# Interpretable CNN Architectures for Rice Leaf Disease Detection: A Grad-CAM and LIME-Based Explainability Approach

OMKAR PRABHU
Manipal Institute of Technology
Department of Computer Science and Engineering
Manipal, Karnataka, 576104
INDIA

*Abstract:* Rice cultivation faces significant challenges from various leaf diseases that substantially impact crop yield and food security. Traditional disease identification methods rely heavily on expert knowledge and visual inspection, which are time-consuming and prone to human error. This study presents a comprehensive deep learning approach for automated rice leaf disease detection, comparing three state-of-the-art convolutional neural network (CNN) architectures: ResNet18, DenseNet121, and EfficientNet-B0. The dataset comprises 5,932 images across four disease categories: Bacterial Blight, Blast, Brown Spot, and Tungro. The experimental results demonstrate exceptional performance across all models, with DenseNet121 achieving the highest test accuracy of 99.83%, followed by ResNet18 at 99.49% and EfficientNet-B0 at 98.99%. To enhance model interpretability and trust, explainable AI (XAI) techniques including Grad-CAM and LIME were integrated, providing visual explanations for model predictions. The implementation of XAI methods enables agricultural practitioners to understand the decision-making process of the deep learning models, thereby increasing confidence in automated diagnosis systems. This research contributes to precision agriculture by offering a reliable, interpretable, and efficient solution for rice disease detection that can be deployed in real-world agricultural settings.

*Key–Words:* Rice leaf disease detection, Convolutional neural networks, Explainable AI, Grad-CAM, LIME.

Received: April 25, 2025. Revised: June 14, 2025. Accepted: June 29, 2025. Published: July 29, 2025.

## 1 Introduction

Rice (Oryza sativa) stands as one of the world's most crucial staple crops, feeding over half of the global population and serving as the primary food source for billions of people worldwide. The agricultural significance of rice extends beyond mere sustenance, as it represents a cornerstone of food security, economic stability, and cultural heritage across numerous nations, particularly in Asia, Africa, and Latin America. However, rice cultivation faces numerous challenges that threaten both yield quality and quantity, with plant diseases representing one of the most significant obstacles to sustainable rice production. Among the most prevalent and destructive diseases affecting rice crops are Bacterial Blight, Blast, Brown Spot, and Tungro, each presenting distinct symptoms and requiring specific management strategies to prevent large-scale crop loss.

Traditional approaches to disease identification in rice cultivation have heavily relied on manual inspection by agricultural experts, extension workers, and experienced farmers. While grounded in centuries of agronomic knowledge, this visual diagnosis method suffers from several limitations, including observer subjectivity, regional inconsistencies, and limited accessibility to trained personnel, especially in remote rural areas where rice farming is prevalent. Moreover, similar visual symptoms across different diseases can lead to frequent misdiagnoses, especially during early infection stages when symptoms are subtle. Manual surveys are also labor-intensive and time-consuming, making them impractical for rapid and scalable disease detection across large cultivation areas. These constraints often lead to delayed intervention, allowing diseases to spread and cause irreversible damage, thereby exacerbating food security challenges.

Recent advancements in artificial intelligence (AI) and deep learning offer promising alternatives through the development of automated, scalable, and accurate plant disease detection systems. Convolutional Neural Networks (CNNs), in particular, have demonstrated exceptional performance in agricultural image analysis due to their ability to learn hierarchical features and extract complex spatial patterns from visual data. When applied to rice leaf disease detection, CNN-based models provide significant advantages: they reduce dependence on human expertise, operate with consistent performance, and are capable of processing large volumes of image data in real time. Additionally, these models show potential in early-stage disease recognition, where subtle visual cues—often imperceptible to the human eye—can be leveraged for timely diagnosis and intervention, thereby mitigating crop losses more effectively.

Despite their potential, deep learning models are often criticized for their ”black box” nature, which hinders their acceptance among agricultural stakeholders who require transparency and trust in decision-making tools. The integration of Explainable AI (XAI) techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM) and Local Interpretable Model-agnostic Explanations (LIME) addresses these concerns by visually highlighting the regions and features influencing model predictions. This study explores the design and evaluation of high-performance CNN architectures for rice leaf disease classification, while also incorporating XAI methods to enhance model interpretability and facilitate human-AI collaboration. Through this dual focus on predictive accuracy and explainability, the research aims to advance practical and trustworthy AI solutions for sustainable rice agriculture.

