyses, which are often labor-intensive, time-consuming, and susceptible to human error (Chitta et al., 2024; Fraiwan et al., 2022; Ghosh et al., 2025; Wang et al., 2025; Zhu et al., 2024).

Recent advances in artificial intelligence (AI) and deep learning (DL) have transformed plant disease diagnosis by enabling automated, rapid, and highly accurate disease detection systems (Colucci et al., 2025; Das et al., 2024; El Sakka et al., 2024; Hamid et al., 2025; Ouamane et al., 2025; Qin et al., 2025; Wang et al., 2025). Early studies predominantly employed convolutional neural networks (CNNs) for plant disease classification. For example, Zhang et al. (2018) evaluated AlexNet, GoogLeNet, and ResNet architectures and reported that ResNet optimized using SGD achieved a classification accuracy of 97.28% for tomato leaf disease identification. Similarly, Liu and Wang (2020) developed a YOLOv3-based framework that improved disease detection accuracy and speed under real-world conditions through optimized feature extraction and multi-scale image analysis.

More recently, transformer-based architectures have demonstrated superior performance in agricultural image classification because of their ability to model both local and global feature relationships. Hossain et al. (2023) compared several transformer architectures, including MaxViT, EANet, CCT, and PVT, and reported that MaxViT achieved the highest classification accuracy when trained with the AdamW optimizer. Likewise, Nishankar et al. (2025) introduced the ViT-RoT benchmarking framework and demonstrated that transformer-based models such as ConvNeXt-Small and Swin-Small exhibited high classification accuracy and robustness using AdamW optimization. Hybrid architectures have also shown promising results. Tiwari et al. (2024) proposed a Vision Transformer–Deep Neural Network framework optimized with Adam and achieved a classification accuracy of 99.74%, while Ghosh et al. (2025) integrated Swin Transformer, ViT, and EfficientNetV2 models and reported approximately 98% classification accuracy under diverse environmental conditions using AdamW.

Despite the rapid advancement of deep learning architectures, model performance depends not only on network design but also on the optimization strategy employed during training. Optimization algorithms directly influence convergence speed, training stability, computational efficiency, and

model generalization. Adam Kingma & Ba (2015) has become one of the most widely used optimizers because of its adaptive learning-rate mechanism based on first- and second-order moment estimation. Several studies have reported that Adam frequently outperforms traditional optimization methods such as SGD and RMSprop in terms of convergence speed and training stability (Osmenaj et al., 2025; Thangaraj et al., 2021). AdamW extends Adam by decoupling weight decay from gradient updates, thereby improving regularization and generalization performance, particularly in transformer-based architectures (Ghosh et al., 2025; Hossain et al., 2023; Nishankar et al., 2025; Shen et al., 2025).

Recently, Liu et al. (2024) proposed Sophia, a scalable second-order optimization algorithm that utilizes a stochastic diagonal Hessian approximation and gradient clipping mechanism to accelerate convergence while maintaining computational efficiency. Sophia demonstrated promising results in large-scale deep learning applications by reducing training costs and convergence time. However, despite its potential advantages, its application in agricultural image classification and plant disease detection remains largely unexplored.

Although previous studies have reported promising results using CNN-based, transformer-based, and hybrid architectures, most investigations have primarily focused on model development while employing a single optimization algorithm. Existing studies rarely perform systematic comparisons of multiple optimization strategies under identical experimental conditions. Furthermore, the effectiveness of recently proposed optimizers such as SophiaG within transformer-based plant disease classification frameworks has not been comprehensively investigated. Therefore, a significant research gap remains regarding the comparative impact of different optimization algorithms on the performance and convergence behavior of transformer-based models for agricultural disease detection. To address this gap, the present study systematically evaluates six optimization algorithms, namely Adam, AdamW, SGD, RMSprop, Adagrad, and SophiaG, using the Multi-Axis Vision Transformer (MaxViT) architecture for tomato leaf disease classification. The study investigates their influence on classification accuracy, convergence characteristics, and training efficiency to identify the most effective optimization strategy for AI-driven disease diagnosis in precision agriculture.

2. MATERIALS AND METHODS

2.1. Dataset and Preprocessing

The PlantVillage Tomato Leaf Disease Dataset served as the dataset for this research. Multiple research studies have utilized this dataset to create automated systems that detect crop leaf diseases using deep learning methods (Dai et al., 2023; Mohameth et al., 2020; Mohanty et al., 2016; Too et al., 2019; Ullah et al., 2023). The researchers found that CNN-based and transfer learning models proved effective on this dataset. The dataset contains 22,930 labeled RGB images of tomato leaves, which researchers organized into 10 classes that include healthy leaves and nine different diseases. The dataset was split into two parts that included 18,345 images for training, which represented approximately 80% of the dataset, and 4,585 images for validation, which formed about 20% of the dataset. The original images came from controlled environments that used standard lighting and a consistent background. The images had a standard resolution of 256×256 pixels, which produced minimal background noise and consistent visual conditions (Mohanty et al., 2016). The preprocessing required all images to be resized to 224×224 pixels for use as input to the MaxViT model, as illustrated in Figure 1.

