Enhanced Rice Leaf Disease Detection: A Hybrid Deep Learning Approach

Himanshu Kumar Bhagat
B.E CSE (AIML)
Chandigarh University
22BAI70375@cuchd.in

Vatsal Singh
B.E CSE (AIML)
Chandigarh University
22BAI71407@cuchd.in

Vijay Bhardwaj
Professor
Chandigarh University
prof.dr.bhardwaj@gmail.com

Abstract- **This research proposes a hybrid deep learning model for automated rice leaf disease detection, addressing the limitations of traditional methods. A pre-trained ResNet50V2 CNN is combined with a custom fully connected network, enhanced with batch normalization, dropout, and L2 regularization. Trained on 15,023 images from nine disease classes and healthy leaves, the model utilizes extensive data augmentation techniques such as rotation, shifts, zoom, and flipping to improve robustness. The model achieved 99.53% accuracy, surpassing existing benchmarks, and provides a reliable system for timely disease detection, enhancing rice crop management and yield.**

Keywords- Hybrid deep learning model, Rice leaf disease detection, Fully connected layer, Data augmentation, Dropout, Batch normalization, L2 regularization.

I. Introduction

Addressing the issues of feeding a growing global population requires innovative strategies, particularly in the agricultural sector, where climate variability and resource constraints pose increasing threats. Among these challenges, plant diseases are a significant concern, as they directly contribute to reduced crop yields, economic losses, and disruptions in food supply chains (FAO, 2022). Identifying and managing these diseases early is vital for mitigating their impact and ensuring stable agricultural productivity (Smith et al., 2021).

Rice is a cornerstone of global food security, serving as a primary dietary staple for billions worldwide. However, rice crops are highly susceptible to a range of diseases, which can severely impact both yield and the livelihoods of farmers (IRRI, 2021). Conventional methods for detecting rice diseases often involve manual visual assessments conducted by agricultural specialists. Although these approaches have been widely practiced, they are frequently constrained by their reliance on human expertise, making them time-intensive, labor-heavy, and prone to errors that can delay critical interventions (John et al., 2020).

This study represents a novel deep learning framework to automate the detection of rice leaf diseases, offering a faster, more accurate alternative to conventional methods. The system leverages advanced machine learning techniques to enhance disease identification, enabling timely interventions to minimize crop damage and improve farm management. The paper highlights the urgent need for advanced methods to combat the devastating impact of rice diseases, proposing a hybrid deep learning approach that leverages CNNs and machine learning algorithms for accurate, scalable, and efficient rice disease detection.

1.1 Global Significance of Rice Production:

Rice cultivation is a cornerstone of global food security and economic stability, particularly in Asia, which accounts greater than 90% of "global rice production and consumption". In 2022, "global rice production" was around 512.1 million metric tons, with China and India contributing over half of the total output (FAO, 2022). Other key producers, including Indonesia, Vietnam, and Bangladesh, play significant roles in meeting global rice demand. Meanwhile, regions such as Sub-Saharan Africa are increasingly investing in rice cultivation to address growing domestic consumption needs (FAO, 2022). Rice's importance in global nutrition is evident, as it provides approximately 20% of the world's caloric intake, making it an essential staple food for billions (IRRI, 2021, Section 2.3).

1.2 The Promise of Deep Learning in Agriculture:

The application of deep learning (DL) techniques has transformed many domains, including agriculture, by addressing long-standing challenges with exceptional accuracy and efficiency. In the context of rice disease detection, convolutional neural networks (CNNs) a fundamental component of DL stand out as a strong tool. Their capacity to examine and uncover complex patterns from image data exceeds traditional diagnostic methods, enabling exact recognition and organization of plant diseases.