## 2 Related Work

The detection and classification of rice leaf diseases have garnered significant attention from researchers due to the critical role of rice as a staple food crop worldwide. Traditional manual inspection methods for disease detection are time-consuming, labor-intensive, and prone to human error, necessitating the development of automated approaches [2], [8]. Early research efforts primarily focused on conventional machine learning techniques combined with image processing methods. Ahmed et al. [2] demonstrated the effectiveness of traditional machine learning algorithms, achieving over 97% accuracy using Decision Tree algorithm with 10-fold cross validation for detecting common rice diseases including leaf smut, bacterial leaf blight, and brown spot. Similarly, comprehensive reviews by Mukherjee et al. [8] and Jackulin and Murugavalli [6] highlighted various machine learning approaches employed between 1999 and 2022, emphasizing the transition from manual feature extraction methods to more sophisticated AI-assisted techniques for early disease detection and surveillance.

Rice cultivation faces significant threats from various diseases that can substantially impact crop yield and quality, necessitating effective detection and management strategies. Bacterial leaf blight (BLB), caused by Xanthomonas oryzae pv. oryzae, represents one of the most significant bacterial diseases affecting rice production, requiring sustainable management approaches that combine biological control methods with traditional chemical treatments to suppress pathogen growth through antagonistic microbes and enhanced plant immunity [16]. Rice blast, caused by the fungal pathogen Pyricularia oryzae, poses a severe threat to global food security as it can infect aboveground tissues at any growth stage, producing characteristic lesions on leaves, culms, and panicles that vary in appearance depending on environmental conditions and varietal resistance [3]. Brown spot disease, caused by Bipolaris oryzae, has emerged as a re-emerging worldwide concern, causing approximately 4% grain yield losses globally with variations ranging from 1% to 34% across different regions in Africa and Asia, while also affecting seed quality with infection rates varying from 0.5% to 76% [**?**]. Additionally, viral diseases such as rice tungro disease present unique challenges in irrigated rice systems where continuous cropping throughout the year facilitates disease spread between fields, with transmission dynamics significantly influenced by vector populations, varietal resistance, and proximity to inoculum sources [4]. These diverse pathogenic threats underscore the critical importance of developing robust automated detection systems capable of accurately identifying and classifying multiple disease types to enable timely intervention and sustainable rice production.

The advent of deep learning has revolutionized plant disease detection, offering superior performance compared to traditional approaches. Ferentinos [5] pioneered the application of convolutional neural networks (CNNs) for plant disease detection, achieving an impressive 99.53% accuracy on a dataset of 87,848 images covering 25 different plant species across 58 distinct disease classes. This breakthrough demonstrated the potential of deep learning models to perform automated disease diagnosis with minimal human intervention. Liu and Wang [7] provided a comprehensive review of deep learning applications in plant disease detection, categorizing approaches into classification networks, detection networks, and segmentation networks, while highlighting the superior performance of deep learning over traditional methods in digital image processing tasks. The success of these early works established deep learning as the preferred approach for automated plant disease identification systems.

Recent advances in rice leaf disease detection have focused on leveraging transfer learning and ensemble methods to improve model performance and generalization capabilities. Simhadri et al. [15] conducted a systematic literature review of 82 high-quality articles published since 2017, identifying transfer learning, ensemble learning, and hybrid approaches as the most effective deep learning strategies for rice disease detection. Krishnamoorthy et al. [9] successfully applied InceptionResNetV2 with transfer learning, achieving 95.67% accuracy for rice leaf disease classification, while Rahman et al. [11] developed an ensemble learning framework combining multiple architectures that achieved 99.78% accuracy, surpassing other state-of-the-art approaches. These studies demonstrate the effectiveness of combining different deep learning architectures and pre-trained models to enhance disease detection performance while reducing computational requirements and training time.

The integration of explainable artificial intelligence (XAI) techniques with deep learning models has emerged as a crucial research direction to enhance model interpretability and user trust in automated disease detection systems. Sagar et al. [13] emphasized the importance of explainable AI in plant disease detection, highlighting the need for transparent and interpretable solutions that can provide meaningful insights to end-users, particularly farmers and agricultural practitioners. Rahman et al. [11] successfully integrated Gradient-weighted Class Activation Mapping (Grad-CAM) into their ensemble models to generate visual explanations for disease predictions, bringing transparency to the diagnostic process. Recent works by Saranya and Kumar [1] and Ning et al. [10] have focused on developing real-time CNN-based systems that not only achieve high accuracy but also provide practical deployment solutions for mobile devices and field applications, incorporating features such as fertilizer recommendation modules and lightweight model architectures for enhanced usability in precision agriculture scenarios.