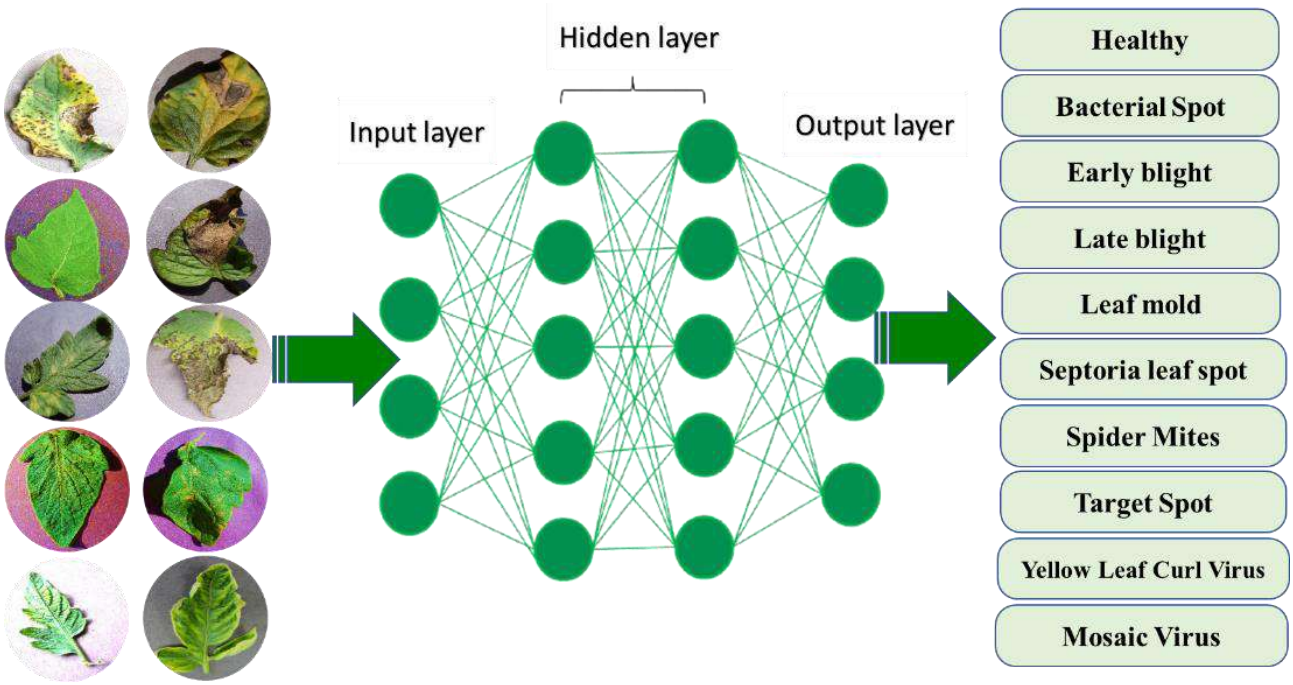


Figure 1. Tomato leaf images are processed through a deep neural network to classify conditions

During training, several preprocessing and data augmentation techniques were applied to improve model generalization and reduce overfitting. Specifically, random horizontal flipping with a probability of 0.5 and random image rotation within ±10° were employed. No augmentation was applied to the validation dataset to ensure unbiased performance evaluation. Subsequently, all images were converted into tensors and normalized using the standard ImageNet normalization statistics with channel-wise mean values of (0.485, 0.456, 0.406) and standard deviation values of (0.229, 0.224, 0.225). These normalization parameters are consistent with the pretrained MaxViT model and facilitate efficient transfer learning by preserving the feature distribution learned from the ImageNet dataset. The processed images were loaded using the PyTorch ImageFolder interface. During training, batches of 32 images were randomly shuffled at each epoch to improve stochastic optimization, whereas the validation dataset was evaluated without shuffling to ensure consistent performance assessment. The methodological workflow of the research is seen in Figure 2.

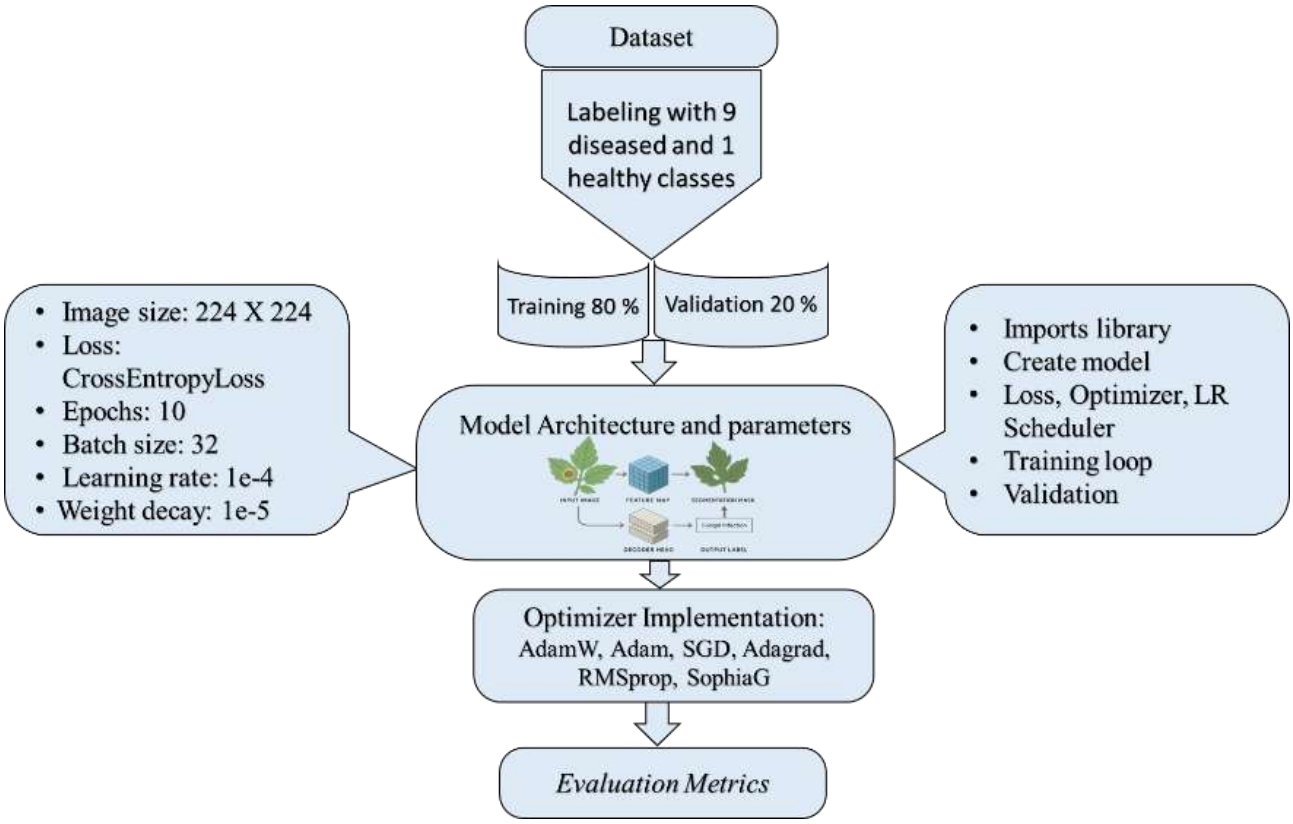


Figure 2. Methodological Framework for Tomato Leaf Disease Classification using MaxViT via various optimizers.

