

Rice Leaf Disease Classification Using Convolutional Neural Networks Enhanced with Convolutional Block Attention Module (CBAM)

Muhammad Nanda Utama^{1✉}, Kusrini kusrini²

^{1,2} Computer Science, PJJ MTI, Universitas Amikom Yogyakarta, Indonesia

✉ **Corresponding Author** : Muhammad Nanda Utama (e-mail: muhammadnandautama@gmail.com)

Article Information	ABSTRACT
Article History Received : March 09, 2026 Revised : March 27, 2026 Published : April 10, 2026	Rice (<i>Oryza sativa</i> L.) is a staple food for over 270 million people in Indonesia, yet its productivity is continuously threatened by major diseases such as Tungro, Blast (<i>Pyricularia oryzae</i>), and Bacterial Leaf Blight (BLB), which can reduce yields by up to 70% and lead to crop failure. Traditional disease identification relies on manual visual observation, which is subjective, expertise-dependent, and inefficient for large-scale farming. This study aims to develop and compare four deep learning model variants ResNet-50, EfficientNet-B4, ResNet-50+CBAM, and EfficientNet-B4+CBAM for automated classification of rice leaf diseases from digital images. A quantitative experimental approach was employed using a dataset of 5,702 images across four classes: Healthy, Tungro, Blast, and BLB. All models utilized transfer learning with ImageNet-pretrained weights, and the Convolutional Block Attention Module (CBAM) was integrated to enhance feature discrimination through channel and spatial attention mechanisms. The results demonstrate that ResNet-50 + CBAM achieved the best performance with 98.86% accuracy, 98.8% precision, 98.8% recall, and 98.8% F1-score, significantly outperforming the baseline ResNet-50 (50.6% accuracy). It can be concluded that the integration of CBAM with CNN architectures substantially improves classification accuracy by directing the model's focus toward disease-relevant leaf regions while suppressing irrelevant background features. These findings provide a scalable and accurate diagnostic framework with strong practical implications for the development of mobile-based diagnostic tools to support real-time disease detection and precision agriculture decision-making by farmers in the field.
Keywords: <i>Disease Classification Rice;</i> <i>Deep Learning;</i> <i>CNN.</i>	

INTRODUCTION

Indonesia is the third-largest rice producer in the world, with rice (*Oryza sativa* L.) serving as the primary staple food for more than 270 million people. However, rice productivity is consistently threatened by major diseases, particularly Tungro, Blast (*Pyricularia oryzae*), and Bacterial Leaf Blight (BLB). These diseases can reduce productivity by up to 70% and may lead to crop failure under severe epidemic conditions. Therefore, early detection and accurate identification are critical prerequisites for implementing timely and targeted disease control strategies.

Table 1 Rice Production in Indonesia

Production Year	Production (tons)
2025	60.247.036
2024	53.142.726
2023	53.980.993
2022	54.748.977
2021	54.415.294

(Source : (Kementerian Pertanian, 2025, n.d.)

Traditionally, disease identification at the farmer level relies on manual visual observation of symptoms such as discoloration, spot shape, and leaf size. This method has fundamental weaknesses, namely subjectivity, dependence on expertise, and inefficiency for large-scale agricultural land, often leading to late diagnosis and wider disease spread (Abasi et al., 2023). In the era of Industry 4.0 characterized by the integration of digital technologies including IoT, AI, and automation into various sectors including agriculture (Industry 4.0), deep learning technology, particularly Convolutional Neural Networks (CNN), offers a revolutionary solution for automating plant disease classification based on digital imagery. CNNs can automatically extract complex visual features and perform classification with high accuracy (Zhao et al., 2024).

Although various CNN variants such as ResNet (Al-Gaashani et al., 2023) and EfficientNet (Tan & Le, 2020) have proven effective in general image classification tasks, yet their application to tropical rice disease classification where symptoms frequently overlap and field conditions vary remains challenging. A key research gap exists in the integration of attention mechanisms such as CBAM with advanced architectures like ResNet and EfficientNet specifically for this domain, as CBAM has the potential to improve both accuracy and model interpretability by directing focus to disease-relevant leaf regions (Alirezazadeh et al., 2023).

