Journal of Science Research International (JSRI)
ISSN: 2456-6365 (Online)

OPEN ACCESS

RICE GUARD: A HYBRID DEEP LEARNING APPROACH FOR RICE LEAF DISEASE DETECTION

Sujal N. Bolke[1], Prathamesh P. Kumbhar[1], Nilesh S. Kulkarni[2] and Laxman G. Mulik[1]

[1]*Ajara Mahavidyalaya, Ajara, Kolhapur, M.S., India*

[3]*Department of Computer Science, Shivaji University, Kolhapur, M.S., India*

**Corresponding author E-mail: bolkesujal@gmail.com, prathameshkumbhar454@gmail.com, meet2nilesh.gad@gmail.com, laxmul@gmail.com*

Received: 13 January 2025	Revised: 08 February 2026	Accepted: 19 March 2026	Published: 27 March 2026

DOI: **https://doi.org/10.5281/zenodo.19256232**

Abstract:

As the primary caloric source for over 3.5 billion people, Oryza sativa productivity is critical to economic stability in South and Southeast Asia. Leaf Blast, Bacterial Leaf Blight and Brown Spot are major yield limiting and income-decreasing foliar diseases. Current practice relies on experts examining crops visually, which is labour-intensive and observer-dependent, with variations among regions. Rice Guard: A New Deep Learning-Based Tool for Monitoring Diseases of Rice Leaves. We developed Rice Guard, an automatic detector of rice leaf diseases based on a hybrid deep learning pipeline. The pipeline combines three pre-trained CNNs: MobileNetV2, DenseNet121, and ResNet50V2-within one framework. (utilizing a dual-pooling fusion step Global Average Pooling) in order to preserve both widespread and spiky activations across scales. Training and testing on 3798 high-resolution rice leaves images for six classes. We get 96% train, 94% on the validation set and 90% on out of sample test set. Healthy class is identified with 99% precision; for diseased classes, both precision and recall are balanced. We present a system that runs on Android and also offers low-latency predictions (<200ms) through a Django-based web application. We also show that it generalizes to real field images beyond curated benchmarks.

Keywords: *CNN, Deep Learning, Hybrid Ensemble, Rice Leaf Disease, Precision Agriculture, Transfer Learning.*

1. Introduction

One of the major foods of more than 3.5 billion people of the world, rice is subject to numerous diseases. For instance, Bacterial Leaf Blight and Leaf Blast can cause a yield loss of 30–50%, thereby threatening food security globally. While early detection is important, there is little manual diagnosis from experts. This can be subjective and error-prone due to different light conditions and symptom stages. Also, manually scouting a large area is not economically viable so we need something automated, scalable, etc. so that we can make quick decisions.

Detecting any pathogens automatically is difficult due to many reasons. One is they have similar symptoms. Other is confounding with abiotic stress. The last one is varied field images such as lighting, background, angles etc. In this paper, we propose Rice Guard, a deep learning hybrid model that balances complexity and feature richness. The system employs a fusion of two CNNs, specifically MobileNetV2 and DenseNet121, to improve robustness and adaptability through a dual pooling approach.

Analysis of paper

i. Hybrid architecture that triples backbone uses 3 CNN backbones to optimize speed and feature capture.
ii. Two Pooling Features are Combined to Extract Feature in order to Improve Discrimination Ability.
iii. A thorough examination of each disease which the researchers studied.
iv. Deploying a web application with Django for real-time prediction of fields
v. In-Situ Validation: Testing generalization on real field images beyond benchmark datasets.

2. Related work

2.1 Traditional machine learning approaches

Prior to deep learning, diagnostic systems relied on engineered features such as color histograms and texture metrics, often failing under variable field lighting (1) (color, texture, shape) and classifiers like SVMs or Random Forestsn. While computationally inexpensive, these methods require extensive domain expertise and often fail to generalize under varying imaging conditions.

2.2 Single CNN model approaches

Deep learning enabled end-to-end feature learning, with architectures like Alex Net, VGG, and ResNet applied to crop disease detection (2,3). These models showed significant gains, especially with limited labeled data, but single networks often plateau in performance by capturing only certain feature types.

2.3 Ensemble and hybrid methods

Ensembles combine predictions from multiple models to reduce variance and improve robustness; hybrids fuse representations from different networks to enhance discrimination. Diversifying features is particularly valuable when disease classes appear visually similar (4), with multi-scale designs effectively capturing both texture and semantic cues.