2.2. Model Architecture and Experimental Setup

The MaxViT-Tiny model functions as the base model architecture because it represents a top-level hierarchical vision transformer system which utilizes multiple attention systems. This model combines both local and global attention mechanisms, which enable it to identify plant leaf disease symptoms that appear in different sizes and patterns. The model was initialized with pretrained weights from the ImageNet-1k dataset, which were obtained through the timm library. The model's final classification head was changed when its adaptive average pooling layer was replaced by a fully connected linear classifier. The model was trained for 10 epochs using a batch size of 32 through the CrossEntropyLoss function, which utilized a learning rate of 1e-4 and a weight decay of 1e-5 for regularization purposes. The model was fully fine-tuned, as all 28.5 million parameters were trainable with no frozen layers. The original classification head of the pretrained MaxViT model was replaced with an Adaptive Average Pooling layer followed by a fully connected linear layer containing 10 output neurons corresponding to the ten tomato leaf disease classes. All experiments were implemented using the PyTorch framework on the Google Colaboratory platform, utilizing an NVIDIA Tesla T4 GPU for computational acceleration. The current setup mirrors recent research studies, which include Google Colaboratory and Tesla T4 GPU systems to train deep learning models for plant disease classification and detection (Jain et al., 2022; Jaiswal et al., 2024).

2.3. Optimizer Implementations

The research evaluated optimization strategies by testing six different optimizers, which included AdamW, Adam, SGD, Adagrad, RMSprop, and SophiaG, as summarized in Table 1. The AdamW optimization method functions as an Adam optimizer variant that separates weight decay from gradient updates to enhance its regularization capabilities. The Adaptive Moment Estimation algorithm computes dynamic learning rates that depend on both the first and second gradient moment

values to optimize performance across diverse tasks. The oldest optimizer in existence, SGD, uses momentum as a method to speed up the process of reaching convergence. Adagrad uses learning rate adjustment based on how often parameters get updated, while RMSprop uses a moving average of previous gradients to handle problems that change over time. SophiaG functions as a new optimization method that uses second-order optimization through a stochastic Hessian approximation and implements a cosine learning rate schedule.

Table 1 summarizes the hyperparameter settings used to ensure a fair comparison among the evaluated optimizers.

Optimizer	Learning Rate	Weight Decay	Momentum	Rho	LR Scheduler	Other Settings
AdamW	1×10^{-4}	1×10^{-5}	–	–	None	PyTorch default parameters
Adam	1×10^{-4}	1×10^{-5}	–	–	None	PyTorch default parameters
SGD	1×10^{-2}	1×10^{-5}	0.9	–	None	Nesterov=False (default)
RMSprop	1×10^{-4}	1×10^{-5}	default	–	None	Centered=False (default)
Adagrad	1×10^{-2}	1×10^{-5}	–	–	None	Initial accumulator=0
SophiaG	1×10^{-4}	1×10^{-5}	–	0.01	CosineAnnealingLR (Tmax=10)	Hessian update every 50 iterations

The evaluation of each optimizer involved measuring its effects on training dynamics through three specific metrics, which included convergence speed, stability, and final performance metrics, which measured accuracy, precision, recall, and F1 score. Previous work by Thangaraj et al. (2021) successfully applied Adam, SGD, and RMSprop in tomato leaf disease classification using CNNs. We performed additional optimization tests for the MaxViT architecture by assessing both traditional optimization methods and their modern optimization counterparts. The complete implementation of the Multi-Axis Vision Transformer model and the optimizer comparison experiments are publicly available for reproducibility on the GitHub[1] repository.

Prior studies successfully applied Adam, SGD, and RMSprop optimizers in tomato leaf disease classification using CNNs. Inspired by this, we expanded the optimizer comparison within the MaxViT architecture to include both conventional and advanced optimizers. The complete implementation of the Multi-Axis Vision Transformer model and the optimizer comparison experiments are publicly available for reproducibility at the repository.

2.4. Evaluation Metrics

The evaluation of model performance used multiple statistical metrics, which included F1 Score, Precision, Recall, and overall Accuracy, to assess the validation dataset. The scikit-learn library was used to calculate these metrics. The researchers used a normalized confusion matrix to study how different classes were incorrectly identified. Training and validation accuracy curves were generated for all epochs to analyze model convergence behavior. The model achieving the best performance for each optimizer was selected and retained for further comparative analysis. The evaluation system used Precision and Recall together with another measurement index after providing formulas for each measurement requirement, which follows:

(1) Categorical Accuracy = Number of Correct Predictions / Total Number of Predictions
(2) Precision=TP/(TP+FP)
(3) Recall=TP/(TP+FN)
(4) F1-Score= (2 × (Precision × Recall)) / ((Precision + Recall))

TP: True Positives; FP: False Positives; and FN: False Negatives.

[1]- https://github.com/RafiullahRafi/TomatoLeaf_Disease_MultiAxisViT_Optimizers

3. EXPERIMENTAL RESULTS AND DISCUSSION

3.1. Performance evaluation.

The comparative evaluation of various optimization algorithms demonstrates that different algorithms produce distinct results when applied to the Multi-Axis Vision Transformer (MaxViT) deep learning model, which detects tomato leaf diseases. AmadW outperforms all other algorithms by achieving the highest results across every measurement category because it attained an F1 score of 0.9989 and recall and precision scores that matched this achievement, and it showed superior performance during training sessions that lasted for multiple epochs, as demonstrated in Figure 3. The optimizer proved to be the best choice for this field because it showed better generalization, faster convergence, and greater system resilience according to test results, while SophiaG provided similar results through its 0.9982 performance across three essential evaluation standards. Adam successfully maintained a high-performance level with a 0.9970 score while RMSprop achieved a 0.9961 score, yet both these optimizers fell short of reaching the top two performance levels. Traditional optimization methods, such as SGD and Adagrad, display lower performance levels when compared to contemporary optimization techniques. Adagrad records the lowest performance across all evaluation categories as it achieved an F1 score of 0.9832 and reached an overall accuracy of 98%, which represents its best performance.