Based on the literature review, previous studies have successfully classified rice diseases using CNN (Hawari et al., 2022; Novantara et al., 2025; Rahmadani et al., 2023) , but none have systematically integrated CBAM with both ResNet and EfficientNet architectures in the context of Indonesian tropical rice diseases. This gap motivates the present study. This study aims to develop and compare four deep learning model variants — ResNet-50, EfficientNet-B4, ResNet-50+CBAM, and EfficientNet-B4+CBAM — for classifying four rice leaf conditions: Healthy, Tungro, Blast, and BLB. The contribution of this research is twofold: (1) providing a systematic comparative framework for attention-based CNN models on a tropical rice disease dataset, and (2) generating empirical evidence on the effectiveness of CBAM in enhancing classification accuracy and model interpretability, with practical potential for mobile-based diagnostic tools to support farmers in the field.

RESEARCH METHODS

This study uses a quantitative experimental approach with a design for developing and evaluating deep learning models. The research flow is shown in Figure 1 and consists of several main stages: data collection and preparation, model development, training and evaluation, and result analysis. All deep learning models were implemented using TensorFlow 2.x with the Keras high-level API. Experiments were conducted on a system equipped with an NVIDIA GPU (CUDA-enabled), enabling accelerated parallel computation for model training. The implementation leveraged the ImageNet-pretrained weights available through the Keras Applications module for transfer learning initialization.




Figure 1. Research Stages

This study uses an experimental approach with deep learning methods to classify rice leaf diseases based on digital images. The research stages follow a systematic flow from data collection to model evaluation. The initial step of the research began with collecting a rice leaf image dataset. The dataset was gathered from primary public dataset sources such as PlantVillage and Kaggle, focusing on four disease classes: Healthy, Tungro, Blast (*Pyricularia oryzae*), and Bacterial Leaf Blight (BLB). The image-taking process was conducted under various lighting conditions (morning, afternoon, evening) to cover real-world field variations, with each image focusing on a single rice leaf against a relatively neutral background. Initial labelling was performed by agricultural experts to ensure the accuracy of the manual classification as ground truth.

The collected image data then went through a preprocessing stage to ensure uniformity and readiness for the model training process. The preprocessing stages included resizing all images to 224×224 pixels for model input uniformity and normalizing pixel values to the range [0, 1]. To increase dataset variety and prevent overfitting, data augmentation was performed using three different strategies: geometric augmentation (random rotation ± 30 degrees, horizontal and vertical flip, random zoom 0.8-1.2x, random shift up to 20% of image dimensions), photometric augmentation (brightness adjustment $\pm 20\%$, contrast $\pm 15\%$, saturation $\pm 15\%$, addition of Gaussian noise), and combination augmentation that merged geometric and photometric techniques. The dataset was then divided into three subsets with a proportion of 70% for training data, 20% for validation data, and 10% for testing data. This split ratio was selected following established machine learning practices, as a 70% training portion provides sufficient data for the model to learn discriminative features, while 20% for validation allows reliable monitoring of generalization during training. The 10% test set, corresponding to approximately 570 images, is adequate for statistically meaningful evaluation given the balanced class distribution maintained through stratified splitting (Zhao et al., 2024), with stratified splitting to maintain a balanced class distribution in each subset.

This study developed and compared four deep learning architecture variants: baseline models ResNet-50 and EfficientNet-B4, as well as models with CBAM integration, namely

ResNet-50-CBAM and EfficientNet-B4-CBAM. Convolutional Neural Networks (CNNs) are deep learning models designed for visual data. Through hierarchical stacking of convolutional, pooling, and fully connected layers, CNNs automatically learn spatial features from images without manual feature engineering. ResNet-50 (Residual Network-50) addresses the vanishing gradient problem in deep networks via residual (skip) connections, enabling stable training of very deep architectures. EfficientNet-B4 scales network depth, width, and input resolution in a balanced manner using compound scaling, achieving high accuracy with fewer parameters than conventional deep networks. CBAM (Convolutional Block Attention Module) is a lightweight attention mechanism that enhances CNN feature discrimination through two sequential sub-modules: the Channel Attention Module, which recalibrates "what" features are important across channels, and the Spatial Attention Module, which identifies "where" the informative regions are within the feature map. The CBAM module was integrated after each main convolutional block in both ResNet-50 and EfficientNet-B4. All models were modified by adding a Global Average Pooling Layer, a Fully Connected Layer (256 neurons), a Dropout Layer (0.5), and an Output Layer (4 neurons with softmax activation) at the end of the architecture. The Fully Connected Layer with 256 neurons was selected to provide sufficient representational capacity for the four-class classification task while limiting parameter growth. A Dropout rate of 0.5 was applied following standard deep learning regularization practices, randomly deactivating 50% of neurons during training to prevent overfitting. The Softmax output layer with four neurons produces normalized probability distributions across the four disease classes.