# 3 Methodology

## 3.1 Dataset Description and Preparation

The foundation of this study rests upon a comprehensive rice leaf disease dataset comprising 5,932 high-quality images distributed across four distinct disease categories. The dataset includes 1,584 images of Bacterial Blight, 1,440 images of Blast disease, 1,600 images of Brown Spot, and 1,308 images of Tungro disease. This dataset represents a diverse collection of rice leaf conditions captured under various environmental conditions, lighting scenarios, and disease severity levels, ensuring comprehensive representation of each disease category for robust model training and evaluation.

The dataset preparation process followed rigorous standards to ensure data quality and consistency. All images underwent careful inspection to verify correct labeling and remove any corrupted or mislabeled samples. The dataset exhibits natural class imbalance, reflecting real-world disease occurrence patterns where certain diseases may be more prevalent than others in specific geographical regions or growing seasons.

To establish reliable model evaluation protocols, the dataset was systematically divided into training, validation, and test sets using stratified sampling to maintain proportional representation of each disease category across all splits. The division followed industry-standard ratios with 75% of images allocated to the training set (4,449 images), 15% to the validation set (889 images), and 15% to the test set (594 images). This partitioning strategy ensures adequate data availability for model training while preserving sufficient samples for unbiased performance evaluation.

The training set serves as the primary learning resource for model parameter optimization, providing diverse examples of each disease category to enable comprehensive feature learning. The validation set functions as an independent evaluation dataset during training, facilitating hyperparameter tuning, early stopping decisions, and model selection processes. The test set remains completely isolated from training and validation procedures, serving as the final benchmark for assessing model generalization capabilities and providing unbiased performance metrics for comparative analysis.

## 3.2 Data Preprocessing and Augmentation

Image preprocessing constitutes a critical component of the deep learning pipeline, ensuring consistent input formats and optimal model performance. All images underwent standardized preprocessing procedures including resizing to 224×224 pixels to match the input requirements of pre-trained CNN architectures. This standardization process maintains aspect ratios while ensuring computational efficiency and compatibility with transfer learning approaches. Normalization procedures followed ImageNet standards, applying mean values of [0.485, 0.456, 0.406] and standard deviation values of [0.229, 0.224, 0.225] across RGB channels. This normalization strategy aligns input data distributions with pre-trained model expectations, facilitating effective transfer learning and stable gradient propagation during training processes. The preprocessing pipeline incorporates tensor conversion operations to transform PIL image objects into PyTorch tensor formats suitable for neural network processing. This conversion process maintains data integrity while enabling efficient batch processing and GPU acceleration during training and inference phases.

## 3.3 Model Architectures

This study evaluates three state-of-the-art convolutional neural network architectures, each representing distinct approaches to deep learning model design and offering unique advantages for image classification tasks.

### 3.3.1 ResNet18 Architecture

The ResNet18 model implements the residual learning framework introduced by He et al., addressing the vanishing gradient problem through skip connections that enable training of deeper networks. The architecture (Figure 1) comprises 18 layers organized into residual blocks, each containing two 3×3 convolutional layers with batch normalization and ReLU activation functions. The skip connections allow gradients to flow directly through the network, enabling stable training and improved convergence characteristics. For this study, the final fully connected layer was modified to output four classes corresponding to the rice disease categories, while all preceding layers were initialized with ImageNet pre-trained weights to leverage transfer learning benefits.

