



RiceGuardAI: An Efficient And LightWeight Deep Learning Framework for Automated Rice Leaf Disease Detection And Agricultural Management

Veena K¹, Darshan Gowda M², Shalom Riona Fernandes³, Shreya Nayak B⁴, Shubham S Taple⁵, Sunil Kumar G⁶, Pushpa C N⁷

Abstract

Rice plays a vital role in global food supply, serving as a primary dietary source for a significant portion of the world's population. However, the occurrence of rice leaf diseases considerably affects agricultural productivity, leading to yield losses of nearly 20–30%. Timely and reliable disease identification is therefore essential for improving crop management practices and ensuring food security. Recent developments in deep learning have enabled automated disease diagnosis systems, but deploying such systems in practical agricultural environments requires a balance between prediction accuracy and computational efficiency. This study presents a comparative evaluation of four deep learning architectures, namely Custom Convolutional Neural Network (CNN), MobileNetV2, ResNet50, and InceptionV3, for the classification of four rice leaf conditions: Bacterial Blight, Brown Spot, Leaf Smut, and Healthy leaves. All models were trained and tested on a curated dataset containing 13,760 images using identical experimental settings to ensure fair performance comparison. Among the evaluated architectures, the proposed lightweight Custom CNN achieved the best overall performance with an accuracy of 99.90%, while maintaining only 2.1 million parameters, a compact model size of 9.8 MB, and a real-time inference speed of 12.13 ms. In comparison, MobileNetV2, InceptionV3, and ResNet50 achieved accuracies of 98.93%, 98.79%, and 88.91%, respectively. In addition to model evaluation, a deployment-oriented web-based agricultural platform was developed, integrating real-time disease diagnosis, ensemble-based yield prediction, expert consultation, and community-driven knowledge exchange. The proposed framework highlights the practical applicability of artificial intelligence in precision agriculture and demonstrates its potential for developing scalable and efficient farming solutions suitable for resource-constrained environments.

^{1,2,3,4,5,6,7} Dept. of Computer Science and Engineering, University of Visvesvaraya College of Engineering, Bengaluru, India.

Email: ¹veena.ise@gmail.com, ²darshugowda200@gmail.com, ³shalomriona@gmail.com,

⁴shreyanayakb26@gmail.com, ⁵shubham.s.taple@gmail.com, ⁶gsuneel.k@gmail.com, ⁷pushpacn@gmail.com

Keywords: Rice leaf disease detection, Convolutional Neural Network (CNN), transfer learning, precision agriculture, model comparison, deep learning, smart farming.

INTRODUCTION

Rice is one of the most widely cultivated crops in the world and serves as a major nutritional source for a large portion of the global population. The importance of rice production is particularly evident in Asian countries, which account for nearly 90% of worldwide rice cultivation and consumption [1]. Despite its agricultural significance, rice plants are frequently affected by several leaf diseases, including Bacterial Blight, Brown Spot, and Leaf Smut, which negatively impact both crop quality and yield. If these diseases are not identified and controlled during the early stages, production losses can reach approximately 20–30%, resulting in substantial economic damage estimated at more than \$10 billion annually [2].

Conventional rice disease diagnosis mainly depends on visual examination performed by farmers or agricultural specialists in the field. Although widely practiced, this method requires significant manual effort, is influenced by human judgment, and may not always be available in rural or economically limited areas where expert assistance is scarce [3]. These limitations have encouraged the development of automated disease detection techniques that can support timely identification of infections, improve decision-making, and minimize agricultural losses.

Recent progress in computer vision and deep learning has significantly improved automated image analysis, making Convolutional Neural Networks (CNNs) highly effective for plant disease recognition tasks [4]. Several research works have reported strong classification performance using both specially designed CNN architectures [5] and transfer learning techniques built on pre-trained networks such as ResNet [6], Inception [7], and MobileNet [8]. Although these approaches produce promising results, a large portion of existing studies primarily concentrates on achieving higher prediction accuracy using individual models, while comparatively fewer works provide a detailed evaluation of multiple architectures under the same experimental conditions.

Despite the progress achieved in deep learning-based disease detection, several important challenges still remain unresolved. First, many existing studies evaluate models using different datasets, preprocessing methods, and training conditions, making direct performance comparison unreliable and difficult to standardize. Second, a majority of the proposed approaches are limited to experimental or laboratory-level evaluations focused mainly on accuracy metrics, while only a few studies address the development of practical deployment-ready systems capable of delivering useful agricultural insights and real-time decision support for farmers.