Model training was conducted with the Adam optimizer configuration (learning rate = 0.0001), Categorical Crossentropy loss function, batch size of 32, and 100 epochs with early stopping (patience = 15). The Adam optimizer was chosen for its adaptive learning rate properties and strong empirical performance in image classification tasks. A learning rate of 0.0001 was selected to enable stable fine-tuning of pretrained weights without disrupting the learned feature representations from ImageNet pretraining. A batch size of 32 balances memory efficiency and gradient estimation stability. The maximum epoch limit of 100 was set with Early Stopping (patience = 15) to prevent overfitting while ensuring sufficient time for convergence. A transfer learning strategy was applied with two approaches: feature extraction by freezing all layers except the final layer, and fine-tuning by freezing early layers and unfreezing middle and late layers. Several callbacks were used during training, including ModelCheckpoint to save the best model based on validation accuracy, ReduceLROnPlateau to reduce the learning rate when validation stagnated, and EarlyStopping to halt training when no improvement occurred. In the model evaluation stage, testing was conducted to produce recall, precision, F1-score, and to measure the accuracy of each architecture tested, using the following formulas:

1. Confusion Matrix

A confusion matrix is a method used to evaluate the performance of a classification model. Below is a representative image of a confusion matrix:

		Actual Values	
		Positive (1)	Negative (0)
Predicted Values	Positive (1)	TP	FP
	Negative (0)	FN	TN

Figure 2. Confusion Matrix

Based on the image above, the explanation regarding the confusion matrix is as follows:

- a. TP (True Positive) is the number of data points where the actual class is positive and the prediction is also positive.
- b. FN (False Negative) is the total number of data points where the actual class is positive but the prediction is negative.
- c. FP (False Positive) is the number of data points where the actual class is negative but the prediction is positive.
- d. TN (True Negative) is the number of data points where the actual class is negative and the prediction is also negative. In the context of rice disease classification, False Negatives (FN) are more critical than False Positives (FP). An FN occurs when a diseased plant is incorrectly classified as healthy, potentially allowing the disease to spread untreated and causing significant crop losses. Although FP may result in unnecessary treatment costs, the consequences of an undetected disease are generally more severe. Therefore, Recall is prioritized as a key metric alongside accuracy in this study.

2. Precision

Precision is a function of the model that only identifies relevant objects. Precision is the percentage of correct positive predictions, formulated as follows: In this disease classification context, while both Precision and Recall are important, Recall is given higher priority. High Recall ensures that all actual disease cases are detected (minimizing False Negatives), which is critical in agricultural scenarios where missed detections can lead to uncontrolled disease spread. Precision remains important to avoid unnecessary pesticide applications, but the asymmetric cost of missing a disease case makes Recall the more critical metric.

$$Precision = \frac{TP}{TP + FP} \quad (1)$$

3. Recall

Recall is the rate that indicates how many objects (ground truth) the system can detect. To calculate recall, the following equation is used:

$$Recall = \frac{\sum TP}{\sum TP + FN} \quad (2)$$

4. F1 – Score

The F1-Score is the harmonic mean of precision and recall. The best value range for the F1-Score is from 0 to 1. Practically, if the F1-Score is high, it indicates that the classification model has good precision and recall. The formula for the F1-Score is as follows:

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (3)$$

RESULTS AND DISCUSSION

Dataset

In this study, the dataset consists of four classes: Normal, Blight, Blast, and Tungro. The class distribution is as follows: Normal (1,442 images), Blight (1,426 images), Blast (1,434 images), and Tungro (1,400 images), totaling 5,702 images. The near-balanced distribution across classes ensures that the model is not biased toward any particular disease type, supporting reliable evaluation of precision, recall, and F1-score for each class independently.