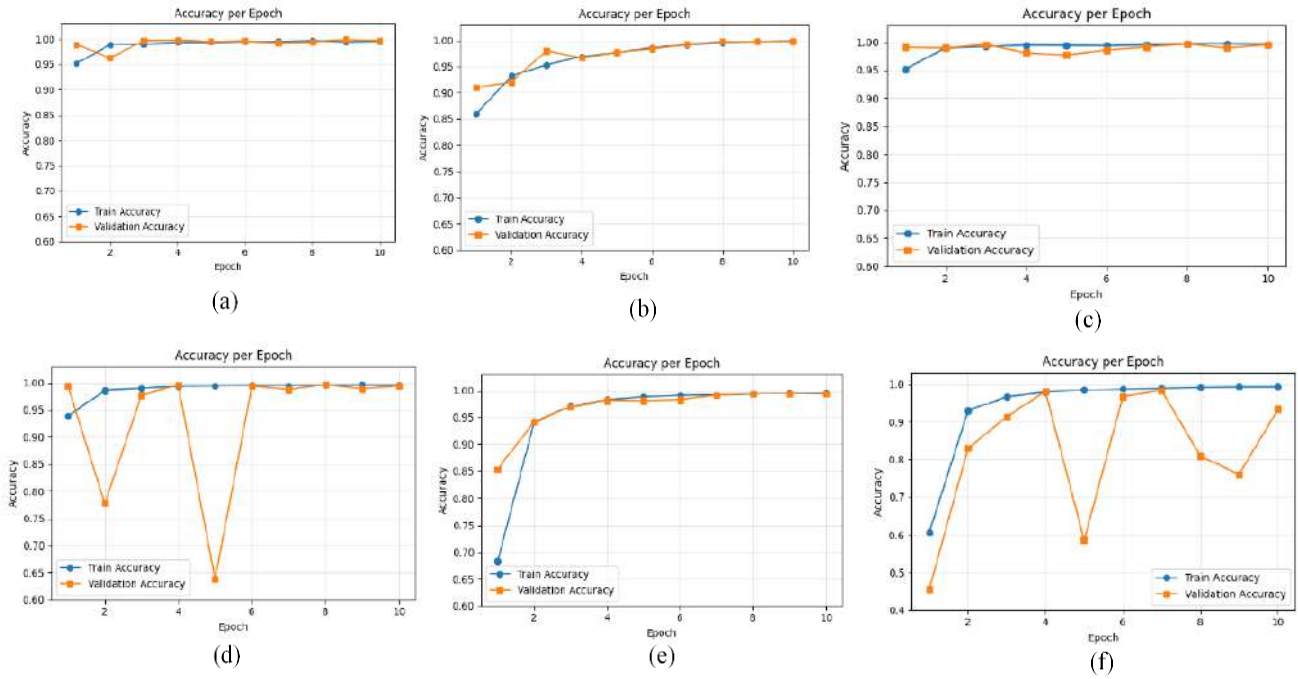


Figure 3. The MaxViT model training and validation accuracy using different optimization algorithms (a) AdamW, (b) SophiaG, (c) Adam, (d) RMSprop, (e) SGD, and (f) Adagrad across training epochs.

The adaptive gradient techniques which developed before modern times become ineffective when they attempt to analyze intricate plant disease detection systems, which require analysis of high-dimensional image data. Modern adaptive optimizers, including AdamW and SophiaG, demonstrate higher performance levels when applied to deep transformer-based models, which include multi-axis vision transformer, because these optimizers deliver better results in high-accuracy requirements of smart agriculture systems, precision agricultural practices, and plant pathology diagnostic processes. Figure 4 illustrates the six different optimizers perform when testing tomato leaf diseases through their metric heatmap assessment, which the MaxViT-Tiny model uses for classification.

