

Original Article

GRAPEVINE DISEASE DETECTION USING AN OPTIMIZED DEEP NEURAL NETWORK FOR SMART AGRICULTURE

Ashraf Atam Mustafa^{1,2,*} , Araz Rajab Abraham³ 

¹Duhok Polytechnic University, Duhok, Kurdistan Region, Iraq.

²Department of Information Technology, Technical College of Informatics, Akre University for Applied Sciences, Akre, Kurdistan Region, Iraq.

³Duhok Technical Institute, Duhok Polytechnic University, Duhok, Kurdistan Region, Iraq.

*Corresponding author, E-mail: ashraf.mustafa@dpu.edu.krd (Tel: +964-7508595109)

ABSTRACT

Received
July 27, 2025

Accepted
September 10, 2025

Published
July 7, 2026

Grapevines are economically vital crops but are highly susceptible to fungal, bacterial, and viral diseases that threaten yield and quality. Traditional detection methods that rely on manual inspection are time consuming, prone to human error, and often delay intervention. This study introduces a lightweight convolutional neural network (CNN) architecture specifically designed for accurately and efficiently detecting grapevine leaf diseases including Black Rot, ESCA, and Leaf Blight based on image classification. The proposed model integrates optimized residual blocks, batch normalization, dropout layers, and global average pooling to maximize accuracy while minimizing computational complexity. This lightweight design makes it well-suited for deployment on edge devices such as drones and mobile systems used in precision agriculture. A comprehensive data augmentation strategy was applied during training to simulate real-world variability and enhance generalization. The model was trained using 9,027 labeled grape leaf images from a publicly available grape disease dataset, achieving 99.8% overall accuracy with near-perfect precision, recall, F1-score, and area under the ROC curve (AUC) across all classes. To further validate its practical utility, we developed a cross-platform mobile application that integrates the trained model for real-time inference. The app was tested using five grapevine leaf images randomly obtained from Google, and all of them were correctly predicted. This preliminary validation demonstrates the feasibility of using the system in real-world agricultural scenarios. These findings highlight the practical potential of the proposed approach for real time, scalable, and sustainable disease monitoring in smart vineyard management systems.

Keywords: Deep Learning, Convolutional Neural Network, Grapevine Disease Detection, Smart Agriculture, Residual Blocks.

1. INTRODUCTION

Ensuring food security remains a critical global priority due to population growth and climate variability, especially considering the increasing world population and the unpredictability of the effects of climate change (Godfray *et al.*, 2010). Sustainable agriculture practices are required to respond to these challenges, prompting new approaches to maximizing crop yield and quality while minimizing the impact to environment. Intelligent agriculture driven by the utilization of advanced artificial intelligence (AI), machine learning (ML), and deep learning (DL) techniques fits in this scenario. Through the use of smart (intelligent) systems, smart farming intends to maximize the management of resources, monitoring of crop and early detection of threats allowing higher productivity and sustainable nature of farms (El-Kenawy *et al.*, 2025; Islam *et al.*, 2025). The economic implication of plant diseases to the food and feed/game production of the world are formidable, with estimated annual losses ranging from 10% to 16% of total crop production, translating into billions of dollars lost worldwide.

Grapevines (*Vitis vinifera*) are one of the most economically im-

portant horticultural crops worldwide with multi-billion-dollar industries based on wine, fresh and dry fruit production. Despite their economic relevance, grapevines are especially vulnerable to several diseases induced by diverse pathogens, including fungus, bacteria, and viruses. Major diseases such as Black Rot (*Guignardia bidwellii*), ESCA, and Leaf Blight can cause significant reduction in grape yield, low quality of fruits, and in worst cases lead to enormous economic loss and total destruction of vineyard (Aher *et al.*, 2025). Grapevine disease identification is traditionally based on the visual manual inspection of the vines by agricultural specialists. However, this process is by nature too much tedious, time-consuming, subjective and error prone, which leads to the delay of the diagnosis beyond a critical time for effective intervention (Elsherbiny, 2024). Impacts of delayed diagnosis are severe with pathogen dissemination, further reinforcing broad-spectrum pesticide dependence, and have growers experiencing significant economic losses. It is therefore necessary to develop rapid and automation approach for accurate diagnosis to promote the sustainable development of the grape industry. Most of the people cannot have specialized treatment in time. Currently, the workflow is largely based on visual and dermoscopic imaging followed by costly

Access this article online



DOI: <https://doi.org/10.25271/sjuoz.2026.14.3.1736>

Printed ISSN 2663-628X;
Electronic ISSN 2663-6298

Science Journal of University of Zakho
Vol. 14, No. 03, pp. 484-493, July -2026

This is an open access article under a CC BY-NC-SA 4.0 license
(<https://creativecommons.org/licenses/by-nc-sa/4.0/>)

and invasive procedures (biopsies) that are read by human doctors. These practices are not only laborious and resource intensive, they are also largely subjective and may, therefore, vary greatly depending on the proficiency and experience of the attending clinician (Malik *et al.*, 2024).

The emergence of digital image processing coupled with deep learning methodologies has revolutionized plant pathology, offering unprecedented opportunities for automated, rapid, and precise disease detection (Shoaib *et al.*, 2023). Specifically, CNN has shown exceptional performance in challenging image classification tasks. Thanks to their natural capacity for automatically learning layered features from raw pixel information (Gupta *et al.*, 2023; Zhong *et al.*, 2020). This capability makes CNNs exceptionally well-suited for identifying the subtle visual symptoms of plant diseases, even at their incipient stages, which may be imperceptible to the unaided human eye.