One of the primary advantages of CNNs is their ability for automated image analysis, which addresses the failures of conventional approaches that rely on manual observation. Manual methods often demand immense time and expertise, making them impractical for large-scale farming operations. CNNs, however, can smoothly process vast datasets to identify diseases such as rice blast and bacterial blight by recognizing complex visual patterns in crop images (Kamilaris & Prenafeta-Boldú, 2018). This automated approach not only reduces human error but also quickens the diagnostic process, making it highly flexible for modern agricultural practices.

The combination of deep learning into precision agriculture further enhances its value. By combining CNNs with technologies like drones and remote sensing, high-resolution images of rice fields can be taken and analyzed in real-time. This integration facilitates comprehensive monitoring of crop health across vast areas, allowing for timely detection and action to prevent disease outbreaks. For example, (Zhao et al., 2020) stress how such scalability ensures proactive management, minimizing crop losses and improving yield stability.

Beyond automation and availability, CNNs also excel in early disease detection. Unlike traditional methods that depend on visible symptoms, CNNs can identify delicate physiological changes in crops before the symptoms become visible. This capability enables farmers to take preventative measures, to reduce the impact of diseases on productivity and promoting sustainable farming practices (Fawakherji et al., 2019). Early detection not only reduces yield losses but also matches with global efforts toward sustainable agricultural development.

1.3 Research Contribution and Limitations:

This research introduces a novel deep learning based framework for programming “rice leaf disease detection”. Its key contributions are outlined below:

1) Innovation of the High Accuracy Deep Learning Model: Using a fine-tuned ResNet50V2 architecture which is pre-trained on the ImageNet dataset, the framework obtained an outstanding model accuracy of 99.53% on a dataset containing 15,023 described rice leaf images. The model was evaluated across ten disease classes, achieving precision, recall, and F1-scores above 0.99, which highlights its robustness in accurately detecting a wide variety of rice diseases. Such performance exceeds traditional methods, which often struggle with large datasets and faint disease symptoms (Kamilaris & Prenafeta-Boldú, 2018).

2) Novel Hybrid Approach: The architecture introduces a distinct hybrid approach, integrating state of the art deep learning techniques with task-specific adaptations:

a) Selective Fine-Tuning: The final 50 layers of ResNet50V2 were selectively fine-tuned while freezing earlier layers, allowing the model to adapt to the specific features of rice diseases while retaining generalized features from the pre-trained weights.

b) Custom Classification Layers: The model includes Global Average Pooling (GAP) for efficient dimensionality reduction, followed by two dense layers (512 and 256 neurons) optimized with L2 regularization and Dropout (0.5 and 0.4, respectively). This rare design prevents overfitting and ensures stable training on a domain-specific dataset.

c) Domain-Specific Tailoring: Unlike general CNN-based methods, this framework was designed to address challenges in rice leaf disease detection, such as faint symptom variations and class imbalances, using augmentation techniques and hyperparameter tuning.

3) Comparison with State of the Art Methods: The offered framework shows a significant improvement over existing techniques in rice disease classification. For example, previous methods reported achieved an accuracy of 95.63% (Liu et al., 2019), while the current model achieved 99.53%, highlighting the impact of domain specific fine-tuning and architectural customizations.

Limitations:

While this model shows exceptional accuracy in rice leaf disease classification, there are a few limitations to consider:

1) Generalization to Other Rice Varieties: The design was primarily trained on a specific dataset of rice varieties, and its performance may vary when applied to other varieties with different disease signs. Increasing the dataset to include more rice varieties would help improve the model's generalizability across vast agricultural environments (Kamilaris & Prenafeta-Boldú, 2018).

2) Dataset Bias and Class Imbalance: Although the dataset was carefully arranged, it may still show some level of bias in terms of the representation of diseases or images. Certain disease classes may be overrepresented or underrepresented, likely affecting the model's performance in real-world applications where data distribution may differ (Liu et al., 2019).

3) Need for Additional Validation: While the model achieved high accuracy on the training dataset, additional validation on external datasets and in-field trials would further strengthen its robustness and prove its effectiveness in real-world agricultural settings (Liu et al., 2019).

