**Islamic University Journal of Applied Sciences (IUJAS)**
https://journals.iu.edu.sa/jesc

Volume VIII, Issue I, July 2026, Pages 189-217

# Explainable Deep Learning Models for Citrus Leaf Disease Identification and Agricultural Decision Support

Sami Saad Albouq [1] *

[1] Department of Computer and Information Systems, Islamic University of Madinah, 42351, Medina, Saudi Arabia, salbouq1@iu.edu.sa

*Corresponding author: (Sami Saad Albouq), *Email Address*: salbouq1@iu.edu.sa

## Abstract

Citrus is a commercially important crop affected by diseases such as anthracnose, bacterial blight, and citrus canker, which damage trees and reduce fruit yield. Classical diagnosis is time-consuming and expert-dependent, limiting scalability. In this paper, we propose explainable deep learning models for multiclass citrus leaf disease detection, incorporating custom pre-processing and extensive data augmentation for robustness. A custom CNN and transfer learning-based VGG16 were fine-tuned on an augmented dataset of nine citrus leaf categories, including both healthy and unhealthy specimens. Performance was evaluated using accuracy, precision, recall, F1-score, and confusion matrices. The CNN achieved 97% accuracy, outperforming VGG16's 95%, demonstrating the effectiveness of our preprocessing approach. For interpretability, Grad-CAM, SHAP, and LIME were uniformly applied to highlight symptom-relevant regions aligned with agronomic knowledge, enhancing transparency and user trust. The framework effectively combines predictability and interpretability, supporting practical, real-time, and sustainable agricultural decision systems.

**Keywords:** Citrus Leaf Diseases; Deep Learning; CNN, VGG16; Explainable AI; Grad-CAM; Agricultural Decision Support.

https://doi.org/10.63070/jesc.2026.023
Received 25 December 2025; 15 February 2026; Accepted 10 April 2026.
Available online 15 April 2026.
Published by Islamic University of Madinah on behalf of *Islamic University Journal of Applied Sciences*.
This is an open access article published under the Creative Commons Attribution 4.0 International license (CC BY 4.0):https://creativecommons.org/licenses/by/4.0/

## 1. Introduction

Citrus belongs to one of the most significant fruit crops, cultivated worldwide with high economic, nutritional, and cultural value. However, its productivity is continually compromised by numerous leaf pathologies, such as anthracnose, bacterial blight, citrus canker, and curvularia virus, among others, which cause drastic losses in productivity and quality [1]. These diseases, in addition to having a significant impact on local food security, also have a substantial impact on international commerce. The existing methods of diagnostic procedures in most instances rely on trained eyesight and lab tests, which are time-consuming, labor-intensive and prone to errors in the agricultural context. Hence, the intelligent prediction model that is scalable, reliable, and achieves early screening and intervention is greatly in need [2].

With the development of deep learning methods, plant disease recognition has also been revolutionised to achieve excellent accuracy in image-based classification. CNNs and transfer learning architectures, such as VGG16, have yielded excellent results on agricultural datasets. Upcoming networks, such as Vision Transformers, continue to challenge the limits of accuracy and the feasibility of heavy computation [3]. Meanwhile, as powerful predictive models, they function as a “black box”, making their decisions completely obscure. This black box challenge results in a trust deficit, as stakeholders in agriculture do not only need accurate predictions but also human-understandable explanations to make informed disease management decisions [4].

Explainable artificial intelligence (XAI) is being employed to address this issue and enhance transparency and trust in deep learning. Gradient-weighted Class Activation Mapping (Grad-CAM), SHapley Additive exPlanations (SHAP), and Local Interpretable Model-agnostic Explanations (LIME) are some of the techniques aimed at providing visual clues to the areas within an image that influence model decisions [5]. Such approaches, when incorporated into citrus leaf disease detection systems, not only reveal the principles underlying the system's decision but also identify patterns of symptoms that are relevant and in agreement with agronomic intuition [6]. This is the connection that is lacking between predictive performance and practical usability, which favours AI-driven systems that can be interpreted, are reliable, and suitable for in-field agricultural decision support. In this research, the following objectives are planned:

- Develop robust CNN and VGG-16-Based models to detect multi-class citrus leaf diseases effectively.
- Implement customized preprocessing and augmentation techniques to enhance the models' generalization capabilities.
- Integrate explainability methods such as Grad-CAM, SHAP, and LIME into the detection pipeline for improved visual interpretation of model predictions.
- Evaluate models based on accuracy, precision, recall, F1-score, and confusion matrices.
- Provide a decision-support framework that translates AI outcomes into easily understandable insights for farmers and policymakers.

The contributions of this study are clarified as follows:

- An explainable citrus leaf disease classification framework is developed by jointly integrating Grad-CAM, SHAP, and LIME with CNN and VGG16 models.
- A lightweight custom CNN (0.33M parameters) is benchmarked against VGG16 under identical training conditions, achieving 97% accuracy with reduced computational cost.
- The study demonstrates that robust preprocessing and multi-XAI validation enable RGB-based models to perform competitively with complex YOLO and hyperspectral baselines.

The rest of this paper follows the following structure. Section 2 conducts a literature review of the gaps in citrus disease detection and interpretability. Section 3 describes the preparation of datasets, preprocessing, and augmentation. The suggested CNN and VGG16 models, which include Grad-CAM, SHAP, and LIME, are depicted. In Section 4, the experimental outcomes are discussed, including performance measures and visual interpretations. Section 5 discusses about the implications of explainable AI on agricultural decision support, and the conclusion is made in Section 6.

## 2. Related Work

### 2.1 Plant Disease Detection: Traditional ML vs Deep Learning

Wang et al. [7] introduced a citrus pest and disease detection system with deep learning models based on YOLOv3 and YOLOv4 and trained on field images with high detection rates of Huanglongbing, citrus canker, and healthy leaves. The framework had a great potential of edge based, real-time monitoring of the orchards. Nonetheless, the method is mainly concerned with the performance of detection of objects and lacks insight that can be explained in regard to disease-specific visual characteristics. Contrastingly, the suggested study highlights the image-level classification of disease along with multi-method explainable AI to aid in greater transparency and diagnostic explainability.

Sujatha et al. [8] proposed an automatic leaf plant disease detection system based on combination of deep learning and machine learning. The research utilised CNN models in the feature extraction process, specifically VGG19 and Inception v3, along with classifiers such as SVM and kNN, to enhance the results. Performance on the Banana, Custard Apple, Fig, and Potato leaf datasets varied: the highest accuracy was obtained for the Custard Apple dataset, at 99.1%. This study demonstrates the flexibility of introducing hybrid AI models to address various types of plant disease detection problems.

### 2.2 Deep Learning Architectures in Agriculture: CNNs and VGG16

CNNs have become the cornerstone of plant disease detection as they capture spatial hierarchies in images. In the meantime, among the CNN-based models, VGG16 is a relatively heavyweight architecture compared to the shallower solutions, which, however, is simple enough to provide strong generalisation in the field of agriculture. Various works support the idea that VGG 16 outperforms shallower CNNs for classification tasks,

albeit at a drastically higher computational load. Generally, multiple researchers use CNN to obtain a low FLOPs baseline, and VGG to implement transfer learning backend. As a result, they are the most popular models in agricultural image analysis.

Dhiman et al. [9] proposed an intelligent citrus detection system that integrates CNN and LSTM models, combined with edge computing. The approach utilises better feature extraction, downsampling, and feature fusion to fit the best models on limited-resource devices. The model was tested on the Kaggle and Plant Village datasets, which included 2,950 citrus fruit images, using pruning and post-quantisation methods. Experimental results demonstrated better accuracy, with 97.18% for pruning and 98.25% for pruning plus post-quantization compared with previous CNN-based techniques.

### 2.3 Explainable AI (XAI) in Limited Agricultural Use

Even though CNNs and VGG16 are accurate, they are frequently referred to as black boxes, which do not remain very interpretable. To address the above issue, various XAI technologies, including Grad-CAM, LIME, and SHAP, have emerged based on the interpretation of model predictions. These techniques are commonly used in the medical and industrial sectors, but are scarcely utilised in agriculture. Current agricultural research often falls short in terms of accuracy and lacks explainability, which undermines farmer trust. Therefore, the instantiation of XAI in CNN and VGG16 predictions presents a new trend in detecting plant diseases.

Kondaveeti and Simhadri [10] proposed a three-stage approach to assess DLMs in rice leaf disease detection through explainable AI. Conventional metrics were combined with a qualitative and quantitative XAI-based analysis, which included LIME visualisations, as well as IoU, DSC, and a new overfitting ratio. ResNet50 was the topmost performing model with a classification accuracy of (99.13%) and higher reliability for feature selection from eight evaluated models. The novel approach enhances the transparency and credibility of the model, both of which have broader applications beyond agriculture.

### 2.4 Beyond Conventional CNNs: Detection, Biochemical, and Graph-Based Methods

Xu et al. [11] focused on detecting Huanglongbing (HLB) infection in citrus plants through the application of volatile organic compound (VOC) profiling with GC 13 MS, GC 13 IMS, and OPLS-DA modelling. Their experiment showed extremely high predictability among various citrus species, which showed that VOC-biochemical signatures are highly discriminating in the early detection of HLB. Nonetheless, the method is based on specific laboratory equipment and chemical analysis, which restricts its applicability and scalability in the field conditions in real-time. In addition, the approach lacks image learning and explainability of visual disease appraisal.

Guo et al. [12] suggested YOLOv8n-DE, a lightweight object detection model with dilated reparameterized feature modules and shared parameter detection head to identify citrus disease. The model offered high detection accuracy and highly decreased the number of parameters and computational cost therefore it is applicable in orchard deployment. Regardless of these benefits, the approach puts emphasis on detection that uses bounding-

boxes and fails to offer model interpretability or pixel-level description of disease symptoms, which becomes more and more significant in precision agriculture to motivate trust and decision support.

Aswini and Vijayakumaran [13] proposed a citrus disease classification system based on a graph neural network that combines feature extraction through FA-ResNet, gated fusion adaptive graph learning with a bio-inspired optimization algorithm. Their hybrid design realised significant performance increase over conventional CNN-based classifiers especially in complicated texture and light conditions. Nevertheless, multiple tightly coupled stages, such as segmentation, graph construction, and evolutionary optimization, take part in the model, and it leads to the complexity of the system and low transparency. Also, clear explainable AI (XAI) evaluation to decode specific visual clues on diseases was not mentioned.

### 2.5 CNN-Based Plant Disease Detection with Explainable AI

Recent studies demonstrated efficiency of convolutional neural networks in classifying plant and leaf diseases with the help of RGB images. Some of these works use such transfer-learning backbones as VGG, ResNet, and MobileNet to enhance classification accuracy in limited data conditions. Nevertheless, a large number of these methods mainly involve predictive performance, and offer little insight on model decisions.