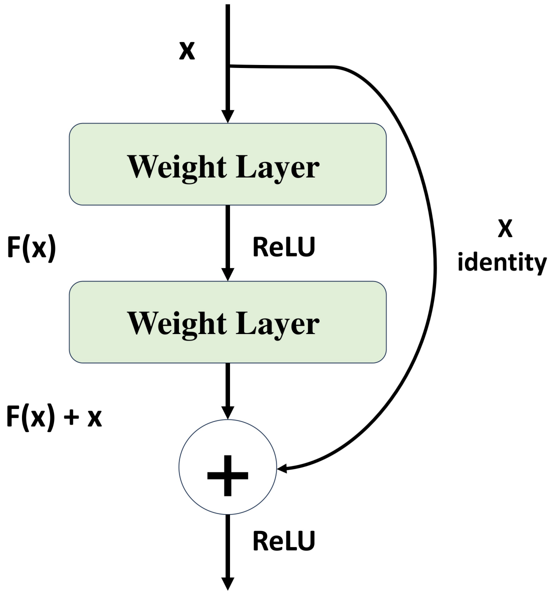

Figure 1: ResNet-18 Architecture denoting Skip Connection

### 3.3.2 DenseNet121 Architecture

The DenseNet121 model implements dense connectivity patterns where each layer receives feature maps from all preceding layers within dense blocks. This architecture (Figure 2) promotes feature reuse and strengthens feature propagation throughout the network, often achieving superior performance with fewer parameters compared to traditional CNN designs. The model comprises four dense blocks with varying numbers of layers, connected through transition layers that perform downsampling operations. The dense connectivity pattern encourages feature diversity and reduces overfitting risks, making it particularly suitable for datasets with limited sample sizes. Similar to ResNet18, the classifier layer was adapted for four-class rice disease classification while maintaining pre-trained feature extraction capabilities.

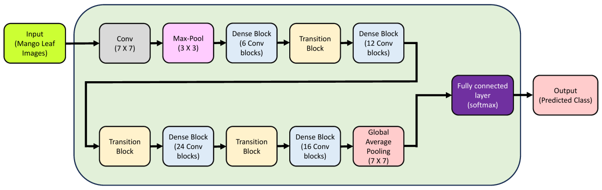

Figure 2: DenseNet-121 Model Architecture

### 3.3.3 EfficientNet-B0 Architecture

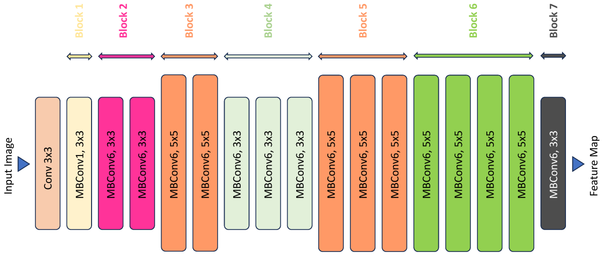

Figure 3: EfficientNet B0 Model Architecture

The EfficientNet-B0 model represents a breakthrough in neural architecture scaling, systematically balancing network depth, width, and input resolution through compound scaling methods. This architecture (Figure 3) achieves optimal trade-offs between model accuracy and computational efficiency, making it suitable for both high-performance applications and resource-constrained deployment scenarios. The model incorporates mobile inverted bottleneck blocks with squeeze-and-excitation optimization, enabling efficient feature extraction and representation learning. The compound scaling approach ensures that network dimensions are optimally balanced, avoiding the inefficiencies associated with arbitrary scaling of individual dimensions.

### 3.4 Explainable AI Integration

To address interpretability requirements and enhance model trustworthiness, this study integrates two complementary explainable AI techniques: Grad-CAM and LIME. These methods provide different perspectives on model decision-making processes, enabling comprehensive understanding of prediction mechanisms.

#### 3.4.1 Grad-CAM

Gradient-weighted Class Activation Mapping (Grad-CAM) generates visual explanations by computing gradients of predicted class scores with respect to feature maps in the final convolutional layer [14]. The implementation registers forward and backward hooks on the target convolutional layer (layer4 for ResNet18) to capture feature activations and gradients during forward and backward passes. The gradient information is used to compute importance weights for each feature map, which are then combined with the feature activations to produce class-specific activation maps. These activation maps highlight image regions that contribute most significantly to the predicted class, providing intuitive visual explanations that align with human understanding of disease symptoms.

#### 3.4.2 LIME

Local Interpretable Model-agnostic Explanations (LIME) provides an alternative interpretability approach by explaining individual predictions through locally faithful approximations [12]. The LIME implementation segments input images into superpixels and systematically perturbs these regions to observe their impact on model predictions. By training a local linear model on the perturbed samples, LIME identifies which image regions contribute positively or negatively to specific class predictions. This approach provides region-based explanations that complement Grad-CAM's activation-based visualizations, offering users multiple perspectives on model decision-making processes.

The integration of both XAI techniques enables comprehensive model interpretability, addressing different user preferences and use cases. Grad-CAM provides smooth, continuous activation maps that highlight important image regions, while LIME offers discrete, region-based explanations that may be more intuitive for users familiar with traditional image analysis approaches.