Contributions:

To overcome these limitations, the present work introduces several key contributions:

1. A detailed comparative evaluation of four deep learning architectures—Custom CNN, MobileNetV2, ResNet50, and InceptionV3—is performed using a uniformly prepared rice leaf dataset containing 13,760 images, ensuring fair and consistent experimental analysis across models. These architectures were selected to represent diverse network designs and learning strategies.
2. A lightweight Custom CNN architecture is proposed, capable of achieving 99.90% classification accuracy while maintaining low computational complexity and reduced model size, making it suitable for real-time applications.
3. This study extends beyond model development by implementing a complete web-based agricultural management platform that combines real-time disease diagnosis, crop yield prediction, expert consultation support, and community-oriented knowledge sharing within a single system. This integration demonstrates a practical transition from research-oriented deep learning models to deployable precision agriculture solutions that emphasize both technical efficiency and real-world usability.

The structure of this paper is arranged as follows: Section II reviews existing studies related to rice leaf disease detection, Section III describes the dataset preparation and proposed methodology, Section IV presents the experimental results along with performance evaluation, Section V discusses the obtained findings and their practical significance, and Section VI summarizes the conclusions and outlines potential directions for future research.

LITERATURE REVIEW

The use of artificial intelligence techniques for automated plant disease identification has grown rapidly in recent years due to advancements in computer vision, deep learning, and the availability of agricultural image datasets [9]. Among various agricultural applications, rice leaf disease detection has received considerable research attention because rice remains one of the most essential food crops worldwide, and diseases such as Bacterial Blight, Brown Spot, and Leaf Smut contribute to major production and economic losses each year [2].

A. Traditional Machine Learning Approaches

Initial research in plant disease analysis mainly focused on conventional image processing and machine learning methods [10]. These approaches generally depended on manually extracting visual characteristics such as color, texture, and shape information from infected leaf regions [11]. Different preprocessing operations, including image segmentation, thresholding, and morphological filtering, were used to isolate diseased portions before classification using algorithms such as Support Vector Machines (SVM), K-Nearest Neighbors (KNN), and Random Forest classifiers [12].

Khirade and Patil [10] proposed a disease detection framework based on handcrafted image features combined with SVM classification and reported an accuracy close to 92% under controlled laboratory environments. Likewise, Mohanraj and Ashokumar [11] utilized Gray-Level Co-occurrence Matrix (GLCM) texture features together with Random Forest classification for rice disease identification. Although these methods produced reasonable performance in constrained settings, their dependency on manually engineered features reduced their adaptability to varying illumination conditions, cluttered backgrounds, and real-field agricultural environments.

B. CNN-Based Deep Learning Approaches

The emergence of Convolutional Neural Networks (CNNs) significantly transformed automated disease classification by enabling direct feature extraction from raw image inputs [13]. Unlike conventional machine learning techniques, CNN-based models can learn hierarchical spatial representations automatically without requiring handcrafted feature design [15]. In addition, recent deep learning architectures such as EfficientNet [14] have improved classification efficiency while reducing computational overhead.

Several researchers have explored customized CNN architectures for rice leaf disease recognition. Singh et al. [16] developed a custom CNN model that achieved approximately 99.83% classification accuracy and demonstrated the influence of optimizer selection, particularly the Adam optimizer, on training performance. Hassam et al. [17] introduced an agricultural image-oriented CNN architecture capable of achieving 98.7% accuracy across multiple crop datasets, outperforming several generalized pre-trained models. Similarly, Sharma et al. [18] evaluated CNN-based systems using rice leaf images collected under field conditions and reported an accuracy of 99.58%, emphasizing the capability of deep learning methods to handle complex environmental variations.

C. Transfer Learning and Pre-trained Architectures

Transfer learning has become a widely adopted strategy in agricultural image analysis, especially when available training data is limited. Instead of training models entirely from scratch, pre-trained architectures such as ResNet [6], Inception [7], VGG, DenseNet, and MobileNet [8] are fine-tuned for disease classification tasks using domain-specific datasets. Studies suggest that transfer learning can considerably improve accuracy and convergence performance when dataset size is relatively small [19].