To overcome this shortcoming, a recent literature on agricultural algorithms has embraced the explainable AI methods, which are mostly Grad-CAM, to visualize disease-relevant regions in leaf images. Although these studies can help make qualitative conclusions on CNN attention patterns, the explainability of their results is commonly limited to analysis of one method and a small sample. Contrary to that, the suggested framework combines several complementary XAI methods (Grad-CAM, SHAP, and LIME) and implements them uniformly in both a custom CNN and a transfer-learning model, which allows obtaining a more holistic and trustworthy view of disease-specific characteristics.

Additionally , in contrast to detection- or hyperspectral-based systems, which rely on complicated architectures or specialized sensors, the suggested design can compete with the others with minimal technical overhead of the standard RGB images and transparent priorities on the reliability of the model and support of real-world agricultural decision making.

### 2.6 Identified Gaps

The literature reveals a few gaps important to this study. First, although CNNs and VGG16 have achieved impressive results, the black-box nature of these models hinders their trustworthiness and applicability in the agricultural sector. Second, the methods of explainability have not been significantly explored in the context of plant diseases, which reduces the relationship between prediction and practical farming activities. Third, there are limited works that combine outputs of AI with real-world applications to assist growers in making decisions. To resolve these problems, CNN and VGG16 need to be employed in conjunction with XAI tools, so that the high accuracy is associated with understandable and farmer-friendly recommendations. Table 1 consolidates recent advances in citrus disease detection, underscoring both high accuracy and the emerging need for explainable, decision-oriented models.

Table 1: Recent Advances in Explainable Citrus Leaf Disease Detection

| Authors / Year | Methodology / Models | Explainability / Analysis | Performance (Accuracy / Metrics) | Key Contribution |
|---|---|---|---|---|
| **Zhu et al., 2024 [14]** | Multi-Model Fusion Network (MMFN) combining AlexNet, VGG, ResNet, Pre-ResNet34, Pre-ResNet50 with Transfer Learning | Comparative model evaluation, transfer learning analysis | 99.72% (healthy vs diseased), 98.68% (multi-class diseases) | Developed MMFN with transfer learning achieving state-of-the-art citrus leaf disease classification |
| **Kaur et al., 2025 [15]** | Nonlinear Fuzzy Rank-Based Ensemble (NL-FuRBE) combining VGG19, AlexNet, Xception with image enhancement (VAD, morphological filtering) | Image quality metrics (PSNR, SSIM, NIQE); ensemble explainability via fuzzy ranking | Avg. accuracy 96.51% (5-fold CV) | Introduced NL-FuRBE with enhanced image preprocessing and fuzzy ensemble, improving robustness |
| **Mo & Wei, 2024** [16] | YOLOv8n-CDDA with Cross-Domain Dynamic Attention (CDDA), ARMS augmentation, pruning + knowledge distillation | Grad-CAM++ for feature visualization | Accuracy 90.89%, Recall 81.12%, mAP50 88.36% | Proposed lightweight YOLOv8n-CDDA for real-time citrus disease detection on edge devices |
| **Min et al., 2024** [17] | Hyperspectral Reflectance + Fluorescence Imaging with ML classifiers (SVM, Discriminant Analysis, PCA feature selection) | Hyperspectral band analysis for interpretability | Accuracy up to 94.5% (SVM & DA classifiers) | Combined hyperspectral reflectance & fluorescence imaging for high-precision citrus disease detection |

## 3. Proposed Methodology

The methodology of this study is structured in a way that explainable deep learning models are created and evaluated to classify various lemon leaf diseases precisely. A comprehensive workflow was constructed, encompassing dataset pre-processing and augmentation for robustness, as well as the design of both a specific CNN model and a transfer learning strategy utilising VGG16. The models were trained with optimal hyperparameters, and the evaluation metrics were carefully chosen to measure performance at both global and class levels. In addition, Explainable AI tools such as Grad-CAM and SHAP were incorporated to improve interpretability by generating visual explanations consistent with agronomic knowledge.

Figure 1 provides a structured methodology for the present study. Using nine classes of lemon leaf diseases images, the pre-processing and augmentations techniques like (normalization, rotation, shift shear zoom flip etc) were applied to obtain a balanced dataset. Two model structures were used as follows: one is a customised CNN with convolution, pooling, dense, and dropout layers, the other is transfer learning based on the VGG16 network with a refining dense layer. Training and validation were performed using the Adam optimiser with Categorical cross-entropy loss, over 30 epochs, and a batch size of 30. Assessment was conducted using accuracy, precision, recall, F1-score, confusion matrices, ROC, and AUC. To enhance interpretability,

explainable AI methods such as Grad-CAM, LIME, and SHAP were incorporated to view results classified diseases with visual explanations and decision-support insights for agriculture.

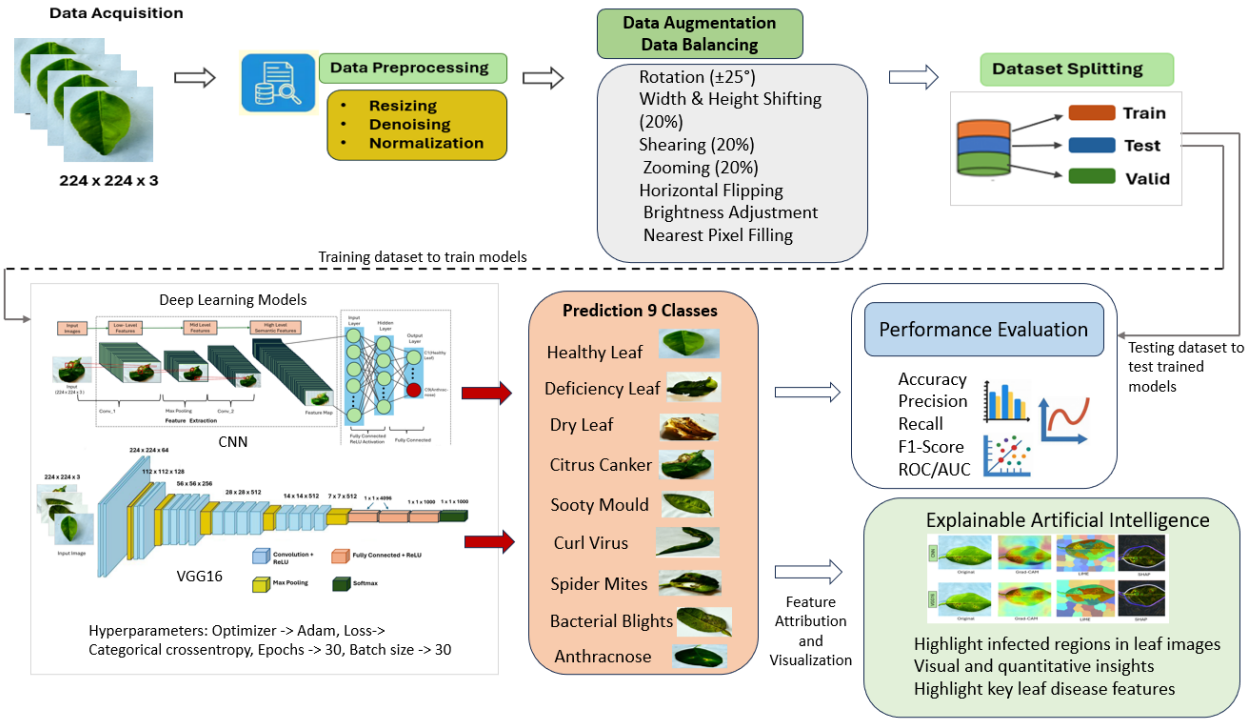



Figure 1: Proposed Methodology for Explainable Citrus Leaf Disease Detection

### 3.1 Dataset Overview

In this work, we utilise the Comprehensive Lemon Leaf Disease Dataset for advanced detection and sustainable agriculture as the reference citrus dataset benchmark for computer vision-based disease recognition. This dataset comprises 1,354 high-resolution raw images (800×800 pixels) of a lemon leaf, captured in Melandaha, Jamalpur, Bangladesh, from July to September 2024. It includes images of nine clearly defined categories — anthracnose, bacterial blight, citrus canker, curl virus deficiency leaf, dry leaf, healthy leaf, sooty mould, spider mites — annotated by agricultural experts to ensure accuracy and consistency in classification. These original images retain natural differences in lighting, orientation, and condition of the leaves, resulting in a dataset that reflects true field conditions.

Figure 2 presents the example images of nine categories from the lemon leaf dataset, which contain both healthy and diseased leaves. The diseases include common illnesses such as Anthracnose, Bacterial Blight, Citrus Canker, Curl Virus, Deficiency Leaf, Dry Leaf, and others, including Sooty Mould and Spider Mites. These visual examples illustrate the diversity in the dataset, which is crucial for training robust deep learning models.




Figure 2: Representative Samples of Lemon Leaf Disease Dataset

### 3.2 Data Preprocessing & Augmentation

All images were normalized and rescaled to a consistent range of pixel intensities, allowing for the convergence of models during training. We employed various augmentation strategies, including random rotation, zoom, width and height shifts, shear, and horizontal flip, to enhance the model's robustness. Although the dataset had a fair proportion of classes, augmentation was applied to increase the intra-class variation, enabling the models to generalise more effectively to unseen conditions (i.e., leaf orientation, lighting, and background variation). Not only does such a preprocessing pipeline alleviate the overfitting problem, but it also enhances the model's ability to cope with real-life agricultural situations where highly non-uniform image acquisition conditions are encountered. The dataset was further split into 80% for training and 20 % for testing for model training and evaluation.

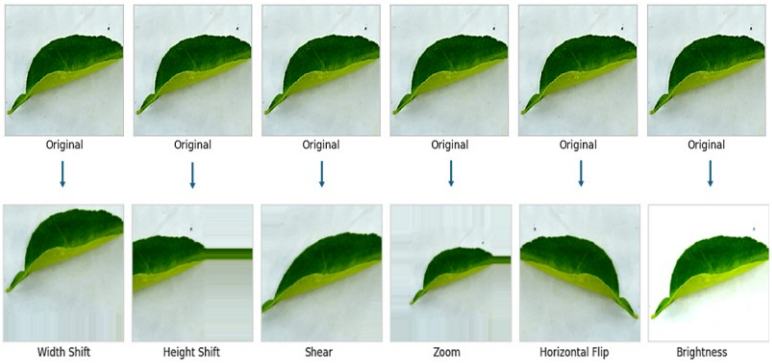



Figure 3: Data Augmentation Techniques for Lemon Leaf Images

Figure 3 illustrates the effectiveness of different types of augmentation employed to diversify dataset and improve model generalization. From the lemon leaf image, several transformations were applied, including width shift, height shift, shear, zoom range (zoom), horizontal flip, and brightness change. Such augmentations mimic possible natural changes in orientation, lighting and scale of the images to ensure that the deep learning models are robust when detecting citrus leaf diseases under different conditions.

