Int J Adv Life Sci Res. Volume 8(4) 138-157
https://doi.org/10.31632/ijalsr.2025.v08i04.011

International Journal of Advancement in Life Sciences Research
Online ISSN: 2581-4877
Journal homepage http://ijalsr.org

Original Article

Enhancing Cassava Leaf Disease Detection through Traditional Segmentation and Attention-driven Deep Learning Approaches

Uddagiri Sirisha[1], Rajeswari Nakka[2], Janaswami Hymavathi[3], S. Lalitha[4], Surapaneni Phani Praveen[1*], Dimitrios A. Karras[5]

[1] *Department of Computer Science and Engineering, Prasad V Potluri Siddhartha Institute of Technology, Kanuru, Vijayawada, A.P, India.*
[2] *Department of Computer Science and Engineering, Seshadri Rao Gudlavalleru Engineering College, A.P, India.*
[3] *Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, A.P, India*
[4] *Department of Computer Science and Engineering (AIML), Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Avadi, Chennai.*
[5] *National and Kapodistrian University of Athens (NKUA), Department of Agricultural Development, Agri-Food and Natural Resources Management (AAADFP), Greece & Computer Engineering Department, EPOKA University, Tirana, Albania.*

**Correspondence E-mail:* phani.0713@gmail.com

Abstract

Diseases of cassava leaves pose a significant threat to food security in tropical regions, as well as to harvests. This paper presents a combined approach for accurate and interpretable cassava leaf disease diagnosis, utilising a deep learning-based ARMUNet architecture alongside conventional image segmentation techniques. Beginning with classical methods—Otsu Thresholding, Distance Transform, and Watershed—the pipeline generates boundary-aware lesion maps that effectively isolate diseased areas. These initial segmentations guide ARMUNet, an enhanced version of UNet with attention gates and residual multi-scale encoders, allowing the model to focus on relevant lesion features while minimising background interference. Featuring various geographical and semantic elements, a Multi-Level Feature Extraction system aids in the correct classification of diseases. Detection accuracy is further improved by an ensemble method that combines ARMUNet predictions with classifiers such as PINN, ResNet50, and EfficientNetB0. The proposed system offers a scalable solution for plant disease diagnostics in precision agriculture, demonstrating high performance, interpretability, and real-time readiness.

Keywords: Distance Transform; Ensemble learning; Multi-Level Feature Extraction; Otsu Thresholding; Segmentation; Watershed

Introduction

A major staple crop in Africa, cassava (Manihot esculenta Crantz), also known as manioc or yucca, is highly beneficial in terms of agriculture and the economy. Although originally from South America, cassava has become a staple crop for over 500 million people due to its ability to thrive in tropical and subtropical regions (Oyewola *et al*., 2021; Emuoyibofarhe *et al*., 2019). Its edible tubers are ranked third in carbohydrates, making them a crucial component of food security. The leaves, grown in Africa, provide a source of protein and vitamins A and B1. Due to its adaptability to different growing conditions and harvest windows (ranging from six months to three years), cassava is a reliable crop for both food security and industrial processing. In 2021, nearly two-thirds of the world's total production, amounting to 324.7 million tonnes, came from Africa. However, yields have been declining in regions such as Sub-Saharan Africa and Thailand due to the evolution of plant diseases like CBSD, CBB, CMD, and CGM over the years.

Received on :17th May 2025; Revised version received on :19th July 2025; Accepted: 3rdSeptember 2025

According to Too *et al*. (2018) and Sangbamrung *et al*. (2020), field inspections have been used by agricultural professionals to identify and manage cassava infections. However, poor nations may lack the resources to implement these methods effectively, even though they would be highly effective in an ideal world. This approach is expensive, slows down the identification process, and reduces yields. Two AI-based technologies, deep learning and machine learning, are powerful tools for automating the diagnosis of plant diseases. The ability of these algorithms to rapidly and efficiently detect diseases in leaf images is invaluable in situations where time is of the essence (Ramcharan *et al*., 2017; Ayu, 2021).

Hybrid deep learning models, which combine various architectures, have recently shown even greater potential (Zhike Zuo, 2019; S *et al*., 2025; Kataoka *et al*., 2003). The feature learning and generalisation capabilities of these models are enhanced by adaptive optimisation and multi-level representation. However, overfitting can still occur with current methods due to factors such as uneven data, poor field setup generalisation, background noise, and difficulties in identifying early or small symptoms. To overcome these challenges, disease diagnostic systems need to be more refined, user-friendly, and accurate (Sharada *et al*., 2016; Rajalakshmi *et al*., 2022).

Kuan filtering and watershed segmentation effectively eliminate background noise and identify disease-relevant zones. UNet and EfficientNet deep learning models have rich hierarchies. When combined with hand-crafted features, such as contour-based measurements, the models can learn more about disease traits, thereby enhancing detection accuracy and generalisability. This method can overcome several image processing limitations. The contributions of this research are as follows:

- To develop a hybrid segmentation method that produces precise and interpretable region-of-interest (ROI) maps for cassava leaf disease diagnosis by integrating conventional techniques—Otsu, Distance Transform, Watershed—with deep learning.
- To enhance the UNet architecture with attention mechanisms and residual multi-scale encoding blocks, enabling targeted learning of diseased regions while minimising background noise.
- To create a Multi-Level Feature Extraction system capable of capturing various spatial and structural characteristics for robust cassava disease categorisation.
- To employ ensemble learning to enhance classification performance by combining ARMUNet predictions with those from other high-performing models such as PINN, ResNet50, and EfficientNetB0.

Literature Review

Many studies have proposed various machine learning (ML) and deep learning (DL) algorithms to improve the accuracy of cassava leaf disease detection. Several recent studies were analysed, focusing on the methods used, the findings achieved, and the significant limitations that remain unresolved, in order to understand the current state of research. The motivation for the proposed study is primarily based on addressing these limitations.

Recent research on cassava disease categorisation has made significant use of both pre-trained and bespoke deep learning models. Although the model showed overfitting, Alford *et al*. (2024) achieved a test accuracy of 80.27% using transfer learning with VGG19. Meanwhile, Singh *et al*. (2024) demonstrated 87% accuracy using InceptionResNetV2, but insufficient training epochs limited the outcomes. Custom lightweight models, such as the one created by Emmanuel *et al*. (2023), which included channel attention processes, reached just 75% test accuracy. These models showed encouraging validation results but struggled to generalise in real-world settings. Despite its resource-intensive design, which restricted its practical use, Zhang *et al*. (2024) developed MAIANet, a multi-attention-based model, achieving 95.83% accuracy.

The authors of Ravi *et al*. (2021) investigated EfficientNet variations using attention mechanisms, thereby achieving notable accuracy on balanced datasets. However, limited dataset diversity remained an issue for real-world scaling. Although the model's small dataset constrained generalisability, the authors of Sholihin *et al*. (2023) combined AlexNet with an SVM classifier to achieve 90.7% accuracy. Using a dataset of 7,000 images spread across two categories, Haq *et al*. (2023) applied CNNs with the Conv2D library and VGG-16 for early pest detection in maize leaves. Using TensorFlow and PyTorch, they also conducted real-time weed detection in wheat, achieving excellent detection accuracy with quick inference on an NVIDIA RTX2070 GPU.

Overall, the common problems identified across all the research were overfitting and poor generalisation to unobserved data (Haq *et al*., 2023). Lightweight models, such as the one developed by Haq *et al*., showed significant declines in test performance, highlighting the difficulty of generalising from simplified designs. On the other hand, Prashanth *et al*. (2024) proposed the MPCSAR-AHH model, which achieved 99.02% accuracy. However, its application in low-resource settings was limited due to its dependence on sophisticated equipment. Similarly, the contributors of Lilhore *et al*. (2022) introduced an Enhanced CNN (ECNN) using depth-wise separable convolutions and gamma correction, achieving 99.3% accuracy on curated datasets. Despite this, it faced deployment restrictions due to the high processing costs.

Another major issue was data quality and variety. Many high-performing models trained on balanced datasets struggled when assessed in various field settings. For example, models described by Sapre *et al*. (2023) performed poorly under real-world complexities, particularly when trained on uneven datasets. When testing ResNet152V2, InceptionResNetV2, DenseNet121, and DenseNet201 for citrus leaf disease identification, they achieved remarkable accuracy and F1 scores on enhanced datasets but raised concerns about real-world dependability.

Hybrid and ensemble methods have attracted interest due to their ability to integrate the strengths of several models. By combining multiple model forecasts, Ravi *et al*. (2021) and Dey *et al*. (2024) demonstrated that an ensemble meta-classifier enhanced overall performance. However, these techniques often require more computing resources, which limits their real-time use. While they may fall short in accuracy and durability, lightweight models such as MobileNetV2 and the one proposed by Tewari and Kumari are more suitable for edge deployment.

Other research (Lakshmi *et al*., 2024) has examined the trade-off between performance and efficiency. Models with attention mechanisms, which required significant processing resources but achieved excellent accuracy, were demonstrated by both Prashanth *et al*. (2024) and Zhang *et al*. (2024). Focusing on cashew disease detection, the authors of Panchbhai *et al*. (2024) developed CAS-CNN and CAS-MODMOBNET, achieving up to 99.8% accuracy with low model size and complexity. Using Densenet-based designs for citrus leaf classification, the authors of Lanjewar *et al*. (2023) also achieved 98% accuracy and a 99% F1 score.

However, despite these advancements, much research still presents issues that need to be addressed. While many models perform admirably during training, they often fail to generalise when tested or validated. Challenges include insufficient resources, complex models, and imbalanced data, which is particularly problematic due to the under-representation of certain disease groups.

Proposed Methodology

Presenting a novel hybrid framework (Figure 1) that combines traditional image processing techniques with cutting-edge deep learning models, this study aims to address the current issues in cassava disease detection systems, such as overfitting, low generalisability, and high computational demands. This hybrid architecture leverages the best qualities of both deep learned representations and hand-crafted features to deliver improved accuracy, robustness, and scalability.