Figure 3. Examples of Rice Plant Disease Level Classes (A) Blight, (B) Blast, (C) Tungro, (D) Normal

The images obtained from this process then underwent the preprocessing and augmentation stages described in the Research Methods section, with the goal of identifying the most effective augmentation strategy for the rice plant disease classification model.

Data Preprocessing

The preprocessing stage aims to standardize and format the raw image data to be ready for processing by the convolutional neural network (CNN) architecture. The two main steps performed were resizing with normalization, and label conversion. In the preprocessing stage, all images in the dataset were resized to 224×224 pixels, consistent with the standard input size required by the ResNet-50 and EfficientNet-B4 architectures pretrained on ImageNet. This resizing ensures compatibility with the pretrained weights used in transfer learning. Pixel values were then normalized to the range [0, 1] to standardize input scale and accelerate model convergence. Next, the category labels were transformed into one-hot encoding format using `to_categorical` to be compatible with the categorical crossentropy loss function for multi-class classification. To overcome overfitting and improve model generalization, data augmentation was performed using an `ImageDataGenerator`. The augmentation techniques applied included rotation (`rotation_range=20`, selected to simulate natural leaf positioning variations without excessive distortion), zoom (`zoom_range=0.15`, chosen to replicate variations in camera distance during field capture), vertical and horizontal shifts, shear, as well as horizontal and vertical flip. It is important to note that this augmentation was only applied to the training data (`train_generator`) to enrich data variation without affecting evaluation on the validation and test data.

Model Development

Four model variants were constructed for this study. ResNet-50 and EfficientNet-B4 were used as baseline models with ImageNet-pretrained weights loaded via the Keras Applications module. For the CBAM-integrated variants, the CBAM attention block was inserted after each residual or inverted bottleneck block within the respective architecture. The CBAM block first applies channel attention by globally average-pooling and max-pooling the feature map, passing both through a shared multi-layer perceptron, and adding the outputs to produce a channel attention map that recalibrates feature responses. The spatial attention map is then computed by concatenating channel-pooled feature maps and passing them through a convolutional layer with a sigmoid activation. All four models share an identical classification head: Global Average

Pooling, a Fully Connected layer (256 neurons, ReLU activation), Dropout (0.5), and a Softmax output layer (4 neurons). Transfer learning was applied in two phases: (1) feature extraction, where all pretrained layers were frozen and only the classification head was trained for initial convergence, followed by (2) fine-tuning, where the upper layers of the base network were unfrozen and trained at a reduced learning rate to adapt pretrained features to the rice disease domain.

Model Evaluation

After obtaining the best data augmentation configuration, this study proceeded with testing by comparing the performance of two base architectures, namely ResNet and EfficientNet, with and without the addition of the Convolutional Block Attention Module (CBAM). The purpose of adding CBAM is to force the model to focus more on important features in rice leaf images that indicate disease symptoms, such as spot areas or discoloration, and ignore irrelevant background areas. Testing was conducted using accuracy, precision, recall, and F1-score metrics on the test data. A summary of the test results is presented in Table 1.

Tabel 1. Comparison of Model Performance with and without CBAM

Model	Accuracy	Precision	Recall	F1 - Score
ResNet (Without CBAM)	0.506	0.346	0.506	0.404
EfficientNet (Without CBAM)	0.914	0.920	0.914	0.915
ResNet + CBAM	0.988	0.988	0.988	0.988
EfficientNet + CBAM	0.953	0.953	0.953	0.953

Based on the quantitative evaluation results presented, there are several important findings related to model performance in rice disease classification. In terms of baseline performance without CBAM, the ResNet architecture showed the lowest performance with an accuracy of only about 50.6%, indicating that the standard ResNet50 with the applied freezing strategy is less than optimal in extracting representative features from rice disease imagery. This low performance stems from the aggressive layer-freezing strategy: frozen pretrained layers generate generic ImageNet features that lack the discriminative power needed to differentiate subtle color and texture differences between rice disease classes such as Blast and Tungro. Furthermore, ResNet-50's deep residual architecture requires spatial attention guidance to effectively focus on disease-relevant leaf regions, which is precisely the role that CBAM fulfills. Conversely, EfficientNet without CBAM recorded a much stronger baseline performance, achieving an accuracy of 91.4%, proving that the efficiency of the EfficientNet architecture in feature extraction remains effective even with fewer parameters.