Table 2: Image Augmentation Parameters for Lemon Leaf Dataset

| Augmentation Technique | Applied Range / Value | Purpose |
|---|---|---|
| **Rescaling** | 1./255 | Normalize pixel values to [0,1] |
| **Rotation** | ±25° | Simulate natural leaf orientation changes |
| **Width Shift** | ±20% (0.2) | Handle horizontal displacement |
| **Height Shift** | ±20% (0.2) | Handle vertical displacement |
| **Shear Transformation** | 0.2 | Distort images for geometric robustness |
| **Zoom** | 0.2 | Simulate different distances to camera |
| **Horizontal Flip** | True | Capture symmetry in leaf orientation |
| **Validation Split** | 20% | Ensure reliable evaluation |
| **Brightness Adjustment** | Range = [0.8, 1.2] | Vary lighting to mimic real-world scenarios |

Table 2 summarises the augmentation methods employed to enhance model generalisation and robustness. All transformations, including rotation, shift, shear, zoom, flipping, and brightness, were properly parameterised to model the actual variation observed in leaf images. Such methods helped create a dataset that represents various states, thereby enhancing the trustworthiness of deep learning models in disease recognition.

Augmentation ranges listed in Table 2 fall within commonly used ImageDataGenerator configurations, but were chosen deliberately because of the size, the class imbalance, and the visual diversity of the citrus leaf data instead of being chosen randomly. The dataset also has moderate variations on leaf orientation, scale, illumination and background conditions hence conservative augmentation ranges were emulated to avoid introducing unrealistic variations which may affect disease texture pattern. Early trials with augmentation levels that were stronger led to lower validation performance and the final parameters are a trade-off between the diversity of the data and retention of features. The augmentation strategy was therefore informed by the dataset characteristics and empirical validation and not arbitrary defaults.

### 3.3 Model Architectures

#### 3.3.1 Convolutional Neural Network (CNN)

For our baseline we used a custom architecture CNN which designed for the task of multi-class lemon leaf disease classification. The architecture consisted of a series of convolutional layers with ReLU activation, followed by max-pooling layers to reduce spatial dimensionality while preserving discriminative information [18]. A flattening layer was inserted between the convolutional backbone and fully connected dense layers, although dropout was employed to prevent overfitting. The model was trained using the Adam optimiser with a learning rate of 1e-4, a batch size of 32, and trained for a maximum of 30 epochs with validation checks. This CNN served as a baseline for the comparison to more sophisticated transfer learning models [19].

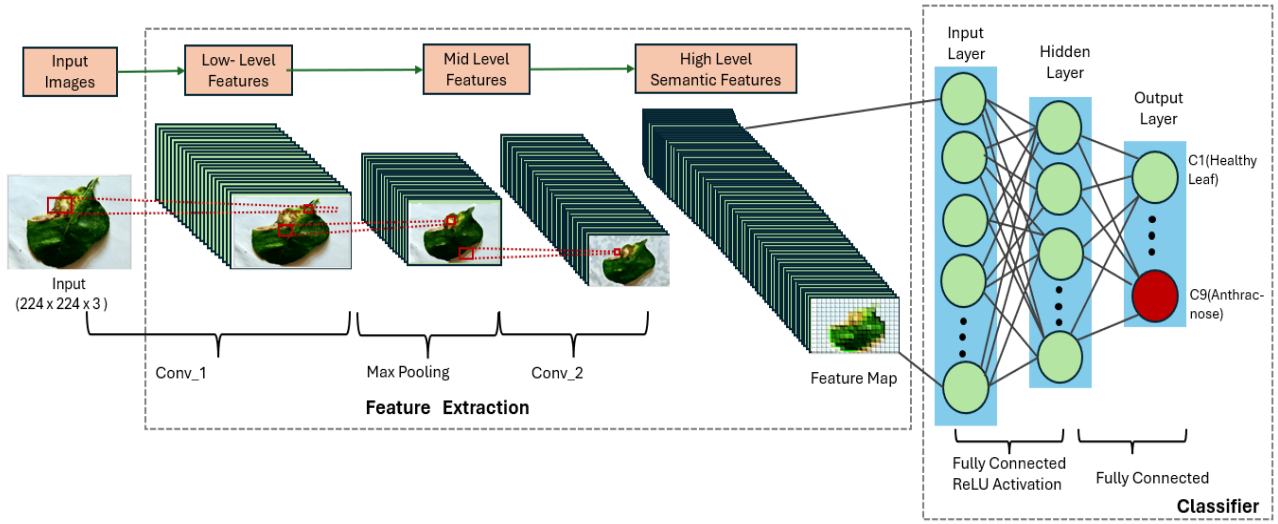



Figure 4: Custom CNN Architecture for Lemon Leaf Disease Detection

The architecture of the specific CNN used in this work is shown in Figure 4. The model leverages hierarchical features from input images through convolutional and pooling operations, capturing different patterns (low-level, mid-level, and high-level). These feature maps are then flattened and fed through two fully connected layers with ReLU activation, which classifies the presence of diseased leaf into nine categories (including both healthy and infected leaves).

### 3.3.2 Transfer Learning with VGG16

To utilise the potential of pre-trained models, we implemented VGG-16 as our backbone network, initialised with ImageNet weights. The convolutional base was partially frozen to retain low-level, generic features; however, the top of the CNN architecture was fine-tuned for domain-specific classification of citrus leaf diseases. On top of the backbone, a Global Average Pooling layer was added, followed by two fully connected dense layers with ReLU activation and dropout for regularisation, and finally a softmax output layer for the nine disease classes.

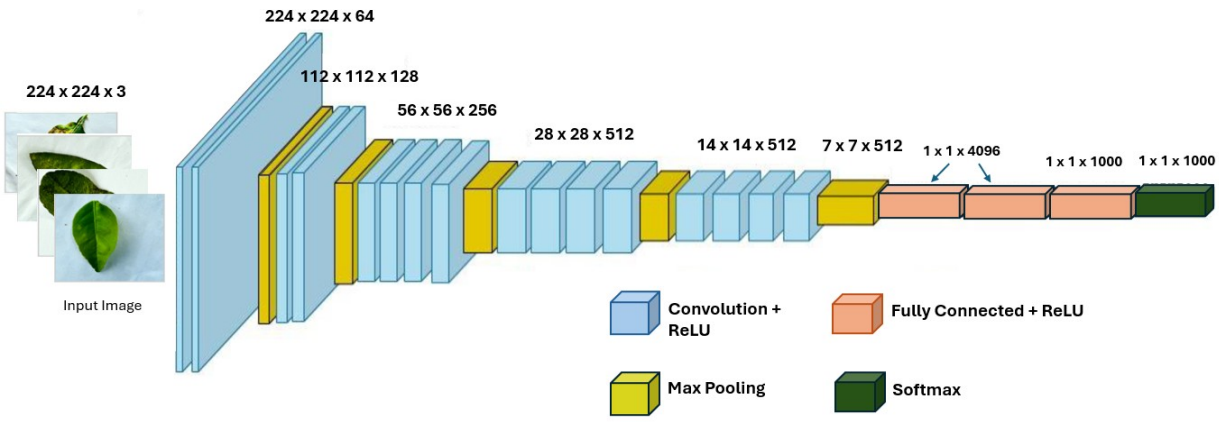



Figure 5: VGG16 Transfer Learning Architecture for Lemon Leaf Disease Detection

This hybrid transfer-learning and fine-tuning method combines the advantages of transfer learning and task-specific fine-tuning, resulting in improved classification performance in both seen and unseen datasets [20].

The VGG-16 transfer learning model used during the experiment is illustrated in Figure 5. Trained on ImageNet, the network utilises several stacked convolutional layers and max pooling to extract complex semantic features. Fully connected layers with ReLU and softmax activation are tuned for nine-class classification of lemon leaf diseases. The model can leverage transfer learning in the form of fine-tuning from pre-trained weights.

Table 3: Hyper Parameters settings for models.

| Parameter | CNN (Proposed) | VGG16 (Transfer Learning) |
|---|---|---|
| Epochs | 30 | 30 |
| Batch Size | 32 | 32 |
| Total Parameters | ≈ 0.33M | ≈ 14.9M |
| Trainable Parameters | ≈ 0.11M | ≈ 0.067M |
| Non-Trainable Parameters | 0 | ≈ 14.7M |
| Learning Rate | 0.0001 | 0.0001 |
| Loss Function | Categorical Crossentropy | Categorical Crossentropy |
| Optimizer | Adam | Adam |

The training configurations of the CNN and VGG16 models are listed in Table 3, demonstrating their divergence in terms of complexity and scale. Although CNN is lightweight (≈0.33M parameters), VGG16 is much deeper (14.9M parameters), and both networks used the Adam optimiser with similar hyperparameters for a fair comparison. Layer-wise implementation for both CNN and VGG16 as follows. The models were trained with the categorical crossentropy loss function together with a softmax output layer since it is a multi-class classification task and the category is mutually exclusive.

**1) Convolution Layer:** Filters are used to capture the local spatial information of the input images. Learnable filter applied to input patches for capturing local features. As shown in (Eq. 1), each output pixel is obtained by a cross-correlation of the input window and the filter weight. This operation is capable of maintaining spatial position and location relationships, as well as capturing hierarchical structures (edges → textures → objects).

$$\text{Output}(i,j) = m\sum n\sum \text{Input}(i+m, j+n) \cdot \text{Filter}(m,n) + \text{Bias} \quad (1)$$

**2) Batch Normalization:** Normalizes the activations to accelerate training and improve stability. Speeds train up for normalization by zero mean and unit variance. The normalization step (Eq. 2a) eliminates internal covariate shift and the scale/shift transformation (Eq. 2b) maintains representational capacity.

Given input xx, BN computes:

$$x = \frac{x-\mu}{\sqrt{a^2}-e} \quad (2a)$$

where $\mu$ and $\sigma^2$ are the mean and variance of the mini-batch. The normalized output is then scaled/shifted:

$$y = \gamma x + \beta \quad (2b)$$

$\gamma$ (scale) and $\beta$ (shift) are learnable parameters; $\epsilon\epsilon$ avoids division by zero.

**3) Clipped ReLU Layer:** Adds non-linearity and avoids activation explosion. To introduce non-linearity but avoiding unbounded activations. Compared to the standard ReLU (Eq. 3), this restriction limits the maximum value (clip_value) to prevent saturation in deep networks.

$$ClippedReLU(x) = min(max(0, x), clip_value) \quad (3)$$

With clip_value = 6 , clip_value = 6, activations are constrained to [0,6].

**4) Grouped Convolution Layer:** Computational efficiency is enhanced by factorizing the filters into groups. As derived in (Eq. 4), which independently operate on a subset of input channels, and results in 1/k parameter reduction (k is the group number). This approach delivers broader networks at the same computation cost.

$$Groupk = Conv(Inputk, Filterk) \quad (4)$$

**5) Fully Connected Layer:** **Purpose**: Projects learned features to class scores via a linear transformation (Eq. 5). The weight matrix W describes combination rules for high-level features, but modern designs now use GAP instead to reduce the number of parameters.

$$y = Wx + b \quad (5)$$

W: Weight matrix (shape: input_dim×output_diminput_dim×output_dim).

x: Flattened input features.

