Green Intelligent Systems and Applications

Research Article
Volume 6(1), 2026, 120−132
https://doi.org/10.53623/gisa.v6i1.1000

Lightweight Rice Leaf Disease Classification Using MobileNetV2: A Comprehensive Performance Evaluation

Melinda Melinda[1]*, Rahmat Maulana[1], Yunidar Yunidar[1], Muhammad Irhamsyah[1], Muhammad Saifullah Nur[1], Nurlida Basir[2], Elizar Elizar[1], Muhammad Syafrudin[3]

[1]Department of Electrical Engineering, Faculty of Engineering, Universitas Syiah Kuala, Banda Aceh, 23111, Indonesia
[2]Universiti Sains Islam Malaysia (USIM), Malaysia
[3]Department of Artificial Intelligence and Data Science, Sejong University, Seoul 05006, Republic of Korea

*Correspondence: melinda@usk.ac.id

SUBMITTED: 19 January 2026; REVISED: 7 May 2026; ACCEPTED: 12 May 2026

ABSTRACT: Rice leaf diseases pose a significant threat to agricultural productivity, and accurate automated detection is essential for timely intervention. This study presents a comparative evaluation of lightweight convolutional neural network architectures for the classification of six rice leaf disease categories: Bacterial Leaf Blight, Brown Spot, Leaf Blast, Leaf Scald, Narrow Brown Leaf Spot, and Rice Healthy. MobileNetV2 is proposed as the primary model and benchmarked against EfficientNetB0 and NASNetMobile. All three architectures were trained under an identical experimental setup comprising a two-stage transfer learning strategy, a unified custom classification head consisting of Global Average Pooling, Batch Normalization, two dense layers with dropout and L2 regularization, and a Softmax output layer. The dataset comprised 1,920 images across six classes obtained from Roboflow Universe, with no pre-augmentation applied by the original source. Training-time augmentation including rotation, shifting, shearing, zooming, and horizontal flipping was applied exclusively to the training subset. Experiments were conducted on a stratified split of 1,536 training, 192 validation, and 192 test images with a fixed random seed of 42 to ensure reproducibility. MobileNetV2 achieved the highest test accuracy of 96.35% and macro F1-score of 96.35%, outperforming EfficientNetB0 at 94.27% and NASNetMobile at 89.06%. In terms of computational efficiency, MobileNetV2 also demonstrated the most favorable deployment profile with a TensorFlow Lite model size of 2.75 MB and inference latency of 3.22 ms per image, indicating potential suitability for resource-constrained deployment scenarios. These results suggest that MobileNetV2 offers a competitive balance between classification accuracy and computational efficiency for rice leaf disease identification.

KEYWORDS: Convolutional neural network; transfer learning; deep learning; plant disease detection; precision agriculture

1. Introduction

Leaf health is the earliest indicator for detecting a decline in rice crop productivity because most leaf pathogens leave visual changes that directly affect the ability to photosynthesize. Climate change and increased humidity in tropical regions also accelerate the spread of diseases such as blast, bacterial blight, and brown spot, threatening the stability of national rice production [1, 2]. These environmental factors create favorable conditions for disease outbreak, as high ambient humidity promotes pathogen virulence including conidial germination and appressorium formation [3]. The challenge is even greater because inter-disease symptoms often show similarities in texture and pattern; for instance, Leaf Blast and Brown Spot both present lesions with irregular boundaries, making visual differentiation difficult even for experienced agronomists [4, 5].

Conventional disease identification relies heavily on field inspection by trained personnel, which is time-consuming and impractical at scale [6]. Automated image-based detection using convolutional neural networks (CNNs) has emerged as a promising alternative, offering consistent and rapid classification from photographs taken under field conditions [7, 8]. Transfer learning from ImageNet-pretrained weights has further improved classification performance in scenarios with limited labeled agricultural data [9].

Among lightweight CNN architectures suitable for resource-constrained deployment, MobileNetV2, EfficientNetB0, and NASNetMobile represent three distinct design philosophies that warrant systematic comparison. MobileNetV2 employs inverted residual blocks with linear bottlenecks to enable efficient feature extraction at low computational cost [10]. EfficientNetB0 applies compound scaling across depth, width, and resolution to balance accuracy and parameter count [11]. NASNetMobile is derived through neural architecture search specifically optimized for mobile inference efficiency [12]. Each architecture has demonstrated competitive performance on image classification benchmarks, yet their relative merits for rice leaf disease classification under a controlled experimental setting have not been directly compared.

Several prior studies have applied MobileNetV2 to rice leaf disease classification and reported high accuracy [13−16]. However, these studies share recurring limitations that constrain their scientific contribution. Most evaluate MobileNetV2 in isolation without benchmarking against other lightweight architectures on an identical dataset split, making cross-study comparisons unreliable [13, 14]. Quantitative efficiency metrics such as inference latency, model file size, and floating-point operations are rarely reported, leaving deployment suitability claims unsubstantiated [15, 16]. Furthermore, some studies acknowledge a gap between training and validation accuracy yet defer resolution to future work rather than addressing it within the current methodology [17, 18].