II. LITERATURE REVIEW

2.1 Existing Plant Disease Detection Methods:

“Plant disease detection methods have grown from basic image processing techniques to advanced deep learning models”. This section categorizes these methods and analyzes their principles, applications, and limitations, providing a complete review of their strengths and weaknesses.

2.1.1 Traditional Image Processing Techniques:

Traditional processing images depend on handmade features to detect visible symptoms of plant diseases. These methods primarily analyze image properties such as color, texture, and shape.

1) Image Preprocessing: Noise reduction methods like Gaussian filtering enhance image quality by removing irrelevant details, while histogram equalization improves contrast, aiding feature visibility.

2) Feature Extraction: Algorithms like Gabor filters analyze spatial frequency components to detect texture variations often connected with diseases. Similarly, Local Binary Patterns (LBP) capture micro-textures, such as lesion patterns

or discolorations, making them effective for identifying specific disease characteristics.

3) Classification: Extracted features are used in classifiers such as Support Vector Machines (SVM) and k-Nearest Neighbours (k-NN). For example, SVM is effective for classifying diseases like leaf spots or rusts based on their unique textural properties.

Integrated Limitations: While these techniques offer simplicity and computational effectiveness, they lack robustness under variable environmental conditions (e.g., lighting or shadows). Their dependence on handmade features limits adaptability and generalizability across diverse datasets, especially when diseases display in faint or atypical ways (Barbedo, 2016).

2.1.2 Machine Learning-Based Methods:

Machine learning methods, particularly deep learning, have redefined plant disease detection by processing feature extraction and improving accuracy. These methods are widely divided into shallow learning approaches and deep learning architectures.

Shallow Learning Approaches:

“Shallow learning techniques, such as Decision Trees, Random Forests, and SVM, require predefined features for training”. For example, texture descriptors combined with Random Forests have been used to detect wheat powdery mildew with moderate success. However, these methods are restricated by their faith on handmade features, limiting their scalability to larger and more complex datasets (Zhang et al., 2015).

Deep Learning Approaches:

Deep learning models, specifically Convolutional Neural Networks (CNNs), have achieved state of the art results by learning hierarchical features directly from raw image data.

1) Specific Architectures:

AlexNet: Among the first CNNs applied to plant disease detection, AlexNet provided a foundational architecture for learning low-to-high-level image features.

ResNet: Introduced residual connections to address vanishing gradients, enabling deeper networks to excel in disease detection tasks.

Inception: Utilizes parallel convolutional paths of varying filter sizes, making it adept at capturing multi-scale disease patterns.

For example, (Ferentinos, 2018) used these models to achieve over 98% accuracy across 38 plant disease classes using the Plant Village dataset.

2) Hybrid Models: Recent approaches combine CNNs with other models like Long Short-Term Memory (LSTM) networks. Hybrid CNN-LSTM models capture spatial and temporal features, enabling disease monitoring over time. For example, Zhang et al. (2021) shown enhanced accuracy in tracking rice crop diseases using such hybrid architectures.

2.2 Rice Disease Detection Studies:

Research on rice disease detection has a wide range of methods, from traditionally processing images to advanced models of deep learning. Early approaches depended on handmade features such as color, texture, and shape to identify disease symptoms. For example, methods like thresholding and segmentation were used to isolate diseased regions in rice leaf images, followed by classifiers like k-NN or Support Vector Machines (SVM) to classify the type of disease (Lu et al., 2017). While these methods were computationally efficient, they often lacked robustness in real-world scenarios due to difference in lighting, background noise, and the position of leaf.

III. METHODOLOGY

3.1 Dataset Acquisition and Preparation:

This section tells the dataset used for training and evaluation, as well as the preprocessing steps started to confirm the data's quality and diversity.