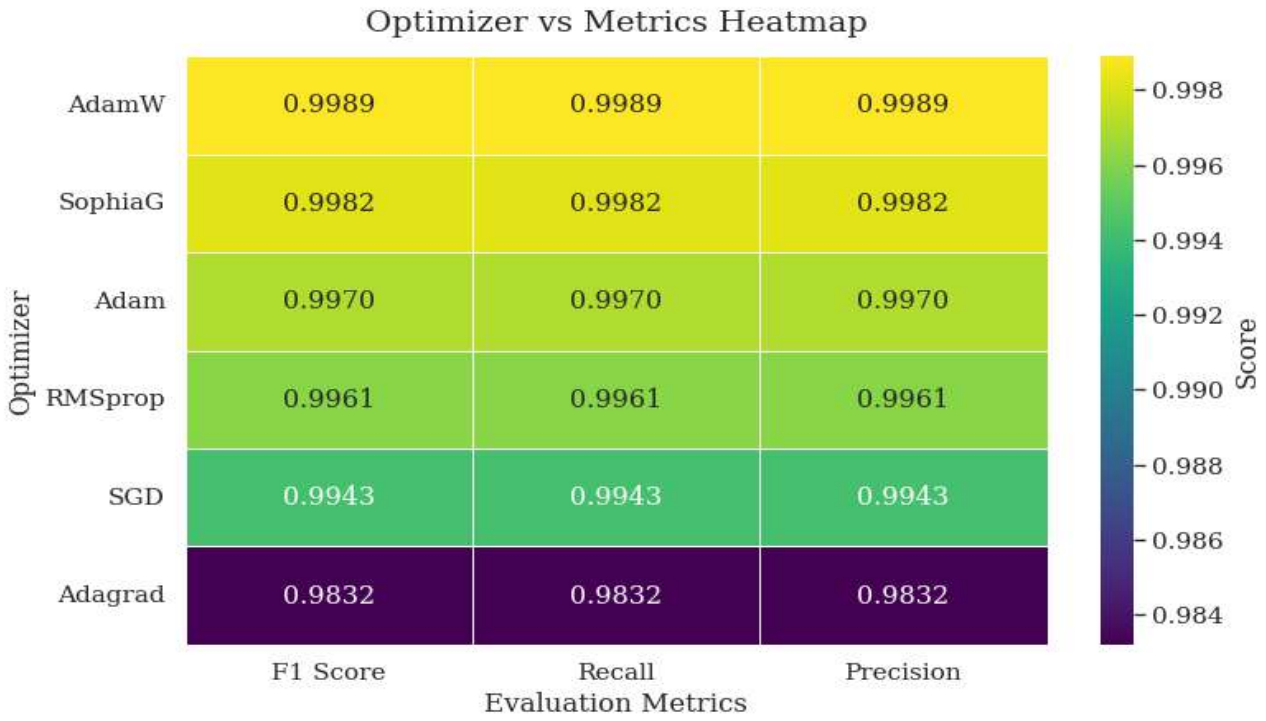
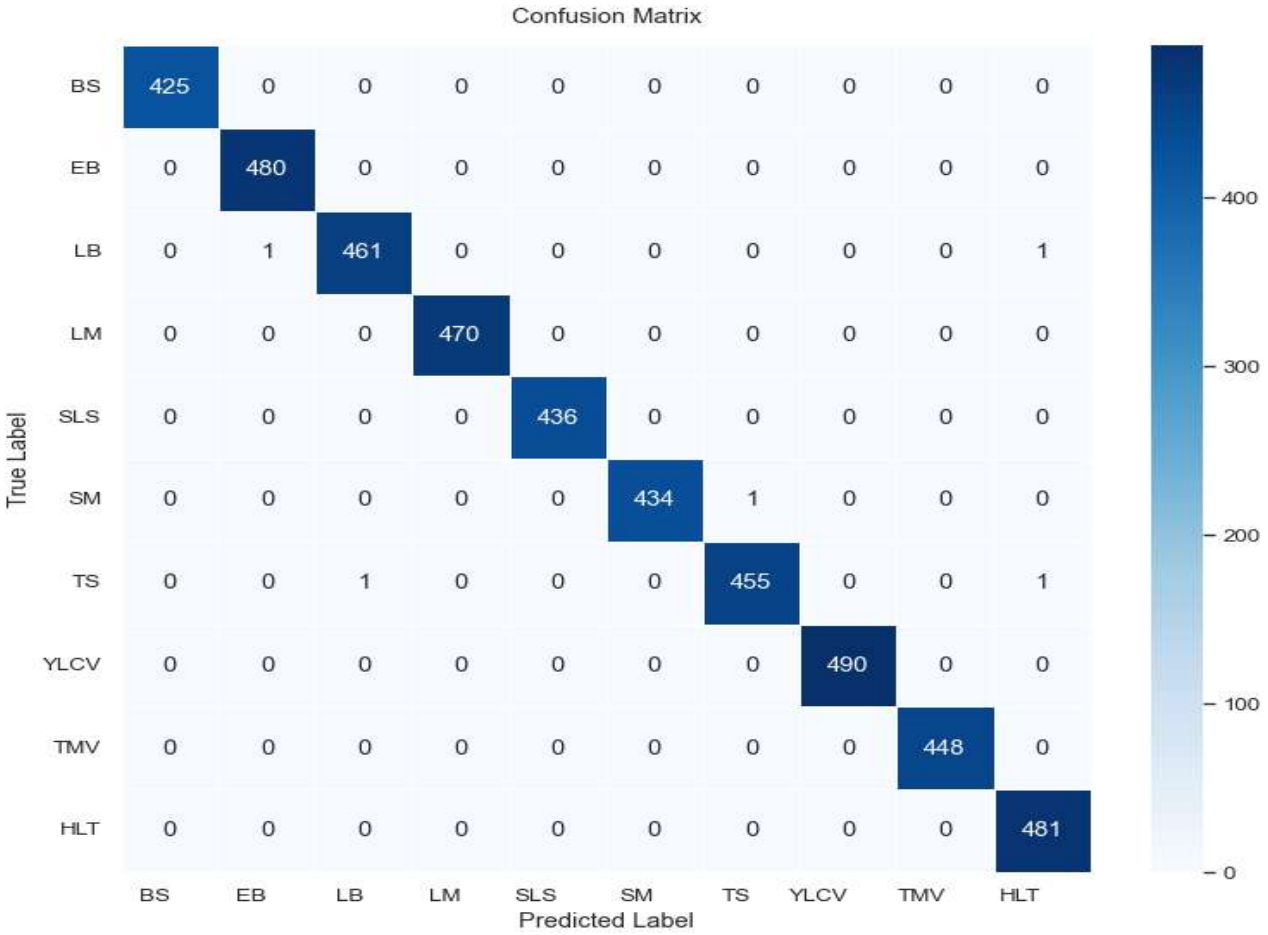


Figure 4. Comparison of various optimizers on tomato leaf disease detection on the MaxViT model.

3.2. Confusion Matrix analysis

Confusion matrix analysis was conducted to assess the classification performance of the MaxViT-Tiny model using different optimization algorithms. The analysis shows how the model predicts every tomato leaf disease category, which demonstrates the strengths and weaknesses of each optimizer to limit classification errors. The AdamW optimizer achieved the lowest number of misclassifications (5), followed closely by SophiaG (8) and Adam (14), indicating superior performance in terms of generalization. The two algorithms, RMSprop and SGD, recorded (18) and (26) misclassifications, while Adagrad showed its highest misclassification count at 77, which indicates that the model struggled with adaptive learning during its training process. The confusion matrix, which Figure 5 displays, shows the classification performance results of the MaxViT-Tiny model, which used the AdamW optimizer for its training. The model demonstrated superior performance compared with the other evaluated optimizers. The AdamW optimizer demonstrates higher correct prediction rates and lower misclassification frequencies, which establish it as the most effective optimization method for this particular situation.

The AdamW optimizer demonstrated superior performance across all metrics, achieving 4580 correct predictions out of 4585 total samples, resulting in only 5 misclassifications, as shown in Figure 5. The system achieved 99.89%accuracy because AdamW effectively reduced loss while maintaining a steady convergence process. The improved performance of the system may be attributed to the efficient second-order gradient approximation method and adaptive learning dynamics, which enhance its capability to handle complex classification tasks.



The abbreviations used in the confusion matrix represent the following full class names: BS – Bacterial Spot, EB – Early Blight, LB – Late Blight, LM – Leaf Mold, SLS – Septoria Leaf Spot, SM – Spider Mites, TS – Target Spot, YLCV – Tomato Yellow Leaf Curl Virus, TMV – Tomato Mosaic Virus, and HLT – Healthy.

Figure 5. Confusion matrix of the MaxViT model for tomato leaf disease classification using the AdamW optimizer, representing the best-performing configuration among all tested optimizers.

The SophiaG optimizer reached 4577 correctly classified samples out of 4585 total samples, with just 8 misclassifications, resulting in an accuracy of 99.82%. The Adam optimizer recorded 4571 correct predictions out of 4585, which resulted in 14 misclassifications and an accuracy of 99.70%. Adam remains an effective optimizer for deep learning workflows because it uses adaptive learning rates together with moment estimation techniques. The RMSprop optimizer produced 4567 correct predictions from 4585 samples, which resulted in 18 misclassifications and achieved an accuracy of 99.61%. The learning rate adjustment mechanism of this optimizer enables efficient handling of non-stationary objectives because it uses recent gradient magnitude data to determine learning rate changes. RMSprop demonstrates strong performance during training because it maintains gradient control even though it falls short of Adam-based optimizers in accuracy.