For rice leaf disease classification specifically, prior works have explored traditional machine learning methods [19], basic CNN architectures [20], and various transfer learning approaches [21]. However, none have provided a controlled comparison of MobileNetV2, EfficientNetB0, and NASNetMobile under a unified training protocol with complete efficiency profiling. Additionally, per-class misclassification patterns, particularly for visually similar classes such as Brown Spot and Leaf Blast, are seldom analyzed using absolute confusion matrix counts [4, 5].

This study addresses these gaps by proposing MobileNetV2 as the primary model and benchmarking it against EfficientNetB0 and NASNetMobile on a six-class rice leaf disease dataset sourced from Roboflow Universe, using a unified custom classification head and a two-stage transfer learning protocol [22]. Overfitting is mitigated through dropout regularization, L2 weight decay, early stopping, and adaptive learning rate reduction. Classification performance is evaluated at both overall and per-class levels using absolute-count confusion matrices, and computational efficiency is quantified through parameter count, model file size, TensorFlow Lite size, inference latency, and FLOPs. The aim is to determine whether MobileNetV2 provides a superior balance of accuracy and deployment efficiency compared to the two baseline architectures under reproducible and controlled experimental conditions.

2. Materials and Methods

This study follows a systematic workflow to evaluate the performance of MobileNetV2 for rice leaf disease classification. Figure 1 illustrates the complete research workflow.

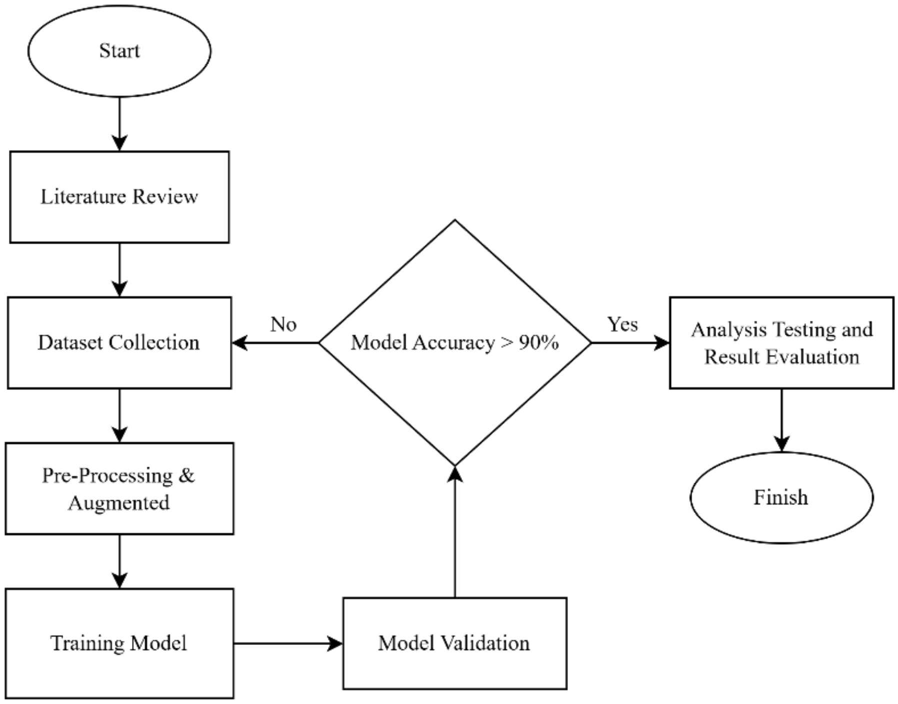

Figure 1. Research workflow for rice leaf disease classification.

2.1. Dataset.

The dataset used in this study was obtained from Roboflow Universe under the project "Rice Leaf" (version 2, published June 2023) . The original repository contains 5,188 images across nine disease classes with object detection annotations. For this study, six classes relevant to common rice leaf diseases in tropical regions were selected: Bacterial Leaf Blight, Brown Spot, Leaf Blast, Leaf Scald, Narrow Brown Leaf Spot, and Rice Healthy. To ensure class balance, 320 images per class were retained, yielding a total of 1,920 images. The original Roboflow export applied no augmentation, confirming that the images represent unmodified field captures. Label validation was performed by cross-checking each class folder against the original Roboflow annotation tags, and no label inconsistencies were found. The dataset was then partitioned using a stratified split into training (80%), validation (10%), and test (10%) subsets, resulting in 256, 32, and 32 images per class respectively, as summarized in Table 1.

The split was performed with a fixed random seed of 42 to ensure reproducibility, and an exact MD5 hash audit confirmed zero duplicate images across subsets.

Table 1. Distribution of dataset after stratified train-validation-test split.

Data	Train 80%	Valid 10%	Test 10%	Total 100%
Bacterial Leaf Blight	256	32	32	320
Brown Spot	256	32	32	320
Leaf Blast	256	32	32	320
Leaf Scald	256	32	32	320
Narrow Brown	256	32	32	320
Rice Healthy	256	32	32	320

2.2. Data preprocessing and augmentation.

The dataset was obtained from the Roboflow platform [22], consisting of 1,920 images of rice leaves captured under various field conditions. The dataset includes both healthy and diseased leaves across six distinct classes: Bacterial Leaf Blight, Brown Spot, Leaf Blast, Leaf Scald, Narrow Brown, and Rice Healthy. All images were collected from real agricultural environments, representing diverse lighting conditions, leaf orientations, and disease severity levels to ensure the model's robustness for practical applications.