This study introduces an innovative deep neural network framework designed for the accurate and efficient diagnosis and monitoring of common grapevine leaf diseases, specifically Black Rot, ESCA, and Leaf Blight, alongside healthy specimens. Our study uses a publicly available, comprehensive dataset of grape leaf photographs. We introduce a tailored CNN framework that systematically integrates a resilient residual block and employs a variety of data augmentation techniques. Residual connections are necessary to solve the vanishing gradient problem, which makes it easier to train deeper and more complex networks while keeping feature information. Simultaneously, our comprehensive data augmentation strategy significantly enhances the model's generalisation capabilities by emulating the inherent variability found in real agricultural environments, thereby reducing the risk of overfitting and increasing resilience to unforeseen fluctuations (Khan *et al.*, 2025). A comprehensive set of established performance metrics, including overall accuracy, class-specific precision, recall (sensitivity), specificity, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), is used to carefully evaluate the proposed model's effectiveness.

The following is a methodical outline of the main contributions of this study:

- **Development of a New and Robust Deep Learning Model:** We design and implement a custom-built CNN architecture, specifically optimized with integrated residual blocks, for highly robust multiclass classification of grapevine diseases directly from leaf images. This architecture is uniquely tailored to extract highly discriminative and subtle pathological features, overcoming limitations of generic architectures.
- **Comprehensive Data Handling and Enhancement with Advanced Augmentation:** A meticulous data preprocessing pipeline is established, encompassing image resizing, normalization, and an extensive, biologically informed data augmentation strategy. This rigorous approach is critical for enhancing model robustness, significantly improving generalization across varied field conditions, and effectively addressing potential class imbalances within the dataset, leading to more resilient models.
- **Rigorous Performance Evaluation and Validation with High Discriminatory Power:** The proposed model's effectiveness was rigorously examined using an extensive set of both conventional and advanced evaluation metrics. The empirical results consistently demonstrate superior classification accuracy and exceptional discriminatory power including all specified illness classifications and the healthy category, validating its practical utility and readiness for real-world deployment. The near-perfect Area Under the Curve (AUC) scores highlight its strong classification boundary.
- **Deployment in a Cross-Platform Mobile Application for Real-Time Execution:** Bridging the gap between laboratory research and real-world practice, the optimized lightweight CNN model was converted using TensorFlow Lite and deployed within a fully functional mobile application for both iOS and Android platforms. This application enables real-time identification of grapevine diseases directly in the field, providing farmers and agricultural stakeholders with a convenient, portable, and effective decision-support tool.

The subsequent sections of this work are structured as follows. Section 2 addresses pertinent literature. Section 3 delineates the materials and procedures, the proposed architecture, and the criteria employed for evaluation. Section 4 presents the experimental outcomes. Section 5 evaluates and analyses findings, juxtaposes them with the prevailing best practices, and identifies the limits of the study. Finally, Section 6 concludes the study and proposes avenues for future research.

2. RELATED WORK

In the field of agricultural image analysis, deep learning methods, particularly Convolutional Neural Networks (CNNs) (Z Tang *et al.*, 2020), have greatly improved disease detection of plants, often with high accuracy. Sophisticated CNN architectures (e.g., VGGNet, GoogLeNet, ResNet, and DenseNet) have been developed and adopted by querying large image datasets and then refining on domain specific agricultural images for tasks such as leaf disease classification. Even with these advances, some challenges including small datasets, class distribution imbalance and image artifacts still remain.

Recent studies have leveraged deep learning for plant disease detection, achieving notable success on different grapevine leaf disease datasets. For example, Z Tang *et al.* (2020) have presented a lightweight CNN model for grape disease identification, obtaining an accuracy of 99.14% over a set of grape leaf images. Their approach was to design an efficient, mobile-friendly image model that was compact and low-latency by leveraging the ShuffleNet architectures along with a channel-wise attention model. Likewise, Liu *et al.* (2020) developed an enhanced CNN called DICNN, utilizing Inception structures and dense connections. Their model achieved a 97.22% accuracy rate on a substantial dataset of synthetic grape leaf images, highlighting the benefits of feature reuse and diverse feature representation.

Some authors have considered alternative building alterations and technologies. Thet *et al.* (2020) modified a VGG16 architecture by using a Global Average Pooling (GAP) layer and obtained an 98.4% accuracy when tested on a dataset of a Myanmar grapevine yard and yielded superior results when compared with ordinary fully connected layers and many more. Lin *et al.* (2022) introduced a CNN model named GrapeNet, which is composed of residual feature fusion blocks and attention modules, and obtained an accuracy of 86.29% on a dataset for different symptom stages, demonstrating the significance of feature fusion in modeling subtle inter-class differences.

More recently, Zinonos *et al.* (2022) demonstrated the feasibility of identifying grape leaf diseases from low-resolution images transmitted over a LoRa network. Their fine-tuned ResNet50 model achieved up to 99.77% accuracy on images with no packet loss. It maintained high accuracy even with significant image corruption, proving the robustness of CNNs in real-world, communication-constrained scenarios.

Prasad *et al.* (2024) formulated a Deep Convolutional Neural Network (DCNN) classifier based on the VGG16 architecture, which they augmented with supplementary layers to elevate its performance. Evaluated on a public Kaggle dataset, their model achieved a test accuracy of 99.06%, with a precision of 0.99, a recall of 0.99, and an F1-score of 0.99, demonstrating that a well-tuned, deeper architecture can achieve high precision in classifying grape leaf diseases. Our work introduces a custom-designed deep CNN for detecting Black Rot, ESCA, and Leaf Blight in grapevines. Our model, which integrates residual blocks, and a novel stage-wise progressive training and advanced data augmentation, was tested on a public grape leaf image dataset. The classifier attained an outstanding accuracy of 99.8%, along with almost perfect precision, recall, and F1-measure for all disease classes. This result is better than many of the above mentioned modern methods, showing the importance of a customized architecture integrated with a flexible data processing pipeline for this specific task.