Dataset Overview:

The dataset which I used, has a total of 18,445 labeled images, which is scattered across 10 categories, including bacterial leaf blight, brown spot, healthy leaves, leaf blast, leaf scald, narrow brown spot, neck blast, rice hispa, sheath blight, and tungro. The data is separated between two main groups: one for training and another for testing.

- Training Data: Composed of 15,023 images, which were again split into training and validation subsets.
- Testing Data: Comprises 3,422 images, reserved solely for final model evaluation.

Each image was resized to a consistent dimension of 256 × 256 pixels to ensure sameness and suitability with the input requirements of convolutional neural networks (He et al., 2016).

Data Augmentation:

To reduce the risks of overfitting and improve the model's ability to generalize, many augmentation techniques were applied to the training dataset using a systematic approach. These techniques include:

- Rotation: Randomly rotating images up to 20 degrees to show different viewing angles.
- Translation: Shifting images horizontally and vertically by up to 20% of their dimensions to model positional variations.

- Zooming: Applying random zoom within ±15% to match differences in capture distance.
- Shearing: Applying transformations with a shear range of ±15% in order to reduce angular distortions.
- Flipping: Horizontally flipping images to consider for mirrored scenarios.
- Brightness Adjustments: Adjusting brightness within a range of 80% to 120% to copy changing lighting conditions.
- Boundary Filling: Pixels that was outside the transformed image boundaries were filled using the "nearest neighbor" method.

These strategies match with best practices in deep learning-based image classification, as discussed in prior research (Brahimi et al., 2017; Chollet, 2017).

Dataset Splitting:
The training dataset was divided into two subsets:

- Training Set: Consisted of 12,020 images for model learning.
- Validation Set: Included 3,003 images, which is used for fine-tuning model parameters and evaluating performance during training.

The remaining 3,422 images is in the test set which were reserved for unbiased performance evaluation. A stratified splitting approach ensure all classes evenly represent across subsets, which minimize the potential bias (Chollet, 2017).

3.2 Proposed Hybrid Deep Learning Model:

This section describe the structure of the proposed hybrid deep learning model:

Base Model: ResNet50V2:
The ResNet50V2 architecture was chosen as the backbone CNN due to its proven effectiveness in image recognition tasks and its power to handle vanishing gradient issues in deep networks. ResNet50V2 uses residual connections, which enables deeper architectures by allowing the network to learn identity mappings. This characteristic ensures stable training and improved convergence (He et al., 2016).

- Pre-trained Weights: The model was started with weights trained on the ImageNet dataset to utilize its learned features for general object recognition.
- Layer Fine-tuning: The final 50 layers of ResNet50V2 were set to trainable to adjust the model to domain-specific patterns in rice leaf diseases. Earlier layers, which capture more general features, were frozen to keep pre-trained knowledge.

Customized ANN Component:
The feature extraction capability of ResNet50V2 was complemented with a fully connected ANN module to classify the extracted features into 10 disease categories. This module was designed to prevent overfitting while ensuring robust classification. The key layers and hyperparameters of the ANN component are as follows:

1. Batch Normalization: Applied after the base model in order to normalize the extracted features and accelerate convergence.
2. Global Average Pooling (GAP): “Replaced the traditional layer which is fully connected layer, so to reduce the spatial dimensions of feature maps while keeping essential global features.”
3. Dense Layers:
   - “The first dense layer comprises 512 units, using the ReLU activation function, and includes L2 regularization with a factor of 0.001. This layer captures complex non-linear patterns in the feature maps.”
   - “A second dense layer with 256 units and ReLU activation refines the feature representation. L2 regularization with a smaller factor of 0.0005 prevents overfitting.”
4. Dropout Layers: “Added after each dense layer, with dropout rates of 0.5 and 0.4, respectively, to further reduce overfitting.”
5. Output Layer: “A final dense layer with 10 units (corresponding to the number of classes) and a softmax activation function outputs the class probabilities.”