**6) Softmax Layer:** The logit-to-probability conversion using the exponential normalization from (Eq. 6). The denominator normalizes so that output values sum to 1 and can be interpreted as a probability. Because of the exponential function.

$$P(yi) = \frac{e^{zi}}{\sum_{j=1}^{c} e^{zj}} \quad (6)$$

C is the number of classes; outputs a probability distribution.

**7) Classification Layer:** Predicts the winning class by the argmax (Eq. 7). This is a deterministic decision rule that transforms the softmax probability distribution into a scalar prediction and then selects the class with the highest softmax probability.

$$\text{Prediction} = \text{argmax}(P(yi)) \quad (7)$$

Table 4. Custom CNN Architecture

| Layer Type | Output Shape | Parameters | Key Notes |
|---|---|---|---|
| Input | 224 × 224 × 3 | 0 | RGB image input |
| Conv2D (32) + MaxPooling | 111 × 111 × 32 | 896 | 3×3 convolution, ReLU activation |
| Conv2D (64) + MaxPooling | 54 × 54 × 64 | 18,496 | Deeper feature extraction |
| Conv2D (128) + MaxPooling | 26 × 26 × 128 | 73,856 | High-level spatial features |
| Global Average Pooling | 128 | 0 | Replaces flattening to reduce parameters |
| Dense (128) + Dropout (0.3) | 128 | 16,512 | Compact fully connected layer |
| Dense (9) + Softmax | 9 | 1,161 | Multi-class classification |
| Total Parameters | — | 332,765 | Lightweight model (< 1M parameters) |
| Trainable Parameters | — | 110,921 | Optimized for efficiency |

Architectural details of the two deep learning models used in this study are provided in Tables 4 and 5. The custom CNN (Table 4) is a lightweight baseline network consisting of three convolutional–pooling blocks, one fully connected layer, and one softmax classifier, total ≈0.33M parameters. On the other hand, the transfer learning model that is built on top of VGG16 (Table 5) uses a pre-trained convolutional backbone from ImageNet with global average pooling and dense layers connected to dropout regularization (total 14.9 million parameters). With these architectures, the tradeoff between efficiency and performance is effectively addressed for citrus leaf disease classification. The VGG16 convolutional base was frozen during training, and only the task-specific dense layers were fine-tuned, resulting in a small number of trainable parameters.

Table 5. VGG16 Transfer Learning Architecture

| Layer Type | Output Shape | Parameters | Key Notes |
|---|---|---|---|
| Input (VGG16 Conv Base, Frozen) | 224 × 224 × 3 | 14,714,688 | Pre-trained ImageNet convolutional layers |
| Global Average Pooling | 512 | 0 | Feature aggregation |
| Dense (128) + Dropout (0.3) | 128 | 65,664 | Task-specific classification head |
| Dense (9) + Softmax | 9 | 1,161 | Multi-class output layer |
| Total Parameters | — | 14,915,165 | Includes optimizer parameters |
| Trainable Parameters | — | 66,825 | Only classification head trained |
| Non-Trainable Parameters | — | 14,714,688 | Frozen convolutional backbone |

### 3.4 Explainable AI Integration

Although CNNs and VGG16 can achieve high prediction quality, deep learning models are black-box systems, which reduces the transparency and trust in their performance in agricultural settings. To address this need, we integrated XAI methods which enable visibility into how models make decisions visually. These approaches further promote interpretability of the areas contributing to the classification and contribute to a more practical useful method for farmers and agronomists by localising the most discriminative areas in citrus leaves.

### 3.4.1 Gradient-weighted Class Activation Mapping (Grad-CAM)

Class-discriminative heatmaps were created using Grad-CAM, which utilises the gradient information flowing into the final convolutional layers of the model. These heatmaps indicate the most important parts of input images for predictions. For citrus leaf disease diagnosis, Grad-CAM consistently drew attention to symptom-specific areas (lesions, discolouration, and fungal spots) as experts observed, yielding supporting evidence for the model's decisions.

### 3.4.2 SHapley Additive exPlanations (SHAP)

To measure the contribution of individual pixels or groups of pixels to a specific prediction, SHAP values were utilized. This framework enables the generation of a consistent and theoretically valid measure of feature importance by approximating Shapley values in cooperative game theory. All explanatory images in this study illustrated how colour intensities and textural methods influenced the predicted probability of a diseased or healthy leaf, providing an interpretable decomposition of model judgments.

### 3.4.3 Local Interpretable Model-Agnostic Explanations (LIME)

The LIME was used to interpret the predictions based on local superpixel segmentation. LIME identifies the super pixels responsible for classification by perturbing image regions and monitoring the impact on prediction. For citrus leaf-based diseases, the overlaid LIMEs effectively highlighted infected patches and suppressed irrelevant background, providing an intuitive and farmer-friendly interpretation of how the model distinguishes different diseases.

In the case of Grad-CAM analysis, feature attribution was calculated based on the last convolutional layer of each model because these layers are the most relevant spatial representations of features of diseases localization. The final Conv2D block was employed in the custom CNN, but, in VGG16, the block5_conv3 layer was applied in Grad-CAM. The representativeness was achieved by the conduct of explainability analysis on a representative half (90 test images) that was chosen randomly among the samples correctly classified in terms of disease to cover classes. Grad-CAM, SHAP, and LIME visualizations were created and compared with each other on a case-by-case basis and across images.

In the reviewed samples, there were recurrent areas of emphasis that were always correlated with the apparent symptoms of the disease like lesions, discoloration, and fungal spots in most instances. Qualitatively, more than

80 percent of the assessed explanations focused on disease-impacted areas as opposed to background or healthy leaf areas, which shows that the models used clinically and agronomically important characteristics to forecast. These results support the validity and clearness of the planned clarifiable model.

## 3.5 Evaluation Metrics

Standard metrics, including accuracy, precision, recall, and the F1-score, were used to evaluate the models' performance both at the overall and class levels. Correct and incorrect predictions across the nine leaf disease categories were presented through confusion matrices. In contrast, ROC curves and AUC values were calculated to assess the models discriminative capacity in a multi-class scenario. Table 6 presents the evaluation metrics adopted in our work, along with their verbal and mathematical definitions. It helps in having clear interpretation of model performance and enables cross comparison with other methodologies being more standardized.

Table 6: Evaluation Metrics of Model

| Metric | Description | Formula |
|---|---|---|
| **Accuracy** | Proportion of total correct predictions among all predictions. | $\frac{TP + TN}{TP + TN + FP + FN}$ |
| **Precision** | Measures how many predicted positives are actually correct. | $\frac{TP}{TP + FP}$ |
| **Recall (Sensitivity)** | Measures how many actual positives are correctly predicted. | $\frac{TP}{TP + FN}$ |
| **F1-Score** | Harmonic mean of precision and recall. Balances false positives/negatives. | $2 * \frac{Precision * Recall}{Precision + Recall}$ |
| **Specificity** | Measures how many actual negatives are correctly predicted. | $\frac{TN}{TN + FP}$ |
| **ROC Curve** | Graphical representation of TPR (Recall) vs. FPR across thresholds. | TPR = TP / (TP + FN), FPR = FP / (FP + TN) |

## 3.6 Experimental setup environment

The data were randomly sampled and classified by a stratified 80% training and 20% validation division in order to maintain the class balance in all nine categories. The randomness was kept by means of a constant seed

(seed = 42). Both CNN and VGG16 models were also trained using the Adam optimizer with a learning rate of 1 x 10 -4, a batch size of 30, and 30 epochs in a workstation with an NVIDIA GPU and TensorFlow 2.x.

## 4. Result Evaluation

The experimental findings indicate a high level of detection and classification by the proposed DNN models in identifying citrus leaf diseases. Both the custom CNN and transfer learning based on VGG16 achieved high performance in prediction, and in all cases, the accuracy on the augmented dataset was above 90%. Large-scale quantitative studies that employ precision, recall, F1-score and class-wise confusion matrices have also indicated the influence of efficacy features and models developed. By incorporating explainable AI methods, such as Grad-CAM, SHAP, and LIME, the visual explanation of how the model arrives at its predictions was significantly enhanced, and it was found that the chosen areas correspond to the areas of disease-affected leaves. Besides the overall accuracy, there were fine differences in the pairwise comparison of the two models. The CNN was faster converging and less expensive to calculate and could therefore be deployed in the field on a lightweight device. Conversely, it was slightly lower than VGG16 in terms of precision and recall across various disease classes, with minimal differences even within disease lookalike similarity. These two competences are complementary, demonstrating the trade-off between efficiency and accuracy, and motivating the use of deep learning models to provide reliable AI in decision support systems that can lead to sustainable farming, provided that explainability is achieved.

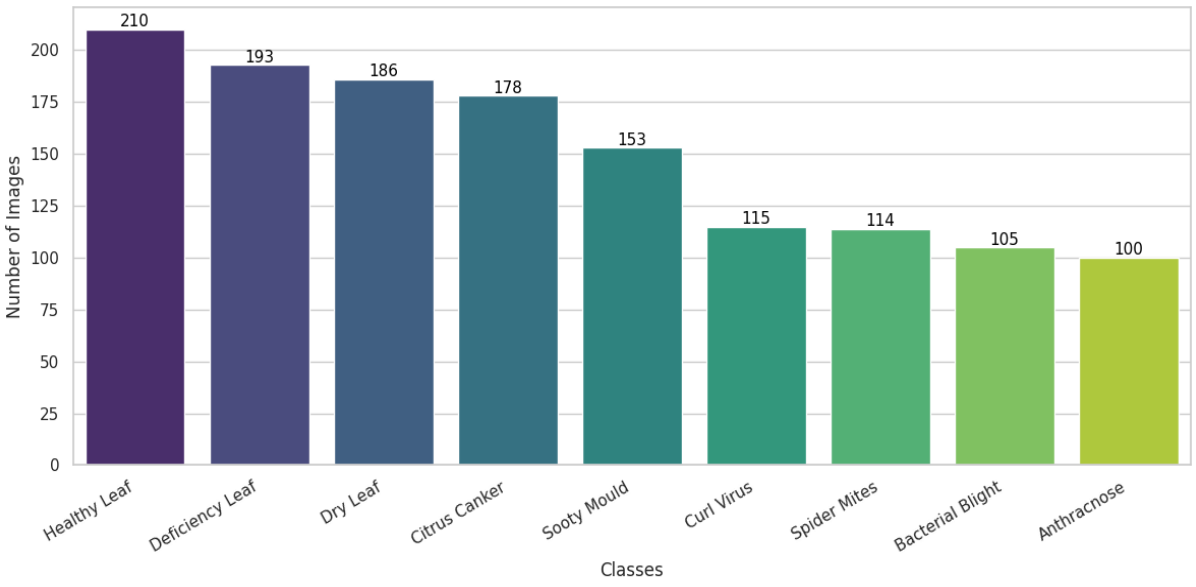



Figure 6: Lemon Leaf Disease Dataset – Class Distribution

Figure 6 shows the class distribution of the original dataset. It is clear from this figure that the categories' Healthy Leaf' and 'Deficiency Leaf' are overrepresented, whereas 'Anthracnose' and 'Bacterial Blight' are

underrepresented. It could harm model training if it is biased towards the majority class. To address the above-mentioned issue and provide even-handedness of all disease types, data augmentation methods were applied to create a well-balanced dataset, resulting in highly robust and generalizable deep learning models.