2.4 Advanced architectural innovations

Recent works integrate attention modules or graph convolutions to emphasize disease-relevant regions while suppressing background noise (5). Vision Transformers have been explored for detection and severity estimation (6) but remain more data- and compute-intensive than CNNs. Studies underscore the importance of model selection, data quality, and hyperparameter tuning for real-world deployment (6).

2.5 Research gap

Key gaps persist: many hybrids use simplistic fusion (averaging or concatenation) without clarifying backbone complementarity; per-class performance analysis and failure modes are rarely detailed; and large model sizes hinder deployment on low-resource devices. Rice Guard addresses these through:

i. Explicit dual-pooling fusion;
ii. Comprehensive class-wise evaluation; and (iii) an end-to-end web deployment validated on real field images.

3. Materials and Methods

3.1 Dataset description and preparation

We provide a curated dataset of 3,798 high-resolution rice leaf images in six classes: Healthy, Brown Spot, Leaf Blast, Leaf Scald, Bacterial Leaf Blight and Sheath Blight. This amount of diversity in the imagery come from scattered growing areas so certain pictures are bright, some are darker or taken from different angles or disease in distinct stages. This version of the dataset labels was all verified by domain experts.

Table 1: Complete dataset distribution:

Class	Total Images	Training (70%)	Validation (15%)	Testing (15%)
Healthy	650	455	97	98
Brown Spot	637	445	95	97
Leaf Blast	619	433	92	94
Leaf Scald	628	439	94	95
Bacterial Blight	636	445	95	96
Sheath Blight	628	439	94	95
Total	3,798	2,656	567	575

Table 1 gives the dataset split by class and split.

Figure 1: Representative samples from each disease category and healthy one

3.2 Process pipeline

In the same manner, all input images processed so training stability is assured and input compatibility with back-bone network pre-trained. The chosen architectures accept images with size 224×224 (height x width). In order to have an RGB color scheme, we will format the pixel intensity of the image to meet data store requirements of Image Net.

We use images in the range of -1, 1. We only use a small amount of data augmentation during our training in order to keep as close as possible to real world deployment. This managed limitation helps the model learn better features for new datasets creation. Thus, this helps improve generalization to the field.

3.3 Design of hybrid architecture

Rice Guard uses three CNNs to balance performance. MobileNetV2 was selected specifically to accommodate the computational constraints of edge devices common in smallholder farming context aDt ~3.5M parameters. DenseNet121 captures fine texture through feature reuse at ~8M parameters. Finally, ResNet50V2 offers deep semantic representations with residual connections at ~25M parameters. Utilizing a dual-pooling fusion strategy, finer features based on global average pooling (GAP) and global maximum pooling (GMP) are finally regrouped. While GAP retrieves contextual information, GMP captures discriminative information. These vectors are appended per backbone and fused into a single descriptor. This vector is passed through two dense layers (512 and 256 units) with ReLU activation in the classification head. After the first layer, we use a dropout of 0.5 to prevent overfitting. Application of 6-way SoftMax layer gives out diseased probability distribution.