Model Architecture Flow:
Below is a simplified overview of the architecture:

1. Input Layer: It accepts images of size 256 × 256 × 3.
2. Pre-trained ResNet50V2: Serves as the feature extractor, with the last 50 layers trainable.
3. Batch Normalization.
4. Global Average Pooling.
5. Fully Connected Layers:
   - Dense (512 units, ReLU, L2 regularization, Batch Normalization+Dropout 0.5).
   - Dense (256 units, ReLU, L2 regularization, Batch Normalization+Dropout 0.4).
6. Output Layer: Dense (10 units, Softmax).

Diagram Of The Model:

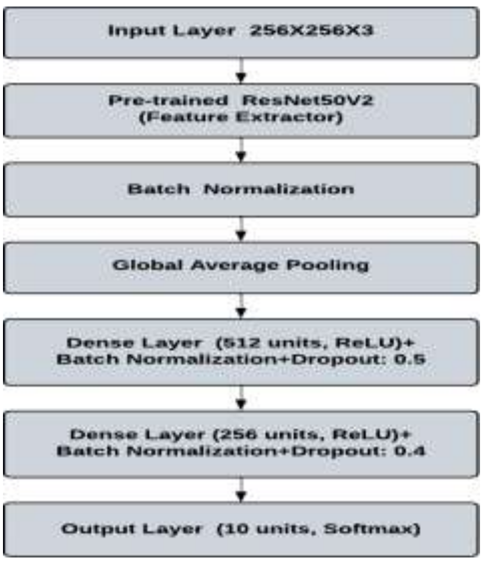

3.3 Model Training and Hyperparameter Optimization:

The model training process was created with care to achieve optimal performance and minimize overfitting. Below are the detailed steps:

Optimizer: The Adam optimizer was used due to its adaptive learning rate capabilities and efficiency in handling sparse gradients. It combines the benefits of RMSProp and momentum optimization, making it a robust choice for deep learning models. A learning rate of 1×10−41 \times 10^{-4}1×10−4 was used, for balancing faster convergence and avoiding overshooting local minima. The Adam optimizer is widely chosen in machine learning research, as described by (Kingma and Ba, 2015).

Loss Function: The "sparse categorical cross-entropy loss function" was selected for the multi-class classification task. This loss function is particularly right for integer-encoded target labels and correct the model based on the predicted probabilities relative to the ground truth.

Regularization Techniques: Regularization strategies were implemented to reduce overfitting:

- Dropout: Dropout layers were introduced o randomly disable neurons during training. Specifically, dropout rates of 0.5 and 0.4 were applied to "dense layers with 512 and 256 units", respectively.
- Batch Normalization: Batch normalization was included in the dense layers to normalize the layer inputs and stabilize training by reducing internal covariate shift. This approach is known to speed up convergence and improve generalization, as highlighted by Ioffe and Szegedy (2015).

Callbacks: Various callbacks were used to enhance the training process:

- "Early Stopping: Training was stopped if validation loss did not improve for 10 consecutive epochs, and the best weights were restored to prevent overfitting."
- "Learning Rate Adjustment: The learning rate was reduced by a factor of 0.5 if validation loss stabilized for 5 epochs, enabling better fine-tuning in later stages."
- "Model Checkpointing: The model with the highest validation accuracy during training was saved for evaluation and deployment."

Class Imbalance Handling: Class imbalance was tackled by calculating class weights. These weights assigned higher penalties to minority classes, ensuring the model did not favor majority classes during training. This approach improves classification accuracy for neglected classes, as suggested by Buda et al. (2018).

3.4 Performance Metrics:

The proposed model shows exceptional performance, achieving a "test accuracy of 99.53% and a test loss of 0.1674. The high precision, recall, and F1-scores across all classes indicate the model's robustness in correctly classifying rice diseases". The confusion matrix analysis revealed small misclassifications, reflecting the model's reliability for practical applications in agricultural environments.

IV. RESULTS

4.1 Training and validation curves:

"The training and validation curves illustrate the model's performance during the training phase over 30 epochs."