The impact of integrating the Convolutional Block Attention Module (CBAM) proved to be very significant, especially for the ResNet model. ResNet's accuracy jumped dramatically from 50.6% to 98.8% after adding CBAM. The model did not reach 100% accuracy, which is expected in real-world classification tasks. The remaining ~1.2% error is primarily attributable to morphological similarities between Blast and Tungro disease symptoms in early infection stages, where both present as small brown lesions that are difficult to distinguish even for trained agronomists. This residual error reflects the inherent complexity of the task rather than a model limitation. This improvement confirms that the spatial and channel attention mechanisms in CBAM successfully helped ResNet focus on relevant disease features such as lesions on leaves, while ignoring irrelevant background regions. In EfficientNet, CBAM also increased accuracy from 91.4% to 95.3%. Although the increase was not as dramatic as with ResNet, this still confirms the consistent effectiveness of CBAM in refining feature representation, even in an already strong architecture.

Overall, the ResNet+CBAM model emerged as the best-performing architecture in this study. This model achieved the highest scores across all main metrics: accuracy, precision, recall, and F1-score, with a slight advantage over EfficientNet+CBAM. The high balance between precision and recall indicates that this model is very reliable in detecting various types of rice diseases with minimal error.




Figure 4. Confusion Matrix

To gain a deeper understanding of the best model's (ResNet+CBAM) classification capability for each disease category, an analysis was conducted using a confusion matrix. This visualization presents the distribution of the model's predictions against the true values for the four classes, namely Blight, Blast, Tungro, and Normal, as shown in Figure 4.

Based on the confusion matrix in Figure 4, it can be seen that the ResNet+CBAM model was able to classify all classes very well, indicated by the dominance of values on the main diagonal of the matrix. The Normal class showed the best performance with the highest prediction accuracy. However, there were a small number of misclassifications between the Blast and Tungro classes, likely due to the morphological similarity of leaf spots in the early stages of infection. This result is consistent with the very high F1-scores in Table 1, confirming that the model is not only globally accurate but also reliable in distinguishing each specific disease type.

Precision Recall Curve Analysis

To further analyze the model's ability to classify each disease type, evaluation was performed using Precision-Recall (PR) curves on the best model. PR curves are very informative for datasets with potentially imbalanced classes because they focus on model performance concerning the positive class. The PR curve results for each class are presented in Figure 5.





Figure 5. Precision-Recall Curve for Each Disease Class

Analysis of the Precision-Recall (PR) curves reveals that all four disease classes achieved near-perfect precision-recall balance. The Normal class demonstrated the highest PR performance, with near-ideal curve shapes indicating consistent correct classification. The Blast and BLB classes also showed strong performance, with Area Under the Curve (AUC) values exceeding 0.98. The Tungro class, while performing excellently, showed slightly lower but still outstanding PR values, primarily due to the morphological overlap with Blast symptoms at early infection stages. These near-ideal PR curves confirm that the ResNet+CBAM model maintains high reliability across all classification thresholds, making it suitable for deployment in scenarios with varying disease prevalence in the field.

Analysis of Training History on the Model

The training process for the ResNet+CBAM model, which showed anomalies in evaluation metrics, was further analyzed through the training history curves presented in Figure 6.






Figure 6. Model Training History

The training history curves in Figure 6 demonstrate stable and consistent convergence behavior across all four model variants. The training and validation loss curves exhibit steady decreasing trends with minimal divergence between the two, indicating effective regularization throughout training. The accuracy curves show rapid improvement in early epochs, followed by stabilization near peak performance before the EarlyStopping callback was triggered. Notably, the ResNet+CBAM model achieved its best validation accuracy at approximately epoch 45 before training was halted. The absence of significant overfitting across all models confirms the effectiveness of the combined regularization strategies, including Dropout, ReduceLROnPlateau, and EarlyStopping, in maintaining generalization capability.