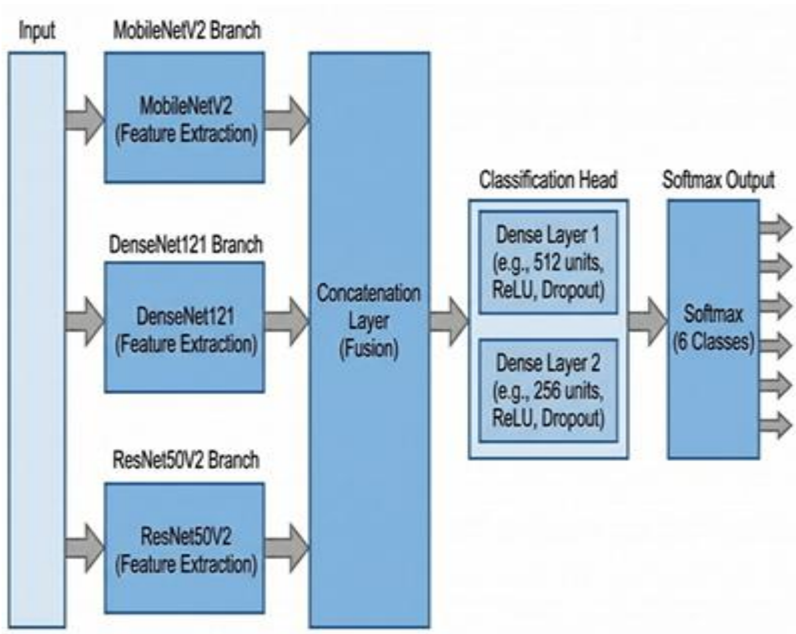

Figure 2: Rice Guard hybrid architecture

3.4 Training configuration

Table 2: Training hyper parameters

Parameter	**Value**	**Justification**
Optimizer	Adam	Adaptive learning rate
Learning Rate	1e-4	Conservative for fine-tuning
Loss Function	Categorical Cross-Entropy	Multi-class classification
Batch Size	32	Memory/stability balance
Epochs	40	Sufficient for convergence
Early Stopping	Patience=5	Prevent overfitting

Table 2 lists the hyperparameters used for training.

4. Experimental results

4.1 Training dynamics

Loss and accuracy flatten out after about 15 epochs. On the training set, we achieve 96% accuracy and the validation set 94%. The training is done with early stopping with patience 5 to avoid overfitting.

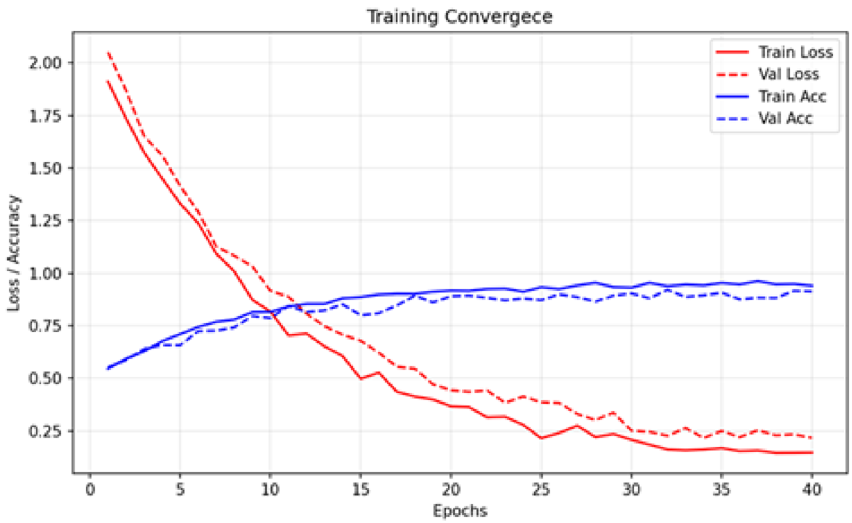

Figure 3: Training and validation accuracy/loss curves.

4.2 Overall performance metrics

The reserved test set of 575 images was used to evaluate Rice Guard performance and reached 90% accuracy. Figure 4 gives the Confusion matrix showing how errors are distributed across the classes.

