

Explainable Hybrid Deep Learning with EfficientNetB3–DenseNet121 Feature Fusion for Banana Leaf Spot Disease Classification

Bilal Bataineh¹, Ghazi Shakah², Bajeszeyadaljunaeidia³, Issa Alsmadi⁴, Mohamed S.Sawah^{5,6*}

¹Department of Computer Science, Jadara University, Irbid, Jordan.

²Department of Software Engineering, Faculty of Information Technology, Ajloun National University, P.O.43, Ajloun-26810, Jordan

³Department of Computer Science, Faculty of Information Technology, Ajloun National University, P.O.43, Ajloun-26810, Jordan

⁴Department of Data Science and Artificial Intelligence, Faculty of Information Technology, Ajloun National University P.O.43, Ajloun-26810, Jordan

⁵Department of Data Science and Artificial Intelligence, Faculty of Information Technology, Ajloun National University, P.O.43, Ajloun-26810, Jordan

⁶Department of Information Systems, Al-Alson Higher Institute, Cairo, Egypt.
(Corresponding author)

Abstract: Banana leaf spot diseases reduce crop quality and yield, making early and accurate diagnosis essential for precision agriculture. This study proposes an accurate and explainable hybrid deep learning framework for multiclass banana leaf spot disease classification. The model combines EfficientNetB3 and DenseNet121 in dual-branch feature-fusion architecture to capture complementary visual features. A stratified 70/15/15 split was used on the BananaLSD dataset (Cordana, Healthy, Pestalotiopsis, and Sigatoka), with class weighting, data augmentation, random erasing, label smoothing, and staged transfer learning to improve generalization. Performance was evaluated using accuracy, precision, recall, F1-score, confusion matrix, and multiclass ROC analysis, while Grad-CAM was used for explainability. The proposed hybrid model outperformed single-backbone baselines (EfficientNetB3 and DenseNet121), achieving 98.58% accuracy with the lowest loss (0.2578) and strong class-wise performance. Grad-CAM visualizations further confirmed that the model focused on disease-relevant lesion regions. These results support the proposed framework as a reliable tool for automated banana disease diagnosis.

Keywords: Banana Leaf Disease Classification; Hybrid Deep Learning; Efficientnetb3; Densenet121; Transfer Learning; Feature Fusion; Explainable Ai; Grad-Cam; Plant Disease Detection; Precision Agriculture.

I. INTRODUCTION

Agriculture is indispensable for the economic development of developing nations, and it is imperative to promptly detect plant diseases. One of the most widely consumed fruits worldwide; bananas are a source of essential nutrients and contribute to the security of the global food supply. Nevertheless, a number of infections that they are communicable and can significantly reduce the productivity and quality of banana trees. Effective disease control and the prevention of infection transmission within plantations necessitate the early identification and accurate categorization of illnesses that affect banana fruit [1].

Traditional methods for the detection of banana maladies, such as field inspections and laboratory testing, are frequently time-consuming, costly, and labor-intensive. In addition, the efficiency of monitoring vast plantations is restricted by ground-based methods, particularly in remote or difficult-to-access regions. Developments in machine learning and remote sensing technologies may present novel opportunities to confront these obstacles [2].

Traditionally, banana disease detection relied heavily on classical methods, primarily expert visual inspection in the field [3]. While expert knowledge is invaluable, it is often inaccessible to small-scale farmers and is inherently subjective, time-consuming, and prone to human error. To overcome these limitations, early computer vision approaches were developed. These classical pipelines typically involved multi-stage image processing: image acquisition, preprocessing to reduce noise, and segmentation using techniques such as K-means clustering to isolate diseased spots. Features based on color, texture, and shape were then manually extracted and fed into traditional machine learning classifiers, such as Support Vector Machines (SVM) [4]. Although these methods marked an improvement over manual inspection, their performance often degraded under varying field conditions, such as inconsistent lighting, complex backgrounds, and overlapping leaves.

In recent years, the field has undergone a paradigm shift toward Artificial Intelligence (AI), specifically Deep Learning (DL). Unlike classical methods that require manual feature engineering, Convolutional Neural Networks (CNNs) automatically learn hierarchical features directly from raw image data, providing superior robustness and accuracy [5][6]. Recent studies have demonstrated the efficacy of architectures like ResNet, Inception, and MobileNet in identifying banana diseases with accuracies exceeding 95% [7].

The main contributions of this study can be summarized as follows:

1. **A novel hybrid deep learning framework** is proposed for multiclass banana leaf spot disease classification by fusing features extracted from EfficientNetB3 and DenseNet121, enabling complementary representation learning for improved discrimination among visually similar disease categories.
2. **A robust end-to-end training pipeline** is designed, including stratified data splitting, class weighting, data augmentation, random erasing, label smoothing, and staged transfer-learning fine-tuning to improve generalization and reduce overfitting.
3. **A comprehensive evaluation framework** is implemented using accuracy, precision, recall, F1-score, confusion matrix, and multiclass ROC-AUC analysis to provide a reliable performance assessment beyond a single metric.
4. **Explainable AI integration** is achieved through Grad-CAM visualizations, which qualitatively verify that the model focuses on biologically meaningful lesion and discoloration regions rather than irrelevant background areas.
5. **A comparative analysis with single-backbone baselines** (DenseNet121 and EfficientNetB3) demonstrates that the proposed hybrid model achieves superior performance, highlighting the effectiveness of feature fusion for banana disease classification.

The remainder of this paper is organized as follows. Section 2 reviews recent related studies on banana leaf disease detection and classification using machine learning, deep learning, and hybrid approaches. Section 3 presents the proposed methodology, including dataset preparation, data splitting and input pipeline construction, hybrid model architecture, staged training strategy, and evaluation/XAI procedures. Section 4 reports and discusses the experimental results using quantitative metrics and qualitative visualizations, including learning curves, ROC analysis, confusion matrix, and Grad-CAM explanations. Finally, Section 5 concludes the paper and outlines potential directions for future work.