Recent studies continue to advance the field of grapevine disease detection through novel deep learning architectures. For instance, Zir-ije Hasani *et al.* (2025) explored various deep learning models and found that a Transfer Learning approach achieved a high accuracy of 98.84% for classifying grape leaf diseases, outperforming their

custom CNN and DNN models. Hu *et al.* (2025) proposed GCS-YOLO, a lightweight model based on YOLOv8, which achieved a mean average precision of 96.2% for grape leaf disease detection while maintaining a small model size of 3.5 MB, making it suitable for deployment on edge devices. Furthermore, Zhang *et al.* (2024) introduced DLVTNet, a lightweight model that combines CNN and Transformer structures, achieving an average recognition accuracy of 98.48%. These studies underscore the trend towards developing lightweight and efficient models for real-world applications.

Other scholars have investigated hybrid models that integrate conventional machine learning classifiers with the feature extraction capabilities of deep learning. For example, Saberi Anari (2022) introduced a hybrid approach where features extracted with a modified dilated ResNet-18 were classified using an ensemble of Support Vector Machines (SVMs), achieving 99.1% accuracy. Such studies highlight the strong potential of combining deep learning and machine learning techniques for plant disease classification.

Although previous studies have achieved remarkable results in grapevine disease detection using CNNs and hybrid models, several challenges remain. Many existing approaches are either computationally intensive, lack robustness, or are unsuitable for real-time deploy-

ment on resource-constrained devices. Lightweight models, while efficient, often compromise accuracy or struggle with class imbalance and image artifacts, and hybrid models combining deep learning with traditional classifiers remain impractical for mobile or edge use. To address these limitations, we propose a custom-designed CNN that builds on the strengths of prior architectures by integrating residual blocks and advanced data augmentation, achieving a state-of-the-art accuracy of 99.8% while remaining lightweight and efficient. The model has been fully deployed in a cross-platform mobile application (Android and iOS), demonstrating its real-time usability, robustness, and practical readiness for agricultural applications.

3. MATERIALS AND METHODS

The current study presents comprehensive methodologies for the classification of grapevine diseases utilising leaf images. The methodology includes characterising the dataset, carefully preparing and adding to the data, designing the architecture of the proposed deep neural network, and performance assessment metrics employed to evaluate the model's efficacy. The subsequent subsections furnish detailed descriptions of each sequential step depicted in Figure 1.




Figure 1: Overarching methodology pursued in this study for classifying grape diseases from leaf images.

Dataset Description:

The experiments in this study were conducted using a publicly accessible dataset sourced from the Kaggle platform (Mandal, 2023). The Grape Disease Dataset (Original) was compiled by researchers specializing in agricultural computer vision and plant pathology, with their effort aimed at establishing a comprehensive, high-quality repository of grape leaf images to support the development and evaluation of deep learning models in precision agriculture and plant disease detection. This collection holds 9,027 grape leaf images distributed among four classes: three disease types (Black Rot, ESCA/Black Measles, Leaf Blight) and one healthy category, with the distribution detailed in Table 1 and examples shown in Figure 2. The original images vary in resolution and are stored in standard formats (JPEG, PNG). The disease classifications are based on expert agricultural and psychopathological assessment, ensuring diagnostic accuracy. Images were captured under diverse field conditions and lighting scenarios, representing natural variability in grape cultivation environments, leaf maturity stages, and disease progression severity.




Figure 2: Overarching methodology pursued in this study for classifying grape diseases from leaf images.

Table 1: Distribution of grape leaf images across different disease classes and the healthy category in the Dataset.

Class Name	Category Code	No. of Samples	Percentage %
Black Rot	br	2,360	26.14
ESCA	esca	2,400	26.59
Healthy	hlth	2,115	23.43
Leaf Blight	lb	2,152	23.84
Total number of samples		9,027	100

Data Preparation:

The input photos were standardised to 224x224 pixels to guarantee equal spatial dimensions across the dataset, facilitating consistent input for the deep learning model. Subsequently, normalisation of picture pixel values was executed, an essential preprocessing measure for improving numerical stability and expediting the convergence of the training process. Pixel values, initially spanning from 0 to 255 in conventional 8-bit pictures, were normalised to the interval [0, 1] by dividing each pixel value by 255.0. This method of Min-Max normalisation guarantees that all input features retain a uniform scale, which is crucial for the robust and efficient training of convolutional neural networks.

The normalized leaf images were then randomly partitioned into three independent sets: training, validation, and testing, as shown in Table 2. This was done to assess the effectiveness of the proposed models for grape disease classification. The data splitting strategy first divided the Dataset into training and testing subsets with a ratio of 70%:30%. From the test subset, a portion equivalent to 50% of the total test subset was subsequently set aside for validation, yielding a total distribution of 70% for training, 15% for validation, and 15% for testing.

Table 2: Distribution of images across training, validation, and test sets before data augmentation, showing the number of samples for each disease type.

Type of diseases	br	esca	hlth	lb
Training	1,652	1,680	1,480	1,506
Validation	354	360	317	323
Testing	354	360	318	323

Data Augmentation:

A comprehensive data augmentation strategy was meticulously implemented to significantly improve the model’s robustness, and its generalization capability to real world variations, and effectively address potential class imbalance within the Dataset. This strategy was exclusively applied during the training phase. The augmentation pipeline was carefully designed to simulate diverse real world imaging conditions of grape leaves while preserving the critical diagnostic features relevant to each disease class. The specific transformations applied as part of this augmentation strategy are detailed in Table 3, encompassing geometric and photometric modifications, as well as boundary handling techniques, to enrich the variability of the training data and prevent overfitting.

All augmentation parameters were carefully selected to maintain biological plausibility and avoid generating unrealistic or non-diagnostic images, while maximizing the training data’s diversity.

Crucially, it should be noted that no augmentation was applied to the validation and testing datasets. This ensures an unbiased and accurate evaluation of the model’s real generalization performance on data directly reflecting real world, untouched conditions.




Figure 3: Data Augmentation Techniques Applied to Grape Leaf Disease Classification Dataset.

Proposed Method:

The cornerstone of this study is the development of a new CNN architecture, specifically engineered for robust and highly accurate classification of grapevine diseases directly from leaf images. The proposed architecture, meticulously detailed in Figure 4, synergistically combines conventional convolutional layers with advanced residual blocks. This structural design efficiently recovers hierarchical picture information and addresses common challenges such as the vanishing gradient problem, which often hinders the training of deep neural networks.