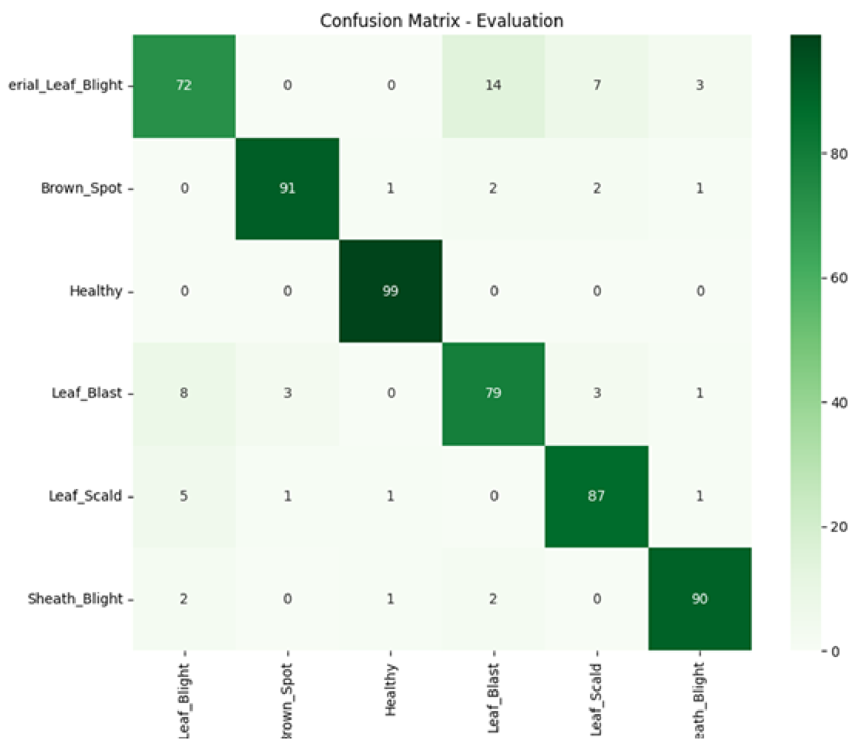

Figure 4: Confusion matrix on test set

4.3 Class-wise performance analysis

Table 3: Class-wise precision, recall, F1-score and support

Class	Precision	Recall	F1-Score	Support
Healthy	0.99	1.00	0.99	98
Brown Spot	0.91	0.92	0.91	97
Leaf Blast	0.88	0.85	0.86	94
Leaf Scald	0.87	0.89	0.88	95
Bacterial Blight	0.89	0.88	0.88	96
Sheath Blight	0.86	0.87	0.87	95
Weighted Avg	0.90	0.90	0.90	575

Table 3 gives precision, recall, F1-score, and support per class.

4.4 Qualitative analysis

Examining the predictions, most samples correctly classified with high confidence. Errors regularly happen between Leaf Blast and Bacterial Blight due to potentially similar color changes. Example predictions are shown in the fig. 5 with confidence to ensure the model focuses on biologically relevant features rather than background noise, we utilized Grad-CAM visualizations.

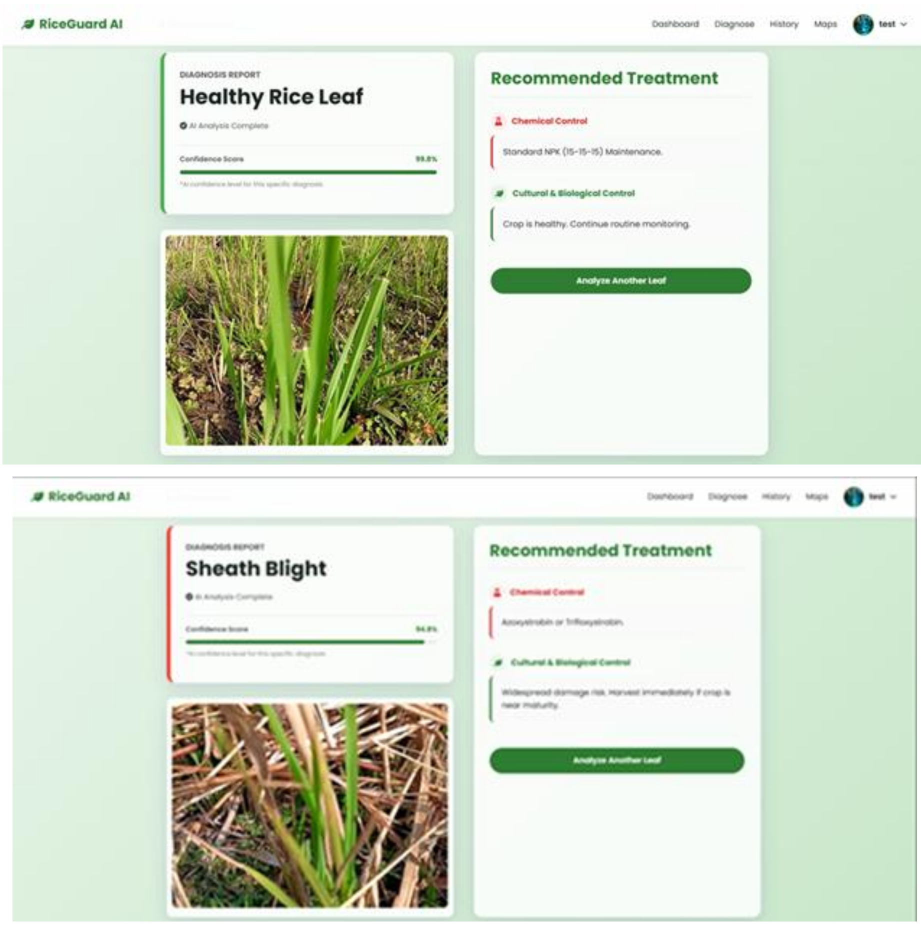

Figure 5: Sample predictions with confidence scores

4.5 Comparison with baseline approaches

Table 4: Comparison with baseline models.