The SGD optimizer achieved 4559 correct predictions from 4585 total predictions, while making 26 incorrect predictions, which resulted in a 99.43% accuracy rate. Adagrad produced 4508 correct predictions from 4585 total samples while making 77 incorrect predictions, which represents the highest number of incorrect predictions across all tested optimizers, resulting in an accuracy rate of 98.32%. Adagrad applies learning rate adjustments based on how often parameters get updated, but the system experiences excessive learning rate decline throughout the training process, which creates obstacles during the later stages of training. The system shows reduced performance for complex

classification tasks, which require handling small differences between classes, because of its high misclassification rate. All optimizers' correct predictions and misclassifications are summarized in Table 2.

Table 2 Performance comparison of optimizers based on accuracy, misclassification, and prediction results.

Optimizer	Total Samples	Correct Predictions	Misclassifications	Accuracy
AdamW	4585	4580	5	99.89%
SophiaG	4585	4577	8	99.82%
Adam	4585	4571	14	99.70%
RMSprop	4585	4567	18	99.61%
SGD	4585	4559	26	99.43%
Adagrad	4585	4508	77	98.32%

3.3 Discussion

This study presented a comparative evaluation of six widely used optimization algorithms—Adam, AdamW, SGD, RMSprop, Adagrad, and SophiaG—within the context of tomato leaf disease classification using the Multi-Axis Vision Transformer (MaxViT) architecture. The study tested multiple optimizers on the PlantVillage dataset because its goal was to find which optimizer achieved the best results for disease detection in precision agriculture. The system achieves its effectiveness because it uses a new stochastic diagonal Hessian approximation method and element-wise gradient clipping system, which leads to rapid learning and effective generalization while using less processing power. AdamW's minimal misclassification rate (only 5 instances) further reinforces its robustness and reliability for high-stakes agricultural applications.

Recent studies have demonstrated that optimizers play a crucial role in improving the performance and stability of deep learning models for plant disease classification. The CNN-based method achieved its highest identification accuracy of 97.28% for tomato leaf diseases through transfer learning, which utilized AlexNet, GoogLeNet, and ResNet with the SGD optimizer (Zhang et al., 2018). The traditional optimizers SGD and RMSprop provide dependable convergence results; however, they need users to spend time tuning their hyperparameters. The Adam optimizer shows quicker convergence times and improved generalization abilities through its automatic learning rate adjustments, which use first and second moment gradient data (Kingma & Ba, 2015). Multiple studies have demonstrated that Adam performs better than SGD and RMSprop because it produces more accurate results and maintains training stability (Osmenaj et al., 2025; Thangaraj et al., 2021). The AdamW variant enables transformers to achieve better performance through its weight decay decoupling feature, which maintains stable optimization processes during transformer-based architecture operations in MaxViT and ViT-RoT (Hossain et al., 2023; Nishankar et al., 2025). The Sophia optimizer developed by Liu et al. (2024) reached its goal of faster convergence through curved second-order approximation, which reduced the needed computation time. The MaxViT model reached its optimal performance through AdamW optimization, which produced an F1-score of 0.9989 and precision and recall scores of 0.9989. Compared to CNN-based models and transformer-based systems, AdamW enables more precise and dependable classification results for tomato leaf diseases.

The superior performance of AdamW can be attributed to its decoupled weight decay mechanism, which improves regularization while preserving adaptive learning rates, making it particularly suitable for transformer-based architectures such as MaxViT. As shown in Figure. 3(a), AdamW

achieved the fastest and most stable convergence, resulting in the highest classification performance. SophiaG produced comparable results because its stochastic second-order Hessian approximation enables efficient optimization and rapid convergence (Figure. 3(b)), although its performance remained slightly below AdamW. In contrast, Adagrad exhibited the weakest performance (Figure. 3(f)) due to its continuously decreasing learning rate, which limited effective parameter updates and caused unstable validation accuracy. Similarly, RMSprop showed noticeable fluctuations during validation (Figure. 3(d)), whereas SGD converged more slowly (Figure. 3(e)) because it relies on a fixed global learning rate despite its stable optimization. These findings indicate that transformer-based models benefit from optimizers that combine adaptive learning rates with effective regularization, explaining the superior performance of AdamW over the other evaluated optimization algorithms.

The findings of this study should be interpreted in light of several limitations. First, the experiments were conducted using the publicly available PlantVillage dataset, which consists of high-quality images captured under controlled lighting conditions with relatively uniform backgrounds. Although these characteristics facilitate consistent model training and fair optimizer comparison, they do not fully represent the variability encountered in real agricultural environments, such as changing illumination, complex backgrounds, occlusions, and varying disease severity. Furthermore, PlantVillage is a curated benchmark dataset and may not adequately capture the diversity of tomato cultivars and field conditions across different geographical regions. Consequently, the generalizability of the proposed MaxViT–AdamW framework may be limited in practical applications. Future research should validate the proposed approach using field-collected datasets acquired under diverse environmental conditions and investigate its robustness across different cultivation systems and geographic locations.

CONCLUSION

This study demonstrates the critical role of optimization algorithms in enhancing the performance, convergence efficiency, and classification accuracy of vision-based deep learning models for agricultural disease detection. The AdamW and SophiaG optimizers prove to be the most effective optimizers for classifying tomato leaf diseases through the MaxViT architecture. The system demonstrates its ability to exceed the performance of both old and new optimization methods, which makes it suitable for use in plant pathology and smart agriculture. The study results provide important knowledge about how artificial intelligence and deep learning technologies link with agricultural diagnostic methods. The research findings lead to better disease detection systems, which help sustainable farming methods to grow and increase agricultural productivity, and improve food security and economic stability. Future research should examine hybrid optimization techniques, validate the proposed approach using field-collected datasets from diverse environmental conditions, and deploy MaxViT-AdamW systems in agricultural field settings to strengthen the connection between precision agriculture research and its practical applications.

CONFLICT OF INTEREST

All authors declare no conflict of interest regarding the publication of this research.

DATA AVAILABILITY

Code used to perform the analyses for this paper is hosted at https://github.com/RafiullahRafi/TomatoLeaf_Disease_MultiAxisViT_Optimizers. All analyses were conducted using Python (v3.x). No additional data have been deposited in a public repository.