## 4 Results and Discussion

### 4.1 Model Performance Comparison

The experimental evaluation of the three CNN architectures reveals exceptional performance across all models, with each achieving remarkably high accuracy rates that demonstrate the effectiveness of deep learning approaches for rice leaf disease detection. The comprehensive performance analysis summarized in Table 1 encompasses training dynamics, validation performance, and final test set evaluation to provide thorough assessment of model capabilities. Acc., Prec., Rec., and F1 stand for accuracy, precision, recall, and F1-score.

Table 1: CNN Model Performance on Rice Leaf Disease Detection

| **Model** | **Acc. (%)** | **Loss** | **Prec.** | **Rec.** | **F1** |
|---|---|---|---|---|---|
| ResNet18 | 99.49 | 0.0271 | 1.00 | 0.99 | 0.99 |
| DenseNet121 | 99.83 | 0.0057 | 1.00 | 1.00 | 1.00 |
| EfficientNet-B0 | 98.99 | 0.0339 | 0.99 | 0.99 | 0.99 |

**DenseNet121** emerged as the top-performing model, achieving an outstanding test accuracy of 99.83% with the lowest test loss of 0.0057. This exceptional performance can be attributed to the dense connectivity patterns that promote feature reuse and strengthen gradient flow throughout the network. The architecture's ability to combine features from multiple scales and depths enables comprehensive representation learning that captures subtle disease characteristics effectively.

**ResNet18** demonstrated highly competitive performance with a test accuracy of 99.49% and test loss of 0.0271. The residual learning framework proves highly effective for rice disease classification, with skip connections enabling stable training and robust feature learning. The slightly lower performance compared to DenseNet121 may be attributed to the less extensive feature reuse capabilities, though the difference remains minimal and practically insignificant.

**EfficientNet-B0** achieved a test accuracy of 98.99% with test loss of 0.0339, representing excellent performance while maintaining computational efficiency advantages. The compound scaling approach balances model complexity and performance effectively, though the slight performance gap compared to the other architectures suggests that the specific scaling configuration may not be optimal for this particular dataset size and complexity.

Based on these results, DenseNet121 is identified as the most effective model for the classification task. The corresponding confusion matrix of DenseNet121 (shown in Figure 4) confirms its robustness, exhibiting minimal misclassifications across all disease categories. This underscores its potential for practical deployment in precision agriculture scenarios.

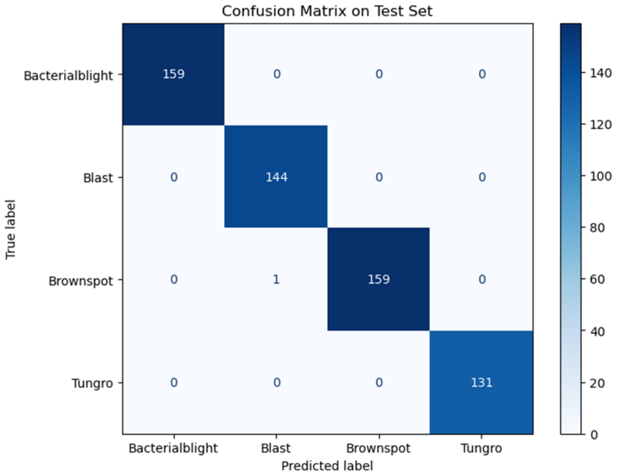

Figure 4: Confusion Matrix of DenseNet121.

## 4.2 Per-Class Performance Analysis

The detailed classification reports reveal consistent high performance across all disease categories, indicating robust model capabilities for discriminating between different rice leaf diseases. The analysis examines precision, recall, and F1-score metrics for each disease class to identify potential biases or weaknesses in model performance. This is depicted in Table 2.

Table 2: Per-Class Performance of CNN Models

| Model | Class | Prec. | Rec. | F1 | N |
|---|---|---|---|---|---|
| ResNet18 | Bact. Blight | 0.99 | 1.00 | 1.00 | 159 |
| | Blast | 0.99 | 1.00 | 0.99 | 144 |
| | Brown Spot | 1.00 | 0.98 | 0.99 | 160 |
| | Tungro | 1.00 | 1.00 | 1.00 | 131 |
| DenseNet121 | Bact. Blight | 1.00 | 1.00 | 1.00 | 159 |
| | Blast | 0.99 | 1.00 | 1.00 | 144 |
| | Brown Spot | 1.00 | 0.99 | 1.00 | 160 |
| | Tungro | 1.00 | 1.00 | 1.00 | 131 |
| EffNet-B0 | Bact. Blight | 1.00 | 0.98 | 0.99 | 159 |
| | Blast | 0.99 | 0.98 | 0.99 | 144 |
| | Brown Spot | 0.97 | 1.00 | 0.98 | 160 |
| | Tungro | 1.00 | 1.00 | 1.00 | 131 |

Prec., Rec., F1, and N stand for precision, recall, F1-score, and number of test samples.