Deb et al. [20] performed a comparative study of multiple transfer learning models and observed that InceptionV3 achieved nearly 96.8% accuracy in paddy disease classification tasks. Islam et al. [21] analyzed the performance of VGG, ResNet, and DenseNet architectures for identifying six rice diseases and found DenseNet121 to provide the best overall accuracy of around 95.2%, though performance varied among disease categories. MobileNet-based approaches have also become popular because of their lightweight architecture and suitability for deployment on mobile and embedded devices. Nayak et al. [22] reported an accuracy of 97.3% using MobileNetV2 for rice disease recognition.

D. Advanced Techniques and Recent Advances

Recent studies have focused on enhancing disease detection systems through advanced deep learning strategies such as attention mechanisms, ensemble learning, and real-time object detection frameworks. Attention-based architectures improve the model's ability to focus on infected regions of leaf images, while ensemble approaches combine multiple models to increase prediction robustness and reduce classification errors. For example, the DEX ensemble framework integrating DenseNet121, EfficientNetB7, and Xception achieved approximately 99.1% accuracy [24]. In parallel, real-time detection approaches based on YOLO architectures [25] and optimization-driven methods inspired by biological processes [26] have also gained importance in precision agriculture applications.

Literature Summary Table

TABLE I. LITERATURE SURVEY SUMMARY

Study	Method	Accuracy	Key Limitation
Khirade & Patil [10]	Handcrafted features + SVM	92%	Limited real-field adaptability
Singh et al. [16]	Custom CNN	99.8%	Absence of comparative analysis
Deb et al. [20]	InceptionV3 (Transfer Learning)	96.8%	Limited architecture diversity
Ahad et al. [24]	DEX Ensemble	99.1%	High Computationally Complexity

D. Research Gap and Motivation

Although previous studies have reported promising results for rice leaf disease classification, several important limitations still remain. Many research works focus primarily on improving the performance of a single architecture rather than conducting standardized comparisons across multiple models. Differences in datasets, preprocessing methods, and evaluation settings also make it difficult to fairly compare reported results across studies. In addition, only limited attention has been given to balancing predictive accuracy with computational efficiency, which is essential for practical deployment in real-world agricultural environments with limited computational resources.

To address these challenges, the present work performs a comparative analysis of four deep learning architectures—Custom CNN, MobileNetV2, ResNet50, and InceptionV3—using a unified dataset of 13,760 rice leaf images, identical preprocessing procedures, and consistent evaluation metrics. Furthermore, this study extends beyond experimental analysis by developing a complete web-based agricultural management platform that integrates disease classification, ensemble-based yield prediction, expert advisory support, and community-oriented information sharing features.

DATASET AND METHODOLOGY

This section presents the dataset used for rice leaf disease classification, along with the preprocessing operations and methodology adopted for model training and performance evaluation. A standardized experimental workflow was maintained throughout the study to ensure consistent and unbiased comparison across all selected deep learning architectures.

E. Dataset Description

The experimental dataset consists of 13,760 rice leaf images categorized into four classes: Bacterial Blight, Brown Spot, Leaf Smut, and Healthy leaves. The images were obtained from publicly accessible agricultural datasets and repositories, including sources such as Kaggle and PlantVillage, followed by careful curation to maintain balanced class representation, image quality, and annotation consistency. The dataset contains considerable diversity in terms of illumination, background conditions, leaf orientation, and disease intensity, enabling effective assessment of model performance under practical field-like scenarios.

To maintain proportional representation of all disease categories, a stratified data partitioning strategy was adopted. The dataset was divided into training, validation, and testing subsets as follows:

- **Training set:** 9,632 images (70%)
- **Validation set:** 2,064 images (15%)
- **Testing set:** 2,064 images (15%)

This distribution ensures that the trained models are evaluated using previously unseen samples, thereby providing a reliable measure of generalization capability and classification performance.

F. Image Preprocessing Pipeline

To enhance classification accuracy and maintain compatibility among all selected deep learning architectures, a uniform preprocessing framework was applied to every input image before model training and evaluation. Standardizing the preprocessing procedure also helped ensure fair performance comparison across different models.