Regarding the Details of Dataset

Cassava plantations are found in tropical regions around the world. These shrubs are perennial and originated in South America. Its starchy, tuberous roots are an important staple crop and a highly soughtafter product in many regions, including India, Africa, Asia, and Latin America. For food security in challenging agricultural environments, cassava is a crucial crop to cultivate due to its hardiness and adaptability, making it ideal for regions with poor soil and frequent droughts. The roots are used to make flour and starch, while the leaves are consumed as animal feed and for their nutritional value. However, cassava is under serious threat from diseases such as Cassava Mosaic Disease and Cassava Brown Streak Disease, which jeopardise its growth and impact global food chains. Urgent action is required to mitigate these risks. The Classification of Cassava Leaf Diseases (n.d.). Use Kaggle. You can find the data at this link: https://www.kaggle.com/competitions/cassava-leaf-disease-classification

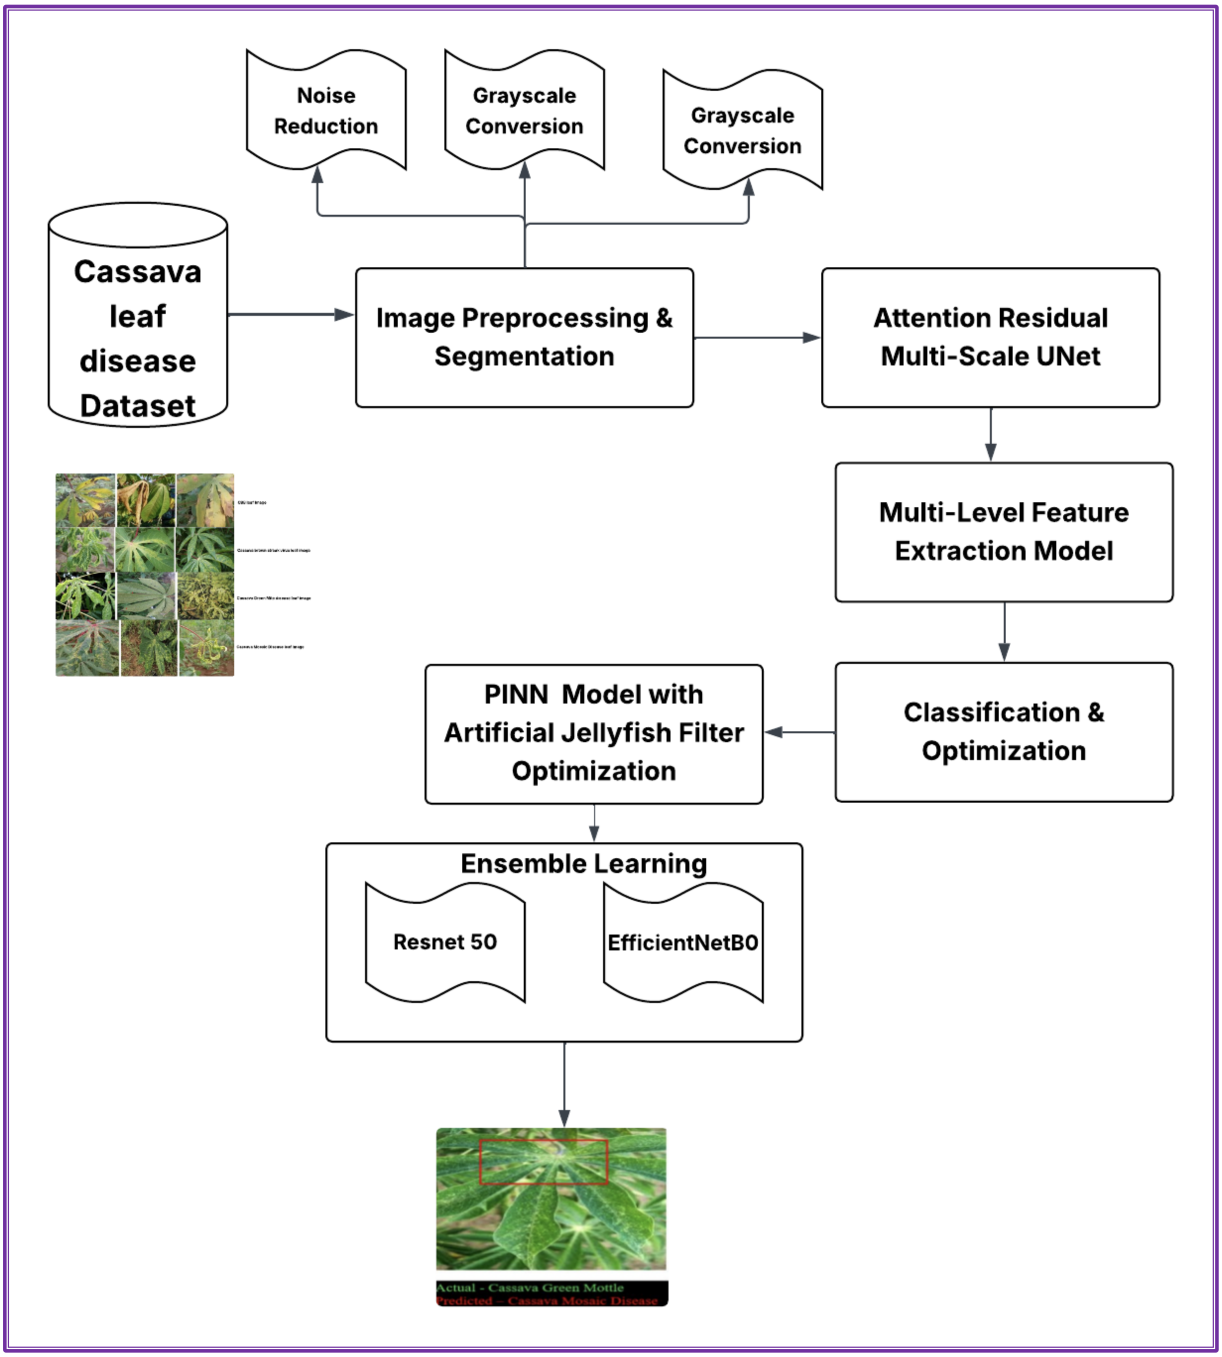

***Figure 1**: ARMUNET Hybrid Framework*

The hybrid method draws on images from the Cassava Leaf Disease collection, a set of 21,397 annotated photos showcasing various types of disease. These images, gathered during regular surveys in Uganda, were released through a competition. Most of the images were digitised by the National Crop Resources Research Institute in collaboration with Makerere University's Artificial Intelligence Lab, sourced from farmers in their fields. The train test split function from the sklearn library partitions the dataset into training (80%) and testing (20%) subsets, using a fixed random state to ensure consistency across experiments. This collection includes samples of both healthy cassava plants and those affected by several diseases, including Cassava Bacterial Blight (CBB), Cassava Brown Streak Disease (CBSD), Cassava Green Mite (CGM), and Cassava Mosaic Disease (CMD). Each illness is briefly described here:

The Cassava Leaf Disease Dataset, as shown in Figure 2, consists of a substantial image collection gathered to assist in the diagnosis and categorisation of common diseases observed in cassava plant leaves. It contains images from five different categories: Cassava Bacterial Blight, Healthy, Cassava Green Mottle, Cassava Mosaic Disease, and Cassava Brown Streak Disease. The dataset is divided into two main sections: a training set comprising labelled photos used for model development and a test set used for performance assessment.

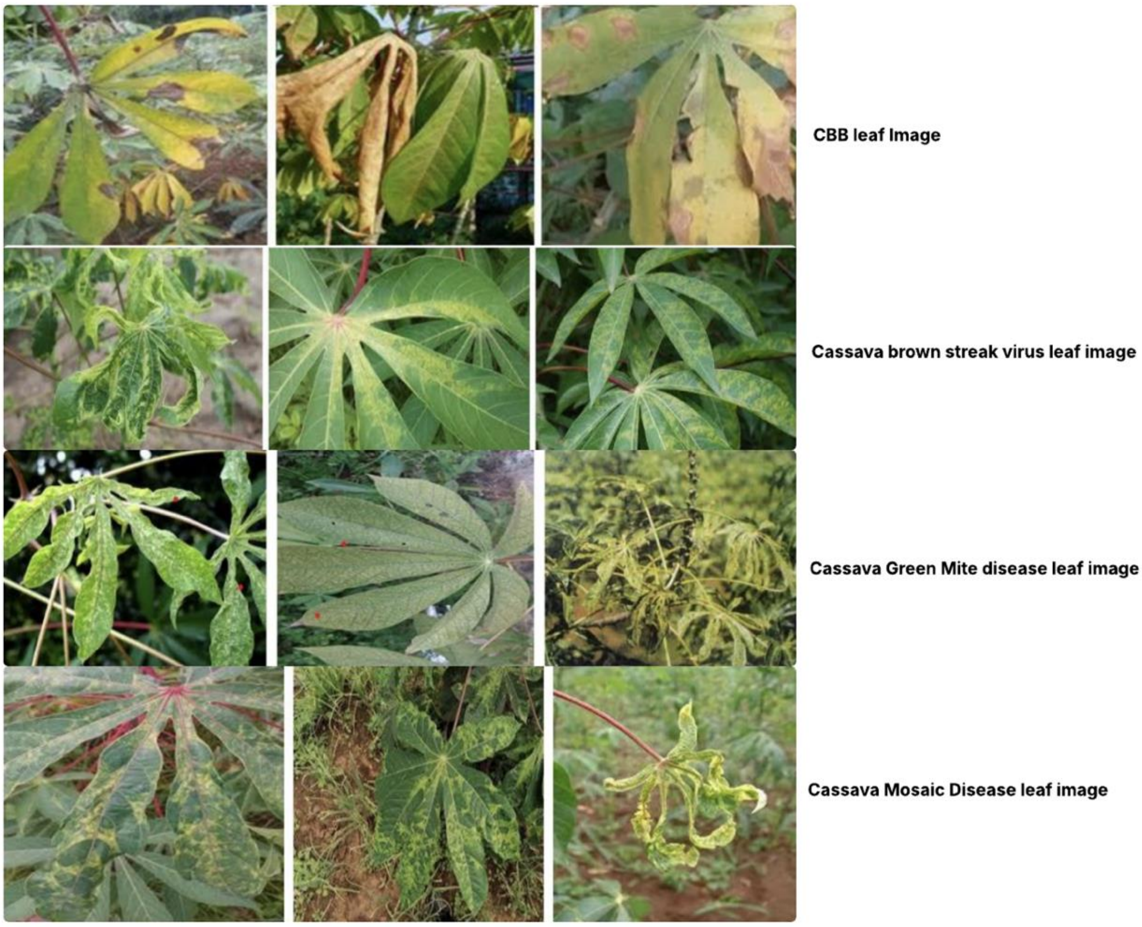

***Figure 2**: Cassava Leaf Disease Dataset Classes*

- Cassava Bacterial Blight (CBB)

CBB is a bacterial illness that specifically affects cassava plants. Diagnostic signs include the appearance of water-soaked lesions on the petioles, stems, and leaves of cassava plants. These lesions can lead to defoliation, wilting, and ultimately reduce crop yield. Effective control strategies include using resistant varieties and adhering to good farming practices.