The proposed CNN architecture is designed to efficiently process input images with 224×224 pixels and 3 color channels (RGB) standardized dimensions. The network’s design comprises several strategically integrated key components, balancing computational efficiency with high feature learning capacity.

Input Layer: The network starts with an input layer that can take images with a 224×224×3 size. This normalization allows for consistent ingestion that may be further analyzed.

Table 3: Data Augmentation Techniques Applied to Grape Leaf Images.

Category	Transformation	Description
Geometric	Rotation	±30° random rotation to simulate varied orientations.
	Shift (X & Y)	Up to 20% image shift to vary leaf position.
	Zoom	±30% zoom to simulate distance variation.
	Flip (H & V)	Random horizontal and vertical flips to reduce bias.
Photometric	Brightness	70%–130% brightness variation to mimic lighting changes.
	Shear	Up to 20% shear to simulate angled views.
Boundary Handling	Fill Mode	‘Nearest’ pixel fills to maintain image consistency after transformation.




Figure 4: Architecture of the proposed DNN for grapevine disease detection.

First Set of Convolutional Blocks (Conv Block 1.1 & Conv Block 1.2): The input image simultaneously undergoes feature extraction through two Conv2D paths with 64 filters of dimensions 3×3. After each Conv2D operation is performed, it applies a Batch Normalization layer to stabilize and speed up the training process and a ReLU activation function to introduce non-linearity. These first blocks take care of the most basic low-level features of the input images, e.g., edges and textures.

A Max-pooling layer with a 2x2 window is used to decrease the spatial size of the feature maps after each convolutional block, and this same layer is applied again following every residual block. This layer forwards down-samples the feature maps, slightly decreasing their spatial dimension. This is important for simplifying the computation and obtaining a more abstract, sparse, and robust feature representation, which further makes the model to be more invariant to small shifts and distortions of input images.

Residual Blocks (Residual Block 2, Residual Block 3): The core of our network lies in two stacked Residual Blocks, which are derived from the ResNet architecture. These are vital for enabling the training of deep networks because they help alleviate the vanishing gradient problem. By stacking two consecutive Conv2D layers with 128 filters each, both followed by Batch Normalization and ReLU activation and a "skip connection" where the input is directly added to the output, a residual block provides a good flow of gradients. Residual Block 2 consists of two sequential Conv2D layers with 128 filters each followed by Batch Normalization and ReLU activation. This enables it to learn a more complex set of features, including pairs of textures and shapes that signal disease. Residual Block 3 is conceptually the same except that it has 256 filters of size 3X3 in each of Conv2D layers. This enables the block to learn an even richer hierarchical feature set at a more abstract level and improves the power of the block to distinguish minor disease characteristics.

Global Average Pooling Layer: Following the last Max Pooling layer, a Global Average Pooling layer creates a 256-dimensional vector representing the salient features. This layer's number of parameters is much smaller than that of traditional fully connected layers, avoids overfitting to some extent, and results in higher robustness.

Dense Blocks (Dense Block 1 & Dense Block 2): The 256-D feature vector is passed through two fully connected (Dense) blocks for classification. A Dense layer with 512 units is used for Dense Block 1 with Batch Normalization, ReLU activation, and a Dropout with a rate of 0.5. Dense Block 2 has similar architecture and consists of Dense layer of 256 units, Batch Normalization, ReLU activation and

a Dropout layer of rate 0.5. These layers provide interpretation of the learned features and their mapping to class probabilities with additional Dropout technique serves as a strong regularizer for mitigating the risk of overfitting enhancing generalization.

Output Layer: The final layer of the network is a Dense layer which consists of 4 units and applicable to the four classes of the grapevine leaf status (Black Rot, ESCA, Healthy, Leaf Blight). The raw output scores are transformed into a probability distribution using a SoftMax activation function, with the most probable class being selected as the model's prediction.

This meticulously crafted architecture synergistically leverages the strengths of CNNs for efficient hierarchical feature extraction and the robustness conferred by residual connections for training deeper models effectively. The comprehensive design, coupled with strategic regularization (Batch Normalization and Dropout), achieves high accuracy, robust generalization, and reliable performance in the challenging domain of grapevine disease detection while maintaining a reasonable computational footprint.

4. RESULT

This section presents a general and complete analysis of all the experimental evidences achieved from training and testing the proposed deep neural network for grapevine disease classification.

Experimental Setup:

A controlled experimental protocol was thoroughly fixed in order to perform a rigorous and reproducible assessment of the proposed deep learning model. This arrangement comprised the description of critical hyperparameters, training parameters, data splitting methods, necessary for training and evaluating models. These hyperparameters were chosen from the best practices in deep learning related research and were usually tuned according to empirical observations for better performance.

Table 4: Hyperparameters used during training.

Hyperparameter	Value
Input Image Size	224 × 224
Learning Rate (Initial)	1e-3 (0.001)
Minimum Learning Rate	1e-7 (0.0000001)
Batch Size	32
Optimizer	Adam
Momentum (β_1)	0.9
Momentum (β_2)	0.999
Optimizer Epsilon	1.00E-07
Epochs	50
Loss Function	Categorical Cross-Entropy
Early Stopping Patience	10 epochs
ReduceLROnPlateau Patience	5 epochs
Data Split	70% train, 15% valid, 15% test
Random State	42

Performance Evaluation Metrics:

To thoroughly assess the performance and reliability of our deep learning model for grape leaf analysis, we used a wide array of established metrics suited for multi-class classification. This evaluation was based on calculating the overall accuracy, class-specific precision, recall (sensitivity), specificity, and the F1-score. These metrics are calculated using common mathematical equations from Eqs. (1) through (5), respectively:

$$\text{Accuracy} = \frac{\sum_{i=1}^n TP_i + TN_i}{\sum_{i=1}^n (TP_i + FP_i + TN_i + FN_i)} \quad (1)$$

$$= \frac{\text{Total Correct Predictions}}{\text{Total Predictions}}$$

$$\text{Precision}_i = \frac{TP_i}{TP_i + FP_i} \quad (2)$$

$$\text{Recall}_i = \text{Sensitivity}_i = \frac{TP_i}{TP_i + FN_i} \quad (3)$$


$$\text{Specificity}_i = \frac{TN_i}{TN_i + FP_i} \quad (4)$$

$$\text{F1-Score}_i = \frac{2 \times \text{Precision}_i \times \text{Recall}_i}{\text{Precision}_i + \text{Recall}_i} \quad (5)$$

In addition to these discrete scores, we employed the Area Under the Receiver Operating Characteristic (ROC) curve (AUC) as a crucial continuous metric. An ROC curve graphically represents a model's performance by plotting the true positive rate (TPR) against the false positive rate (FPR) across all classification thresholds. The resulting AUC value is highly informative because it reflects the model's overall capability to distinguish between classes, independent of any single threshold choice.

Training and Validation Performance:

As shown in Figure 5, the convergence of both the training and validation curves is smooth and steady, which indicates that the model learned features effectively and generalized well. On the left, training accuracy rapidly rises toward a near-perfect 1.00, while validation accuracy steadily climbs and levels off around 0.98 after roughly 20 epochs, the the consistently small gap between training and validation accuracy indicates that the model generalized well without signs of overfitting. On the right, training and validation loss start relatively high as is typical in early training then sharply decrease and converge toward zero, signifying that the model makes almost no errors on seen data and maintains a very low error rate on unseen validation images. The stable, parallel trajectory of these curves highlights a consistent training process and validates the chosen hyperparameters and regularization methods, confirming the model's capacity to generalize to new data.

**Figure 5:** Evolution of the training and validation accuracy, alongside their corresponding loss curves, across the training duration.

Performance Analysis:

For a detailed evaluation of each class (Black Rot, ESCA, Leaf Blight, and Healthy), we calculated the specific precision, recall, and F1-score. These metrics provide a deeper understanding of the model's ability to correctly identify positive cases (precision) while also capturing all actual positives (recall). Figure 6 visually displays these crucial performance metrics for each class on the independent test dataset.


**Figure 6:** Performance Metrics (Precision, Recall, F1-Score) by Class, demonstrating high and balanced classification across all categories on the test set.

Figure 6 illustrates that the suggested model exhibits outstanding and equitable performance across all four separate categories. The precision, recall, and F1-score for each class are repeatedly seen to be exceptionally high, nearing unity (1.000).

These outstanding results robustly indicate that the model possesses an exceptionally high degree of accuracy in its positive predictions (reflected by high precision scores) and is highly effective in comprehensively identifying all relevant instances for each class (demonstrated by high recall scores). The constantly elevated F1-scores, indicative of the harmonic mean of precision and recall, further validate the model's general robustness, equilibrium, and dependability across all categories. The negligible variation noted in these scores among the various classes strongly indicates that the model is not biased towards any specific class, an essential attribute for practical agricultural applications where the precise and timely identification of all disease types is equally vital, irrespective of their frequency in the dataset.

Confusion Matrix Analysis:

The confusion matrix is a valuable tool that offers an in-depth study of a classifier's predictions, outlining true positives, false positives, true negatives, and false negatives. It offers a detailed, cell-by-cell evaluation of the proposed model's efficacy, highlighting the sites of misclassifications. Figure 7 illustrates the confusion matrix for the

proposed model, based on its predictions from the independent test dataset.




Figure 7: Confusion Matrix of classification results on the test dataset- Limitations and Future Work.

The confusion matrix in Figure 7 visually and quantitatively confirms the exceptionally high accuracy of the proposed model, with the diagonal elements representing correctly classified instances for each class being overwhelmingly dominant across all categories. Specifically, for Black Rot, 353 instances were correctly classified, with only a single misclassification as ESCA. ESCA and Leaf Blight both achieved perfect classification, with all 360 and 323 instances correctly identified, respectively. For the Healthy class, 317 instances were accurately classified, with only one misclassified as Leaf Blight. These results show very low elements outside of the diagonal, meaning few false positives and negatives. Achieving this level of classification precision is vital in agriculture. A minimal incidence of erroneous positives helps avoid the needless application of pesticides, which in turn protects the environment and reduces operational costs. Similarly, a low rate of false negatives ensures that diseases are caught early, allowing for timely treatment to protect crop yields. The model's excellent ability to differentiate subtle disease symptoms emphasizes its applicability of being reliable and robust under real world smart agriculture systems.

Analysis of ROC Curves and AUC:

The ROC and AUC are very important tools for a full evaluation of the model's classification accuracy because they measure how well it can tell the difference between classes at all possible thresholds. In one picture, Figure 8 illustrates the ROC curves for each class alongside the microaverage curve, indicating the model's efficacy in distinguishing between them.




Figure 8: Class-wise and Micro-Average ROC Curves showing model performance.

Figure 8 presents the individual and micro-average ROC analyses for the four classes Black Rot, ESCA, Healthy, and Leaf Blight demonstrating the model's exceptional discriminatory power. The graph on the left demonstrates that each class-specific ROC curve is positioned close to the top-left corner. This positioning signifies a very high true positive rate and a near-zero false positive rate, resulting in a perfect AUC score of 1.000 for each category and indicating flawless class separation. On the right, the micro-average ROC curve aggregates true and false positives across all classes, follows the top-left boundary, and yields a Micro-average AUC of 1.000. Together, these results underscore the proposed model's robust, near-perfect ability to identify and differentiate grapevine leaf conditions, validating its reliability for practical applications in smart agriculture.