The preprocessing steps include:

1. **Image Resizing:** All images were resized to a resolution of 224×224 pixels using bilinear interpolation to match the standard input dimensions required by architectures such as MobileNetV2, ResNet50, and InceptionV3.
2. **Pixel Value Normalization:** Image pixel intensities were normalized to a range between 0 and 1 to improve numerical stability during training and accelerate model convergence.
3. **Colour Space Standardization:** Every image was converted into RGB format to maintain consistency in colour representation across the dataset.
4. **Noise Suppression:** Bilateral filtering techniques were employed to minimize unwanted background noise while preserving important disease-related texture and edge information within the leaf regions.
5. **Label Representation:** Disease categories were transformed into one-hot encoded vectors to support multi-class classification using the categorical cross-entropy loss function.

C. Data Augmentation

To enhance model generalization and prevent overfitting, extensive data augmentation was applied during training using TensorFlow's ImageDataGenerator as shown in Fig 1.

```
train_datagen = ImageDataGenerator(
    rescale=1./255,
    rotation_range=20,
    width_shift_range=0.2,
    height_shift_range=0.2,
    shear_range=0.2,
    zoom_range=0.2,
    horizontal_flip=True,
    fill_mode='nearest'
)
```

Fig. 1. ImageDataGenerator function

The augmentation techniques include: Random rotations ($\pm 20^\circ$) Horizontal flipping, Width and height shifts ($\pm 20^\circ$), shearing transformations ($\pm 20^\circ$) and zoom (0.8-1.2 $\times$). These transformations stimulate real-world variations such as different leaf orientations, camera positioning and lighting conditions.

D. Deep Learning Models

Four deep learning architectures were selected for comparative evaluation:

I. Proposed Custom Convolutional Neural Network (CNN): A lightweight CNN architecture was specifically developed for this study with nearly 2.1 million trainable parameters, aiming to achieve high classification performance while maintaining computational efficiency. The proposed network is composed of four convolutional stages with gradually increasing filter sizes of 32, 64, 128, and 256, respectively. Each convolutional layer is followed by batch normalization, ReLU activation, and 2×2 max-pooling operations to improve feature extraction and training stability. Following the convolutional feature extraction process, global average pooling is applied to reduce dimensional complexity before passing the extracted representations through two fully connected layers containing 128 and 64 neurons. ReLU activation and dropout regularization values of 0.5 and 0.3 are incorporated to minimize overfitting and improve model generalization. Finally, a softmax output layer with four neurons is used to perform multi-class rice leaf disease classification. The architecture was intentionally designed to balance predictive accuracy, reduced model size, and faster inference suitable for real-time agricultural applications.

II. MobileNetV2: MobileNetV2 is a lightweight deep learning architecture designed for resource-efficient image classification tasks. The model utilizes depthwise separable convolutions and inverted residual blocks to significantly reduce computational cost and memory usage while maintaining strong predictive performance. Due to its compact structure and faster inference capability, MobileNetV2 is particularly suitable for deployment on mobile devices and edge-based agricultural monitoring systems.

III. ResNet50: ResNet50 is a deep residual learning architecture that incorporates skip connections, also known as residual connections, to overcome issues related to vanishing gradients during training. These connections enable the network to learn deeper feature representations effectively and improve overall classification capability. Although ResNet50 demonstrates strong performance in extracting complex disease-related patterns, it requires comparatively higher computational resources and training time.

IV. InceptionV3: InceptionV3 is a deep convolutional architecture designed to capture features at multiple spatial scales using parallel convolutional operations with different filter sizes. This multi-scale feature extraction capability allows the model to effectively identify intricate and irregular disease symptoms present in rice leaf images. The architecture also incorporates factorized convolutions and dimensionality reduction techniques to improve computational efficiency while maintaining high classification accuracy..

E. Training Strategy

To ensure a fair comparison, all models were trained under identical conditions:

- **Image Size:** 224 X 224 X 3
- **Optimizer:** Adam (lr=0.001)
- **Loss Function:** Categorical Cross-Entropy
- **Batch size and Epochs:** 32 with maximum 50 epochs
- **Early stopping:** Patience=10(monitor:val_accuracy)
- **Learning Rate Schedule:** ReduceLROnPlateau (factor=0.5, patience=5)
- **Regularization:** L2 weight decay (0.0001), Dropout (0.3-0.5)

For transfer learning models, a two-phase training strategy was employed: **Phase 1:** Base layers frozen, only the new classification head trained. **Phase 2:** Last 10-20% of base layers unfrozen and fine-tuned with a reduced learning rate (0.0001). The overall Methodology is as shown in Fig 2.