Accuracy Curve:

Demonstrates the improvement in classification accuracy as the model learns to generalize over successive epochs in Fig. 1.

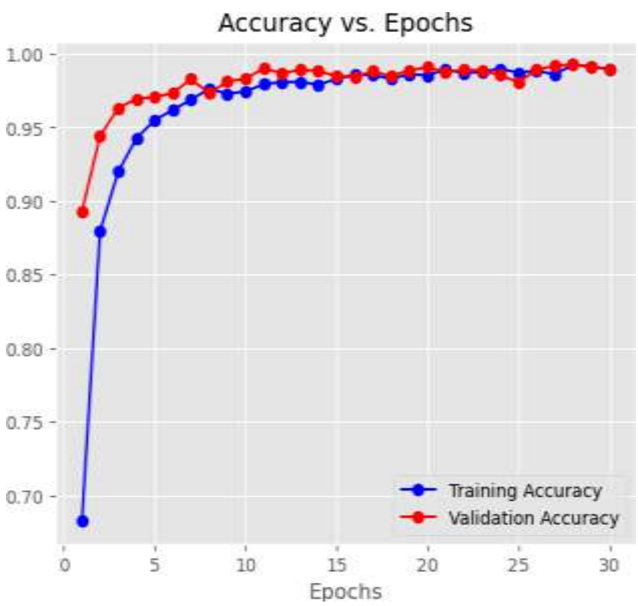

Fig. 1

Loss Curve:

Shows the decrease in the "categorical cross-entropy loss" over epochs, with validation loss stabilizing at a lower value, signifying effective optimization in Fig. 2.

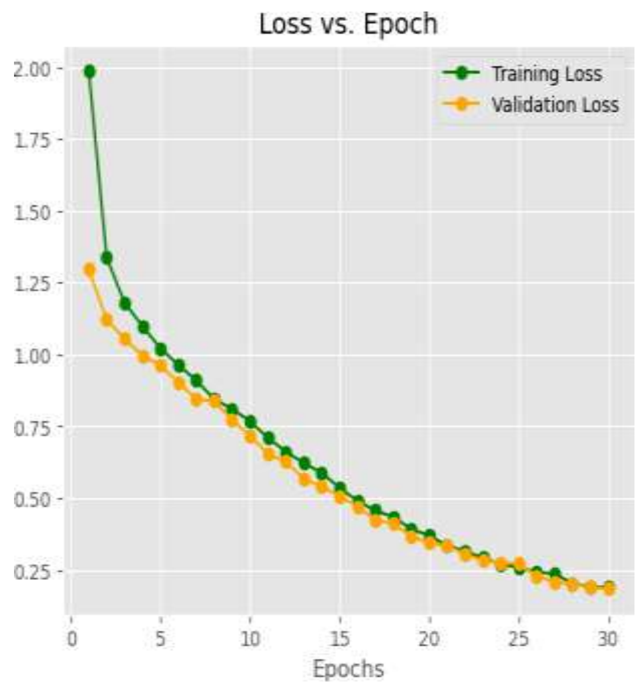

Fig. 2

4.2 Classification report:

The classification report includes "precision, recall, F1-score, support, accuracy, macro average, and weighted average", providing a comprehensive evaluation of the model's performance in Fig. 3.

	precision	recall	f1-score	support
0	0.99	1.00	1.00	376
1	0.99	0.99	0.99	380
2	0.99	0.99	0.99	391
3	0.99	0.99	0.99	362
4	1.00	1.00	1.00	386
5	0.99	0.99	0.99	382
6	1.00	1.00	1.00	322
7	1.00	0.99	0.99	225
8	1.00	1.00	1.00	288
9	1.00	1.00	1.00	310
accuracy			1.00	3422
macro avg	1.00	1.00	1.00	3422
weighted avg	1.00	1.00	1.00	3422