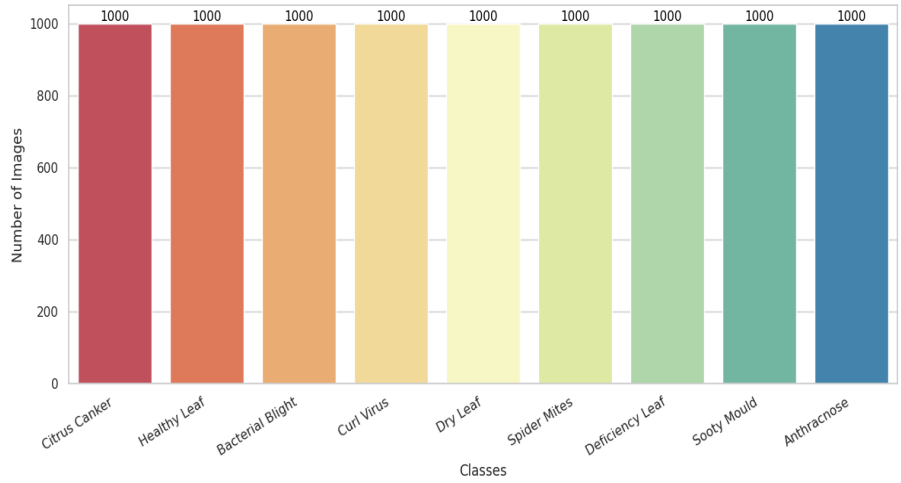



Figure 7: Balanced Lemon Leaf Disease Dataset – Post Augmentation

Figure 7 shows the balanced class distribution after employing systematic data augmentation, where there are 1,000 images per each of the total nine classes. This balanced data generation prevents bias in finding disease-specific features caused by an imbalanced original dataset, and thus permits the learning of diverse category-dependent characteristics equally. We enhanced the generalisation of models, lowered misclassification risks, and increased the reliability of citrus disease detection in practical applications with real agricultural scenes by balancing the dataset.

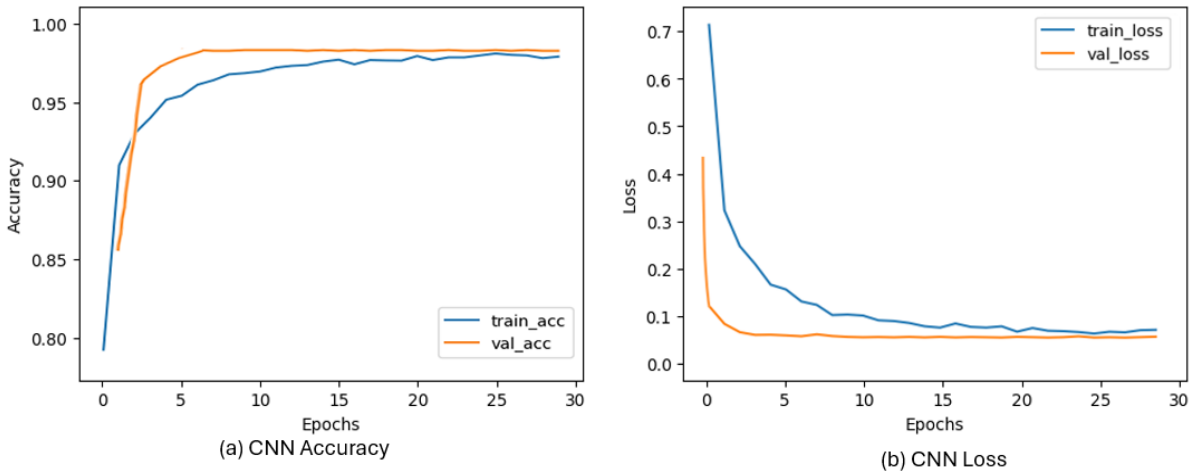



Figure 8: Training and Validation Performance of CNN

Figure 8 displays the accuracy and loss curves for the training and validation sets (30 epochs) of the CNN model. The accuracy curve increases rapidly in the early epochs, then levels off around 98%, with a slightly higher validation set accuracy than the training set accuracy, suggesting good generalisation. Similarly, the loss curves show a steep drop in training and validation losses, that smoothly converge at low values powering overfitting. These results indicate that the CNN model is capable of learning discriminative representations for the identification of citrus leaf diseases.

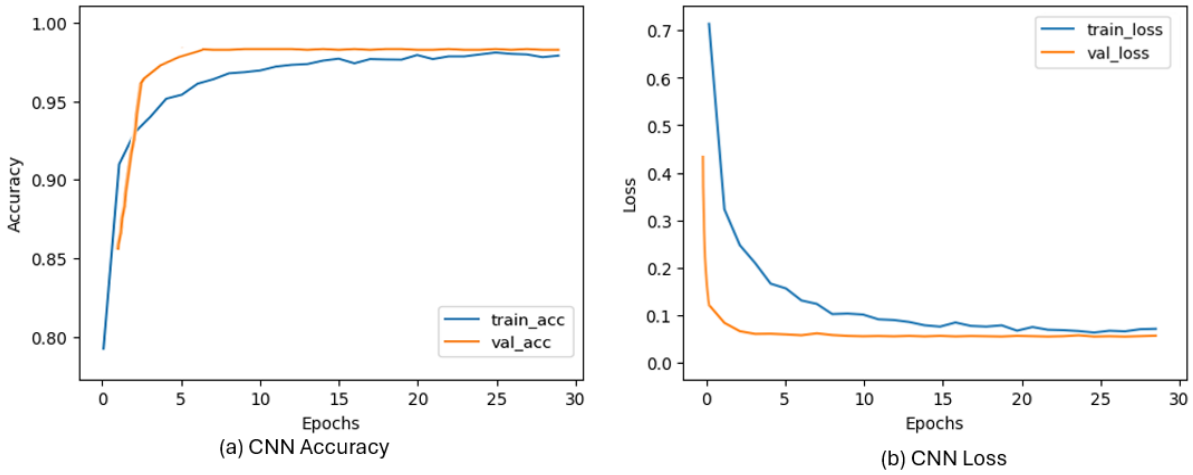



Figure 9: CNN Training and Validation Accuracy and Loss

Figure 9 show the training process of the CNN model. Subfigure (a) depicts the curves of accuracy, showing that both training accuracy and validation accuracy rise rapidly and converge above 97%, indicating that the model consistently learns. Subfigure (b) shows the loss curves, which both decrease gradually to low values before stabilising, indicating satisfactory convergence with minimal overfitting. Such phenomena also proved the generalization ability of the CNN in citrus leaf disease diagnosis.

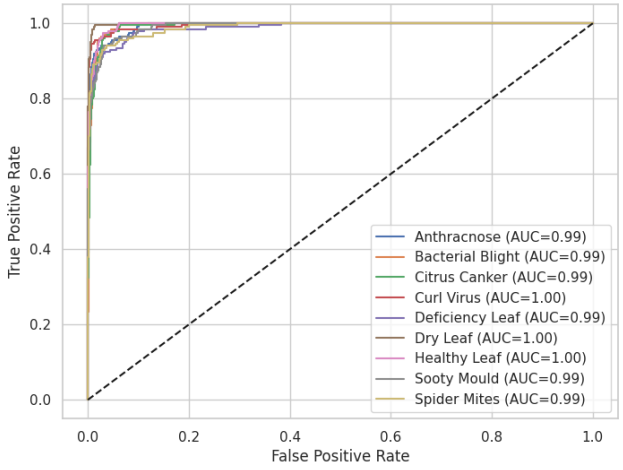



Figure 10: ROC Curve for CNN Model

As shown in Figure 10, the ROC curve highlights the class-wise performance of the CNN model for citrus leaf disease identification. The ROC analysis shows that the performance of the discriminative analysis is strong in all the nine classes with AUC values of 0.99 in Anthracnose, Bacterial Blight, Citrus Canker, Deficiency Leaf, Sooty Mould, and Spider Mites and 1.00 perfect AUC values in Curl Virus, Dry Leaf, and Healthy Leaf. Most classes have true positive rates of more than 0.95 at false positive rates less than 0.05 and thus have high sensitivity and low misclassification. These findings provide a quantitative support of the strength and stability of the proposed model in the process of multi-class citrus leaf disease detection.

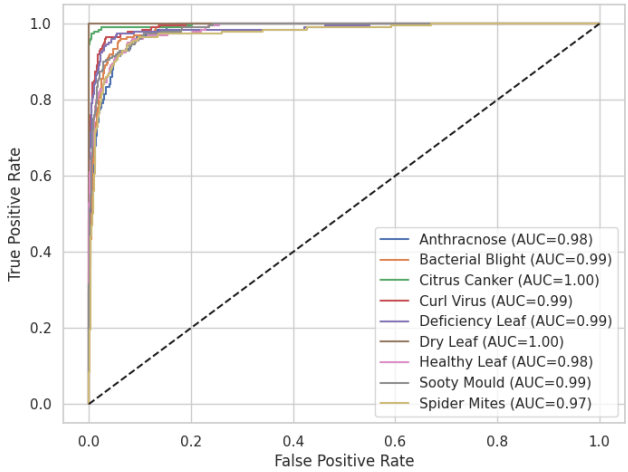



Figure 11 Class-Wise ROC Curve for VGG16

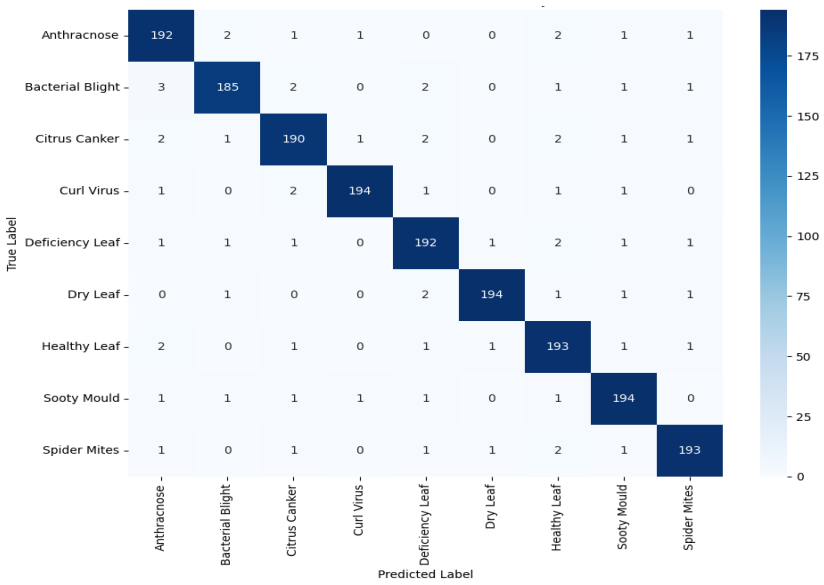



Figure 12: Confusion Matrix for CNN Model

The ROC curve shown in Figure 11 represents the classification performance of VGG16 model for all categories of citrus leaf diseases. The AUC values are between 0.97 and 1.00 with Citrus Canker and Dry Leaf having the highest AUCs of 1.00, Anthracnose having AUC of 0.98 and the lowest AUC of 0.97 of Spider Mites. Most classes achieve the true positive rates of more than 0.90 when the false positive rate is less than 0.10 meaning that it has high but slightly lower discriminative power as compared to the proposed model. Altogether, the findings validate the classification performance of all disease classes with slight and limited differences in the sensitivity of the visually similar conditions.