Fig. 2. Custom CNN Architecture

EXPERIMENTAL SETUP AND RESULTS

Experiments were conducted under a unified hardware and software environment to ensure fairness.

G. Experimental Setup

All experimental evaluations were carried out within a standardized hardware and software setup to maintain consistency and ensure fair comparison among the selected deep learning architectures. The implementation of all

models was performed using Python 3.10 along with TensorFlow 2.10.0 and Keras 2.13.1 libraries. Model training and evaluation were accelerated using an NVIDIA RTX 3080 GPU. In order to preserve experimental uniformity, every architecture was trained using the same dataset partitioning strategy consisting of 70% training data, 15% validation data, and 15% testing data, together with the identical preprocessing and augmentation procedures described in Section III.

Training Configuration: To ensure unbiased performance evaluation, all models were trained using a common set of hyperparameters. The Adam optimization algorithm was employed with an initial learning rate of 0.001, while categorical cross-entropy was used as the loss function for multi-class classification. Training was conducted using a batch size of 32 for a maximum of 50 epochs. Early stopping based on validation accuracy was incorporated to reduce overfitting and prevent unnecessary training iterations. Additionally, adaptive learning rate adjustment was implemented using the ReduceLROnPlateau strategy with a reduction factor of 0.5 and a patience value of 5 epochs. Regularization techniques such as L2 weight decay and dropout were also applied according to the requirements of each architecture to improve model generalization performance.

H. Model Performance Comparison

The performance of the four models—Custom CNN, MobileNetV2, ResNet50, and InceptionV3—was evaluated on the test dataset of 2,064 images. Table II summarizes the classification results based on accuracy, precision, recall, and F1-score.

TABLE II. PERFORMANCE COMPARISON OF DEEP LEARNING MODELS

MODEL	ACCURACY	PRECISION	RECALL	F1-SCORE
CUSTOM CNN	99.90%	99.91%	99.90%	99.90%
MOBILENETV2	98.93%	98.94%	98.93%	98.93%
INCEPTIONV3	98.79%	98.79%	98.79%	98.79%
RESNET50	88.91%	89.00%	88.91%	88.93%

The experimental findings demonstrate that the proposed Custom CNN architecture achieved the best overall performance among all evaluated models. The model attained a classification accuracy of 99.90% while maintaining a low inference time of 12.13 ms and a compact model size of only 9.8 MB. These results indicate that a carefully designed lightweight architecture tailored for rice leaf disease classification can outperform several widely used pre-trained deep learning models when optimized specifically for agricultural image patterns and disease characteristics. Both MobileNetV2 and InceptionV3 also produced strong classification results with accuracies exceeding 98%, highlighting the effectiveness of transfer learning approaches for rice disease recognition tasks. In contrast, ResNet50 exhibited comparatively lower performance while requiring significantly higher computational resources. This behavior suggests that very deep residual architectures may not always be the most suitable choice for this specific application, possibly due to excessive model complexity, over-parameterization, or insufficient adaptation to the characteristics of the available dataset.

I. Computational Efficiency Analysis

Beyond classification accuracy, computational efficiency was critically evaluated to assess deployment suitability for resource-constrained agricultural environments. Table III presents detailed computational characteristics.

TABLE III. COMPUTATIONAL CHARACTERISTICS OF MODELS

Model	Parameters (Millions)	Model Size (MB)	Inference Time (ms)	Efficiency Score
CUSTOM CNN	~2.1	9.8	12.13	8.24
MOBILENETV2	~3.4	14.2	21.80	4.54
INCEPTIONV3	~23.8	92.1	40.41	2.44
RESNET50	~25.6	98.4	50.23	1.77

The proposed Custom CNN achieved the highest Efficiency Score, calculated using the ratio of classification accuracy to inference time multiplied by 100, with a value of 8.24. In comparison, MobileNetV2, InceptionV3, and ResNet50 achieved efficiency scores of 4.54, 2.44, and 1.77, respectively. These results demonstrate that the Custom CNN provides a highly favorable balance between predictive accuracy and computational efficiency. The lightweight architecture, reduced inference latency, and smaller model size make it particularly well-suited for real-time agricultural deployment scenarios where limited hardware resources, faster response time, and low power consumption are critical considerations.