Model	Parameters	Test Accuracy	Training Time
MobileNetV2	3.5M	85%	1.2 hrs.
DenseNet121	8M	87%	2.1 hrs.
ResNet50V2	25M	86%	2.8 hrs.
Simple CNN	2M	78%	3.5 hrs.
Rice Guard (Ours)	36.5M	90%	~4.2 hrs.

Table 4 contrasts Rice Guard with each backbone alone and with a simple CNN.

4.6 Ablation study

To validate the contribution of the dual-pooling fusion strategy, we conducted an ablation study comparing three model variants: (i) MobileNetV2 backbone only, (ii) RiceGuard without pooling fusion (simple concatenation), and (iii) the full RiceGuard architecture with dual-pooling. As shown in Table V, the dual-pooling mechanism improves test accuracy by 3% over simple concatenation and 5% over the single backbone, confirming that preserving both global context (GAP) and discriminative regions (GMP) is critical for distinguishing visually similar diseases.

Table 5: Ablation study on fusion strategies

Model Variant	Accuracy	Parameters
MobileNetV2 Only	85%	3.5M
RiceGuard (No Pooling)	87%	36.5M
RiceGuard (Dual-Pooling)	**90%**	**36.5M**

4.7 Deployment and latency analysis

To assess practical viability, the model was deployed via a Django web application and tested on mid-range hardware. The Django backend processes inference requests in <200ms on average. Furthermore, an Android prototype achieved 15 FPS on a device with 4GB RAM (Snapdragon 678), demonstrating capability for real-time field scouting without requiring high-end computational resources.

Discussion

Utilizing three distinct backbones results in enhanced uniqueness of features, consequently bolstering the robustness of the system. Additionally, the dual-pooling step enables our model to effectively employ both broad and local cues. The complete framework comprises 36.5M parameters which may be too much for very low-end phones; pruning or quantization are natural next steps. Despite strong performance, a limitation of Rice Guard is the 36.5M parameter count, which may strain devices with <4GB RAM. Future work will explore knowledge distillation to compress the ensemble into a single lightweight student network. Additionally, the current model predicts disease presence but not severity; integrating regression heads for lesion quantification is a planned extension.

Conclusion and future work

This hybrid deep learning system Rice Guard is designed to detect rice leaf disease and is built using MobileNetV2, DenseNet121, and ResNet50V2 with dual pooling. It reaches 90% accuracy on the test, exceeding each backbone.

Deployment tests confirm low-latency inference suitable for field use. The next steps will focus on lighter models which are mobile deployable and severity estimation for wider crop health support.

References

1. Baiju, B. V., et al. (2023). Automatic diagnosis of rice leaves diseases using hybrid deep learning model with self-attention module. *Journal of Advances in Information Technology, 14*(3), 418–425. https://doi.org/10.12720/jait.14.3.418-425
2. Mehnaz, S., & Islam, M. T. (2025). Rice leaf disease detection: A comparative study between CNN, transformer and non-neural network architectures. *arXiv*. https://doi.org/10.48550/arXiv.2501.06740
3. Naresh Kumar, B., & Sakthivel, S. (2025). Rice leaf disease classification using a fusion vision approach. *Scientific Reports, 15*, Article 8692. https://doi.org/10.1038/s41598-025-87800-3
4. Raman, R., & Jayaraman, S. (2025). Artificial intelligence for sustainable farming with dual branch convolutional graph attention networks in rice leaf disease detection. *Scientific Reports, 15*, Article 10595. https://doi.org/10.1038/s41598-025-94891-5
5. Rimi, S. A., et al. (2025). Empowering agricultural insights: Rice leaf DB a new dataset and optimal model selecting for diagnosis of rice leaf disease through transfer learning technique. *arXiv*. https://doi.org/10.48550/arXiv.2501.08912
6. Roy, P. S., & Kukreja, V. (2025). Vision transformers for rice leaf disease detection and severity estimation: A precision agriculture approach. *Journal of the Saudi Society of Agricultural Sciences, 24*(3), Article 3.