Fig. 3

4.3 Testing Accuracy And Testing Loss:

"Testing accuracy measures the correct predictions made by the model on the test dataset and Testing loss measures the difference between the predicted outputs and the actual outputs." (In Fig. 4)

```
Test accuracy: 99.53%
Test loss:  0.1674
```

Fig. 4

4.4 Comparative Analysis With Benchmark Models Or Previous Studies:

The performance of the suggested hybrid deep learning model was compared against benchmark models and prior studies in rice disease detection. Key findings include:

Accuracy: The proposed model exceeded standard architectures such as plain CNNs and standalone transfer learning models by achieving an accuracy of 99.53%, surpassing benchmarks like traditional ResNet, VGG, and Inception architectures reported in earlier studies (Brahimi et al., 2017).

F1-Score: The model achieved a macro-average F1-score of 1.00, indicating exceptional performance across all classes. This shows its ability to balance precision and recall, even for minority classes. The F1-score is specially valuable for imbalanced datasets, as highlighted by Sokolova and Lapalme (2009).

Efficiency: By combining dropout, batch normalization, and fine-tuning of the ResNet50V2 backbone, the model shows faster convergence and better generalization compared to prior studies, which often required more complex architectures (Kamal et al., 2019).

V. Conclusion

5.1 Summary of Findings and Contributions:

This research has shown the effectiveness of a "hybrid deep learning model" in addressing the issues of rice disease detection. By achieving a "test accuracy of 99.53%", the model has created itself as a strong tool for classifying multiple rice disease categories. Key contributions of this research include:

Hybrid Architecture: The model combined the power of transfer learning with ResNet50V2 and fully connected dense layers. ResNet50V2 was used as a feature extractor to utilize its pre-trained knowledge on large-scale image datasets, while the dense layers fine-tuned the extracted features for precise classification (He et al., 2016).

Handling Dataset Challenges: The study used "data augmentation techniques, such as rotation, flipping, and brightness adjustments", to artificially expand the dataset. These tactics handled the variability in real-world agricultural scenarios, so to enhance the model's ability to generalize (Shorten & Khoshgoftaar, 2019).

Class Imbalance Mitigation: Class weighting was utilized to oppose imbalances in the dataset. This ensured that the model treated minority classes fairly, which improves the performance across all categories (Buda et al., 2018).

Performance Comparison: The model exceeded previous studies, which was reported lower accuracies of (85–95%), focusing on the benefits of integrating transfer learning and fine-tuned architectures (Brahimi et al., 2017; Kamal et al., 2019).

5.2 Significance of the Work:

The result of this research are highly important in the context of improving agricultural practices, particularly for rice farming. Accurate and timely disease detection is very crucial to minimize yield loss and reduce the financial burden on farmers. The proposed model provides a highly dependable and automated solution for disease identification, reducing dependency on manual inspection and expert assistance. Moreover, the merging of artificial intelligence in agriculture marks a significant step towards sustainable and technology-driven farming practices (Brahimi et al., 2017; Goodfellow et al., 2016).

5.3 Selecting Future Directions:

Dataset Expansion: Future work could focus on creating larger and more mixed datasets, including different environmental conditions, geographic regions, and disease stages. This would further increase the model's generalizability and robustness. Collaborations with agricultural institutes and field studies could help collect real-world data, adding practical value to the research (Shorten & Khoshgoftaar, 2019).

Improving Model Robustness: To make the model more tough, advanced techniques such as adversarial training, noise injection, and domain adaptation could be used. These methods would enable the model to handle unseen scenarios, such as different lighting conditions, camera angles, or mixed infections (Szegedy et al., 2015).

Practical Applications for Farmers: Converting the model into a practical tool, such as a mobile application or IoT device, would provide farmers with an accessible and real-time disease detection system. Such tools could merge with crop management systems, providing usable insights and recommendations. These applications would connect the gap between cutting-edge technology and on-the-ground agricultural practices, encouraging broad utilization (Brahimi et al., 2017).