Grad-CAM visualization:

To guarantee that our proposed model predicts using pertinent features, we utilised Gradient-weighted Class Activation Mapping (Grad-CAM), a method that offers a visual elucidation of the model's decision-making process by producing heatmaps that emphasise the most significant areas in input images for particular predictions. Understanding why the model classifies a leaf in a certain way is crucial for building trust and verifying that it has learned correct pathological features rather than relying on irrelevant background artifacts. Figure 9 presents the Grad-CAM visualization for Black rot classification, showing the source specimen, class activation map, and pathological analysis. The analysis reveals that for Black rot, the heatmap correctly localizes the dark, circular lesions characteristic of the disease, with the model's attention concentrated directly on the symptomatic areas, confirming it has learned to identify these specific visual cues rather than background artifacts. The visualization demonstrates that the model focuses precisely on the pathological features the distinctive brown, circular spots scattered across the leaf surface that are the primary indicators of Black rot disease. In summary, the Grad-CAM results demonstrate that the trained model is not a "black box" but has successfully learned to identify and localize relevant, disease-specific features, providing high interpretability that increases trust in the model's predictive accuracy and its potential for reliable, real world deployment in smart agriculture.




Figure 9: Grad-CAM visualization.

Mobile Application Design And Development:

To further examine the generalizability and practicality of the proposed CNN model beyond offline testing, a dedicated mobile application named Grape was developed for real-time grapevine disease detection in vineyard environments. The app was designed as part of an end-to-end deployment pipeline, enabling direct evaluation of the model under realistic agricultural conditions.



Figure 10: End-to-End Deployment Pipeline From Research To Field App.

The deployment process began with model conversion, where the trained CNN model was exported into TensorFlow Lite format using the tf.lite.TFLiteConverter framework. This lightweight representation significantly reduced model size and optimized inference speed while preserving classification accuracy, making it suitable for execution on resource-limited mobile devices.

Following model optimization, the application was implemented using Flutter and Dart to ensure cross-platform compatibility. The mobile app integrates a camera interface that allows farmers and specialists to capture grapevine leaf images directly in the field. Once an image is captured, the TensorFlow Lite model performs on-device inference to classify the leaf as healthy or affected by Black Rot, ESCA, or Leaf Blight. The results are immediately displayed on the screen along with additional information describing disease symptoms and suggested management practices, thereby providing instant decision support to end users.

To validate the app in practice, we randomly obtained five grapevine leaf images from Google and tested the mobile application on them, as shown in Figure 11. The app correctly predicted the

disease status for all five samples, further confirming the robustness of the model when applied to real world images outside the training dataset. This preliminary experiment demonstrated the capability of the mobile deployment to maintain the accuracy of the original model even under variable conditions and sources.

The application was deployed on both Android and iOS devices to maximize accessibility. To further improve usability for the target community, the interface was designed with simple navigation and support for local languages. Figure 10 presents screenshots of the mobile application, demonstrating its key functionalities including image acquisition, automated classification, and disease information display.




Figure 11: Google Images as Tested by Mobile Application.

5. DISCUSSION

Discussion of Results:

An ablation study was systematically carried out to isolate and quantify the contribution of each primary component to our network’s superior performance. This analysis involves selectively removing specific architectural elements or data processing techniques from the full model and evaluating the resulting performance degradation. This analysis highlights that the combined use of residual blocks, data augmentation, batch normalization, and dropout layers was crucial for attaining the reported accuracy and robustness, with dropout layers playing an indispensable role rather than a coincidental one. The findings of this ablation study are detailed in Table 5, which illustrates how the removal of specific components affected the model’s accuracy, precision, recall, and F1-score.

Table 5: Results of the ablation study demonstrating the impact of removing or modifying key components on the model’s performance metrics.

Experiment Setup		Removed Component(s)		Accuracy (%)	Precision	Recall	F1-Score	Observation / Impact
Baseline (Full Model)		None (original model)		0.998	0.997	0.997	0.996	Highest performance, stable training/validation.
	Without Residual Blocks	Residual Blocks 2 and 3		0.984	0.984	0.983	0.983	Notable accuracy drop indicates the importance of residual connections for robust feature learning in deeper networks.
Without Data Augmentation		All augmentation techniques		0.986	0.986	0.986	0.986	Performance decrease, suggesting increased overfitting and reduced generalization to unseen data without diverse augmentation.
SGD optimizer		Replaced with Adam		0.792	0.8466	0.8013	0.792	Substantial performance drop indicates Adam’s superiority for this task and faster convergence.
Batch Normalization (Batch 64)		Changed from 32 to 64		0.9918	0.991	0.992	0.991	Slight performance decrease compared to optimal batch size (32), indicating sensitivity to batch size for peak performance.

The experimental results confirm the proposed deep neural network’s outstanding performance in classifying grapevine diseases. The model exhibited high accuracy, with rapid and stable convergence of the training and validation curves. It effectively learned distinctive pathological features while generalizing remarkably well to unseen data. The small gap between training and validation performance demonstrates the efficacy of the integrated regularization techniques specifically batch normalization and dropout and the comprehensive data augmentation strategy in supporting strong generalization.

The model demonstrated strong and consistent performance across every category, attaining nearly perfect precision, recall, and F1-scores (close to 1.000) for the Black Rot, ESCA, Leaf Blight, and healthy classes. This consistency is critical for real world deployment, ensuring reliable differentiation between diseased and healthy leaves.

Analysis of the confusion matrix revealed exceptionally few misclassifications, underscoring the model’s potent discriminatory power. The extremely low rates of false positives and negatives are vital for agricultural applications; they help prevent unnecessary pesticide use while ensuring timely disease detection to mitigate crop loss, promoting sustainable and cost-effective practices.

Finally, the perfect Area Under the ROC Curve (AUC) scores of 1.000 for all classes signify an ideal capacity to separate the different leaf conditions within this Dataset. This result confirms the model’s robust classification ability and establishes the high reliability of its predictions for use in autonomous smart agriculture systems.

Comparative Analysis with State-of-the-Art:

To benchmark the efficacy of our proposed model, its performance is compared against several recent state-of-the-art approaches in grapevine disease detection. While direct comparisons can be challenging due to variations in datasets, computational environments, and specific disease types, Table 6 provides a summarized overview of our model’s performance about published works.