II. RELATED WORK

A hybrid deep learning model, DenseMobilenetV2, was developed for banana leaf spot disease classification, achieving 98.42% accuracy and high precision and recall scores on a dataset of 937 images [8]. The authors evaluated three transfer-learning models (EfficientNetB0, ResNet50, and VGG19), reporting accuracies of 88.33%, 88.90%, and 87.22%, respectively, with ResNet50 achieving the highest performance [9]. A hybrid fuzzy C-means procedure is used for automatic segmentation and classification of banana leaf diseases, with comparison to existing deep learning methods for diseases like black sigatoka and yellow sigatoka[10].

An improved agro deep learning model for detecting Panama wilt disease in banana leaves achieved 91.56% accuracy, 91.61% precision, 88.56% recall, and 81.56% F1-score, using color and shape changes for prediction [11]. A CNN model achieved 96.41% accuracy in identifying Sigatoka leaf spot disease in bananas using a dataset of 2008 images [12]. A hybrid approach using Ant Colony Optimization and CNNs for banana leaf disease classification provides greater accuracy [13].

A hybrid Xception-RF model achieved 100% accuracy in banana leaf disease classification using a combined dataset and merging deep learning and machine learning techniques [14]. A deep learning model, BananaSqueezeNet, was developed to diagnose banana leaf diseases with high accuracy (0.9625) and outperformed other models like EfficientNetB0, MobileNetV3, and ResNet-101, using Bayesian optimization for image-based diagnosis without human intervention [15]. ResNet50 model achieves 96% accuracy in classifying banana leaf spot diseases using CNN architecture [16].

A hybrid CNN-SVM model achieves 94.77% accuracy in banana leaf wilt disease severity assessment, outperforming previous techniques with robust feature extraction and classification capabilities [17]. A hybrid deep learning model combining region-based separation and CNN with Transfer Learning and vision transformer (TLMViT) for banana disease detection and classification [18]. A hybrid deep learning model combining SCCNN and SVM achieves 99.1% training accuracy and 92.8% testing accuracy for classifying healthy and diseased banana leaves [19].

A CNN model was developed and deployed in a mobile application to identify Fusarium Wilt and Black Sigatoka in bananas with 91.17% accuracy using a dataset of 27,360 images [20]. A novel hybrid deep learning model, OEDTN, is developed for banana leaf disease classification, showing improved performance over state-of-the-art techniques through experiments on Banana Leaf and BBW-BBS datasets [21]. Two new approaches are introduced in [22] for disease detection in banana leaves using hybrid models that combine ANN with SIFT or HOG and LBP to address rotation and scale invariance issues.

A deep learning model combining CNN with spatial attention mechanism achieves 96.37% accuracy in classifying banana leaf diseases using a dataset of 5,583 images [23]. A deep learning model using MobileNetV2 architecture was proposed for banana disease classification, trained on 3012 annotated images, showing high accuracy after fine-tuning and image augmentation [24]. Three deep learning models (AlexNet, VGG16, GoogleNet) were introduced in [25] that used to classify banana leaf diseases, with GoogleNet achieving 95.31% accuracy through fine-tuning pre-trained networks and CNN feature extraction.

A DCNN model for early detection of banana leaf diseases achieves 99% accuracy, advancing agricultural practices with reliable disease prediction [26]. A deep CNN model using transfer learning achieved near-perfect accuracy of 0.99 and an F-1 score of 0.98 for identifying Fusarium wilt in banana crops, outperforming existing models by at least 0.07 in accuracy [27]. ResNet-50 is proposed for leaf disease classification in banana and sunflower, showing maximum efficiency compared to other algorithms. The methodology involves image pre-processing and ResNet-50 classification. Empirical evidence is based on a dataset collected from farms and villages [28].

Despite the promising progress reported in recent banana leaf disease classification studies, several limitations remain in the current literature. First, many studies rely on single-model CNN architectures or report high accuracy without a rigorous comparison against feature-fusion hybrid backbones that can capture complementary disease representations. Second, several works emphasize overall accuracy while providing limited analysis of class-wise performance, macro/weighted metrics, and threshold-independent evaluation (e.g., multiclass ROC-AUC), which are important for reliable multiclass assessment. Third, although visualization and interpretability are occasionally discussed, explainable AI (XAI) is not consistently integrated with high-performing hybrid architectures within a complete end-to-end framework for banana leaf spot disease classification. As summarized in Table 1, prior studies report strong results using both single and hybrid models; however, they differ substantially in disease scope, evaluation depth, and explainability integration, making direct comparison difficult. Therefore, this study addresses these gaps by proposing an explainable hybrid EfficientNetB3–DenseNet121 model with staged fine-tuning, comprehensive multiclass evaluation, and Grad-CAM-based interpretation, while also benchmarking against single-backbone baselines.

TABLE 1 Comparative Summary of Selected Related Work on Banana Leaf Disease Classification.