The confusion matrix represents the performance of CNN model on nine classes of citrus leaves as presented in Figure 12. Many predictions are in the top left quadrant, indicating high TP rates for all categories; Anthracnose, Curl Virus, Dry Leaf and Sooty Mould had particularly high precision. Some false assignments are distracting for classes like Bacterial Blight and Citrus Canker but those are still few. These results show that the CNN exhibits excellent classification performance and is well balanced between different disease pattern recognitions.

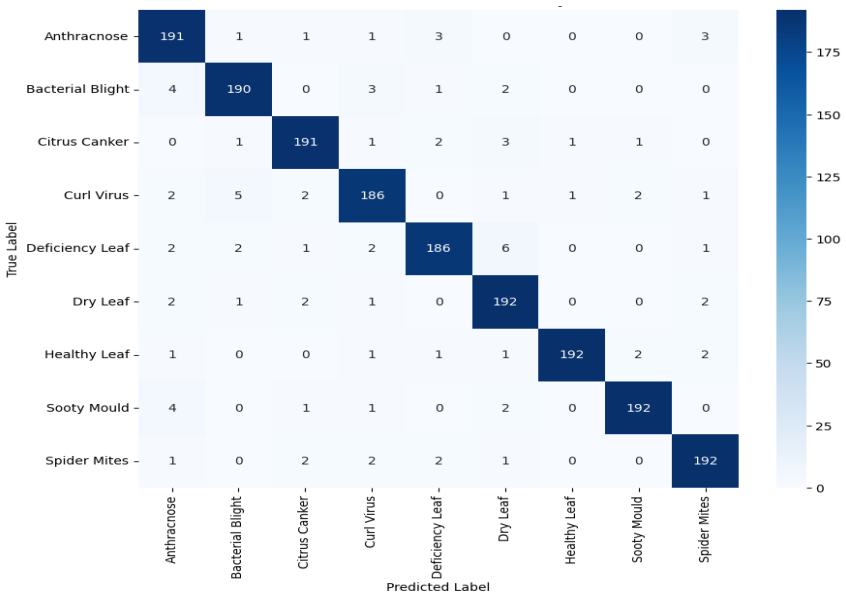



Figure 13: Confusion Matrix for VGG16 Model

The confusion matrix in Figure 13 shows the classification performance of the VGG16 model for nine classes of citrus leaf diseases. Although most predictions fall on the diagonal (showing good recognition for most of the categories), misclassifications tend to be more common than CNN. Curl Virus and Deficiency Leaf, etc., have higher overlaps with their neighbours, under sensitivity to visual closeness. Nevertheless, the model performs well in classification, achieving a general accuracy of 95%, which supports the effectiveness of our proposed transfer learning method.

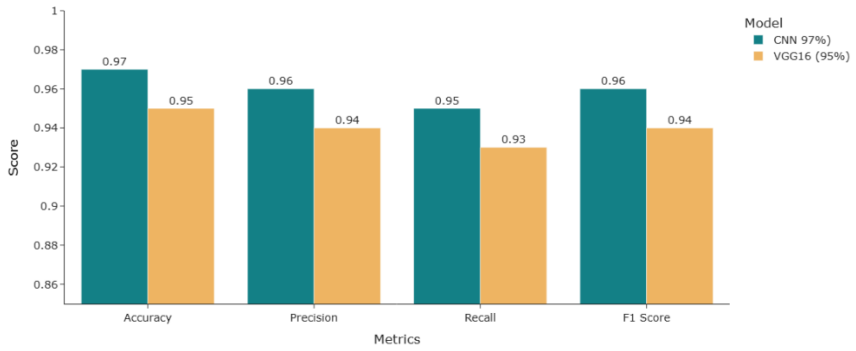



Figure 14: Model Comparison on Performance Metrics

Figure 14 illustrates the comparative performance of CNN (97%) and VGG16 (95%) on four key metrics: Accuracy, Precision, Recall, and F1-Score. The results in all categories are almost always higher for the CNN model, with an accuracy of 0.97, precision of 0.96, recall of 0.95, and an F1-score of 0.96, while VGG16 lags behind with values of 0.95, 0.94, 0.93, and 0.94, respectively. Findings demonstrate that the customised CNN architecture exhibits better generalisation and robustness than the transfer learning-based VGG16 for the automated classification of citrus leaf diseases.

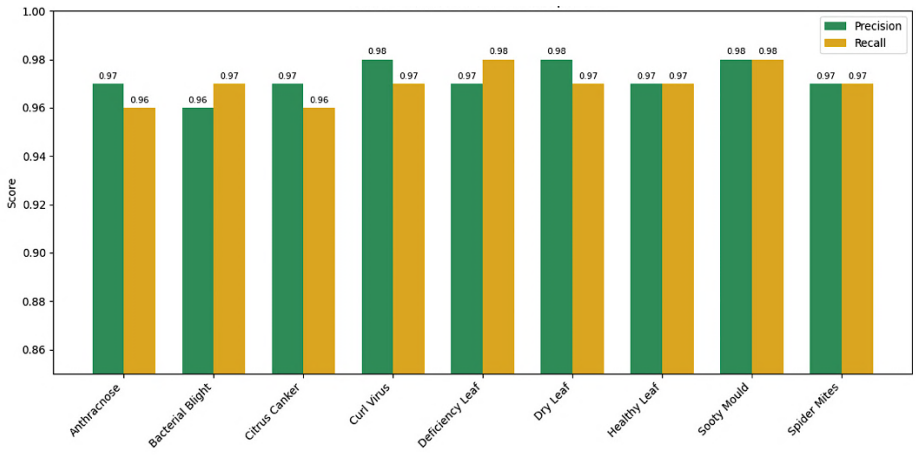



Figure 15: Precision and Recall per Class of CNN Model

Figure 15 shows the average class-wise precision and recall of the CNN model, with an overall accuracy rate of 97%. The obtained results, indicative of consistent behaviour in all nine classes with precision and recall values between 0.96 and 0.98, are tabulated. The categories of Dry Leaf, Curl Virus, and Sooty Mould receive the highest scores (0.98), which demonstrates that this model performs well for different types of diseases affecting citrus tree leaves. The near coincidence of precision and recall implies that there is a strong compromise between false positive and false negative minimisations, guaranteeing reliable disease detection.

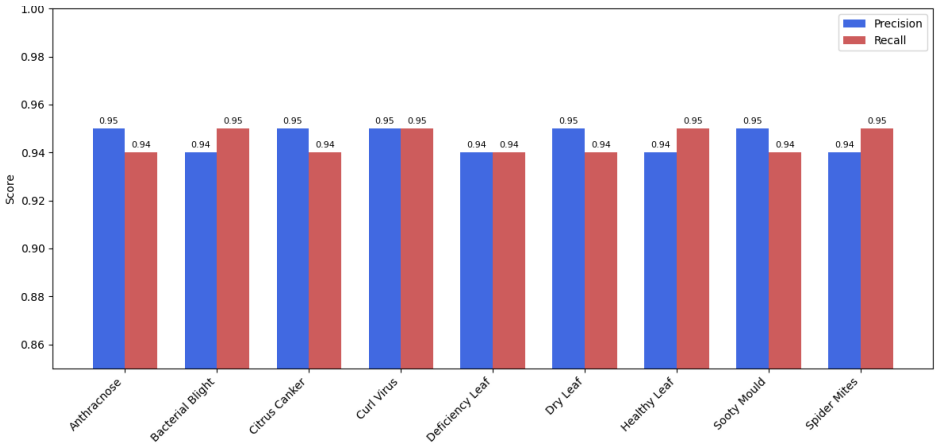



Figure 16: Precision and Recall per Class of VGG16 Model

Figure 16 illustrates the class-wise precisions and recalls of our VGG16 model, which achieved 95% accuracy. Both precision and recall fall within a tight range of 0.94–0.95, indicating that the method exhibits consistent performance across all nine citrus leaf disease classes. Although smaller than CNN, VGG16 possesses an equally good detection ability, thus guaranteeing that it avoids being biased in any one disease category. The results reveal that VGG16 is a very consistent transfer learning approach, but with slightly lower sensitivity than the CNN model.

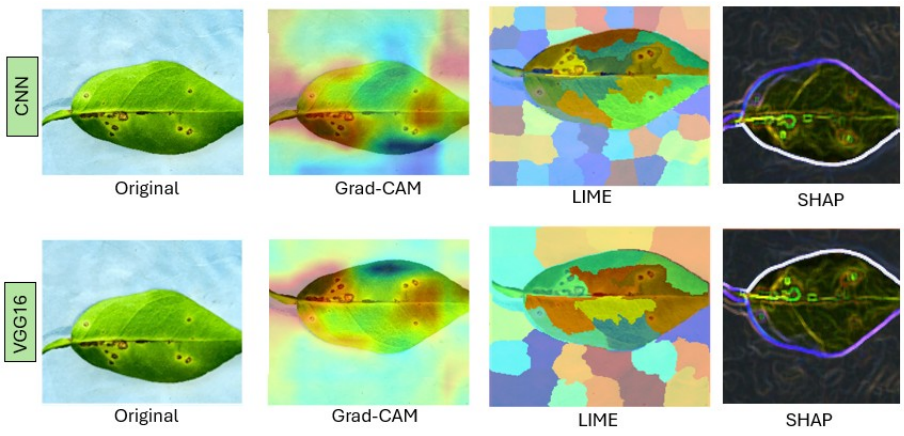



Figure 17: Comparative Explainable AI Visualizations for Citrus Leaf Disease Detection

Figure 17 presents a comparison of explanation outputs for CNN and VGG16, which use Grad-CAM, LIME, and SHAP, on a single sheet with an identical diseased leave image as input. Grad-CAM visually attains the broader lesion region of interest for classification, whereas LIME segments and accentuates superpixels contributing to predictions. The SHAP provides attribution of pixels that identify important symptomatic areas, consistent with expert diagnosis. In combination, these visualisations corroborate the models' logic, thereby increasing interpretability and confidence in real-world agricultural decision support.