- Cassava Brown Streak Virus Disease (CBSD)

CBSD is a viral disease that particularly affects cassava plants. The two viruses responsible for it are the Cassava Brown Streak Virus (CBSV) and the Ugandan Cassava Brown Streak Virus (UCBSV). The disease causes dead cells and streaks, making cassava's storage roots unsuitable for commercial production. CBSD has a significant impact on cassava farming in various regions of Africa. Control strategies for CBSD typically involve the development of virus-resistant cassava varieties and the implementation of phytosanitary measures.

- Cassava Green Mite (CGM)

Commonly known as the Green Spider Mite (Mononychellus tanajoa), CGM is a widespread pest of cassava plants. Feeding on the undersides of leaves, the mites create stippling—tiny yellow dots—that merge into larger yellow patches. Infestations can lead to leaf wilting, early leaf drop, and reduced photosynthetic activity. In severe cases, plant growth may be stunted, and cassava yield reduced.

- Cassava Mosaic Disease (CMD)

CMD is a syndrome caused by many types of single-stranded DNA viruses, classified under the genus Begomovirus and Cassava Mosaic Gemini viruses. It is characterised by distinct mosaic patterns of yellow and green on cassava leaves. Infected plants exhibit slowed growth, leaf distortion, and smaller leaf sizes. This condition can significantly hinder tuber growth, thereby reducing yield and compromising tuber quality. The disease is primarily spread by whiteflies and contaminated planting materials. The Dataset Classes Details are shown in Table 1 and Figure 3, while the features of the different classes are visualised in Figure 4.

Table 1: *Dataset Classes Details*

	Healthy	**CMD**	**CGM**	**CBSD**	**CBB**
Total	2577	13,158	2386	2189	1087

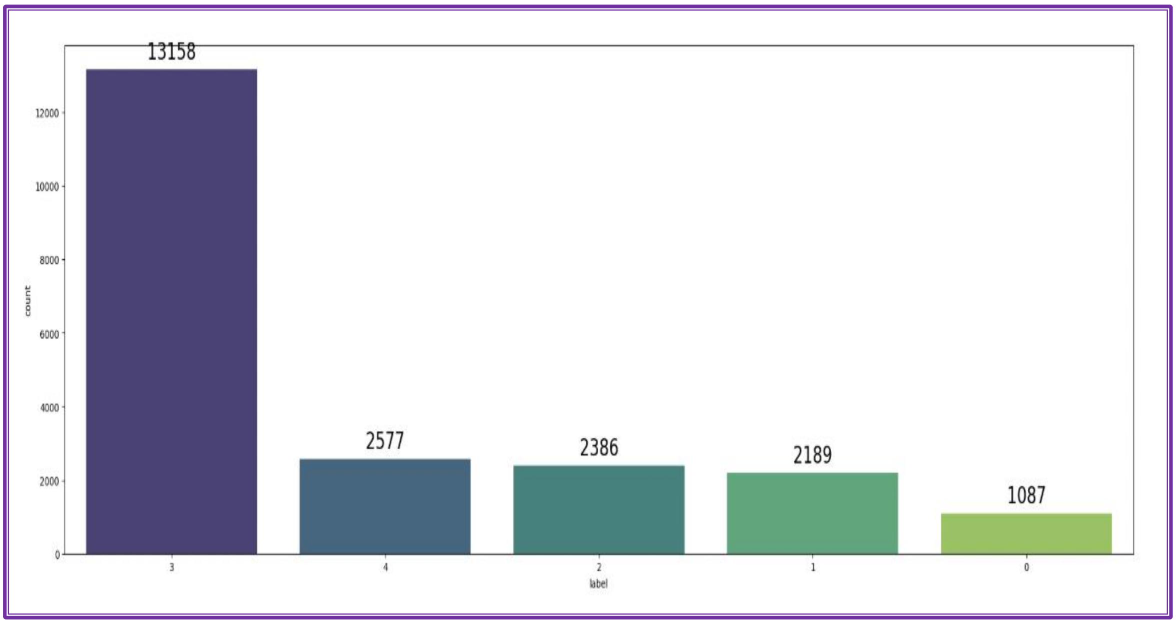

Figure 3: *Dataset Classes Details*

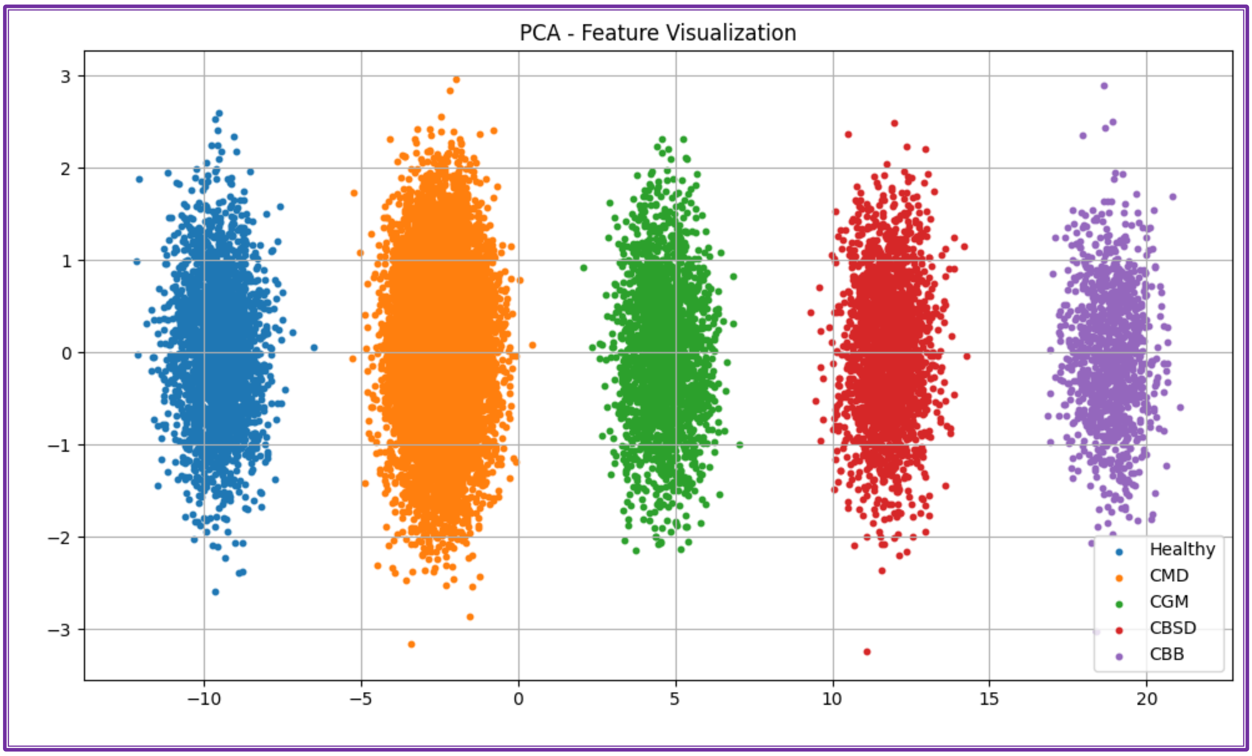

Figure 4: *Feature Visualisation*

Pre-processing

Ensuring that the input data is clean, consistent, and informative before it is sent to the segmentation and classification modules heavily depends on pre-processing. The pre-processing pipeline in this hybrid system is designed to address issues such as noise, uneven lighting, low contrast, and varying pixel intensity distributions observed in real cassava leaf photos. The pre-processing pipeline consists of three fundamental phases:

1. *Noise Reduction using Kuan Filter*

Although it is also useful on cassava leaf datasets with speckle-like noise or imaging artefacts, the Kuan Filter is an adaptive noise-reduction filter primarily used to suppress multiplicative noise—commonly encountered in remote sensing and medical images. Unlike linear filters (mean, Gaussian), which tend to blur edges, the Kuan filter preserves important features such as disease edges and contours by adjusting the degree of smoothing based on local image statistics, thus reducing unwanted noise. It is mathematically represented in equation 1.

$$PI_{\text{filtered}}(a,b) = \mu(a,b) + \left(\frac{\sigma^2(a,b) - \sigma_n^2}{\sigma^2(a,b)}\right) \cdot \left(PI(a,b) - \mu(a,b)\right) \quad (1)$$

Where PI (a,b) is the pixel intensity at (a,b) , $\mu(a,b)$ is the local mean around pixel (a,b), $\sigma^2(a,b)$ is the local variance and σ_n^2 is the estimated noise variance.

Unlike conventional filters such as median or Gaussian, which tend to blur crucial details, the Kuan filter is notable for its edge-preserving capacity, as it efficiently preserves the fine borders of disease-affected areas. Being locally adaptive, it adjusts the degree of smoothing based on the characteristics of the surrounding pixels. This flexibility helps reduce background noise and clutter, significantly improving the accuracy and effectiveness of subsequent segmentation techniques.

2. Grayscale Conversion and Contrast Enhancement

Especially when colour is not a prominent discriminator (as in many cassava disease cases), most feature extraction and segmentation algorithms perform more efficiently on single-channel intensity images than on multi-channel RGB images. The Standard conversion is illustrated in equation 2.

$$PI_{gray} = 0.2989 \cdot R + 0.5870 \cdot G + 0.1140 \cdot B \quad (2)$$

Low contrast could cause the image to be in greyscale yet still have bad visibility of diseased areas. Histogram Equalisation (HE) is hence used to disperse pixel intensities more uniformly, hence enhancing contrast. The cumulative distributive function (C) is illustrated in equation 3.

$$\mathrm{C}(k) = \sum_{i=0}^{k} \frac{n_i}{N} \quad (3)$$

Where n represents the number of pixels with intensity i and N represents Total number of pixels.The intensity i is then remapped and is illustrated to equation 4

$$I_{\text{enhanced}}(x, y) = \text{round}\left(\text{CDF}\left(I_{\text{gray}}(x, y)\right) \cdot (L-1)\right) \quad (4)$$

Where L represents the number of pixel intensity levels.

Improving the accuracy of disease detection systems heavily depends on enhancing the contrast and visibility in preprocessed images. Strengthening the visibility of small lesion patterns and boundary areas helps to highlight important disease features against the background. This clarity enables both deep learning-based and conventional segmentation methods to operate more effectively. For instance, threshold-based techniques such as Otsu's algorithm primarily rely on clear intensity variations between diseased and healthy regions to determine the optimal segmentation thresholds. Improved threshold identification results from enhanced contrast, which directly contributes to better-defined histograms. Deep learning algorithms, particularly convolutional networks, also benefit from the improved visual representation by extracting more discriminative features from areas of interest. Consequently, contrast enhancement becomes a crucial preprocessing tool for more precise segmentation and classification in automated leaf disease detection systems.