Ref.	Year	Method / Model	Task / Scope	Performance	Key Limitation / Observation
[8]	2024	DenseMobileNetV2 (hybrid)	Banana leaf spot classification	98.42% accuracy	Strong result, but limited interpretability discussion
[9]	2024	EfficientNetB0, ResNet50+ VGG19	Multiclass banana leaf disease classification	88.90% accuracy	High computational cost for ResNet50/VGG19
[10]	2022	Hybrid fuzzy C-means + classification	Banana leaf disease segmentation/classification	Comparative performance reported	Focus includes segmentation; not an end-to-end explainable hybrid CNN fusion framework
[11]	2023	Improved agro deep learning model	Panama wilt detection	91.56% acc., 91.61% prec., 88.56% rec., 81.56% F1	Lower performance; disease-specific focus
[12]	2023	CNN	Sigatoka identification	96.41% accuracy	Targets a specific disease rather than multiclass leaf spot set
[13]	2023	Hybrid ACO + CNN	Banana leaf disease classification	Improved accuracy (qualitative claim)	Limited standardized multiclass metric
[14]	2023	Hybrid Xception-RF	Multiclass plant leaf disease classification	100% accuracy	require careful reproducibility and dataset generalization analysis
[15]	2023	BananaSqueezeNet	Three banana leaf diseases	0.9625 accuracy	Lightweight and fast; limited XAI integration
[16]	2024	ResNet50	Banana leaf spot disease detection	96% accuracy	Single backbone baseline
[17]	2024	Hybrid CNN-SVM	Banana wilt severity assessment	94.77% accuracy	Severity assessment task differs from multiclass spot-disease classification
[18]	2025	CNN + TL + ViT (hybrid/TLMViT)	Banana disease detection/classification	88.90% accuracy	Hybrid design exists, but explainability and standardized benchmarking details vary
[19]	2023	SCCNN + SVM (hybrid)	Healthy vs diseased banana leaves	99.1% train / 92.8% test	Binary or simplified setting; lower test generalization
[20]	2024	Mobile CNN app-based	Fusarium Wilt + Black Sigatoka	91.17% accuracy	Mobile deployment strength, but fewer classes and lower accuracy
[21]	2025	OEDTN (ensemble transfer network)	Banana leaf disease classification	Improved over SOTA	Ensemble strength, but XAI not emphasized in summary

Ref.	Year	Method / Model	Task / Scope	Performance	Key Limitation / Observation
[22]	2024	Hybrid ANN + SIFT HOG + LBP	Banana leaf disease detection	95.00% accuracy	handcrafted features + ANN rather than end-to-end deep feature fusion
[23]	2024	CNN + Spatial Attention	Banana leaf disease classification	96.37% accuracy	Strong DL baseline; attention
[24]	2024	MobileNetV2 (fine- tuned)	Banana disease classification	Achived high accuracy	Single backbone; exact metric details should be standardized
[25]	2021	AlexNet / VGG16 / GoogleNet	Banana leaf disease classification	GoogleNet 95.31% accuracy	Earlier baselines; lower performance than recent hybrid methods
[26]	2025	DCNN	Early banana leaf disease detection	99% accuracy	Needs more comparison of other datasets
[27]	2023	Transfer learning- based DCNN	Fusarium wilt detection	~0.99 accuracy, F1 = 0.98	Disease-specific focus
[28]	2023	ResNet-50	Banana + sunflower disease classification	96.5 accuracy for banana and 96.8 accuracy for sunflower	Multi-crop scope; not directly equivalent task setting

III. METHODOLOGY

This study proposes a hybrid deep learning framework for multiclass banana leaf disease classification using transfer learning, feature fusion, and explainable artificial intelligence (XAI). The methodology is designed as a four-phase pipeline as shown in Figure .1 that starts with structured data preparation, proceeds through a robust input pipeline and stratified data splitting, then trains a hybrid model composed of EfficientNetB3 and DenseNet121, and finally evaluates the model using comprehensive classification metrics and visual interpretability analysis. The overall design emphasizes not only predictive performance, but also generalization, class balance handling, and model transparency. To improve stability and reproducibility, fixed random seeds are used, and the training process is conducted in staged fine-tuning steps with regularization, adaptive learning-rate control, and checkpointing. The final framework outputs standard performance indicators (accuracy, precision, recall, F1-score), graphical evaluation summaries (learning curves, confusion matrix, ROC curves), and a 3×3 Grad-CAM explainability figure that highlights disease-relevant leaf regions used by the model during decision-making.

A. Dataset

The experiments are conducted using the Banana Leaf Spot Diseases image dataset [29]. The dataset is organized into class-specific folders and contains four categories representing one healthy condition and three disease conditions: cordana, healthy, pestalotiopsis, and sigatoka. This structure is suitable for supervised multiclass image classification because each image is associated with a single ground-truth label.

During dataset ingestion, all image file paths are collected and paired with their corresponding class labels, then stored in a structured dataframe containing the image path, class label, and an integer-encoded class index. This tabular representation simplifies downstream processing, including stratified splitting, batch generation, and evaluation. The label encoding step ensures compatibility with neural-network training while preserving the mapping between integer labels and human-readable

class names for reporting and visualization. The dataset is processed at an input resolution of 300×300 pixels, which is selected to preserve disease texture patterns and lesion structures while maintaining feasible computational cost on Kaggle GPU resources. Since the classes are not perfectly balanced, class-weight computation is later incorporated during training to reduce bias toward majority classes and improve minority-class sensitivity.




Fig.1. Proposed Methodology

B. Phase 1: Data Preparation

The first phase focuses on data preparation and experiment configuration, ensuring that the learning process starts from a clean and reproducible setup. Initially, the dataset is loaded, and the original images are selected as the primary source for model development. Choosing the original images is an important methodological decision because it allows augmentation to be applied only during training, thereby reducing the risk of leakage between training and test partitions and improving the fairness of evaluation.

After loading the dataset, the framework scans the class folders and extracts image file paths and labels for all supported image extensions. These pairs are stored in a dataframe with three key columns: *filepath*, *label*, and *label_idx*. The *label* field stores the categorical disease/health class name, while *label_idx* stores the numeric encoding used by the model. In addition, a bidirectional mapping (*class_to_idx* and *idx_to_class*) is created to ensure consistent encoding during training and accurate label decoding during result interpretation (e.g., classification report, confusion matrix labels, and Grad-CAM figure titles).

This phase also defines the experimental hyperparameters and reproducibility settings. A fixed random seed is applied to Python's random module, NumPy, and TensorFlow, which helps reduce variability across runs and improves experiment repeatability. Core hyperparameters such as image size (*IMG_SIZE* = 300), batch size (*BATCH_SIZE* = 16), and stage-specific training epochs are set at this stage. The methodology optionally enables mixed-precision training (when supported by the hardware), which can accelerate model training and reduce memory usage without changing the conceptual pipeline. Overall, Phase 1 establishes a standardized and reproducible foundation for the subsequent training and evaluation stages.