As evident from Table 6, our proposed model consistently achieves superior performance metrics, particularly in overall accuracy and AUC, compared to other robust deep learning approaches in similar contexts. This exceptional performance can be attributed to the combined strength of our custom residual CNN architecture’s ability to effectively extract nuanced, deep features and the sophisticated, biologically informed data augmentation strategy, which collectively enhance generalization and ensure robust performance. Many comparative studies rely solely on transfer learning from pre-trained ImageNet models with less tailored augmentation or simpler architectures. This may not capture the specific pathological nuances and inter-class differences as effectively as our bespoke design. The highly consistent metrics across all classes also indicate a robust and unbiased classification ability, which is critical for practical deployment. This research offers a highly robust, accurate and dependable automated approach for the diagnosis of grapevine diseases, paving a significant pathway towards more efficient, sustainable, and proactive vineyard management practices, ultimately benefiting crop yield and quality globally.

While the proposed model achieves near-perfect accuracy, precision, recall, and F1-scores, it is important to note that these results

are based on a single public dataset. Thus, there is still risk of overfitting, where the model may capture dataset-specific features rather than generalizable disease indicators. Although augmentation and regularization were employed to mitigate this risk, broader validation using multiple datasets, field-collected images, and cross-validation remains necessary to ensure real-world reliability.

Limitations and Future Work:

The study’s dependence on a single public dataset and a fixed train/validation/test split restricts its generalizability. To overcome this, comprehensive k-fold cross-validation (both 5 and 10-fold) will be carried out, along with formal statistical analyses (e.g., McNemar’s test or bootstrap hypothesis testing) when comparing models. Broader generalization will be assessed using multiple independent public datasets, as well as field-collected images that reflect a variety of cultivars, imaging devices, and environmental conditions, with per-class performance measured under domain shifts. Different foreground-segmentation methods (such as color-threshold masking, GrabCut, and learned models like U-Net or Mask R-CNN) will be tested, and ablation studies will evaluate the role of background removal. To address dataset bias, domain-adaptation and synthetic-data approaches (adversarial domain adaptation, style transfer, and GAN-based augmentation) will be explored. Furthermore, uncertainty estimation (via MC-dropout and deep ensembles) and interpretability techniques (Grad-CAM, SHAP) will be included to enhance reliability. Finally, model optimization strategies (pruning, quantization, knowledge distillation) will be examined for use on resource-limited devices, with latency and memory usage reported on representative edge hardware.

6. CONCLUSION

Fungal infections are a major and expanding global health concern, because of their diversity, changing pathogenic patterns, and growing resistance to existing antifungal therapies. To improve diagnosis, treatment, and prevention strategies, a thorough understanding of fungal biology, virulence factors, and host interactions is essential. Advances in antifungal drugs, vaccines, and diagnostic tools offer promising directions. Vaccines could be an effective approach to reduce the impact of fungal diseases but complex fungal biology, immune response against fungal infections, diagnostic barriers and similarity to human host are current challenges in development of fungal vaccines. So, there is a dire need to work on emerging fungal pathogens, resistance patterns, and development of diagnostic tools. Scientists must use different approaches to develop fungal vaccines particularly against diseases having high burden and mortality, keeping in view the challenges associated with fungal vaccines to lessen the burden of medical mycology.

Acknowledgment:

The authors acknowledge all individuals and institutions that supported the completion of this study.

Table 6: Comparative analysis of the proposed model’s performance against recent state-of-the-art plant and grapevine disease detection methods.

Author (Year)	Model/Approach	Accuracy (%)	F1-score	Recall	Precision
(Zhe Tang <i>et al.</i> , 2020)	Lightweight CNN (ShuffleNet with CA)	99.14	99.14	99.14	99.15
(Saber Anari, 2022)	Hybrid ResNet-18 + SVM ensemble	99.1	99.05	98.9	99.2
(Liu <i>et al.</i> , 2020)	DICNN (Inception & Dense Connectivity)	97.22	97.20	97.19	97.21
(Lin <i>et al.</i> , 2022)	GrapeNet	86.29	79.05	88.43	77.7
(Zinonos <i>et al.</i> , 2022)	Fine-tuned ResNet50 with LoRa	99.77	98.5	98.5	98.5
(Thet <i>et al.</i> , 2020)	Fine-tuned VGG16 with GAP	98.4	98.4	98.4	98.4
(Prasad <i>et al.</i> , 2024)	DCNN Classifier (VGG16-based)	99.06	98.8	99.3	99.0
(Zhang <i>et al.</i> , 2024)	DLVTNet	98.48	98.49	98.48	98.51
(Zirije Hasani <i>et al.</i> , 2025)	Transfer Learning	98.84	-	-	-
(Hu <i>et al.</i> , 2025)	GCS-YOLO	96.2	-	93.2	95.1
Proposed Method	Lightweight CNN	99.8	99.6	99.7	99.7

Ethical Statment:

This study did not involve any experiments on humans or animals and did not require ethical approval. All data used in this research were obtained from publicly available international datasets and used in accordance with the terms and conditions stated by the dataset providers.

Conflict of Interests:

The authors declare no competing interests.

Funding:

This research did not receive any specific funding from public, commercial, or non-profit organizations.

Author Contributions:

A.A.M.: conceptualization, methodology, software implementation, data curation, formal analysis, visualization, and original draft preparation. A.R.A.: supervision, validation, review and editing. All authors have read and agreed to the published version of the manuscript.