AUTHORS CONTRIBUTIONS:

Rafiullah Rafi: Conceptualization, investigation, formal analysis, resources, original draft preparation, review and editing, visualization, supervision, project administration, and funding acquisition. Assistant Professor Alireza Soleimanipour and Associate Professor Abbas Rezaei: Supervision, consultation, review support, academic consultation, and guidance throughout the research. Sarajudin and Sediqullah Hammas: Methodology development and technical assistance.
All authors have read and agreed to the published version of the manuscript.

Funding
This research received no external funding.

AI Statement

Artificial intelligence tools were used only for language improvement and grammatical editing.

REFERENCES

Agarwal, M., Singh, A., Arjaria, S., Sinha, A., & Gupta, S. (2020). ToLeD: Tomato Leaf Disease Detection using Convolution Neural Network. *Procedia Computer Science*, *167*, 293–301. https://doi.org/10.1016/j.procs.2020.03.225

Amin, H., Darwish, A., Hassanien, A. E., & Soliman, M. (2022). End-to-End Deep Learning Model for Corn Leaf Disease Classification. *IEEE Access*, *10*, 31103–31115. https://doi.org/10.1109/ACCESS.2022.3159678

Barman, U., Sarma, P., Rahman, M., Deka, V., Lahkar, S., Sharma, V., & Saikia, M. J. (2024). ViT-SmartAgri: Vision Transformer and Smartphone-Based Plant Disease Detection for Smart Agriculture. *Agronomy*, *14*(2), 327. https://doi.org/10.3390/agronomy14020327

Chitta, S., Yandrapalli, V. K., & Sharma, S. (2024). Deep Learning for Precision Agriculture: Evaluating CNNs and Vision Transformers in Rice Disease Classification. *2024 OPJU International Technology Conference on Smart Computing for Innovation and Advancement in Industry 4.0, OTCON 2024*, 1–6. https://doi.org/10.1109/OTCON60325.2024.10687983

Collins, E. J., Bowyer, C., Tsouza, A., & Chopra, M. (2022). Tomatoes: An Extensive Review of the Associated Health Impacts of Tomatoes and Factors That Can Affect Their Cultivation. *Biology*, *11*(2), 239. https://doi.org/10.3390/biology11020239

Colucci, G. P., Battilani, P., Camardo Leggieri, M., & Trinchero, D. (2025). Algorithms for Plant Monitoring Applications: A Comprehensive Review. In *Algorithms* (Vol. 18, Number 2, pp. 1–47). https://doi.org/10.3390/a18020084

Dai, D., Xia, P., Zhu, Z., & Che, H. (2023). MTDL-EPDCLD: A Multi-Task Deep-Learning-Based System for Enhanced Precision Detection and Diagnosis of Corn Leaf Diseases. *Plants*, *12*(13), 2433. https://doi.org/10.3390/plants12132433

Das, P. K., Rupa, S. S., Pumrin, S., Das, U. C., & Hossen, Md. K. (2024). Deep Learning for Plant Disease Detection and Classification: A Systematic Analysis and Review. *CURRENT APPLIED SCIENCE AND TECHNOLOGY*, e0259016. https://doi.org/10.55003/cast.2024.259016

El Sakka, M., Mothe, J., & Ivanovici, M. (2024). Images and CNN applications in smart agriculture. In *European Journal of Remote Sensing* (Vol. 57, Number 1, pp. 1–30). Taylor & Francis. https://doi.org/10.1080/22797254.2024.2352386

Fraiwan, M., Faouri, E., & Khasawneh, N. (2022). Classification of Corn Diseases from Leaf Images Using Deep Transfer Learning. *Plants*, *11*(20), 1–14. https://doi.org/10.3390/plants11202668

Ghosh, H., Rahat, I. S., Emon, Md. M. R., Mashrafi, Md. J., Al Arafat Tanzin, M. A., Mohanty, S. N., & Kant, S. (2025). Advanced neural network architectures for tomato leaf disease diagnosis in precision agriculture. *Discover Sustainability*, *6*(1), 312. https://doi.org/10.1007/s43621-025-01149-1

Hamid, M. K., Shah, S. K., Husnain, G., Ghadi, Y. Y., Al Maaytah, S. A., & Qahmash, A. (2025). Enhancing Agricultural Disease Detection: A Multi-Model Deep Learning Novel Approach. *Engineering Reports*, *7*(1). https://doi.org/10.1002/eng2.13113

Hossain, S., Tanzim Reza, M., Chakrabarty, A., & Jung, Y. J. (2023). Aggregating Different Scales of Attention on Feature Variants for Tomato Leaf Disease Diagnosis from Image Data: A Transformer Driven Study. *Sensors*, *23*(7), 3751. https://doi.org/10.3390/s23073751

Jain, S., Sahni, R., Khargonkar, T., Gupta, H., Verma, O. P., Sharma, T. K., Bhardwaj, T., Agarwal, S., & Kim, H. (2022). Automatic Rice Disease Detection and Assistance Framework Using Deep Learning and a Chatbot. *Electronics (Switzerland)*, *11*(14), 2110. https://doi.org/10.3390/electronics11142110

Jaiswal, M., Nagwanshi, K. K., Kumar, R., Gaurav, S., & Watti, Y. (2024). Classification of Plant Disease using a State-of-the Art Deep learning Algorithm on a Tesla GPU. *International Journal of Computing*, 240–246. https://doi.org/10.47839/ijc.23.2.3542

Javidan, S. M., Ampatzidis, Y., Banakar, A., Asefpour Vakilian, K., & Rahnama, K. (2025). An Intelligent Group Learning Framework for Detecting Common Tomato Diseases Using Simple and Weighted Majority Voting with Deep Learning Models. *AgriEngineering*, *7*(2), 31. https://doi.org/10.3390/agriengineering7020031

Kingma, D. P., & Ba, J. L. (2015, December 22). Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*.