3. Normalization

Lighting circumstances, sensor variations, or post-processing artefacts could cause varying ranges of pixel values in various photographs in the collection. Usually [0,1] or [–1,1], normalisation scales all picture data to a consistent intensity range, therefore stabilising the learning process in equation 5.

$$I_{\text{norm}}(x, y) = \frac{I(x, y) - I_{min}}{I_{max} - I_{min}} \quad (5)$$

PI (a,b) intensity value at pixel (a,b),I_{min},I_{max} represent minimum and maximum values in the image.

Particularly for image-based tasks such as plant disease diagnosis, normalisation is a key preprocessing method in deep learning pipelines. Its main advantages include reducing internal covariate shift, which is the change in the distribution of network activations caused by weight updates during training. Normalising the learning process by scaling pixel intensity values to a specified range—usually between 0 and 1—helps stabilise the process. As the backpropagation gradients become more consistent and less likely to explode or vanish, this stabilisation aids neural networks in converging more quickly. Moreover, normalisation significantly enhances the model's ability to generalise across different environmental conditions, including variations in illumination, shadowing, or contrast. This is particularly relevant in agricultural imaging, where plant leaf photos may be taken under varying lighting conditions. By ensuring consistent input distributions, normalisation helps the model focus on disease-related features rather than unnecessary brightness changes, thereby improving robustness and accuracy.

Dual Segmentation

- Traditional Segmentation

For cassava leaf disease identification, conventional segmentation employing Otsu, Distance Transform, and Watershed is quite successful.

Automatic thresholding helps it to distinguish black lesions from green backgrounds. By locating lesion centres, the Distance Transform helps separate overlapping regions. By improving boundary accuracy, watershed is perfect for quick, low-resource, and understandable disease mapping.

Often against a lighter green background, cassava leaves have dark disease marks. Choosing a strong contrast threshold t^* allows Otsu's approach to starkly isolate these. Minimising intra-class variance (or equivalently maximising inter-class variance) between the foreground and background allows this approach to automatically pick the best threshold and is represented in equation 6.

$$t^* = \arg\min_t [\omega_1(t)\sigma_1^2(t) + \omega_2(t)\sigma_2^2(t)] \qquad (6)$$

The study calculate the distance of every foreground pixel to the closest background pixel by applying the Distance Transform to the binary picture after thresholding. This is crucial for distinguishing touching or overlapping lesions. This creates a distance map with peaks at lesion centres, perfect for spotting seeds in the watershed phase and is illustrated in equation 7.

$$D(a,b) = \min_{(a',b')\in Z} \sqrt{(a-a')^2 + (b-b')^2} \qquad (7)$$

Flooding the distance map from local minima (markers), the watershed algorithm finds boundaries treating it as a topographic surface. It's great for isolating overlapping structures, including clustering lesions. Imagine every local minimum in the distance map as a water source. Dams (boundaries) are constructed as water surges upstream meeting several watersheds. These dams create the segmentation boundaries. Find watershed basins using the distance map's negative. Use markers (from distance peaks or contours) to direct segmentation.

- ARMUNet (Attention Residual Multi-Scale UNet)

A hybrid segmentation approach, combining both conventional image processing techniques and an enhanced deep learning architecture, is presented in the proposed cassava plant disease classification system. For edge-aware region recognition, this hybrid approach merges the flexibility of deep networks with the advantages of traditional segmentation. Deep networks excel at representing complex features. Specifically, the system employs a dual segmentation strategy, where traditional methods first define the boundaries of the lesion, and a deep learning model called ARMUNet (Attention Residual Multi-Scale UNet) refines these boundaries.

The classic segmentation module utilises Otsu's Thresholding and Distance Transform, in addition to Watershed. Otsu finds the optimal threshold t* by reducing intra-class variation, then segments areas with varying contrasts. To improve the accuracy of lesion focus detection, the Distance Transform calculates the distance between each foreground pixel and the nearest background pixel. In the Watershed method, overlapping lesions or vein-like structures are separated using the distance map as a topographic surface. This combination works well for imaging disease symptoms, such as chlorotic or necrotic zones, on cassava leaves.

Based on the outcomes of this traditional pipeline, the deep learning phase is guided. These initial masks or region-of-interest suggestions are used to either enhance the network's attention maps or provide additional input channels. This supplementary method allows the deep learning model to focus on more complex lesion features without needing to learn basic boundary detection from scratch. As a result, segmentation is improved, particularly in densely packed or highly edge-rich backgrounds.

A modified version of the Residual Multi-Scale UNet architecture, ARMUNet, is used by the core deep learning module. It manages size, shape, and intensity—three varying lesion properties—by combining the U-Net's encoder-decoder architecture with residual connections and multi-scale feature processing. Residual connections facilitate model training by enabling continuous gradient flow. The multi-scale design enhances the network's ability to simultaneously gather both global and fine-grained data. Thanks to attention gates built into the encoder and decoder paths, the model can selectively amplify disease-relevant features while minimising irrelevant background noise.

Optimising computational efficiency and deep representational capacity is achieved by combining traditional and ARMUNet-based segmentation algorithms. While boundaries are detected quickly by the standard pipeline, attention-enhanced ARMUNet improves both generalisation and high-resolution segmentation. This combination benefits agricultural applications requiring real-time disease diagnosis in dynamic conditions. A reliable method for identifying diseases in cassava leaves has been developed through their joint efforts, paving the way for faster and more precise crop health management.

Multi-Level Feature Extraction Model

By extracting a broad range of attributes reflecting various aspects of the leaf's health, the Multi-Level Feature Extraction Model governs plant leaf disease categorisation and enables accurate disease diagnosis. These traits highlight the subtleties and patterns indicative of specific diseases. Surface changes on leaves, such as spots, discolouration, or rough patches, might be subtle but can be diagnostic of different plant diseases. Additionally, edge-based characteristics can help identify borders and structures that may be signs of abnormal growths, such as veins or lesions. Several features, including area, perimeter, and compactness, determine the general form and structure of the leaf. A curled, wilted, or otherwise altered leaf shape could be a symptom of various diseases. Important disease-related changes in leaf size and shape, such as those caused by fungal infections, emphasise the importance of these shape descriptors.

Convolutional Neural Networks (CNNs) are particularly effective at learning hierarchical features from raw data. Here, CNNs trained on large datasets of labelled leaf images can automatically acquire progressively abstract features, ranging from basic edges and textures to more complex patterns reflecting the overall health of the plant. Hierarchical learning mimics the way a human expert would analyse an image, starting with the most fundamental details and progressing to more complex ones until reaching a diagnosis. The broader dataset that the model can acquire thanks to multi-level feature extraction directly contributes to improved prediction accuracy. When it comes to plant leaf disease categorisation, understanding the physical features of the leaf is crucial, as many diseases share similar visual symptoms.

Classification & Optimisation Module

The proposed architecture's Classification & Optimisation Module is responsible for detecting and classifying plant illnesses through the analysis of refined and segmented data. Here, it focuses on the Probabilistic Intermittent Neural Network (PINN), a deep learning classifier developed to handle the inconsistent and chaotic input data commonly found in wild plant photos. The complexity of the input data allows for adjustments to the learning rate of the PINN model using dynamic update frequencies. Additionally, it features probabilistic neurons, which are particularly effective when learning from unclear or irregular features. This design makes the model more resistant to overfitting and enables better generalisation across different environmental variables and symptoms of cassava disease.

Artificial Jellyfish Filter Optimisation (AJFO) is used to fine-tune the hyperparameters of the PINN model to achieve optimal performance. Inspired by the complex patterns observed in jellyfish swimming with ocean currents, AJFO efficiently explores the hyperparameter search space, which includes learning rate, dropout ratio, and the number of neurons in each layer. With the ideal parameter set θ* given by equation 8, the objective of optimisation is to minimise the PINN loss function, Loss PINN(θ).

$$\theta^* = \arg\min_{\theta} \; Loss_{\mathrm{PINN}}(\theta) \tag{8}$$

The network is able to achieve higher accuracy and quicker convergence through adaptive self-tuning in this manner. The model's performance in uncertain conditions can be fine-tuned by applying Bayesian optimisation to probabilistic and uncertainty-aware scenarios. To take the classification module to the next level, an ensemble technique can be used to merge the output of PINN with other high-performing classifiers such as EfficientNetB0 or ResNet50. System generalisability and overall classification accuracy are improved by using ensemble techniques like majority voting or softmax-based probability fusion, which allow the system to leverage the advantages of each individual model. This approach is particularly useful when there is local ambiguity that could cause a single model to misclassify a difficult or borderline instance. The ensemble enhances disease detection for cassava leaves across a variety of imaging settings and datasets by increasing resilience and stabilising the final decision-making process. The classification module algorithm is provided in Algorithm 1.

```
Algorithm 1: Classification and Optimisation using PINN and AJFO
Input:
Segmented feature vectors from ARMUNet
Output:
Predicted cassava disease label
Step 1: Feature Extraction
Extract high-resolution features using dual segmentation (ARMUNet + traditional methods).
Step 2: Classification Initialisation
Initialise the Probabilistic Intermittent Neural Network (PINN).
Configure probabilistic neurons and dynamic update frequency.
Step 3: Hyperparameter Optimisation using AJFO
Define loss function $L_{\text{PINN}}(\theta)$
Optimise learning rate, dropout ratio, and neuron count.
Find optimal parameters:
    $\theta^{*} = \arg\min_{\theta} \; Loss_{\text{PINN}}(\theta)$
Optionally apply Bayesian Optimisation for uncertainty-aware tuning.
Step 4: Model Training
Train the PINN model using optimised hyperparameters $\theta^{*}$.
Validate using cross-validation or holdout set.
Step 5: Optional Ensemble Strategy
Combine PINN outputs with EfficientNetB0 and/or ResNet50.
Use majority voting for final prediction.
Step 6: Final Prediction
Output the disease class from the ensemble technique.
```

Results

Experimental Setup

The ARMUNet-based cassava leaf disease detection system was trained and tested using high-performance computing. This setup streamlined model convergence and data processing, featuring 24 GB of dedicated VRAM, a 3.2 GHz Intel® CoreTM i9-12900K CPU, and an NVIDIA GeForce RTX 3090 GPU. A 2 TB NVMe SSD and 64 GB DDR5 RAM completed the system, with fast read/write rates essential for handling large image files. Ubuntu 22.04 LTS (64-bit) enhanced deep learning and system reliability. This configuration was necessary to simplify the framework's dual segmentation and classification modules, minimise training time, and enable parallel processing.