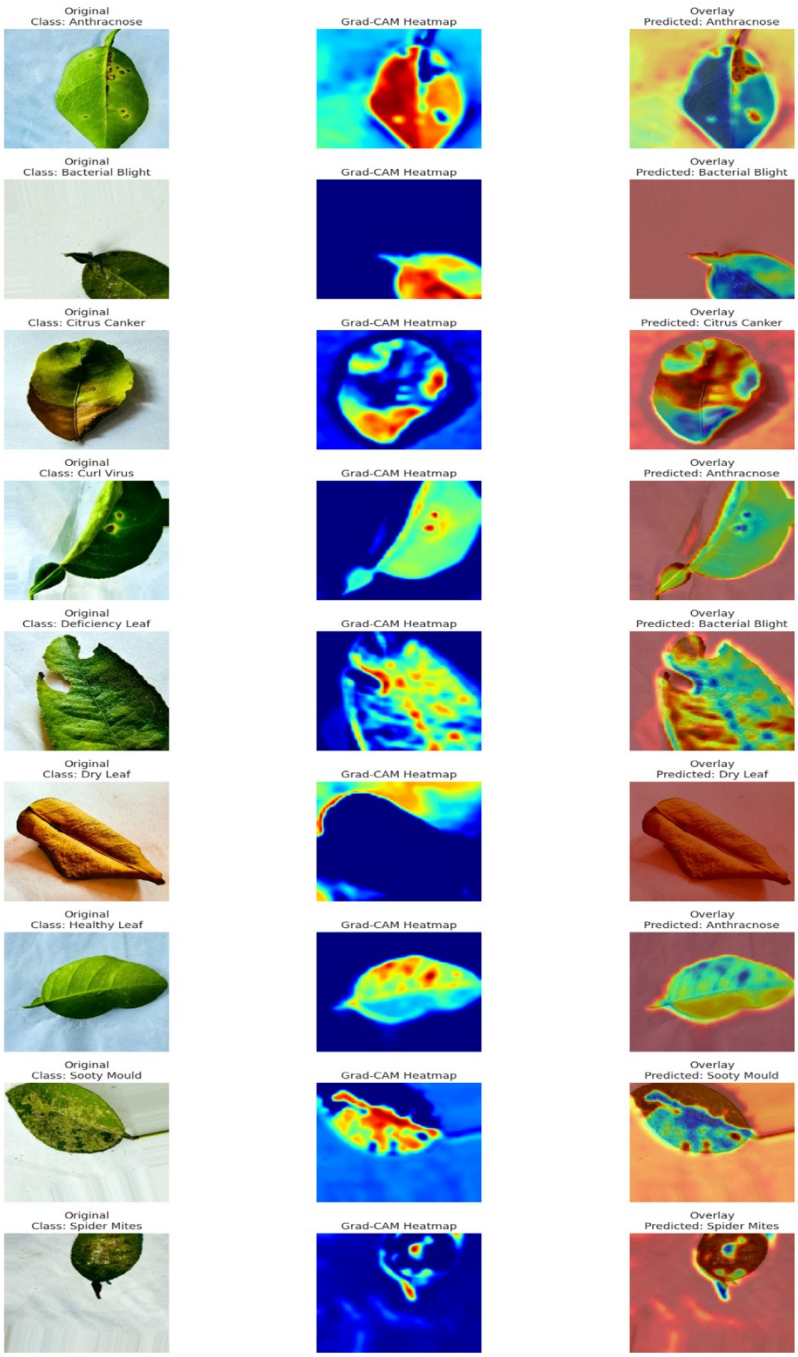



Figure 18: Grad-CAM–Based Visual Explanation of CNN

Figure 18 depicts the visual explanations (Grad-CAM) of the last convolutional layer of the CNN model to classify citrus leaf diseases. The disease-relevant areas including lesions, discoloration, and texture abnormalities are identified by the heatmaps and employed to make predictions with the models. Having proper classifications demonstrate the high coincidence between activated areas and observed pathological phenomena, whereas wrongly classified samples demonstrate the coincidence of visual cues belonging to different classes. These findings indicate that CNN makes decisions basing on significant leaf features and not background facts.

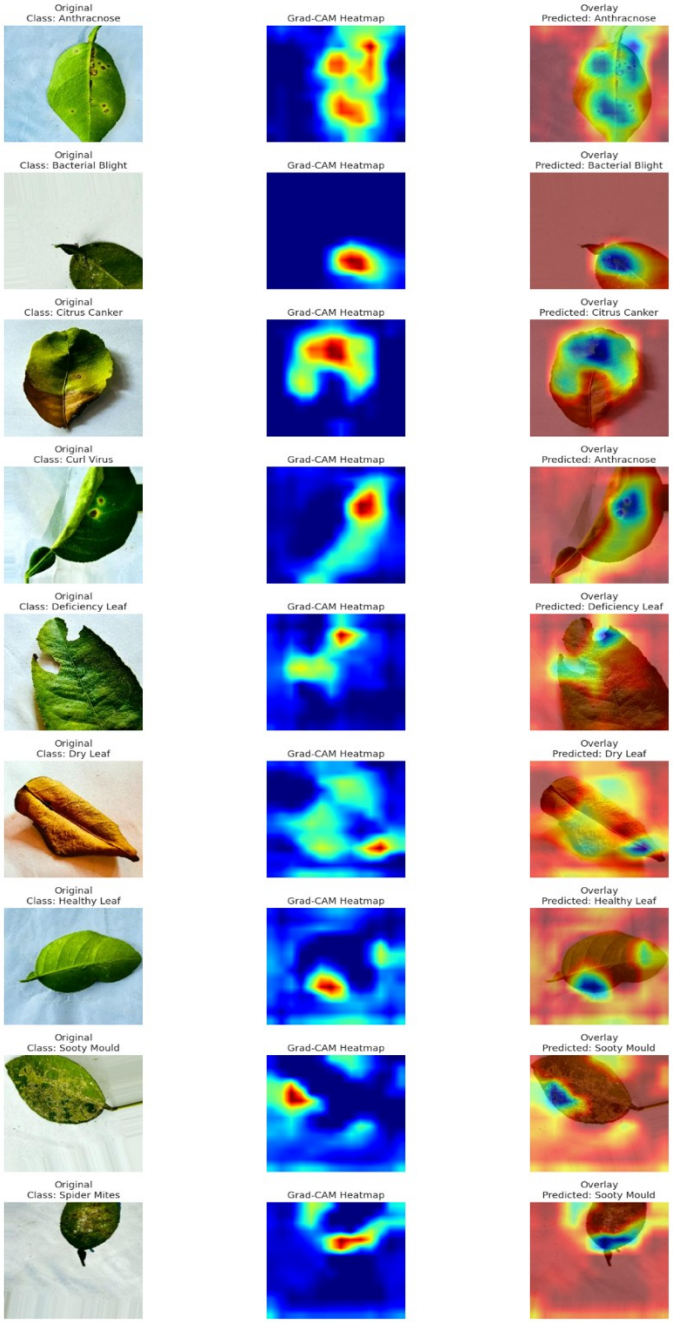



Figure 19: Grad-CAM–Based Visual Explanation of VGG16

Figure 19 demonstrates Grad-CAM visualizations generated on the basis of the last convolutional layer of VGG16 model to understand citrus leaf disease predictions. The salient areas of the disease identified with visual salience are the lesions, discoloration, and surface irregularities that were trained in ImageNet-pretrained features. Right predictions show high similarities between the activated and pathological areas whereas the misclassified ones show overlapping features activation across visually similar classes. Through these illustrations, it is possible to understand how transfer learning facilitates VGG16 to learn discriminative disease information and retain model interpretability.

The results have shown that, on average, CNN and VGG16 performed well in modelling for citrus leaf disease classification, though CNN slightly outperformed VGG16 with regard to overall accuracy, precision, recall, and F1-score. Explainable AI techniques, namely Grad-CAM, LIME, and SHAP, were incorporated to identify regions of interest as indicated by the models, which aligned well with observable disease symptoms. These observations not only confirm the robustness of the models but also enhance their interpretability, thereby improving their potential for agricultural adoption. In summary, the results suggest that integrating deep learning and XAI can provide well-performing, interpretable decision support systems for sustainable crop production.

## 5. Discussion

The results of this paper reinforce the transformational nature of deep learning in precision agriculture, particularly in the automated detection of citrus leaf disease. We were able to experimentally demonstrate that the custom CNN was more accurate than the VGG16 transfer learning model, highlighting the importance of domain-specific model design. Although both models worked efficiently in various disease classes, the CNN was consistently able to produce superior recall and F1-scores, which ensured a more reliable detection of underlying disease symptoms. Notably, a balanced and augmented dataset was used, which ensured resistance to class imbalance, a phenomenon frequently reported in studies of plant diseases. The findings suggest that the accuracy and reliability of agricultural artificial intelligence systems can be significantly enhanced when specifically designed architectures are implemented with the aid of various training datasets.

In addition to performance, this work focuses on the explainability aspect to ensure that AI in agriculture is reliable and can be effectively applied. Grad-CAM, LIME, and SHAP were not only effective in revealing the essential visual symptoms but also in good agreement with agronomic knowledge of experts, facilitating interpretability between the precision of the computer and the discernibility by the farmer. Our models can inform decisions by interpreting them, which can serve as a guide for practice during the early phases of disease monitoring and specific interventions. This combined high-performance classification with easy interpretability makes the proposed framework a viable and scalable solution for sustainable agriculture, benefiting direct farmers, researchers, and authorities.

Table 7 provides a comparison of the proposed explainable CNN-VGG16 framework with the newest state-of-the-art citrus leaf disease detection methods on the basis of model structure, data type, support of explainability and reported performance. In contrast to the YOLO-based and hyperspectral approaches that either focus on the accuracy of the detection, or use special imaging tools, the proposed study focuses on lightweight classification, multi-method explainability, and their practical deployability on normal RGB images. As noted by the comparison, the proposed approach is competitive in its accuracy, and it is the only one that offers transparent and understandable predictions that can be applied in the real-world support of agricultural decisions.

To contextualize the achieved accuracies of 97% for the custom CNN and 95% for VGG16, Table 7 provides the comparative analysis of the given results with the representative citrus and leaf disease detection methods in the literature. The comparison demonstrates that the suggested models perform competitively in comparison to detection-based and hyperspectral systems and is the only one that includes explainable AI practices, and their computational complexity is lower with the usage of standard RGB images.