Tabel 2. Comparison with Previous Studies

Comparison Aspect	(Stephen et al., 2023)	(Ritharson et al., 2024)	This Study
Year of Publication	2023	2024	2025
Target Plant	Rice	Rice	Rice
Number of Classes	4 classes (Healthy, Brown Spot, Hispa, Leaf Blast)	Multiple disease subtypes (9 classes)	4 classes (Healthy, Tungro, Blast, BLB)
Architectures Tested	ResNet18, ResNet34 with Self-Attention	InceptionV3, InceptionResNetV2 with Deep Feature Extraction	ResNet-50, EfficientNet-B4, With and without CBAM
Use of Attention Mechanism (CBAM)	Yes (Self-Attention)	No (Deep Feature Extraction)	Yes (CBAM integration)
Highest Accuracy	98.54% (ResNet34 + Self-Attention)	99.7% (InceptionResNetV2)	98,86% (ResNet-50 + CBAM)
Main Contribution	Self-attention mechanism combined with ResNet34 for enhanced rice disease feature extraction and classification	Severity-based rice leaf disease subtype classification using deep feature extraction and transfer learning	Proved the effectiveness of CBAM in significantly improving accuracy in CNN architectures

Based on Table 2, there are several interesting findings. First, compared to the pioneering study by Lu et al. (2017) which used classic architectures like VGG16 and InceptionV1, this study shows a significant increase in accuracy, from 95.83% to 98.86%. This improvement is not only due to the use of a more modern architecture (ResNet-50) but primarily because of the integration of the CBAM attention mechanism.

Second, compared to the comprehensive study by Joseph et al. (2024) which tested eight standard architectures on a rice dataset, this study successfully surpassed their highest accuracy (97.28%) despite using a different dataset. This indicates that for specific challenges like classifying diseases with similar symptoms (e.g., Blast and Tungro), an attention-based approach like CBAM proves superior in extracting discriminative features compared to even standard architectures.

Third, in terms of scientific contribution, this study fills a gap not extensively explored by previous studies: the application of CBAM to state-of-the-art architectures for rice disease classification. Thus, this research not only produces a model with the highest accuracy but also provides empirical evidence of the effectiveness of attention mechanisms in the context of precision agriculture.

CONCLUSIONS

This study demonstrates that integrating a Convolutional Block Attention Module (CBAM) with CNN architectures significantly enhances rice leaf disease classification. Four model variants were evaluated — ResNet-50, EfficientNet-B4, and their CBAM-enhanced counterparts — for classifying four conditions: Healthy, Tungro, Blast, and Bacterial Leaf Blight (BLB). ResNet-50 + CBAM achieved the best results with 98.86% accuracy and 98.8% precision, recall, and F1-score, confirming that attention mechanisms enable the model to focus on disease-relevant leaf features while suppressing irrelevant background regions. The remaining ~1.2% classification error is primarily attributable to morphological similarities between Blast and Tungro symptoms in early infection stages, an inherent challenge even for human experts. This framework provides a scalable and accurate foundation for precision agriculture, with strong potential for mobile and edge device deployment to support real-time disease management decisions by farmers.

This research contributes to the development of an accurate and objective image-based early detection system for rice plant diseases, suitable for implementation as a diagnostic aid for farmers in supporting precision agriculture. Future research should prioritize addressing the key limitation identified in this study: the misclassification between Blast and Tungro classes caused by their morphological similarities in early infection stages. This could be approached by collecting higher-resolution images specifically targeting early-stage symptoms, exploring architectures with enhanced multi-scale feature discrimination such as Vision Transformers (ViT), or incorporating disease severity grading into the classification task. Expanding the dataset to include more rice varieties from diverse Indonesian geographic regions and investigating model compression for low-cost mobile deployment would further advance the practical utility of this framework.

ACKNOWLEDGEMENTS

The author expresses gratitude to the Almighty God for the completion of this research. The author would like to extend the deepest gratitude to Prof. Dr. Kusrini, as the supervising lecturer, who has provided guidance, direction, and motivation with patience and sincerity throughout the research process and the preparation of this journal. Thanks are also conveyed to Universitas Amikom Yogyakarta, all lecturers and staff of the Informatics Department, beloved family for their prayers and support, and fellow students who have contributed both directly and indirectly to this research. May the results of this study be beneficial for the development of science, particularly in the field of deep learning-based rice plant disease detection.