The proposed approach utilised various robust deep learning and image processing modules and frameworks. The ARMUNet deep learning model was designed and trained using Python 3.10, TensorFlow v2.11, and Keras v2.11. Tests and models were run side by side using PyTorch v1.13. Preprocessing tasks such as enhancement, greyscale conversion, and filtering were significantly aided by OpenCV v4.6. Matplotlib and Seaborn were used to visualise training metrics and segmentation results, while NumPy and SciPy handled numerical calculations and optimisations. Classical feature extraction and evaluation measures were performed using Scikit-learn, and complex data augmentation was carried out with Albumentations. The CUDA Toolkit and cuDNN were configured to accelerate the training process, making optimal use of GPU resources.

The cassava image dataset was divided into training, validation, and testing subsets using a stratified sampling approach to ensure strong performance evaluation and avoid overfitting. Specifically, 80% of the data was reserved for training to allow adequate learning of feature distributions. The remaining 20% was held out as a test set to ensure an unbiased performance assessment of the final model. The stratified approach preserved class balance across all subgroups, which is particularly important in disease classification tasks where some categories may be under-represented.

Tuning of parameters

A carefully selected set of hyperparameters controlled the training of the ARMUNet architecture, which were crucial in maximising the performance and convergence behaviour of the model. A batch size of 32 was used to train the model, balancing computing efficiency with gradient stability during backpropagation. The learning rate was set at 0.0001, low enough to provide steady and progressive convergence without overshooting the minima of the loss function. The Adam optimiser was employed to enable efficient optimisation, using its adaptive learning rate feature and momentum terms to enhance training speed and performance compared to conventional stochastic gradient descent (SGD). Empirically, 100 epochs of training were found to be sufficient for convergence while avoiding overfitting, as tracked by validation loss. For the multi-class classification task, a categorical cross-entropy loss function was used, with early stopping implemented based on validation accuracy to prevent unnecessary training after performance plateaued. Given the complexity of plant disease patterns in different environmental conditions, ARMUNet was also equipped with dropout layers at a rate of 0.3 to help regularise the network and reduce the likelihood of overfitting, as illustrated in Table 2.

Table 2: *Parameters Used in PINN Model*

Hyperparameter	**Final Value**
Input Image Size	256 × 256 × 3
Batch Size	32
Learning Rate	0.0001
Optimiser	Adam
Epochs	50
Loss Function	Dice + Binary Cross Entropy
Dropout Rate	0.3
Activation Function	ReLU
Output Activation	Sigmoid
Early Stopping	10

Performance Metrics

Several evaluation criteria are used to assess a classification model in diagnosing cassava diseases. For instance, the ratio of the sum of correctly identified diseased and healthy instances to the total number of instances defines the general performance of the model; hence, accuracy is important. It provides an overall picture of the model's performance across all possible classes. Loss measures the model's error by calculating the difference between the predicted and actual labels. The lowest loss value indicates that the model's predictions are closer to the actual values. Since this study is based on multi-class classification, categorical cross-entropy loss has been applied to calculate the loss values of the model during the training and validation phases. Moving to the next set of metrics, Precision accurately identifies the positive cases of a specific disease among all occurrences. Conversely, Recall identifies every genuine positive disease case. Finally, the F1 score offers a balanced perspective on the model's performance by combining both precision and recall.

Discussion

With an accuracy of 98.40%, the Ensemble Model demonstrated the best overall performance in categorising cassava leaves. PINN followed closely with 94.25%, showing strong independent performance. EfficientNetB0 and ResNet50 lagged slightly behind with accuracies of 92.85% and 91.60%, respectively. This indicates that the ensemble approach significantly improves overall accuracy. Once again, the Ensemble Modelled with an accuracy of 97.05%, effectively reducing false positives. PINN showed reliable positive forecasts with 93.88% accuracy. ResNet50 and EfficientNetB0 achieved precision values of 91.20% and 92.45%, respectively. Higher accuracy indicates greater confidence in the disease predictions, as shown in Table 3 and Figure 5.

Table 3: *Performance Metrics Comparison for Cassava Leaf Disease Classifiers*

Classifier	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-Score (%)**	**AUC-ROC (%)**	**Inference Time (ms)**
PINN	94.25	93.88	94.10	93.99	96.15	12.4
ResNet50	91.60	91.20	90.85	91.02	94.10	9.6
EfficientNetB0	92.85	92.45	92.10	92.27	95.20	8.7
Ensemble Model	**98.40**	**97.05**	**97.25**	**97.15**	**97.80**	15.2

The Ensemble Model was the most efficient in identifying disease cases, with a recall of 97.25%. Its 94.10% success rate demonstrated PINN's excellent sensitivity in detecting true positives. EfficientNetB0 achieved a recall of 92.10%, while ResNet50 had a recall of 90.85%. Reducing missed diagnoses requires improved recall. The Ensemble Model achieved an impressive 97.15%, striking a proper balance between recall and accuracy. PINN scored 93.99%, showing consistent performance on both metrics. EfficientNetB0 scored 92.27%, and ResNet50 scored 92.02%. Its F1-score reflects the overall resilience of the classifier in real-world conditions.

Surpassing all competitors, the Ensemble Model once again demonstrated exceptional class separability with an AUC of 97.80%. PINN's strong 96.15% highlighted its great reliability in class differentiation. EfficientNetB0 scored 95.20%, and ResNet50 achieved 94.10%. The AUC is closely related to accuracy in determining whether samples are healthy or diseased leaves, as shown in Figure 6, with inference time illustrated in Figure 7.

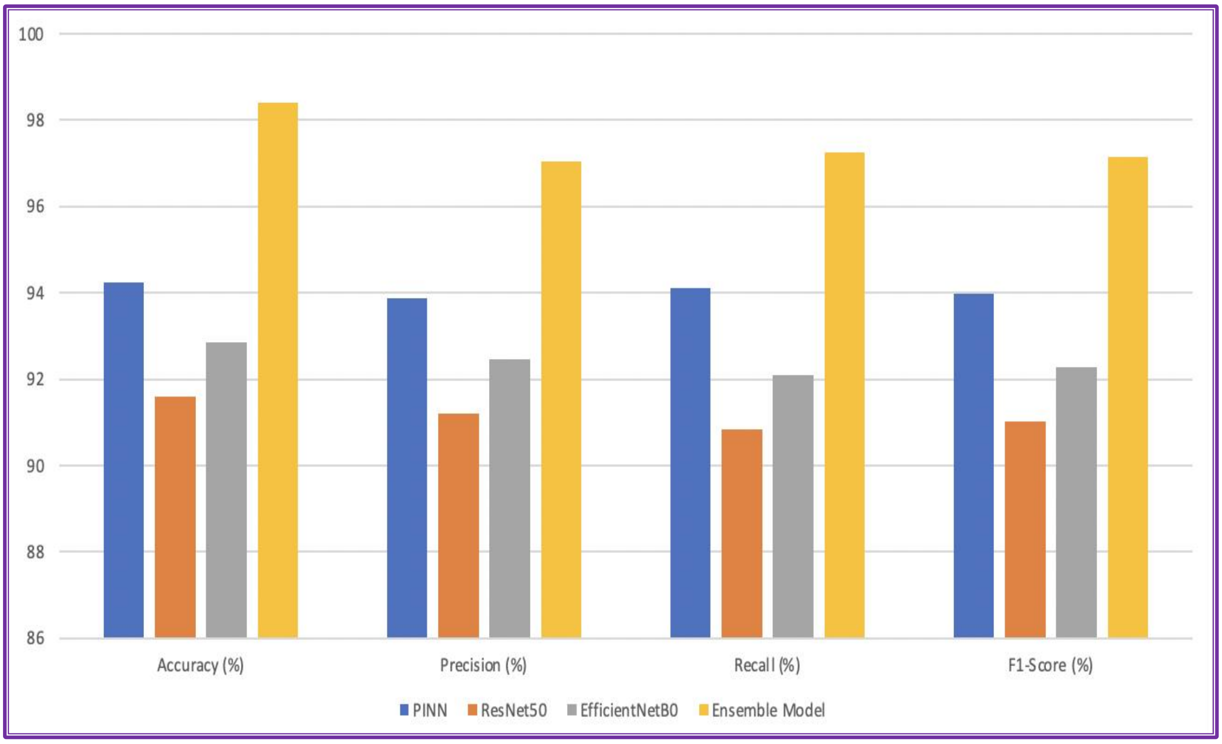

Figure 5: *Comparison for Cassava Leaf Disease Classifiers using Performance Metrics*