Table 7. Comparative Positioning of the Proposed Method Against Recent Citrus Leaf Disease

| Study | Core Methodology | Data Modality | Explainability | Performance | Key Limitation vs. This Work |
|---|---|---|---|---|---|
| Zhu et al., 2025 [21] | YOLOv8-CMS with CSPPC, attention modules + segmentation | RGB images | Not integrated | Accuracy: 97.9%, mAP50: 98.2% | Heavy detection model; no interpretability or decision transparency |
| Frederick et al., 2025 [22] | YOLOv8 with hyperspectral band selection | Hyperspectral (HSI) | Not integrated | Weighted F1: 0.896 | Requires expensive sensors; limited field scalability |
| Zhu et al., 2025 [23] | CBACA-YOLOv5 with CBAM + coordinate attention | RGB images | Not integrated | Accuracy: 96.1%, mAP: 92.1% | Focused on detection; lacks interpretability and decision support |
| Frederick et al., 2024 [24] | YOLOv8 with supervised spectral band selection | Hyperspectral (adaxial/abaxial) | Not integrated | Accuracy: 87.1% | Complex imaging setup; lower accuracy than proposed CNN |
| Xiao et al., 2026 [25] | Spectral–physiological fusion with ML classifiers | Spectral + physiological | Not integrated | Accuracy: up to 98.9% | Single-disease focus (anthracnose); not image-based, no XAI |
| This Work | Custom CNN + VGG16 with unified XAI (Grad-CAM, SHAP, LIME) | RGB images | Multi-XAI integrated | Accuracy: 97%, F1: 0.96 | Lightweight, interpretable, decision-oriented |

Existing approaches often apply either no explainability strategy or use one visualization method, which is problematic when it comes to the assessment of trust. This research shows much stronger evidence that the learned features reflect the real disease symptoms as opposed to spurious correlations by comparatively analyzing Grad-CAM, SHAP, and LIME outputs of two different model families. The multi-XAI validation is a realistic development on the conventional CNN based classification pipelines.

## 6. Conclusion

This study proposed a deep learning model to explainable detect citrus leaf disease, which integrates the predictive power of the CNN and VGG16 models with interpretability methods, including Grad-CAM, LIME, and SHAP. The experimental findings revealed that the custom CNN achieved state-of-the-art accuracy (97%),

surpassing VGG16 (95%) in terms of accuracy, precision, recall, and F1-score, and was made robust by utilising advanced preprocessing and augmentation techniques. Explainable AI techniques enabled the clear visualisation of the disease-impacted areas of the land, which increased confidence in the results received by the models and matched the predictions with professional agronomic observations. These results demonstrate the need to strike a very delicate balance between predictive accuracy and transparency in the development of trustworthy AI-based systems in agriculture.

The broader impact of this study is that it can help develop sustainable and intelligent agricultural practices. The proposed framework can provide an early-warning system for crop health management, as well as a powerful decision-support tool, empowering farmers, agronomists, and policymakers by combining high-performance disease detection with explainability. This paper provides a basis for incorporating AI into real-time agricultural surveillance frameworks, which will enhance food security, minimise financial losses, and ensure environmentally friendly farming practices. Future directions include generalising the framework to other crops, utilising multimodal data (e.g., hyperspectral and environmental data), and running lightweight models on edge devices to implement the framework on a large scale in the field.

**Funding**

The authors declare that this research received no external funding.

**Data Availability Statement**

Not applicable.

**Conflicts of Interest**

The authors declare no conflict of interest.

**References**

[1] P. Sharma and P. Abrol, “Multi-component image analysis for citrus disease detection using convolutional neural networks,” Crop Protection, vol. 193, p. 107181, Jul. 2025, doi: 10.1016/j.cropro.2025.107181.
[2] A. Goyal and K. Lakhwani, “Integrating advanced deep learning techniques for enhanced detection and classification of citrus leaf and fruit diseases,” Scientific Reports, vol. 15, no. 1, p. 12659, Apr. 2025, doi: 10.1038/s41598-025-97159-0.
[3] K. Yan et al., “Three-dimensional fluorescence spectral characteristic of flavonoids for citrus Huanglongbing disease early detection,” Microchemical Journal, vol. 208, p. 112263, Jan. 2025, doi: 10.1016/j.microc.2024.112263.

[4] S. A. Higgins et al., “Plant-derived, nodule-specific cysteine-rich peptides as a novel source of biopesticides for controlling citrus greening disease,” Phytopathology, vol. 114, no. 5, pp. 971–981, May 2024, doi: 10.1094/PHYTO-09-23-0322-KC.
[5] P. K. Yadav et al., “Automated classification of citrus disease on fruits and leaves using convolutional neural network generated features from hyperspectral images and machine learning classifiers,” Journal of Applied Remote Sensing, vol. 18, no. 1, p. 014512, Jan. 2024, doi: 10.1117/1.JRS.18.014512.
[6] P. Dinesh and R. Lakshmanan, “Deep learning-driven citrus disease detection: A novel approach with DeepOverlay L-UNet and VGG-RefineNet,” International Journal of Advanced Computer Science and Applications, vol. 15, no. 7, pp. 1023–1041, Jul. 2024.
[7] C. Wang et al., “A simple and effective VIGS system facilitates the control of citrus canker by silencing CsLOB1,” Phytopathology Research, vol. 6, no. 1, p. 12, Mar. 2024, doi: 10.1186/s42483-024-00234-z.
[8] R. Sujatha, S. Krishnan, J. M. Chatterjee, and A. H. Gandomi, "Advancing plant leaf disease detection integrating machine learning and deep learning," Scientific Reports, vol. 15, Art. no. 11552, Apr. 2025, doi: 10.1038/s41598-024-72197-2.
[9] P. Dhiman, A. Kaur, Y. Hamid, E. Alabdulkreem, H. Elmannai, and N. Ababneh, "Smart Disease Detection System for Citrus Fruits Using Deep Learning with Edge Computing," Sustainability, vol. 15, no. 5, p. 4576, Mar. 2023, doi: 10.3390/su15054576.
[10] H. K. Kondaveeti and C. G. Simhadri, "Evaluation of deep learning models using explainable AI with qualitative and quantitative analysis for rice leaf disease detection," Scientific Reports, vol. 15, Art. no. 31850, Aug. 2025, doi: 10.1038/s41598-025-14306-3.
[11] Q. Xu, L. Ma, R. Zhou, C. Wang, J. Bai, L. Sun, and J. Cai, “Detection of huanglongbing infection in citrus using compositional analysis of volatile organic compounds,” Plant Pathology, vol. 73, no. 8, pp. 2084–2100, Oct. 2024, doi: 10.1111/ppa.13964.
[12] X. Guo, X. Wang, W. Zhu, S. X. Yang, L. Song, P. Li, and Q. Li, “Citrus disease detection based on dilated reparameterized feature enhancement and shared parameter head,” Sensors, vol. 25, no. 7, p. 1971, Mar. 2025, doi: 10.3390/s25071971.
[13] E. Aswini and C. Vijayakumaran, “Citrus plant classification using gated fusion adaptive graph neural network with harbor seal whiskers optimization algorithm,” Knowledge-Based Systems, vol. 329, p. 114093, Nov. 2025, doi: 10.1016/j.knosys.2025.114093.
[14] H. Zhu, D. Wang, Y. Wei, X. Zhang, and L. Li, "Combining Transfer Learning and Ensemble Algorithms for Improved Citrus Leaf Disease Classification," Agriculture, vol. 14, no. 9, p. 1549, Sep. 2024. doi: 10.3390/agriculture14091549. [Online]. Available: https://doi.org/10.3390/agriculture14091549
[15] B. Kaur, S. K. Gupta, M. Janarthan, D. M. Alsekait, and D. S. AbdElminaam, "Precision diagnosis of citrus leaf diseases using image enhancement and nonlinear fuzzy ranking ensemble approach NLFuRBe," Scientific Reports, vol. 15, Art. no. 32296, 2025. doi: 10.1038/s41598-025-16923-4. [Online]. Available: https://doi.org/10.1038/s41598-025-16923-4
[16] H. Mo and L. Wei, "Lightweight citrus leaf disease detection model based on ARMS and cross-domain dynamic attention," Journal of King Saud University – Computer and Information Sciences, vol. 36, no. 10, p. 102133, Oct. 2024. doi: 10.1016/j.jksuci.2024.102133. [Online]. Available: https://doi.org/10.1016/j.jksuci.2024.102133

[17] H. J. Min, J. Qin, P. K. Yadav, Q. Frederick, T. Burks, M. Dewdney, I. Baek, and M. Kim, "Classification of Citrus Leaf Diseases Using Hyperspectral Reflectance and Fluorescence Imaging and Machine Learning Techniques," Horticulturae, vol. 10, no. 11, p. 1124, Oct. 2024. doi: 10.3390/horticulturae10111124. [Online]. Available: https://doi.org/10.3390/horticulturae10111124
[18] A. Bora and P. Sarma, "A Visionary Analysis on Ways and Techniques Used for Identification of Citrus Plant Disease," 2025 3rd International Conference on Intelligent Systems, Advanced Computing and Communication (ISACC), Silchar, India, Feb. 2025. doi: 10.1109/ISACC65211.2025.10969180.
[19] M. Franco-Hernández, V. Mejía-Domínguez, and Y. Rodriguez-Gallo, "Advanced Image Processing Methods for Early Detection and Monitoring of Aphid Infestations," 2024 IEEE Central America and Panama Student Conference (CONESCAPAN), Panama, Panama, Sep. 2024. doi: 10.1109/CONESCAPAN62181.2024.10891106.
[20] N. Valarmathi, M. Bharathi Kannan, M. Elankumaran, M. Gowri Shankar, and B. Gowtham, "Fruit Disease Prediction with Fertilizer Recommendation for Citrus Family using Deep Learning," 2023 International Conference on Sustainable Computing and Data Communication Systems (ICSCDS), Erode, India, Mar. 2023. doi: 10.1109/ICSCDS56580.2023.10104601.
[21] H. Zhu, D. Wang, Y. Wei, P. Wang, and M. Su, "YOLOV8-CMS: A high-accuracy deep learning model for automated citrus leaf disease classification and grading," Plant Methods, vol. 21, no. 1, p. 88, Jun. 2025, doi: 10.1186/s13007-025-01396-3.
[22] Q. Frederick, T. Burks, J. A. Watson, P. K. Yadav, J. Qin, M. Kim, and M. M. Dewdney, "Supervised hyperspectral band selection using texture features for classification of citrus leaf diseases with YOLOv8," Sensors, vol. 25, no. 4, p. 1034, Feb. 2025, doi: 10.3390/s25041034.
[23] J. Zhu, J. Chen, H. He, W. Bai, and T. Zhou, "CBACA-YOLOv5: A symmetric and asymmetric attention-driven detection framework for citrus leaf disease identification," Symmetry, vol. 17, no. 4, p. 617, Apr. 2025, doi: 10.3390/sym17040617.
[24] Q. Frederick, T. Burks, P. K. Yadav, J. Qin, M. Kim, and M. Dewdney, "Classifying adaxial and abaxial sides of diseased citrus leaves with selected hyperspectral bands and YOLOv8," Smart Agricultural Technology, vol. 9, p. 100600, Dec. 2024, doi: 10.1016/j.atech.2024.100600.
[25] Z. Xiao, C. Yang, Z. Wei, X. Fu, C. Zhao, and G. Li, "A spectral–physiological feature fusion model for the early detection of anthracnose in citrus leaves," Computers and Electronics in Agriculture, vol. 243, p. 111410, Mar. 2026, doi: 10.1016/j.compag.2026.111410.