C. Phase 2: Data Splitting and Input Pipeline

The second phase constructs a robust data partitioning strategy and an efficient input pipeline tailored for hybrid deep learning. The dataset is divided using a stratified split into training, validation, and test sets with a ratio of 70% / 15% / 15%, ensuring that each class maintains approximately the same relative distribution across all subsets. Stratification is particularly important in multiclass disease classification because it prevents class imbalance in the validation or test partitions, which could otherwise produce misleading performance estimates.

To further address class imbalance during optimization, class weights are computed from the training set using a balanced weighting strategy. These weights are later passed to the training routine so that misclassification of underrepresented classes contributes proportionally more to the loss function. This improves the model's ability to recognize all disease categories rather than overfitting to the dominant class.

Next, the methodology defines a training-time augmentation pipeline to improve generalization and robustness to real-world image variability. The augmentation sequence includes random horizontal flipping, random rotation, random zoom, random translation, and random contrast adjustment. These transformations simulate natural variations in camera orientation, leaf positioning, scale, and illumination conditions. In addition, an optional random erasing (cutout-like augmentation) operation is applied with a controlled probability. This operation masks a randomly selected image region, forcing the model to rely on distributed disease cues instead of memorizing a single local pattern or background artifact.

After augmentation, images are decoded, converted to RGB, resized to 300×300 , and cast to floating-point format. Since the proposed classifier is a hybrid model with two pretrained backbones, the same image is transformed into two preprocessed versions: one using the preprocessing function of EfficientNet and the other using the preprocessing function of DenseNet. This dual preprocessing aligns input pixel distributions with the expectations of each backbone's ImageNet-based pretrained weights. The class labels are then converted to one-hot encoded vectors, matching the softmax output layer used in multiclass classification. Finally, the methodology builds optimized tf.data pipelines for train, validation, and test sets, using shuffling (for training only), parallel mapping, batching, and prefetching to improve throughput and training efficiency.

D. Phase 3: Hybrid Deep Learning Modeling and Training

The third phase introduces the core contribution of this work: a hybrid feature-fusion deep learning model that combines the strengths of two complementary pretrained convolutional neural networks, EfficientNetB3 and DenseNet121. The model is designed with dual inputs, where each branch receives the same banana leaf image but in the preprocessing format expected by its respective backbone. Both backbones are initialized with ImageNet pretrained weights and configured with *include_top=False*,

meaning that their original classification heads are removed and only their feature extraction blocks are retained.

In the first branch, EfficientNetB3 is used to capture high-quality hierarchical visual features with strong parameter efficiency and scaling behavior. In the second branch, DenseNet121 contributes densely connected feature reuse patterns that can improve gradient flow and capture complementary texture and lesion representations. For each branch, the output feature map is compressed using Global Average Pooling (GAP), followed by Batch Normalization and Dropout to reduce internal covariate shift and overfitting. The resulting branch-level embeddings are then concatenated in a feature-fusion layer, creating a unified representation that aggregates information learned by both backbones.

The fused features are passed through a multi-layer classification head composed of fully connected layers ($512 \rightarrow 256 \rightarrow 128$ neurons) with ReLU activation and progressively applied dropout. The final layer is a softmax classifier that outputs class probabilities for the four disease categories. To ensure numerical stability when mixed precision is enabled, the output layer is explicitly set to float32.

Training is performed in a three-stage transfer learning and fine-tuning strategy. **In Stage 1**, both pretrained backbones are frozen and only the newly added classification head is trained. This allows the fusion head to learn task-specific decision boundaries without immediately disturbing pretrained feature representations. The model is compiled with the AdamW optimizer (Adam with decoupled weight decay), which improves regularization and optimization behavior compared to standard Adam. The loss function is categorical cross-entropy with label smoothing (0.05), which reduces overconfidence and often improves generalization in multiclass classification. Training callbacks include Early Stopping (to prevent overfitting), ReduceLROnPlateau (to lower the learning rate when validation loss stagnates), and Model Checkpointing (to save the best-performing model based on validation loss).

In Stage 2, both backbones are unfrozen, but only the upper layers are allowed to train while the lower layers remain frozen (approximately the top 25% trainable, depending on indexing). The model is then recompiled with a lower learning rate ($1e-4$) to perform controlled fine-tuning of high-level convolutional features. This stage helps adapt pretrained representations to the visual characteristics of banana leaf diseases while preserving general low-level visual filters.

In Stage 3, the fine-tuning depth is increased by unfreezing additional layers (approximately the top 40% trainable), and the model is again recompiled with an even smaller learning rate ($2e-5$) for stable deeper adaptation. This final stage refines the feature extractors more extensively while minimizing catastrophic forgetting. After training is completed, the histories from all three stages (head training + two fine-tuning stages) are merged to produce unified learning curves for accuracy and loss across the full optimization process. The best saved model checkpoint is then reloaded for final evaluation to ensure that the reported results correspond to the minimum validation-loss model rather than the final epoch state.

E. Phase 4: Evaluation, Reporting, and Explainable AI

The fourth phase provides a comprehensive evaluation and reporting framework for assessing both the predictive performance and interpretability of the proposed hybrid classifier. First, the best checkpointed model is evaluated on the held-out test set, yielding the final test loss and test accuracy. Next, the model generates class-probability predictions for all test samples, from which the predicted class labels are extracted using the maximum softmax probability.

To summarize performance rigorously, the methodology computes a metric summary including accuracy, precision, recall, and F1-score using both macro averaging and weighted averaging. Macro-averaged metrics provide an equal contribution from each class and are useful for assessing class-balanced performance, while weighted-average metrics account for class frequency and provide a population-level view. In addition, a classification report is generated to present per-class precision, recall, and F1-score, enabling detailed analysis of strengths and weaknesses across the four banana leaf conditions.