J. Confusion Matrix Analysis

Detailed class-wise evaluation using confusion matrix analysis revealed consistent classification trends across all evaluated models. The proposed Custom CNN achieved near-perfect prediction performance, producing only

4 misclassifications out of 2,064 testing samples. The observed errors were limited to confusion between the Brown Spot and Leaf Smut categories, indicating that certain disease symptoms share similar visual characteristics that may occasionally challenge the classification process.

MobileNetV2 produced a total of 22 misclassified samples, with errors distributed relatively evenly across different disease categories. InceptionV3 demonstrated comparable behaviour, resulting in 43 classification errors overall. In contrast, ResNet50 showed significantly weaker class-level performance with 229 misclassifications, most of which occurred between Brown Spot and Leaf Smut classes, as well as between Bacterial Blight and Brown Spot categories. These findings suggest that although deeper architectures can learn complex feature representations, they may not always effectively distinguish visually overlapping disease patterns within this dataset.




Fig. 2. Confusion Matrix Analysis

Healthy leaves were consistently identified with 100% accuracy by Custom CNN and MobileNetV2, and $\geq 96.9\%$ accuracy by all models, indicating robust discrimination between diseased and healthy tissue—a critical requirement for practical agricultural applications.

K. Convergence Analysis

The Custom CNN demonstrated optimal training behavior with rapid convergence (98% accuracy within 10 epochs) and minimal overfitting (validation loss stabilized at 0.0987). MobileNetV2 showed typical transfer learning patterns with initial rapid gains but developed a validation gap (training loss: 0.0044 vs. validation loss: 0.7780) in later epochs, indicating potential overfitting to training specifics. ResNet50 exhibited unstable convergence requiring extensive hyperparameter tuning, while InceptionV3 showed steady but slower convergence consistent with its architectural complexity.

DISCUSSION

The comparative evaluation of deep learning models for rice leaf disease classification provides several important insights into model performance and practical usability.

Although all evaluated models achieved high classification accuracy, noticeable differences were observed in terms of computational efficiency, model complexity, and deployment suitability.

L. Architectural Suitability Analysis

The Custom CNN's high performance stems from its tailored design with only 2.1M parameters, balancing capacity without overfitting. Its progressive convolutional layers capture hierarchical disease features efficiently. By contrast, ResNet50's depth may overly smooth fine-grained textures, harming discrimination of similar diseases like Brown Spot and Leaf Smut.

M. Transfer Learning Effectiveness

MobileNetV2 and InceptionV3 achieved $>98\%$ accuracy via transfer learning, yet their generic ImageNet features slightly lagged the domain-specific Custom CNN. This finding underscores that architectural suitability is critical, a factor often overlooked in prior comparisons that used inconsistent experimental settings.

N. Practical Deployment Considerations

The Custom CNN excels in field deployment due to its 99.90% accuracy, 12.13 ms inference, 9.8MB size, and simpler interpretability. MobileNetV2 suits highly constrained environments, while InceptionV3 offers multi-scale analysis at higher computational cost.

O. Error Pattern Implications

Persistent confusion between Brown Spot and Leaf Smut across models highlights a diagnostic challenge, guiding future dataset enrichment with more diverse and early-stage samples—addressing the multi-disease classification gap.

CONCLUSION AND FUTURE WORK

- This study demonstrates that a lightweight and domain-focused Convolutional Neural Network (CNN) can achieve highly effective performance for rice leaf disease classification while maintaining computational efficiency. The proposed Custom CNN attained a classification accuracy of 99.90% with an inference time of 12.13 ms and approximately 2.1 million trainable parameters. Experimental evaluation conducted under identical training and testing conditions showed that the proposed architecture outperformed established pre-trained models such as MobileNetV2, InceptionV3, and ResNet50. These findings emphasize that architectures specifically designed for agricultural image analysis can provide better efficiency and generalization compared to deeper generic networks originally developed for large-scale image recognition tasks.
- The results further suggest that models containing nearly 2 million parameters offer an effective balance between learning capacity, computational cost, and generalization ability for rice leaf disease classification. Although the overall classification performance was nearly perfect, limited confusion was still observed between Brown Spot and Leaf Smut categories, mainly due to visual similarity in certain disease stages. This indicates the importance of incorporating more diverse and early-stage disease samples in future datasets to improve model robustness.
- Future research can focus on real-world field validation, integration of explainable artificial intelligence techniques for improving model interpretability, and deployment within IoT-enabled precision agriculture systems for automated crop monitoring. Overall, this work contributes both a systematic framework for fair deep learning model comparison and a lightweight deployment-ready solution that helps bridge the gap between laboratory-based experimental performance and practical agricultural applications.