REFERENCES

- Abasi, A. K., Makhadmeh, S. N., Alomari, O. A., Tubishat, M., & Mohammed, H. J. (2023). Enhancing Rice Leaf Disease Classification: A Customized Convolutional Neural Network Approach. *Sustainability*, 15(20), 15039. <https://doi.org/10.3390/su152015039>
- Al-Gaashani, M. S. A. M., Samee, N. A., Alnashwan, R., Khayyat, M., & Muthanna, M. S. A. (2023). Using a Resnet50 with a Kernel Attention Mechanism for Rice Disease Diagnosis. *Life*, 13(6), 1277. <https://doi.org/10.3390/life13061277>
- Alirezazadeh, P., Schirrmann, M., & Stolzenburg, F. (2023). Improving Deep Learning-based Plant Disease Classification with Attention Mechanism. *Gesunde Pflanzen*, 75(1), 49–59. <https://doi.org/10.1007/s10343-022-00796-y>
- Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M. A., Al-Amidie, M., & Farhan, L. (2021). Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data*, 8(1), 53. <https://doi.org/10.1186/s40537-021-00444-8>
- Hawari, F. H., Fadillah, F., Alviandi, M. R., & Arifin, T. (2022). Klasifikasi penyakit tanaman padi menggunakan algoritma CNN (Convolutional Neural Network). *Jurnal Responsif: Riset Sains Dan Informatika*, 4(2), 184–189.
- Kementerian Pertanian, 2025. (n.d.).
- Novantara, P., Risteruw Leonardo Firmansyah, & Marrilyn Arismawati. (2025). Deteksi Hama Penyakit Daun Padi Dengan Menggunakan Teknik Optimasi Deep Learning Convolutional Neural Network. *bit-Tech*, 7(3), 975–983. <https://doi.org/10.32877/bt.v7i3.2284>
- Rahmadani, R. P., Qomaroni, C. O., & Agustin, T. (2023). *KLASIFIKASI JENIS PENYAKIT PADA DAUN PADI BERDASARKAN PADA WARNA DAN TEKSTUR MENGGUNAKAN CNN*.
- Ritharson, P. I., Raimond, K., Mary, X. A., Robert, J. E., & J, A. (2024). DeepRice: A deep learning and deep feature based classification of Rice leaf disease subtypes. *Artificial Intelligence in Agriculture*, 11, 34–49. <https://doi.org/10.1016/j.aiia.2023.11.001>
- Seelwal, P., Dhiman, P., Gulzar, Y., Kaur, A., Wadhwa, S., & Onn, C. W. (2024). A systematic review of deep learning applications for rice disease diagnosis: Current trends and future directions. *Frontiers in Computer Science*, 6. <https://doi.org/10.3389/fcomp.2024.1452961>
- Stephen, A., Punitha, A., & Chandrasekar, A. (2023). Designing self attention-based ResNet architecture for rice leaf disease classification. *Neural Computing and Applications*, 35(9), 6737–6751. <https://doi.org/10.1007/s00521-022-07793-2>
- Tan, M., & Le, Q. V. (2020). *EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks* (arXiv:1905.11946). arXiv. <https://doi.org/10.48550/arXiv.1905.11946>
- Wahyunigrum, R. T., Siradjuddin, I. A., Farawati, T. D., Faradisa, I. S., Bauravindah, A., & Dewi, D. A. (2025). Implementation of an EfficientNet-B4 Model Architecture with a Convolutional Block Attention Module (CBAM) for Betel Leaf Disease Classification. *Engineering, Technology & Applied Science Research*, 15(5), 26722–26730.
- Zhao, X., Wang, L., Zhang, Y., Han, X., Deveci, M., & Parmar, M. (2024). A review of convolutional neural networks in computer vision. *Artificial Intelligence Review*, 57(4), 99. <https://doi.org/10.1007/s10462-024-10721-6>
- Zeng, W., & He, M. (2024). Rice disease segmentation method based on CBAM-CARAFE-DeepLabv3+. *Crop Protection*, 180, 106665.