To quantitatively assess the classification performance, the proposed model was evaluated using Accuracy, Precision, Recall, and F1-score. The evaluation metrics (Accuracy, Precision, Recall, and F1-score) were computed from the confusion matrix. For the multiclass setting, Precision, Recall, and F1-score were calculated using a one-vs-rest scheme for each class, and then summarized using macro and weighted averages. The corresponding metric definitions are given in Eqs. (1) – (4) [30].

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (2)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3)$$

$$\text{F1 - score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4)$$

The training dynamics are visualized using learning curves, specifically the accuracy evolution curve and loss evolution curve, generated from the merged multi-stage training history. These plots help assess convergence behavior, detect potential overfitting, and show the impact of staged fine-tuning. A confusion matrix is also plotted to reveal class-wise error patterns, such as which disease classes are most often confused with one another. This is especially valuable in plant disease diagnosis, where symptom overlap may exist between disease categories.

To evaluate the discriminative performance beyond hard class predictions, the methodology performs multiclass ROC analysis. The ground-truth labels are binarized, and ROC curves are computed for each class individually, along with corresponding Area Under the Curve (AUC) values. A micro-average ROC curve is also reported to summarize overall multiclass separability. These analyses provide threshold-independent evidence of classifier quality and complement the confusion matrix and classification report.

Finally, to improve transparency and support explainable decision-making, the framework applies Grad-CAM (Gradient-weighted Class Activation Mapping) to produce a 3×3 explainability figure. Grad-CAM is computed from the final convolutional representations, and the resulting heatmaps are overlaid on selected test images. Each visualization includes the true label, predicted label, and prediction confidence, allowing qualitative inspection of whether the model is focusing on biologically meaningful disease regions (e.g., spots, lesions, discoloration patterns) rather than irrelevant background areas. This XAI component is important for increasing trust in the proposed system and supporting its potential use in practical agricultural screening scenarios.

IV. RESULT AND DISSCUSSION

This section presents the experimental results and discussion for the proposed hybrid deep learning framework on the BananaLSD multiclass dataset. The analysis begins with a quantitative comparison between the proposed hybrid model and the single-backbone baselines (DenseNet121 and EfficientNetB3), followed by a detailed evaluation of the hybrid model using class-wise metrics, learning curves, confusion matrix, and multiclass ROC analysis. In addition to predictive performance, this section also examines model interpretability through Grad-CAM visualizations to verify whether the learned decisions are based on meaningful disease-related leaf regions. Together, these results provide both quantitative and qualitative evidence of the effectiveness, robustness, and practical reliability of the proposed approach.

Table 2 presents a comparative performance analysis of the three evaluated models, namely DenseNet121, EfficientNetB3, and the proposed hybrid model (EfficientNetB3 + DenseNet121), for banana leaf spot disease classification. The results show that both DenseNet121 and EfficientNetB3

achieved the same overall accuracy of 96.45%, although their loss values and class-wise metrics differed slightly, indicating different prediction confidence and error distributions across classes. In contrast, the proposed hybrid model achieved the best overall performance, with a higher accuracy of 98.58% and the lowest loss value of 0.2578, demonstrating improved generalization and more reliable classification. At the class level, the hybrid model also produced consistently high precision, recall, and F1-scores across all disease categories (Cordana, Healthy, Pestalotiopsis, and Sigatoka), outperforming the single-backbone models in most cases. These findings confirm that feature fusion from EfficientNetB3 and DenseNet121 provides complementary representations that enhance multiclass discrimination performance.

Table 3 presents a comparative evaluation of the proposed hybrid model against several state-of-the-art methods for banana leaf disease classification. As shown, traditional and earlier deep learning approaches, including K-Means + SVM and a conventional CNN, achieved accuracies of 85.00% and 90.62%, respectively, while more advanced deep learning models such as Optimized EfficientNetB2 and SqueezeNet reached 95.31% and 96.25%. In comparison, the proposed hybrid model (EfficientNetB3 + DenseNet121) achieved the highest accuracy of 98.58%, outperforming all referenced methods in this comparison. These results indicate that the proposed feature-fusion strategy effectively enhances discriminative learning by combining complementary representations from the two backbone networks, leading to superior classification performance on banana leaf disease images.

TABLE 2 COMPARATIVE PERFORMANCE OF DENSENET121, EFFICIENTNETB3, AND THE PROPOSED HYBRID MODEL

<i>Model</i>	<i>Accuracy %</i>	<i>Loss</i>	<i>Class Name</i>	<i>Precision</i>	<i>Recall</i>	<i>F1-Score</i>
DensNet121	96.45	0.2999	<i>Cordana</i>	0.96	0.96	0.96
			<i>Healthy</i>	0.95	0.95	0.95
			<i>Pestalotiopsis</i>	0.93	0.96	0.94
			<i>Sigatoka</i>	0.99	0.97	0.98
EfficientNetB3	96.45	0.3249	<i>Cordana</i>	1.00	0.96	0.98
			<i>Healthy</i>	0.95	0.95	0.95
			<i>Pestalotiopsis</i>	0.89	0.96	0.93
			<i>Sigatoka</i>	0.99	0.97	0.98
Hybrid Model (EfficientNetB3+ DensNet121)	98.58	0.2578	<i>Cordana</i>	1.00	1.00	1.00
			<i>Healthy</i>	0.95	1.00	0.98
			<i>Pestalotiopsis</i>	0.96	0.96	0.96
			<i>Sigatoka</i>	1.00	0.99	0.99