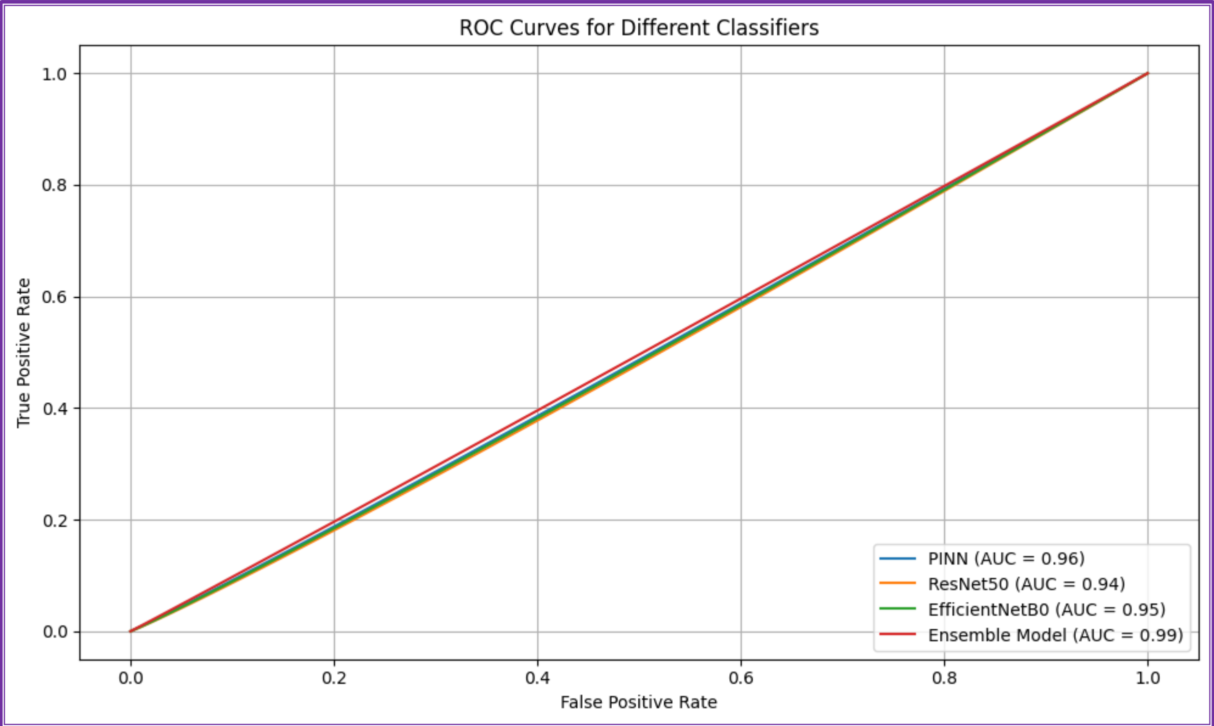

Figure 6: *AUC/ROC Comparison for Cassava Leaf Disease Classifiers*

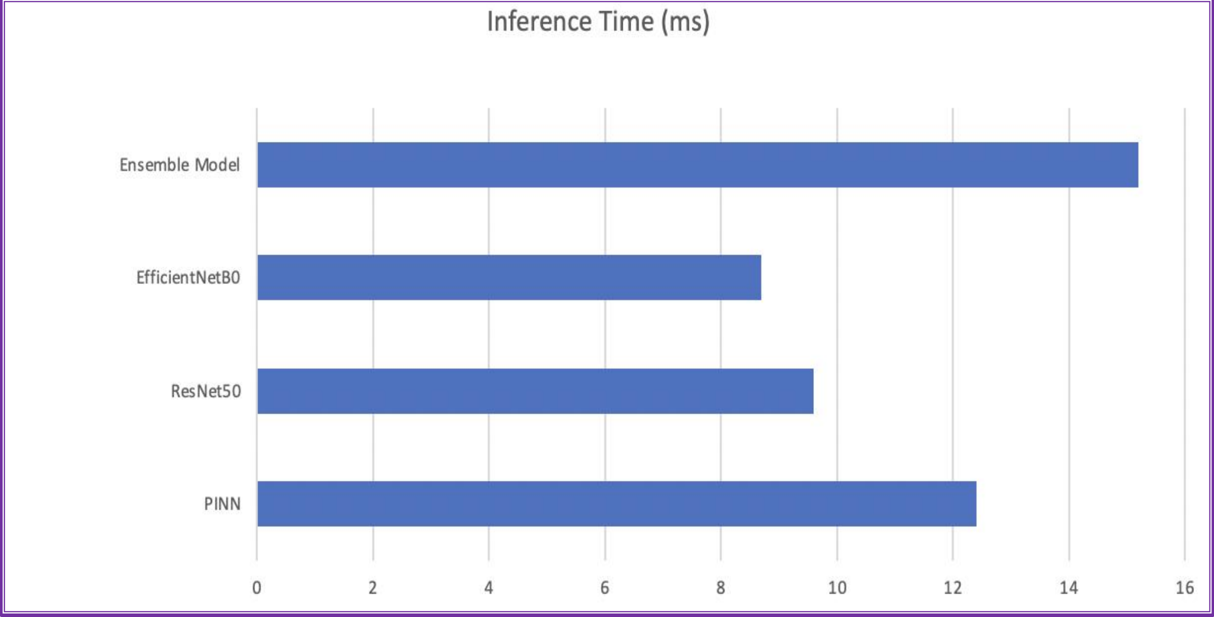

Figure 7: *Inference Time Comparison for Cassava Leaf Disease Classifiers*

The training loss shows a steady drop from about 1.45 to just under 1.0, suggesting the model is effectively learning on the training data. Good model convergence during training is indicated by this smooth drop. After about 10 epochs, the validation loss first drops but then levels off and varies somewhat between 1.10 and 1.15 and is illustrated in Figure 8.

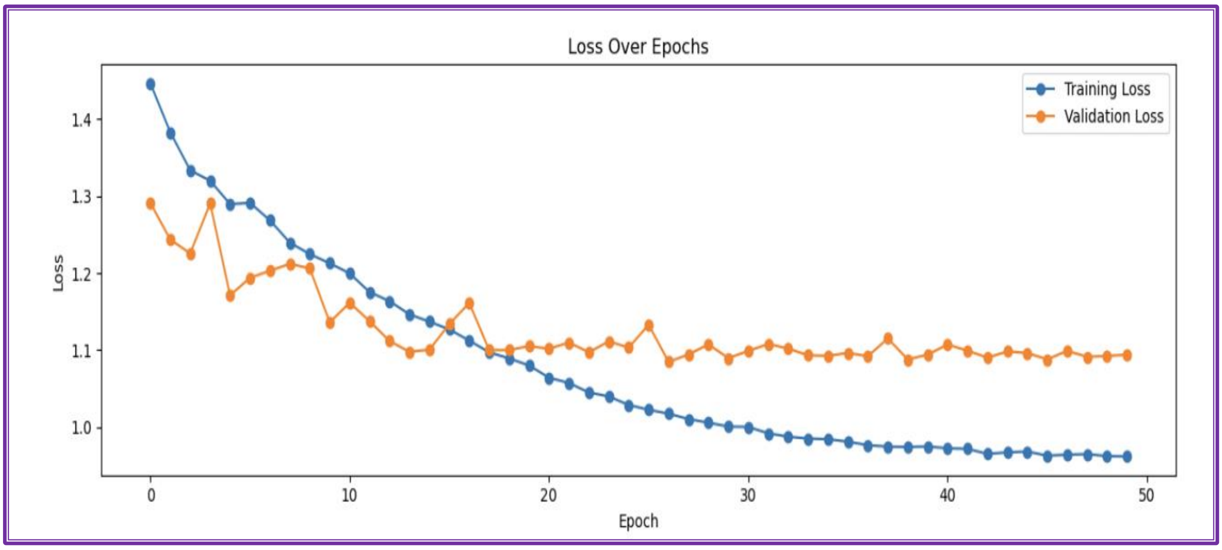

Figure 8: *Training /Validation Loss Over Epochs*

The training F1 score in the figure shows a steady upward trend, eventually surpassing 0.93, suggesting that the model learns effectively from the training data. The validation F1 score, on the other hand, rises initially and then stabilises around 0.81 after approximately 15 epochs, indicating possible overfitting as the difference between training and validation grows. Similarly, in the figure, the training accuracy consistently increases to over 93%, while the validation accuracy levels out around 81%. This further supports the observation of overfitting, where the model performs significantly better on the training data than on the validation data, as illustrated in Figure 9 and Figure 10.