7. REFERENCES

- Aher, P. G., Sabnis, V., & Jain, J. K. (2025). Deep learning for grape leaf disease detection: A review. *Malque Publishing*. DOI:<https://doi.org/10.31893/multirev.2025364>
- El-Kenawy, A. M., Ezzat, A. A., Khedr, S. M., Ali, A. H., Abdelaziz, S. I., & El-Mashad, M. G. (2025). Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. *Artificial Intelligence Review*. DOI:<https://doi.org/10.1007/s10462-024-11100-x>
- Elsherbiny, O. (2024). Rapid Grapevine Health Diagnosis Based on Digital Imaging and Deep Learning. *Plants*, 13(1), 135. DOI:<https://doi.org/10.3390/plants13010135>
- Godfray, H. C. J., Beddington, J. R., Crute, I. R., Haddad, L., Lawrence, D., Muir, J. F., & Toulmin, C. (2010). Food security: the challenge of feeding 9 billion people. *Science*, 327(5967), 812–818. DOI:<https://doi.org/10.1126/science.1185383>
- Gupta, H., Jha, A., Baluni, A., & Suryavanshi, R. (2023). Smart Agriculture through Convolutional Neural Networks for Plant Disease Classification. *International Journal of Research - GRANTHAALAYAH*, 11(12), 1–9. DOI:<https://doi.org/10.29121/granthaalayah.v11.i12.2023.4912>
- Hu, Q., & Zhang, Y. (2025). GCS-YOLO: A Lightweight Detection Algorithm for Grape Leaf Diseases Based on Improved YOLOv8. *Applied Sciences (Switzerland)*, 15(7). DOI:<https://doi.org/10.3390/app15073910>
- Islam, M. I. S., Sultana, S., Al Farid, F., Islam, M. N., Rashid, M., & Bari, B. S. (2025). Deep Learning Approaches for Plant Disease Diagnosis Systems: A Review and Future Research Agendas. *Journal of Applied Agricultural Science and Technology*, 9(1). DOI:<https://doi.org/10.55043/jaast.v9i2.308>
- Khan, H., Khan, M. I., Hamza, M. A., Tariq, M., & Al-Gaashani, A. (2025). Enhancing plant disease detection through deep learning: a Depthwise CNN with squeeze and excitation integration and residual skip connections. *Frontiers in Plant Science*, 15, 1505857. DOI:<https://doi.org/10.3389/fpls.2024.1505857>
- Lin, J., Chen, X., Pan, R., Cao, T., Cai, J., Chen, Y., Peng, X., Cernava, T., & Zhang, X. (2022). GrapeNet: A Lightweight Convolutional Neural Network Model for Identification of Grape Leaf Diseases. *Agriculture (Switzerland)*, 12(6). DOI:<https://doi.org/10.3390/agriculture12060887>
- Liu, B., Ding, Z., Tian, L., He, D., Li, S., & Wang, H. (2020). Grape Leaf Disease Identification Using Improved Deep Convolutional Neural Networks. *Frontiers in Plant Science*, 11. DOI:<https://doi.org/10.3389/fpls.2020.01082>
- Malik, S. G., Jamil, S. S., Aziz, A., Ullah, S., Ullah, I., & Abohashrh, M. (2024). High-Precision Skin Disease Diagnosis through Deep Learning on Dermoscopic Images. *Bioengineering*, 11(9), 867. DOI:<https://doi.org/10.3390/bioengineering11090867>
- Mandal, R. (2023). Grapevine Disease Dataset (Original). *Kaggle*. DOI:<https://doi.org/10.34740/KAGGLE/DS/2928682>
- Prasad, K. V., Vaidya, H., Rajashekhar, C., Karekal, K. S., Sali, R., & Nisar, K. S. (2024). Multiclass classification of diseased grape leaf identification using deep convolutional neural network(DCNN) classifier. *Scientific Reports*, 14(1). DOI:<https://doi.org/10.1038/s41598-024-59562-x>
- Saberi Anari, M. (2022). A Hybrid Model for Leaf Diseases Classification Based on the Modified Deep Transfer Learning and Ensemble Approach for Agricultural AIoT-Based Monitoring. *Computational Intelligence and Neuroscience*, 2022. DOI:<https://doi.org/10.1155/2022/6504616>
- Shoaib, I., Shah, M., El-Sappagh, S., Ali, M., Ullah, M., Alenezi, F., & Ali, J. (2023). An advanced deep learning models-based plant disease detection: A review of recent research. *Frontiers in Plant Science*, 14, 1158933. DOI:<https://doi.org/10.3389/fpls.2023.1158933>
- Tang, Z., Yang, J., Li, Z., & Qi, F. (2020). Grape disease image classification based on lightweight convolution neural networks and channelwise attention. *IEEE Access*, 8, 102789–102801. DOI:<https://doi.org/10.1109/ACCESS.2020.2997382>
- Tang, Zhe, Yang, J., Li, Z., & Qi, F. (2020). Grape disease image classification based on lightweight convolution neural networks and channelwise attention. *Computers and Electronics in Agriculture*, 178. DOI:<https://doi.org/10.1016/j.compag.2020.105735>
- Thet, K. Z., Htwe, K. K., & Thein, M. M. (2020). Grape Leaf Diseases Classification using Convolutional Neural Network. *Proceedings of the 4th International Conference on Advanced Information Technologies, ICAIT 2020*, 147–152. DOI:<https://doi.org/10.1109/ICAIT51105.2020.9261801>
- Zhang, E., Qi, G., Li, F., Lv, C., & Zhang, N. (2024). Lightweight Grape Leaf Disease Recognition Method Based on Transformer Framework. DOI:<https://doi.org/10.21203/rs.3.rs-4858639/v1>
- Zhong, Y., & Zhao, M. (2020). Research on deep learning in apple leaf disease recognition. *Computers and Electronics in Agriculture*, 168, 105146. DOI:<https://doi.org/10.1016/j.compag.2019.105146>
- Zinonos, Z., Gkelios, S., Khalifeh, A. F., Hadjimitsis, D. G., Boutalis, Y. S., & Chatzichristofis, S. A. (2022). Grape Leaf Diseases Identification System Using Convolutional Neural Networks and LoRa Technology. *IEEE Access*, 10, 122–133. DOI:<https://doi.org/10.1109/ACCESS.2021.3138050>