Table 3 Comparison with Prior Studies Using the BananaLSD Dataset

<i>Reference</i>	<i>Model</i>	<i>Accuracy</i>
[4]	K-Means + SVM	85.00%
[31]	CNN	90.62%
[32]	OPTIMIZED EfficientNetB2	95.31%
[15]	SqueezeNet	96.25%
Proposed Approach	Hybrid Model (EfficientNetB3+DensNet121)	98.58%

Figure 2 illustrates the accuracy evolution curve of the proposed hybrid model (EfficientNetB3 + DenseNet121) across the training epochs, showing both training and validation accuracy trends. The figure indicates a rapid improvement during the early epochs, where training accuracy increases sharply from approximately 0.70 and validation accuracy rises from about 0.90, reflecting effective feature learning from the transfer-learning backbones and the hybrid fusion head. As training progresses, both curves continue to improve gradually and converge near 0.99, with only minor fluctuations, particularly in the training curve around the middle epochs. The close alignment between the training and validation accuracy curves suggests good generalization performance and limited overfitting, while the high final accuracy values confirm the stability and effectiveness of the staged fine-tuning strategy adopted in the proposed methodology.

Figure 3 presents the loss evolution curve of the proposed hybrid model (EfficientNetB3 + DenseNet121) over the training epochs, showing both training and validation loss trajectories. The figure demonstrates a substantial reduction in loss during the initial epochs, where the training loss decreases sharply from approximately 0.89 and the validation loss drops from around 0.58, indicating rapid convergence and effective learning of discriminative features for banana leaf disease classification. As training continues, both curves decline more gradually and stabilize at low values (approximately 0.26–0.28), with only minor oscillations in the training loss, particularly in the mid-training stage. Notably, the validation loss remains slightly lower than or close to the training loss across most epochs, which suggests that the regularization strategies used in the model (dropout, label smoothing, augmentation, and staged fine-tuning) helped prevent overfitting and improved generalization. Overall, the consistent downward trend and late-stage stabilization confirm the training stability and convergence quality of the proposed hybrid framework.

Figure 4 shows the multiclass ROC curves for the proposed hybrid model (EfficientNetB3 + DenseNet121), along with the micro-average ROC curve, demonstrating the model's excellent discriminative capability across all banana leaf classes. The class-wise AUC values are extremely high, with *cordana* = 1.000, *sigatoka* = 1.000, *healthy* = 0.999, and *pestalotiopsis* = 0.996, while the micro-average AUC = 0.999, indicating near-perfect overall separability. All ROC curves are concentrated near the top-left corner, far above the diagonal reference line (random classifier), which confirms that the model achieves a very high true positive rate at very low false positive rates. These results strongly support the robustness of the proposed hybrid feature-fusion architecture and are consistent with the high accuracy and F1-scores reported in Table 2, confirming that the model maintains strong classification performance not only at a single threshold but across different decision thresholds as well.

Figure 5 presents the confusion matrix of the proposed hybrid model (EfficientNetB3 + DenseNet121) on the test set, providing a detailed view of class-wise prediction performance. The matrix shows that the model correctly classified all *cordana* samples (25/25) and all *healthy* samples (19/19), indicating perfect recognition for these two classes in this test split. Similarly, *pestalotiopsis* achieved 25 correct predictions out of 26, with only one sample misclassified as *healthy*, while *sigatoka* achieved 70 correct predictions out of 71, with only one sample misclassified as *pestalotiopsis*. The strong concentration of values along the main diagonal and the very limited off-diagonal errors confirm the model's high classification reliability and strong class separability. The small confusion observed between *pestalotiopsis* and *healthy*, and between *sigatoka* and *pestalotiopsis*, is likely due to partial visual similarity in lesion appearance, but these errors remain minimal and do not significantly affect the overall performance.

Figure 6 presents the 3×3 Grad-CAM explainability visualization for representative test images, illustrating how the proposed hybrid model (EfficientNetB3 + DenseNet121) localizes disease-relevant regions during classification. In each panel, the heatmap overlay highlights the areas that contributed most strongly to the prediction, while the titles report the true class (T), predicted class (P), and prediction confidence. The visualizations show that the model generally focuses on biologically meaningful regions such as lesions, necrotic spots, discoloration patterns, and infected leaf tissues, rather than relying on unrelated background information. For correctly classified samples (e.g., *sigatoka*, *cordana*, *pestalotiopsis*, and *healthy* leaves), the high-activation regions are concentrated around characteristic symptom patterns, supporting the reliability of the learned features. Even in

lower-confidence cases, the highlighted regions remain close to visible leaf structures and symptom areas, indicating that the model's decisions are guided by relevant visual cues. Overall, this figure provides qualitative evidence that the proposed framework is not only highly accurate but also interpretable, which enhances confidence in its potential use for practical banana disease screening and decision support.




Fig.2. Accuracy Evolution for the Proposed Hybrid Model.




Fig.3. Loss Evolution for the Proposed Hybrid Model.




Fig.4. ROC-Curve for the Proposed Hybrid Model.




Fig.5. Confusion Matrix for the Proposed Hybrid Model.




Fig.6. Grad-CAM Explainability Visualization (3×3) for the Proposed Hybrid Model.

IV. CONCLUSION AND FUTURE WORK

This paper proposed an accurate and explainable hybrid deep learning framework for banana leaf spot disease classification by combining EfficientNetB3 and DenseNet121 in a dual-branch feature-fusion architecture. The framework was designed as a complete pipeline covering dataset preparation, stratified splitting, augmentation-based input processing, staged transfer learning, comprehensive multiclass evaluation, and explainable AI analysis using Grad-CAM. Experimental results on the BananaLSD dataset demonstrated that the proposed hybrid model outperformed the two single-backbone baselines, achieving the best overall performance with 98.58% test accuracy and the lowest test loss, while also providing strong class-wise precision, recall, and F1-scores. The multiclass ROC results further confirmed the high discriminative power of the model, and the Grad-CAM visualizations showed that the classifier focused on disease-relevant lesion regions, supporting the trustworthiness and interpretability of the predictions.