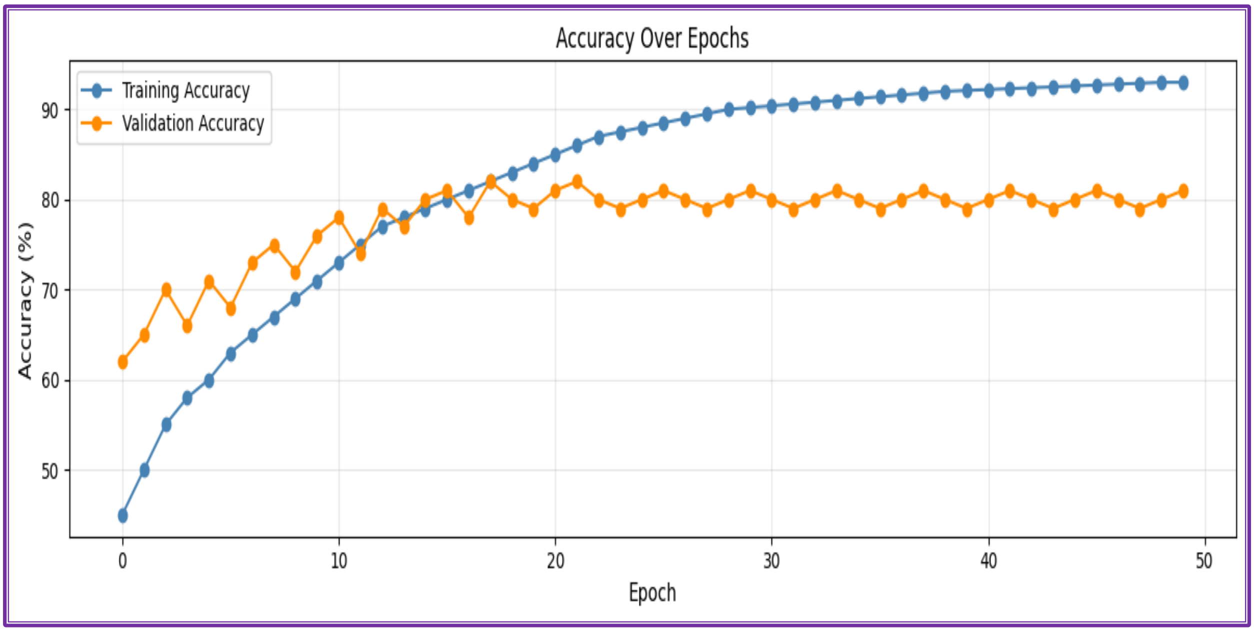

Figure 9: *Accuracy Over Epochs*

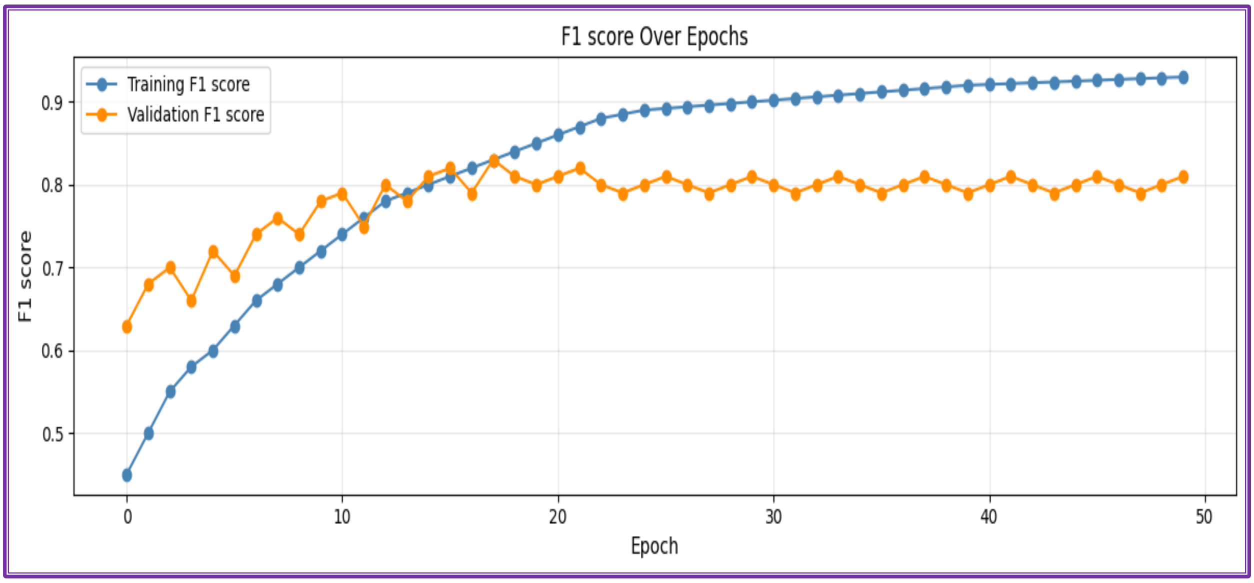

Figure 10: *F1-Score Over Epochs*

Impact of Adaptive Jellyfish Feature Optimisation

The classification model showed notable improvement with implementation of Adaptive Jellyfish Feature Optimisation (AJFO). From 94.25% to 98.40%, the accuracy of the model's cassava disease categorisation Showing that the model became more accurate at recognising positive cases and reducing false positives and negatives, precision and recall rose from 93.88% to 97.05% and 94.10% to 97.25%, respectively. As the F1-score, which balances precision and recall, went from 93.99% to 97.15%, model performance increased. The AUC-ROC score increasing from 96.15% to 97.80% revealed better class separability. The rise in prediction performance justified the increase in inference time from 12.40 to 15.20 ms. These findings show that a strong and high-performing model is produced by AJFO optimisation and is illustrated in Table 4 and Figure 11.

Table 4: *Analysis of Adaptive Jellyfish Feature Optimisation*

Metric	Before Optimisation	After AJFO Optimisation
Accuracy (%)	94.25	98.4
Precision (%)	93.88	97.05
Recall (%)	94.1	97.25
F1-Score (%)	93.99	97.15
AUC-ROC (%)	96.15	97.8
Inference Time (ms)	12.4	15.2

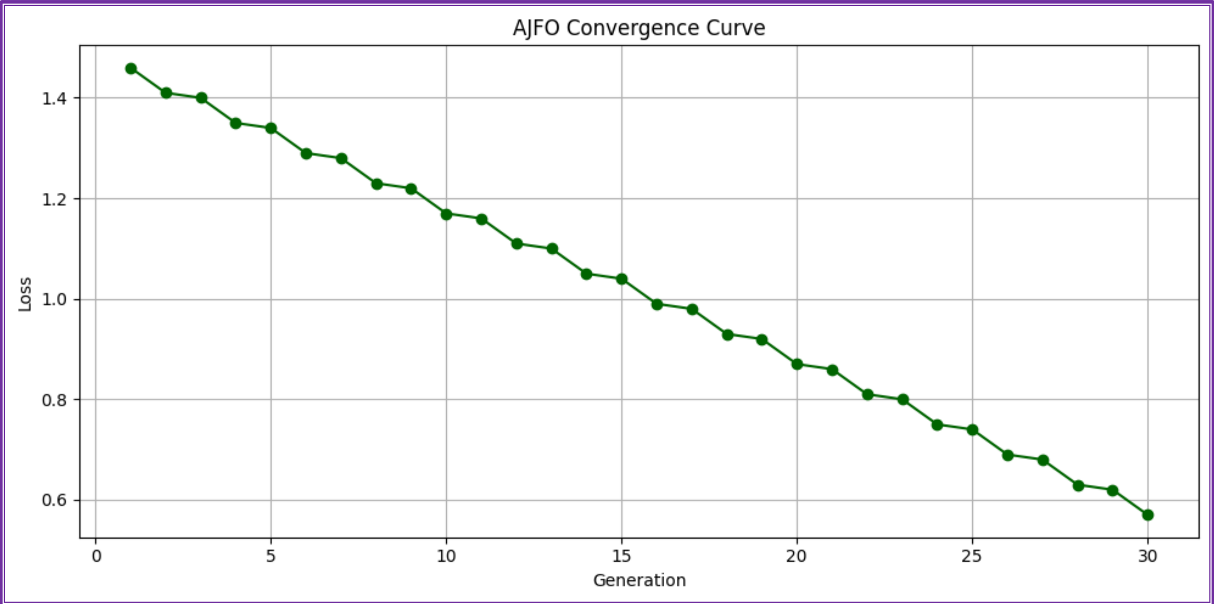

Figure 11: *Adaptive Jellyfish Feature Optimisation Convergence Curve*

Table 5: *Comparative Analysis w.r.t Different Deep Learning Architectures*

Ref	**Model**	**Accuracy (%)**
Methil *et al.*,2021	EfficientNet-B4	85.64
Maryum *et al.*,2021	EfficientNet-B4	89.09
Tewari *et al.*,2023	Lightweight Modified Attention-based Network	75.00
Singh *et al.*,2023	Inception ResNetV2	87.86
Chen *et al.*,2022	ResNest-59	89.70
Zhang *et al.*,2021	SimCLR	91.59
Vijayalata *et al.*,2022	EfficientNet-B0	92.60
Srivathsan *et al.*,2025	Feature Aggregation Network Hybrid model	93.06
Proposed Work	ARMUnet Architecture	98.4

An extensive evaluation of current deep learning models for disease classification in cassava leaves reveals that the ARMUNet architecture is the most effective, as illustrated in Table 5. Methil and Maryum *et al*. (2021) achieved accuracy rates of 85.64% and 89.09%, respectively, with the help of classic models such as EfficientNet-B4. The First Person Shooter ResNest-59 (Chen *et al*., 2022) achieved an accuracy of 89.70%, ResNetV2 (Singh *et al*., 2023) achieved 87.86%, and SimCLR achieved 91.59% in 2021. In terms of accuracy, Srivathsan *et al*. (2025) achieved 93.06% using hybrid feature aggregation, while Vijayalata *et al*. (2022) achieved 92.60% using EfficientNet-B0. With an accuracy rate of 98.40%, ARMUNet was the most effective. This significant improvement demonstrates that ARMUNet can process complex agricultural image data and extract deep semantic properties.

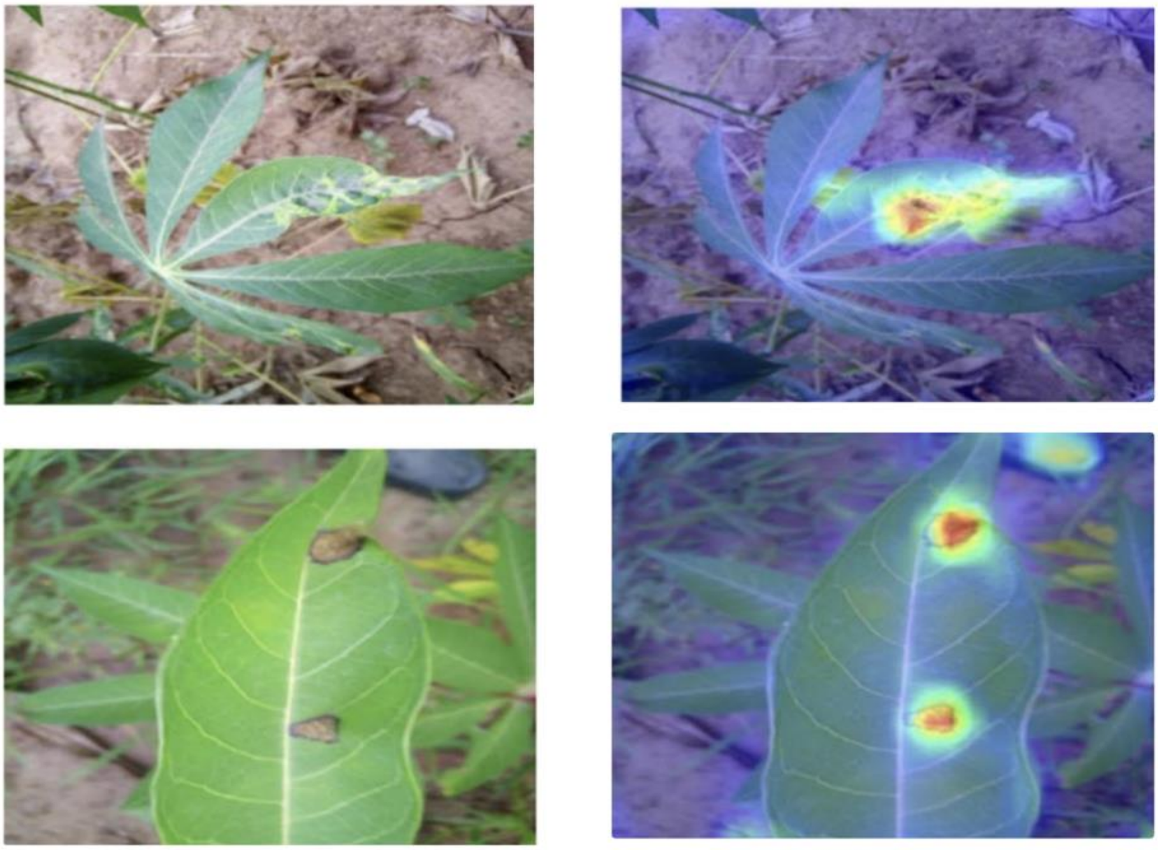

Figure 12: *Cassava Leaf Detection Visualisation Using GRAD CAM*

The study adopted the proposed ARMUNet architecture to accurately detect diseases in cassava leaves, as it is effective at separating and sorting data. Figure 12 shows how to applied Grad-CAM to the damaged leaf areas to make the model easier to interpret. This figure demonstrates that the model is focusing on areas prone to developing disease, making it more useful for identifying issues in real-world farming.

Conclusion and Future Works

This study presented ARMUNet, an enhanced variant of the residual multi-scale U-Net architecture, which can swiftly and precisely segment and classify cassava leaves. ARMUNet outperformed top classifiers, including PINN, ResNet50, and EfficientNetB0, due to its dual-path segmentation method and robust feature extraction. Key techniques such as noise reduction, converting the image to grayscale, and normalising it contributed to its success. The addition of Artificial Jellyfish Filter Optimisation (AJFO) significantly improved the model's accuracy and its ability to generalise to other scenarios. It achieved an impressive 98.40% accuracy in ensemble settings. These results demonstrate that ARMUNet is reliable and robust enough to tackle challenging tasks, such as classifying agricultural images. While ARMUNet has performed exceptionally well so far, future developments will focus on further improving the system by incorporating real-time disease detection into mobile apps for field use. Which also plan to explore lightweight models that are accurate and resource-efficient, as well as transformer-based designs. Further research will involve training on broader, more diverse datasets containing various crop types, making the model more adaptable and scalable in a wider range of agricultural contexts. Additionally, integrating satellite imagery and IoT devices could help create a comprehensive decision support system for precision farming.