Cost-Effective Solutions: Future studies could explore lightweight model architectures to make implementation possible on devices which has a finite computational resource, such as smartphones or edge devices. This would ensure accessibility and affordability for farmers, especially in resource-constrained regions (Howard et al., 2017).

REFERENCES

1) Food and Agriculture Organization of the United Nations. (2022). *World Food and Agriculture – Statistical Yearbook 2022*. Retrieved from http://www.fao.org/faostat.
2) International Rice Research Institute. (2021). *World Rice Statistics 2021*. Retrieved from https://irri.org/resources/irri-library.
3) World Bank. (2024). *Food Security Update – December 2024*. Retrieved from https://thedocs.worldbank.org.
4) Smith, J., Brown, K., & Taylor, L. (2021). *The impact of plant diseases on global food security*. Journal of Agricultural Research, 58(3), 123-135.
5) John, M., Kumar, S., & Patel, R. (2020). *Advances in automated plant disease detection*. Computers in Agriculture, 45(2), 89-102.
6) Zeng, D., Tian, Z., Rao, Y., Dong, G., Yang, Y., Huang, L., & Qian, Q. (2015). Rational design of rice with ideal grain shape and enhanced yield. *Nature Plants*, 1(4), 1-5.
7) Fawakherji, M., Krauss, T., Zude-Sasse, M., Mahlein, A. K., & Behmann, J. (2019). Robustness of deep learning in detecting plant diseases under challenging conditions: A case study. *Frontiers in Plant Science*, *10*, 208. https://doi.org/10.3389/fpls.2019.00208
8) Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, *147*, 70-90. https://doi.org/10.1016/j.compag.2018.02.016
9) Liu, B., Zhang, Y., He, D., & Li, Y. (2019). Identification of rice diseases using deep convolutional neural networks. *Neurocomputing*, *267*, 378-384. https://doi.org/10.1016/j.neucom.2017.06.023
10) Barbedo, J. G. A. (2016). A review on the main challenges in automatic plant disease identification based on visible range images. *Biosystems Engineering*, *144*, 52–60. https://doi.org/10.1016/j.biosystemseng.2016.01.017
11) Zhang, S., Huang, W., & Zhang, C. (2015). A framework for citrus disease detection based on machine learning. *Precision Agriculture*, *16*(6), 642–659. https://doi.org/10.1007/s11119-015-9391-y
12) Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
13) He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. https://doi.org/10.1109/CVPR.2016.90
14) Lu, Y., Yi, S., & Wang, J. (2017). A survey of traditional methods for rice leaf disease detection. *Computational Agriculture*, 19(4), 45–56.
15) Zhang, Z., Li, W., & Zhao, Y. (2021). Development of field-specific datasets for rice disease detection in natural conditions. *Remote Sensing in Agriculture*, 12(5), 457.
16) Chollet, F. (2017). Deep Learning with Python. *Manning Publications*.
17) Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*.
18) Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proceedings of the 32nd International Conference on Machine Learning (ICML)*.
19) Buda, M., Maki, A., & Mazurowski, M. A. (2018). A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106, 249-259.
20) Brahimi, M., Boukhalfa, K., & Moussaoui, A. (2017). Deep learning for plant diseases: Detection and saliency map visualization. *Computers and Electronics in Agriculture*, *137*, 46–58.
21) Kamal, K., Singh, M., & Singh, K. (2019). Rice disease classification using convolutional neural networks. *International Journal of Agricultural Sciences*, *12*(2), 1–6.
22) Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. *MIT Press*.
23) Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, *6*(1), 1–48.
24) Szegedy, C., Ioffe, S., Vanhoucke, V., & Alemi, A. (2015). Inception-v4, Inception-ResNet and the impact of residual connections on learning. *Proceedings of the AAAI Conference on Artificial Intelligence*.
25) Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations* (ICLR