Although the results are highly promising, several directions remain for future work. First, the proposed framework can be validated on larger and more diverse field datasets collected under real plantation conditions with varying illumination, occlusion, and background complexity to further assess generalization. Second, future studies may investigate additional hybrid backbones (e.g., CNN-transformer fusion or lightweight mobile-oriented hybrids) to balance performance and computational efficiency for edge deployment. Third, cross-dataset evaluation and domain adaptation techniques could be explored to improve robustness across regions and acquisition devices. Finally, the current Grad-CAM-based explanation can be extended with multiple XAI methods (e.g., Score-CAM, LIME, SHAP) and expert agronomist validation to provide deeper clinical/agricultural interpretability and stronger deployment readiness in precision agriculture applications.

Acknowledgment: “The authors extend their appreciation to JADARA University in Jordan to support this work.”

REFERENCES

- [1] H. Patel, N. Patil, and D. Patel, “An Intelligent Grading System for Automated Identification and Classification of Banana Fruit Diseases Using Deep Neural Network,” *International Journal of Computing and Digital Systems*, vol. 15, no. 1, pp. 2210–142, doi: <https://doi.org/10.12785/ijcds/150155>.
- [2] R. Retkute, K. S. Crew, J. E. Thomas, and C. A. Gilligan, “Detection of Banana Diseases Based on Landsat-8 Data and Machine Learning,” *Remote Sensing*, vol. 17, no. 13, pp. 2308–2308, Jul. 2025, doi: <https://doi.org/10.3390/rs17132308>.
- [3] C. Ünal, “Incorporating Deep Learning into the Diagnosis of Banana Leaf Spot Diseases for the Protection of Banana Crops,” *Tarım Bilimleri Dergisi*, vol. 31, no. 3, pp. 780–794, Jul. 2025, doi: <https://doi.org/10.15832/ankutbd.1579442>.
- [4] V. Chaudhari and M. Patil, “Banana leaf disease detection using K-means clustering and Feature extraction techniques,” *2020 International Conference on Advances in Computing, Communication & Materials (ICACCM)*, Dehradun, India, 2020, pp. 126–130, doi: 10.1109/ICACCM50413.2020.9212816.
- [5] C. Kanimalar and M. Karthikeyan, “Advances in Machine Learning and Deep Learning for Automated Banana Disease Detection: A Comprehensive Survey,” *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)*, Gobichettipalayam, India, 2024, pp. 1160–1168, doi: 10.1109/ICUIS64676.2024.10866142.
- [6] P. R. and V. A., “Disease diagnosis in banana leaves: a review on AI powered techniques,” *PeerJ Computer Science*, vol. 11, p. e3310, Dec. 2025, doi: <https://doi.org/10.7717/peerj-cs.3310>.
- [7] M. S. Devi, B. Vamshikrishna, J. A. Pandian, P. V. Rao and S. R. Yaseen, “Eight Convolutional Layered Deep Convolutional Neural Network based Banana Leaf Disease Prediction,” *2022 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)*, Greater Noida, India, 2022, pp. 313–317, doi: 10.1109/ICCCIS56430.2022.10037628.

- [8] M. K. Haq, K. Tanvir, M. S. Kabir, N. S. B. Alam, S. K. Dey and V. Raju, "Enhancing Banana Leaf Spot Disease Classification using Dense Mobilenet V2," 2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI), Chennai, India, 2024, pp. 1-6, doi: 10.1109/ACCAI61061.2024.10602228.
- [9] N. Jiménez, S. Orellana, B. Mazon-Olivo, W. Rivas-Asanza, and I. Ramírez-Morales, "Detection of Leaf Diseases in Banana Crops Using Deep Learning Techniques," AI, vol. 6, no. 3, p. 61, Mar. 2025, doi: <https://doi.org/10.3390/ai6030061>.
- [10] K. V. Gokula, D. J., R. Pinagadi Venkateswara, D. V., and K. S., "An automated segmentation and classification model for banana leaf disease detection," Journal of Applied Biology & Biotechnology, Oct. 2022, doi: <https://doi.org/10.7324/jabb.2021.100126>.
- [11] R. Sangeetha, J. Logeshwaran, J. Rocher, and J. Lloret, "An Improved Agro Deep Learning Model for Detection of Panama Wilts Disease in Banana Leaves," AgriEngineering, vol. 5, no. 2, pp. 660–679, Mar. 2023, doi: <https://doi.org/10.3390/agriengineering5020042>.
- [12] S. Samridhi, M. Kalpana, R. Parimalarangan, and N. V. Palanichamy, "Identification of Sigatoka Leaf Spot Disease in Banana Using Convolutional Neural Network (CNN)," Asian Journal of Agricultural Extension, Economics & Sociology, vol. 41, no. 9, pp. 931–936, Aug. 2023, doi: <https://doi.org/10.9734/ajaees/2023/v41i92123>.
- [13] D. Thomas and J. M. David, "Banana Plant Disease Detection-Hybrid Machine Learning Approach," 2023 Annual International Conference on Emerging Research Areas: International Conference on Intelligent Systems (AICERA/ICIS), Kanjirapally, India, 2023, pp. 1-5, doi: 10.1109/AICERA/ICIS59538.2023.10420299.
- [14] A. A and B. N., "An Hybrid Feature Extraction and Classification using Xception-RF for Multiclass Disease Classification in Plant Leaves," Applied Artificial Intelligence, vol. 37, no. 1, Feb. 2023, doi: <https://doi.org/10.1080/08839514.2023.2176614>.
- [15] A. Baki, Hasan Muhammad Abdullah, S. E. Arman, Sayed Saminur Rahman, and Kaies Al Mahmud, "BananaSqueezeNet: A very fast, lightweight convolutional neural network for the diagnosis of three prominent banana leaf diseases," Smart Agricultural Technology, vol. 4, pp. 100214–100214, Aug. 2023, doi: <https://doi.org/10.1016/j.atech.2023.100214>.
- [16] P. Nasra and S. Gupta, "ResNet50: Deep Learning Method for Automated Detection of Banana Leaf Spot Diseases," 2024 First International Conference on Innovations in Communications, Electrical and Computer Engineering (ICICEC), Davangere, India, 2024, pp. 1-6, doi: 10.1109/ICICEC62498.2024.10808847.
- [17] M. Kumar and A. Kumar, "Deep Learning Meets Support Vector Machines: An Effective Hybrid Model for Banana Leaf Wilt Disease Severity Assessment," 2024 2nd International Conference on Disruptive Technologies (ICDT), Greater Noida, India, 2024, pp. 386-390, doi: 10.1109/ICDT61202.2024.10489081.
- [18] Muhammad Saifuddeen, "LEAF DISEASE DETECTION IN BANANA PLANTS FOR PRECISION AGRICULTURE," International Journal of Applied Mathematics, vol. 38, no. 7s, pp. 1251–1279, Oct. 2025, doi: <https://doi.org/10.12732/ijam.v38i7s.585>.
- [19] J. B. Kakati and T. K. Das, "Classification of healthy and unhealthy banana leaves using deep learning approach: A comparative assessment," 2023 4th International Conference on Computing and Communication Systems (I3CS), Shillong, India, 2023, pp. 1-5, doi: 10.1109/I3CS58314.2023.10127309.
- [20] C. A. Elinisa and N. Mduma, "Mobile-Based convolutional neural network model for the early identification of banana diseases," Smart Agricultural Technology, vol. 7, p. 100423, Mar. 2024, doi: <https://doi.org/10.1016/j.atech.2024.100423>.
- [21] N. B. Raja and P. S. Rajendran, "An efficient banana plant leaf disease classification using optimal ensemble deep transfer network," Journal of Experimental & Theoretical Artificial Intelligence, pp. 1–24, Aug. 2023, doi: <https://doi.org/10.1080/0952813x.2023.2241867>.
- [22] J. D. Thiagarajan et al., "Analysis of banana plant health using machine learning techniques," Scientific Reports, vol. 14, no. 1, Jul. 2024, doi: <https://doi.org/10.1038/s41598-024-63930-y>.