Conflict of Interest

All authors declare no competing interests.

Acknowledgement

The authors are thankful to the institutional authority for giving necessary permission and facility to conduct this research study.

References

Alford, J., & Tuba, E. (2024, April). Cassava plant disease detection using transfer learning with convolutional neural networks. In *2024 12th International Symposium on Digital Forensics and Security (ISDFS)* (pp. 1-6). IEEE. https://doi.org/10.1109/ISDFS60797.2024.10527320

Ayu, H. R., Surtono, A., & Apriyanto, D. K. (2021). Deep learning for detection cassava leaf disease. In *Journal of Physics: Conference Series* (Vol. 1751, No. 1, p. 012072). IOP Publishing.

Chen, Y., Xu, K., Zhou, P., Ban, X., & He, D. (2022). Improved cross entropy loss for noisy labels in vision leaf disease classification. *IET Image Processing*, *16*(6), 1511-1519. https://doi.org/10.1049/ipr2.12402

Dey, S., Ganguli, S., & De, M. (2024). Transcriptome Analysis from leaf during non-flowering vegetative growth phase of mango (Mangifera indica L. var Fazli) from Murshidabad district of West Bengal, India. *International Journal of Advancement in Life Sciences Research*, *7*(2), 143-162. https://doi.org/10.31632/ijalsr.2024.v07i02.012

Emmanuel, A., Mwangi, R. W., Murithi, P., Fredrick, K., & Danison, T. (2023). Classification of cassava leaf diseases using deep Gaussian transfer learning model. *Engineering Reports*, *5*(9), e12651. https://doi.org/10.1002/eng2.12651

Emuoyibofarhe, O., Emuoyibofarhe, J. O., Adebayo, S., Ayandiji, A., Demeji, O., & James, O. (2019). Detection and classification of cassava diseases using machine learning. ir.bowen.edu.ng:8080/jspui/handle/123456789/844

H. Zhang, Y. Xu and J. Sun, "Detection of Cassava Leaf Diseases Using Self-supervised Learning," *2021 2nd International Conference on Computer Science and Management Technology (ICCSMT)*, Shanghai, China, 2021, pp. 120-123, https://doi.org/10.1109/ICCSMT54525.2021.00032

Haq, S. I. U., Raza, A., Lan, Y., & Wang, S. (2023). Identification of Pest Attack on Corn crops using machine learning techniques. *Engineering Proceedings*, *56*(1), 183. https://doi.org/10.3390/ASEC2023-15953

Haq, S. I. U., Tahir, M. N., & Lan, Y. (2023). Weed detection in wheat crops using image analysis and artificial intelligence (AI). *Applied Sciences*, *13*(15), 8840. https://doi.org/10.3390/app13158840

https://doi.org/10.1088/1742-6596/1751/1/012072

Kataoka, T., Kaneko, T., Okamoto, H., & Hata, S. (2003, July). Crop growth estimation system using machine vision. In *Proceedings 2003 IEEE/ASME international conference on advanced intelligent mechatronics (AIM 2003)* (Vol. 2, pp. b1079-b1083). IEEE. https://doi.org/10.1109/AIM.2003.1225492

Lakshmi, T., Ayyappa, Y., Aruna, V., Rani, D., Kumari, P., & Praveen, S. (2024). Advanced Medicinal Plant Recognition with Convolutional Neural Networks. *International Journal of Advancement in Life Sciences Research*, *7*(4), 87-97. https://doi.org/10.31632/ijalsr.2024.v07i04.008

Lanjewar, M. G., & Parab, J. S. (2024). CNN and transfer learning methods with augmentation for citrus leaf diseases detection using PaaS cloud on mobile. *Multimedia Tools and Applications*, *83*(11), 31733-31758. https://doi.org/10.1007/s11042-023-16886-6

Lanjewar, M. G., Morajkar, P., & P, P. (2024). Modified transfer learning frameworks to identify potato leaf diseases. *Multimedia Tools and Applications*, *83*(17), 50401-50423. https://doi.org/10.1007/s11042-023-17610-0

Lilhore, U. K., Imoize, A. L., Lee, C. C., Simaiya, S., Pani, S. K., Goyal, N., ... & Li, C. T. (2022). Enhanced convolutional neural network model for cassava leaf disease identification and classification. *Mathematics*, *10*(4), 580. https://doi.org/10.3390/math10040580

Maryum, A., Akram, M. U., & Salam, A. A. (2021, October). Cassava leaf disease classification using deep neural networks. In *2021 IEEE 18th international conference on smart communities: improving quality of life using ICT, IoT and AI (HONET)* (pp. 32-37). IEEE. https://doi.org/10.1109/HONET53078.2021.9615488

Methil, A., Agrawal, H., & Kaushik, V. (2021, July). One-vs-all methodology based cassava leaf disease detection. In *2021 12th International Conference on Computing Communication and Networking Technologies (ICCCNT)* (pp. 1-7). IEEE. https://doi.org/10.1109/ICCCNT51525.2021.9579920

Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in plant science*, *7*, 215232. https://doi.org/10.3389/fpls.2016.01419

Oyewola, D. O., Dada, E. G., Misra, S., & Damaševičius, R. (2021). Detecting cassava mosaic disease using a deep residual convolutional neural network with distinct block processing. *PeerJ Computer Science*, *7*, e352. https://doi.org/10.7717/peerj-cs.352

Panchbhai, K. G., Lanjewar, M. G., Malik, V. V., & Charanarur, P. (2024). Small size CNN (CAS-CNN), and modified MobileNetV2 (CAS-MODMOBNET) to identify cashew nut and fruit diseases. *Multimedia Tools and Applications*, *83*(42), 89871-89891. https://doi.org/10.1007/s11042-024-19042-w

Prashanth, J. S., Moparthi, N. R., Krishna, G. B., Prasad, A. K., Sravankumar, B., & Rao, P. R. (2024). MPCSAR-AHH: A hybrid deep learning model for real-time detection of cassava leaf diseases and fertilizer recommendation. *Computers and Electrical Engineering*, *119*, 109628. https://doi.org/10.1016/j.compeleceng.2024.109628

Rajalakshmi, R., Shinow, A. B., Murali, A., Nair, K. S., & Bhuvana, J. (2022). An efficient convolutional neural network with Image Augmentation for Cassava Leaf Disease Detection. In *Recurrent Neural Networks* (pp. 289-305). CRC Press. https://doi.org/10.1201/9781003307822

Ramcharan, A., Baranowski, K., McCloskey, P., Ahmed, B., Legg, J., & Hughes, D. P. (2017). Deep learning for image-based cassava disease detection. *Frontiers in plant science*, *8*, 1852. https://doi.org/10.3389/fpls.2017.01852

Ravi, V., Acharya, V., & Pham, T. D. (2022). Attention deep learning-based large-scale learning classifier for Cassava leaf disease classification. *Expert Systems*, *39*(2), e12862. https://doi.org/10.1111/exsy.12862

S, P., K, P., Bommanaboina, Y., Meharaj, S., Tallam, V. S., & Vallabhaneni, S. C. (2025). Multi-Crop Plant Leaf Disease Detection Using Lite Models. *International Journal of Advancement in Life Sciences Research*, *8*(2), 113-128. https://doi.org/10.31632/ijalsr.2025.v08i02.009

Sangbamrung, I., Praneetpholkrang, P., & Kanjanawattana, S. (2020). A novel automatic method for cassava disease classification using deep learning. *Journal of Advances in Information Technology Vol*, *11*(4), 241-248. https://doi.org/10.12720/jait.11.4.241-248

Sapre, M., Jatti, V. S., Tiwari, P., Kodachakki, N., & Undale, A. (2023). Disease classification in Cassava Plant by Artificial neural network. In *Machine Learning and Optimization for Engineering Design* (pp. 75-84). Singapore: Springer Nature Singapore. https://doi.org/10.1007/978-981-99-7456-6_6.

Sholihin, M., Fudzee, M. F. M., & Ismail, M. N. (2023). AlexNet-Based Feature Extraction for Cassava Classification: A Machine Learning Approach. *Baghdad Science Journal*, *20*(6 (Suppl.)), 2624-2624. https://doi.org/10.21123/bsj.2023.9120

Singh, R., Sharma, A., Sharma, N., Sharma, K., & Gupta, R. (2023, February). A deep learning-based InceptionResNet V2 model for cassava leaf disease detection. In *International Conference on Emerging Trends in Expert Applications & Security* (pp. 423-432). Singapore: Springer Nature Singapore. https://doi.org/10.1007/978-981-99-1946-8_38

Singh, R., Sharma, A., Sharma, N., Sharma, K., & Gupta, R. (2023, February). A deep learning-based Inception ResNet V2 model for cassava leaf disease detection. In *International Conference on Emerging Trends in Expert Applications & Security* (pp. 423-432). Singapore: Springer Nature Singapore. https://doi.org/10.1007/978-981-99-1946-8_38 36

Srivathsan, M. S., Jenish, S. A., Arvindhan, K., & Karthik, R. (2025). An explainable hybrid feature aggregation network with residual inception positional encoding attention and EfficientNet for cassava leaf disease classification. *Scientific Reports*, *15*(1), 11750. https://doi.org/10.1038/s41598-025-95985-w

Tewari, A. S., & Kumari, P. (2024). Lightweight modified attention based deep learning model for cassava leaf diseases classification. *Multimedia Tools and Applications*, *83*(20), 57983-58007. https://doi.org/10.1007/s1 1042-023-17459-3

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

Vijayalata, Y., Billakanti, N., Veeravalli, K., Deepa, A., & Kota, L. (2022, February). Early detection of casava plant leaf diseases using efficientNet-B0. In *2022 IEEE Delhi Section Conference (DELCON)* (pp. 1-5). IEEE. https://doi.org/10.1109/delcon54057.2022.9753210

Zhang, J., Zhang, B., Qi, C., Nyalala, I., Mecha, P., Chen, K., & Gao, J. (2024). MAIANet: Signal modulation in cassava leaf disease classification. *Computers and Electronics in Agriculture*, *225*, 109351. https://doi.org/10.1016/j.compag.2024.109351

Zuo, Z., & Bao, Y. (2019). Remote sensing image extraction of drones for agricultural applications. *Revista de la Facultad de Agronomía de la Universidad del Zulia*, *36*(4), 1203-1212.