All models demonstrate excellent precision and recall values, indicating both accurate positive predictions and comprehensive detection of disease cases. The consistent high performance across different disease types suggests robust feature learning capabilities that capture distinctive characteristics of each disease category.

## 4.3 Training Dynamics and Convergence Analysis

The training process analysis reveals fascinating insights into model learning behaviors and convergence characteristics across the seven-epoch training period. Each model exhibited distinct learning patterns that reflect their architectural differences and optimization characteristics.

### 4.3.1 ResNet18 Training Progression

The ResNet18 model demonstrated rapid initial learning with training accuracy improving from 90.78% in epoch 1 to 95.84% in epoch 2, accompanied by significant validation accuracy improvement from 98.43% to 99.21% (Figure 5). This rapid convergence indicates effective transfer learning utilization and stable gradient propagation through residual connections. The training loss decreased consistently from 0.2667 to 0.0349 over seven epochs, while validation loss fluctuated between 0.0091 and 0.0626, indicating good generalization capabilities with minimal overfitting. The final training accuracy reached 98.85% with validation accuracy of 99.10%, demonstrating excellent model performance and appropriate regularization.

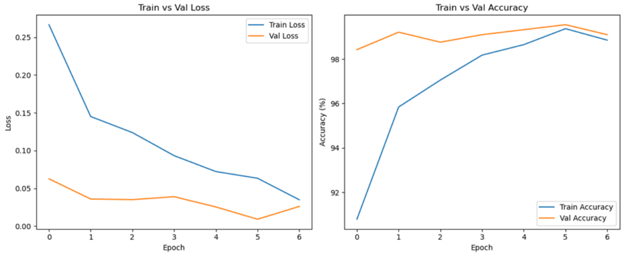

Figure 5: Training vs validation curves for ResNet18.

### 4.3.2 DenseNet121 Training Progression

DenseNet121 exhibited more gradual but steady improvement patterns, with training accuracy progressing from 93.46% in epoch 1 to 99.24% in epoch 7

(Figure 6). The model showed remarkable validation performance improvement, achieving 99.89% validation accuracy in the final epoch with minimal validation loss of 0.0029. This exceptional validation performance, combined with the steady training progression, indicates optimal feature learning and excellent generalization capabilities. The dense connectivity patterns appear to provide superior feature reuse and gradient flow, resulting in the most stable and effective learning trajectory among the evaluated models.


Train vs Val Loss
Train vs Val Accuracy

Figure 6: Training vs validation curves for DenseNet121.

### 4.3.3 EfficientNet-B0 Training Progression

EfficientNet-B0 demonstrated rapid early convergence with training accuracy reaching 95.80% in the first epoch and achieving consistently high validation accuracy above 99% from epoch 1 onwards (Figure 7). The model maintained stable performance throughout training, with final training accuracy of 98.22% and validation accuracy of 99.78%. The compound scaling approach appears to provide effective initialization and balanced model capacity, enabling rapid convergence while maintaining excellent generalization performance.

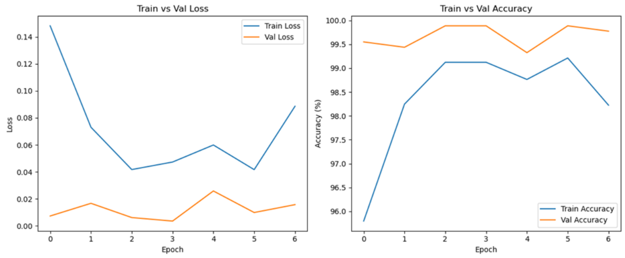
Train vs Val Loss
Train vs Val Accuracy

Figure 7: Training vs validation curves for EfficientNet B0.