- [23] T. Hoque, N. A. Abdullah, M. A. Hasan and N. M. Sadali, "Banana Disease Recognition Using Deep Learning Based Method Combined With Attention Mechanism," 2024 27th International Conference on Computer and Information Technology (ICCIT), Cox's Bazar, Bangladesh, 2024, pp. 2852-2857, doi: 10.1109/ICCIT64611.2024.11022158.
- [24] S. Singh, P. Prakash, G. Baghel, A. Singh, D. Raj and A. K. Agrawal, "Banana Crop Health: A Deep Learning-Based Model for Disease Detection and Classification," 2024 27th International Symposium on Wireless Personal Multimedia Communications (WPMC), Greater Noida, India, 2024, pp. 1-6, doi: 10.1109/WPMC63271.2024.10863138.
- [25] P. Sahu, A. Chug, A. P. Singh, D. Singh, and R. P. Singh, "Classification and activation map visualization of banana diseases using deep learning models," in *International Conference on Innovative Computing and Communications*, A. Khanna, D. Gupta, S. Bhattacharyya, A. E. Hassanien, S. Anand, and A. Jaiswal, Eds., *Advances in Intelligent Systems and Computing*, vol. 1394. Singapore: Springer, 2022, doi: 10.1007/978-981-16-3071-2_61.
- [26] N. R. Rajalakshmi, S. Saravanan, J. Arunpandian, Sandeep Kumar Mathivanan, and G. Qin, "Early Detection of Banana Leaf Disease Using Novel Deep Convolutional Neural Network," *Journal of Data Science and Intelligent Systems*, Aug. 2024, doi: <https://doi.org/10.47852/bonviewJDSIS42021530>.
- [27] K. Yan, M. Kamran, and Y. Sun, "A Transfer Learning-Based Deep Convolutional Neural Network for Detection of Fusarium Wilt in Banana Crops," *AgriEngineering*, vol. 5, no. 4, pp. 2381-2394, Dec. 2023, doi: <https://doi.org/10.3390/agriengineering5040146>.
- [28] C. H R, R. G S, S. BR, and G. N. Yallappa, "Plant Disease Detection using Deep Learning in Banana and Sunflower," *Journal of Advanced Zoology*, vol. 44, no. 3, pp. 93-106, Oct. 2023, doi: <https://doi.org/10.17762/jaz.v44i3.369>.
- [29] S. E. Arman, Md. A. B. Bhuiyan, H. M. Abdullah, S. Islam, T. T. Chowdhury, and Md. A. Hossain, "BananaLSD: A banana leaf images dataset for classification of banana leaf diseases using machine learning," *Data in Brief*, vol. 50, p. 109608, Oct. 2023, doi: <https://doi.org/10.1016/j.dib.2023.109608>.
- [30] M. S. Sawah, "A novel feature selection-driven framework for sustainable and efficient rainfall classification using machine learning models," *Unconventional Resources*, vol. 9, p. 100286, Jan. 2026, doi: <https://doi.org/10.1016/j.uncres.2025.100286>.
- [31] N. Helmawati and E. Utami, "Analysis for Detecting Banana Leaf Disease Using the CNN Method," *JUITA: Jurnal Informatika*, pp. 29-36, Mar. 2025, doi: <https://doi.org/10.30595/juita.v13i1.24514>.
- [32] A. Jadhaw and H. R. Shah, "Hybrid EfficientNetB2 model for banana leaf disease classification using PCM-segmented and augmented images with stratified K-fold validation," *IAENG International Journal of Computer Science*, vol. 53, no. 2, pp. 706-720, Feb. 2026.