REFERENCES

- [1] FAO, "Rice Market Monitor," Food and Agriculture Organization, Rome, Italy, vol. 26, no. 1, 2023.
- [2] S. Savary et al., "The global burden of pathogens and pests on major food crops," *Nature Ecology & Evolution*, vol. 3, pp. 430–439, 2019.
- [3] P. K. Sethy et al., "Deep feature-based rice leaf disease identification using support vector machine," *Computers and Electronics in Agriculture*, vol. 175, p. 105527, 2020.
- [4] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.
- [5] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [7] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2818–2826.
- [8] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4510–4520.
- [9] J. G. A. Barbedo, "A review on the main challenges in automatic plant disease identification based on visible range images," *Biosystems Engineering*, vol. 144, pp. 52–60, 2016.
- [10] S. D. Khirade and A. B. Patil, "Plant disease detection using image processing," in *Proc. International Conference on Computing, Communication, Control and Automation (ICCUBEA)*, 2015, pp. 768–771.
- [11] I. Mohanraj and K. Ashokumar, "Automated rice plant disease detection using texture feature analysis," *International Journal of Engineering and Advanced Technology*, vol. 9, no. 1, pp. 5332–5336, 2019.
- [12] S. Dhar, M. K. Ghosh, and R. Sen, "Comparative evaluation of SVM, KNN, and AdaBoost for rice foliar disease detection," *International Journal of Agricultural and Environmental Information Systems*, vol. 13, no. 1, pp. 1–17, 2022.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2012, pp. 1097–1105.
- [14] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. International Conference on Machine Learning (ICML)*, 2019, pp. 6105–6114.
- [15] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [16] A. Singh, V. Kumar, and P. R. Bisen, "Custom CNN for rice leaf disease classification using optimized learning strategies," *IEEE Access*, vol. 11, pp. 23456–23467, 2023.
- [17] S. M. Hassan, A. K. Maji, and M. Jasinski, "A novel CNN architecture for agricultural disease recognition across multiple datasets," *Expert Systems with Applications*, vol. 206, p. 117870, 2022.
- [18] P. Sharma, R. Singh, and S. Gupta, "Plant disease detection in field images using CNN and classical machine learning methods," *Computers and Electronics in Agriculture*, vol. 198, p. 107012, 2022.
- [19] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [20] M. Deb, S. Roy, and A. Dutta, "Comparative evaluation of deep learning models for paddy leaf disease classification," *Journal of Plant Diseases and Protection*, vol. 128, no. 3, pp. 789–802, 2021.

- [21] M. M. Islam, M. A. Hossain, and T. A. T. Ali, "Deep learning architectures for multi-class rice disease detection using leaf images," *IEEE Transactions on AgriFood Electronics*, vol. 2, no. 3, pp. 215–226, 2020.
- [22] A. Nayak, D. K. Behera, and M. R. Panda, "Mobile-based plant disease detection using transfer learning with optimized image acquisition," *Multimedia Tools and Applications*, vol. 82, no. 10, pp. 15247–15268, 2023.
- [23] N. Abinaya, S. P. Kumar, and R. Venkatesan, "CAAR U-Net: Cascading autoencoder with attention residual U-Net for plant leaf disease segmentation and classification," *Engineering Applications of Artificial Intelligence*, vol. 123, p. 106241, 2023.
- [24] Md. T. Ahad, Y. Li, and B. Zhang, "DEX: A hybrid ensemble framework for rice disease detection using DenseNet121, EfficientNetB7 and Xception," *Computers and Electronics in Agriculture*, vol. 204, p. 107532, 2023.
- [25] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.
- [26] T. Daniya and S. Vigneshwari, "Rider Water Wave Neural Network for early rice leaf disease detection with SegNet-based segmentation," *Soft Computing*, vol. 27, no. 14, pp. 9871–9885, 2023.