## 4.4 Explainable AI

The integration of explainable AI techniques provides valuable insights into model decision-making processes, enhancing interpretability and building user confidence in automated disease detection systems. As DenseNet121 achieved the highest classification performance among the evaluated models, it was selected for interpretability analysis using Grad-CAM and LIME. The visualizations from both techniques consistently reveal that the model focuses on clinically relevant image regions, aligning well with human expert understanding of disease symptoms. This interpretability not only supports the model's reliability but also facilitates its adoption in real-world agricultural settings.

### 4.4.1 Grad-CAM Analysis

The Grad-CAM visualizations demonstrate that all three models consistently focus on diseased leaf regions when making classification decisions. For Bacterial Blight cases, the activation maps highlight water-soaked lesions and yellowing areas that correspond to characteristic disease symptoms. Blast disease visualizations show concentrated attention on diamond-shaped lesions with gray centers and brown borders, accurately identifying the distinctive morphological features of this fungal disease. Brown Spot cases reveal activation patterns focused on oval brown spots with yellow halos, demonstrating the model's ability to recognize these characteristic symptoms. Tungro disease visualizations highlight yellowing and discoloration patterns associated with viral infection symptoms.

The consistency of Grad-CAM results across different model architectures suggests robust feature learning that captures fundamental disease characteristics rather than spurious correlations or background artifacts. The activation maps provide intuitive visual explanations that agricultural experts can readily interpret and validate against their domain knowledge.

### 4.4.2 LIME Analysis

LIME explanations complement Grad-CAM visualizations by providing region-based interpretations that segment images into discrete areas of influence. The LIME results consistently identify diseased leaf regions as positive contributors to disease classification decisions while showing neutral or negative contributions from healthy leaf areas and background regions. This behavior aligns with expected decision patterns and provides confidence in model reliability.

The superpixel-based explanations reveal that models make decisions based on multiple image regions rather than single distinctive features, indicating comprehensive pattern recognition capabilities. The explanations show appropriate weighting of different symptom characteristics, such as lesion morphology, color patterns, and spatial distribution of disease symptoms.

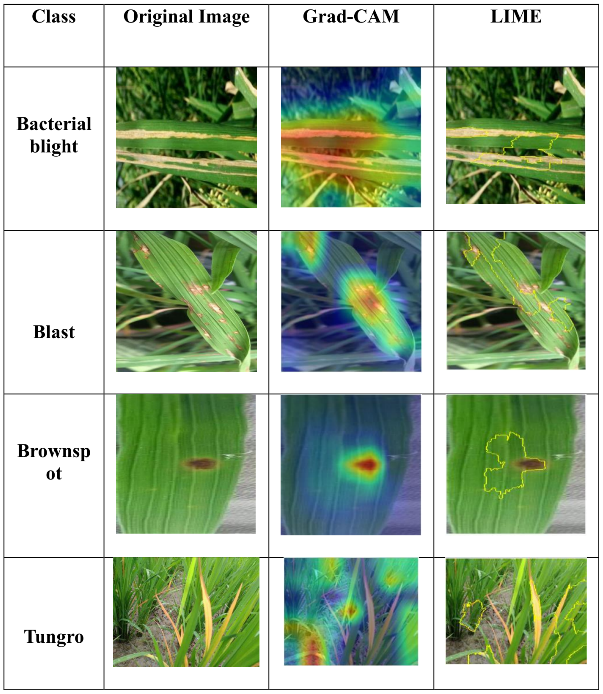

Figure 8: Visual Explanations using Grad-CAM and LIME for DenseNet121 Predictions

The combination of Grad-CAM and LIME provides comprehensive interpretability coverage, addressing different user preferences and analytical needs. Agricultural practitioners can utilize these explanations to verify model decisions, understand disease identification criteria, and build confidence in automated diagnosis systems for practical field applications, as illustrated in Figure 8.

## 5 Conclusion

This study presents a comprehensive investigation of deep learning methods for rice leaf disease detection using state-of-the-art CNN architectures—ResNet18, DenseNet121, and EfficientNet-B0—achieving exceptional test accuracies of 99.49%, 99.83%, and 98.99%, respectively. The integration of explainable AI techniques such as Grad-CAM and LIME effectively addresses the "black box" nature of deep learning models, enhancing interpretability and user trust. Visual explanations consistently aligned with human expert understanding, reinforcing the credibility of automated decisions and supporting their adoption in real-world agricultural environments. These results demonstrate that CNN-based systems can offer both accuracy and transparency, meeting the dual needs of high-performance classification and human-centered interpretability.