Krishna, M. S., Machado, P., Otuka, R. I., Yahaya, S. W., Neves dos Santos, F., & Ihianle, I. K. (2025). Plant Leaf Disease Detection Using Deep Learning: A Multi-Dataset Approach. *J*, *8*(1), 4. https://doi.org/10.3390/j8010004

Lee, S., & Lee, S. (2025). Efficient Data Augmentation Methods for Crop Disease Recognition in Sustainable Environmental Systems. *Big Data and Cognitive Computing*, *9*(1), 1–19. https://doi.org/10.3390/bdcc9010008

Liu, H., Li, Z., Hall, D., Liang, P., & Ma, T. (2024, May 23). SOPHIA: A SCALABLE STOCHASTIC SECOND-ORDER OPTIMIZER FOR LANGUAGE MODEL PRE-TRAINING. *12th International Conference on Learning Representations, ICLR 2024*.

Liu, J., & Wang, X. (2020). Tomato Diseases and Pests Detection Based on Improved Yolo V3 Convolutional Neural Network. *Frontiers in Plant Science*, *11*. https://doi.org/10.3389/fpls.2020.00898

Liu, W., & Zhang, A. (2025). Plant Disease Detection Algorithm Based on Efficient Swin Transformer. *Computers, Materials and Continua*, *82*(2), 3045–3068. https://doi.org/10.32604/cmc.2024.058640

Mohameth, F., Bingcai, C., & Sada, K. A. (2020). Plant Disease Detection with Deep Learning and Feature Extraction Using Plant Village. *Journal of Computer and Communications*, *08*(06), 10–22. https://doi.org/10.4236/jcc.2020.86002

Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using Deep Learning for Image-Based Plant Disease Detection. *Frontiers in Plant Science*, *7*(September). https://doi.org/10.3389/fpls.2016.01419

Nishankar, S., Pavindran, V., Mithuran, T., Nimishan, S., Thuseethan, S., & Sebastian, Y. (2025). ViT-RoT: Vision Transformer-Based Robust Framework for Tomato Leaf Disease Recognition. *AgriEngineering*, *7*(6), 185. https://doi.org/10.3390/agriengineering7060185

Osmenaj, Z., Tseliki, E.-M., Kapellaki, S. H., Tselikis, G., & Tselikas, N. D. (2025). From Pixels to Diagnosis: Implementing and Evaluating a CNN Model for Tomato Leaf Disease Detection. *Information*, *16*(3), 231. https://doi.org/10.3390/info16030231

Ouamane, A., Chouchane, A., Himeur, Y., Miniaoui, S., Atalla, S., Mansoor, W., & Al-Ahmad, H. (2025). Optimized Vision Transformers for Superior Plant Disease Detection. *IEEE Access*, *13*, 48552–48570. https://doi.org/10.1109/ACCESS.2025.3547416

Panno, S., Davino, S., Caruso, A. G., Bertacca, S., Crnogorac, A., Mandić, A., Noris, E., & Matić, S. (2021). A Review of the Most Common and Economically Important Diseases That Undermine the Cultivation of Tomato Crop in the Mediterranean Basin. *Agronomy*, *11*(11), 2188. https://doi.org/10.3390/agronomy11112188

Qin, Y. M., Tu, Y. H., Li, T., Ni, Y., Wang, R. F., & Wang, H. (2025). Deep Learning for Sustainable Agriculture: A Systematic Review on Applications in Lettuce Cultivation. In *Sustainability (Switzerland)* (Vol. 17, Number 7, p. 3190). https://doi.org/10.3390/su17073190

Sakkarvarthi, G., Sathianesan, G. W., Murugan, V. S., Reddy, A. J., Jayagopal, P., & Elsisi, M. (2022). Detection and Classification of Tomato Crop Disease Using Convolutional Neural Network. *Electronics*, *11*(21), 3618. https://doi.org/10.3390/electronics11213618

Shanthi, D. L., Vinutha, K., Ashwini, N., & Vashistha, S. (2024). Tomato Leaf Disease Detection Using CNN. *Procedia Computer Science*, *235*, 2975–2984. https://doi.org/10.1016/j.procs.2024.04.281

Shen, Y., Yang, Z., Khan, Z., Liu, H., Chen, W., & Duan, S. (2025). Optimization of Improved YOLOv8 for Precision Tomato Leaf Disease Detection in Sustainable Agriculture. *Sensors*, *25*(5), 1398. https://doi.org/10.3390/s25051398

Thangaraj, R., Anandamurugan, S., & Kaliappan, V. K. (2021). Automated tomato leaf disease classification using transfer learning-based deep convolution neural network. *Journal of Plant Diseases and Protection*, *128*(1), 73–86. https://doi.org/10.1007/s41348-020-00403-0

Tiwari, M., Kumar, H., Prakash, N., Kumar, S., Neware, R., Tripathi, S., & Agarwal, R. (2024). Tomato Disease Detection Using Vision Transformer with Residual L1-Norm Attention and Deep Neural Networks. *International Journal of Intelligent Engineering and Systems*, *17*(1), 679–688. https://doi.org/10.22266/ijies2024.0229.57

Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture*, *161*, 272–279. https://doi.org/10.1016/j.compag.2018.03.032

Ullah, Z., Alsubaie, N., Jamjoom, M., Alajmani, S. H., & Saleem, F. (2023). EffiMob-Net: A Deep Learning-Based Hybrid Model for Detection and Identification of Tomato Diseases Using Leaf Images. *Agriculture*, *13*(3), 737. https://doi.org/10.3390/agriculture13030737

Upadhyay, A., Chandel, N. S., Singh, K. P., Chakraborty, S. K., Nandede, B. M., Kumar, M., Subeesh, A., Upendar, K., Salem, A., & Elbeltagi, A. (2025). Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. *Artificial Intelligence Review*, *58*(3), 1–64. https://doi.org/10.1007/s10462-024-11100-x

Wang, S., Xu, D., Liang, H., Bai, Y., Li, X., Zhou, J., Su, C., & Wei, W. (2025). Advances in Deep Learning Applications for Plant Disease and Pest Detection: A Review. In *Remote Sensing* (Vol. 17, Number 4, pp. 1–31). https://doi.org/10.3390/rs17040698

Zhang, K., Wu, Q., Liu, A., & Meng, X. (2018). Can Deep Learning Identify Tomato Leaf Disease? *Advances in Multimedia*, *2018*, 1–10. https://doi.org/10.1155/2018/6710865

Zhu, T., Yan, F., Lv, X., Zhao, H., Wang, Z., Dong, K., Fu, Z., Jia, R., & Lv, C. (2024). A Deep Learning Model for Accurate Maize Disease Detection Based on State-Space Attention and Feature Fusion. *Plants*, *13*(22), 1–24. https://doi.org/10.3390/plants13223151