Beyond classification, this work contributes meaningfully to the domain of precision agriculture by laying out a framework for designing interpretable and deployable AI solutions. The comparative analysis of CNN models provides actionable guidance for selecting architectures based on accuracy, explainability, and computational efficiency. Future work can explore integrating drone-based imagery and temporal analysis for large-scale monitoring, or building lightweight models suitable for mobile deployment in low-resource farming regions. Additionally, expanding datasets to include more rice varieties and rare disease cases will help improve generalization across diverse agricultural conditions. Ultimately, this study establishes a foundation for scalable, intelligent systems capable of not only disease

*References:*

[1] *Real time rice leaf diseases detection using cnn with deep learning*, International Journal of Emerging Technologies and Innovative Research **12** (2025), no. 4, ppd477–d487, UGC and ISSN Approved, Available at www.jetir.org.

[2] Kawcher Ahmed, Tasmia Rahman Shahidi, Syed Md. Irfanul Alam, and Sifat Momen, *Rice leaf disease detection using machine learning techniques*, 2019 International Conference on Sustainable Technologies for Industry 4.0 (STI), 2019, pp. 1–5.

[3] Aziiba Emmanuel Asibi, Qiang Chai, and Jeffrey A. Coulter, *Rice blast: A disease with implications for global food security*, Agronomy **9** (2019), no. 8.

[4] T.C.B. Chancellor, J. Holt, S. Villareal, E.R. Tiongco, and J. Venn, *Spread of plant virus disease to new plantings: A case study of rice tungro disease*, Advances in Virus Research, vol. 66, Academic Press, 2006, pp. 1–29.

[5] Konstantinos P. Ferentinos, *Deep learning models for plant disease detection and diagnosis*, Computers and Electronics in Agriculture **145** (2018), 311–318.

[6] C Jackulin and S. Murugavalli, *A comprehensive review on detection of plant disease using machine learning and deep learning approaches*, Measurement: Sensors **24** (2022), 100441.

[7] Jun Liu and Xuewei Wang, *Plant diseases and pests detection based on deep learning: a review*, Plant Methods **17** (2021), no. 1, 22.

[8] Rashmi Mukherjee, Anushri Ghosh, Chandan Chakraborty, Jayanta Narayan De, and Debi Prasad Mishra, *Rice leaf disease identification and classification using machine learning techniques: A comprehensive review*, Engineering Applications of Artificial Intelligence **139** (2025), 109639.

[9] Krishnamoorthy N, L.V. Narasimha Prasad, C.S. Pavan Kumar, Bharat Subedi, Haftom Baraki Abraha, and Sathishkumar V E, *Rice leaf diseases prediction using deep neural networks with transfer learning*, Environmental Research **198** (2021), 111275.

[10] H. Ning, S. Liu, Q. Zhu, and T. Zhou, *Convolutional neural network in rice disease recognition: accuracy, speed and lightweight*, Frontiers in Plant Science **14** (2023), 1269371.

[11] Md Mokshedur Rahman, Zhang Yan, Mohammad Tarek Aziz, MD Abu Bakar Siddick, Tien Truong, Md. Maskat Sharif, Nippon Datta, Tanjim Mahmud, Renzon Daniel Cosme Pecho, and Sha Md Farid, *Explainable deep transfer learning framework for rice leaf disease diagnosis and classification*, International Journal of Advanced Computer Science and Applications **15** (2024), no. 12.

[12] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, *"why should i trust you?": Explaining the predictions of any classifier*, 2016.

[13] Saurav Sagar, Mohammed Javed, and David S Doermann, *Leaf-based plant disease detection and explainable ai*, 2023.

[14] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra, *Grad-cam: Visual explanations from deep networks via gradient-based localization*, International Journal of Computer Vision **128** (2019), no. 2, 336–359.

[15] Chinna Gopi Simhadri, Hari Kishan Kondaveeti, Valli Kumari Vatsavayi, Alakananda Mitra, and Preethi Ananthachari, *Deep learning for rice leaf disease detection: A systematic literature review on emerging trends, methodologies and techniques*, Information Processing in Agriculture **12** (2025), no. 2, 151–168.

[16] Bestha Sai Teja, Gayatri Jamwal, Vishal Gupta, Mansi Verma, Ayushi Sharma, Akash Sharma, and Vinod Pandit, *Biological control of bacterial leaf blight (blb) in rice–a sustainable approach*, Heliyon **11** (2025), no. 2, e41769.