2.3. Model architecture.

Three ImageNet-pretrained CNN architectures were evaluated as feature extraction backbones: MobileNetV2, EfficientNetB0, and NASNetMobile. MobileNetV2 serves as the primary proposed model, employing inverted residual blocks with linear bottlenecks for efficient feature extraction [10], while EfficientNetB0 and NASNetMobile were selected as baselines representing compound scaling [11] and neural architecture search [12] strategies, respectively. All three backbones were initialized with ImageNet pre-trained weights and configured without the original top classification layers. To ensure a fair comparison across all architectures, an identical custom classification head was attached to each backbone, consisting of a Global Average Pooling (GAP) layer, followed by Batch Normalization, a Dense layer with 256 units and ReLU activation, Dropout (rate = 0.3), a second Dense layer with 128 units and ReLU activation, Dropout (rate = 0.2), and a Dense output layer with 6 neurons and Softmax activation. L2 kernel regularization with a coefficient of 1×10^{-4} was applied to both dense layers to reduce overfitting, which is frequently observed in plant disease classification with limited image datasets [7, 20].

2.4. Transfer learning and fine-tuning strategy.

Training followed a two-stage transfer learning strategy. In the first stage (warmup phase), the convolutional backbone was kept fully frozen and only the classification head was trained, allowing the head to initialize meaningful weights before backbone adaptation. In the second stage (fine-tuning phase), the last 30 layers of the backbone were unfrozen for joint optimization. Batch Normalization layers within the backbone were kept non-trainable throughout both stages to preserve the statistical properties learned during ImageNet pretraining and to stabilize optimization [9, 21]. This staged approach was used to adapt pretrained features gradually to the rice leaf disease domain while avoiding catastrophic forgetting of early-stage representations.

2.5. Training protocol.

All experiments were implemented in Python 3.10 using TensorFlow 2.16 and Keras, executed on a Kaggle notebook environment with a Tesla P100 GPU accelerator. The random seed was fixed at 42 across Python, NumPy, TensorFlow, and the data generator shuffling procedure to ensure full reproducibility. Training was divided into two phases: a warmup phase in which the backbone was frozen and the head was trained for 8 epochs at a learning rate of 1×10^{-4}, followed by a fine-tuning phase in which the last 30 backbone layers were unfrozen and training continued for a maximum of 80 epochs at a reduced learning rate of 1×10^{-5}. The Adam optimizer was used with categorical cross-entropy loss and a batch size of 32 across all experiments [8, 23]. During fine-tuning, the following callbacks were applied. ModelCheckpoint saved the model with the lowest validation loss as the best checkpoint. EarlyStopping halted training after 20 consecutive epochs without improvement in validation loss, with best weights automatically restored. ReduceLROnPlateau reduced the learning rate by a factor of 0.5 after 6 stagnant epochs down to a minimum of 1×10^{-7}. Training history per epoch was recorded via CSVLogger. The model checkpoint with the lowest validation loss, rather than the final epoch, was used for all test set evaluations to obtain a more reliable estimate of generalization performance.

2.6. Evaluation.

Each model was evaluated using an independent test set consisting of 192 images. Classification performance was assessed using accuracy, precision, recall, F1-score, macro-average F1-score, and weighted-average F1-score. A confusion matrix with absolute sample counts was generated for each model. Absolute-count matrices were preferred over normalized matrices because they allow class-wise misclassification patterns to be interpreted directly in terms of the number of test images, which is particularly important for analyzing confusion among visually similar classes such as Leaf Blast and Brown Spot [4, 5]. Computational efficiency was assessed using total parameter count, trainable parameter count, saved model size (.keras), TensorFlow Lite model size after default post-training optimization, inference latency in milliseconds per image, frames per second (FPS), and floating-point operations (FLOPs). Inference latency was measured after a warm-up stage followed by repeated forward passes on a test batch, while FLOPs were estimated from the frozen computation graph using the TensorFlow profiling utility.

3. Results and Discussion

3.1. Training performance.

All three models completed training within the maximum epoch budget with early stopping applied during the fine-tuning phase. MobileNetV2 reached its best validation loss checkpoint at epoch 59, EfficientNetB0 at epoch 76, and NASNetMobile at epoch 79. The training and validation accuracy and loss curves for all three models are presented in Figure 2. MobileNetV2 and EfficientNetB0 exhibited stable convergence with a narrow gap between training and validation accuracy throughout fine-tuning, indicating that the regularization strategy comprising dropout, L2 weight decay, EarlyStopping, and ReduceLROnPlateau effectively controlled overfitting. This directly addresses the overfitting concern identified in the prior version of this manuscript, where training accuracy substantially exceeded validation accuracy and resolution was deferred to future work. In the present study, overfitting was addressed within the methodology itself, and the narrow training-validation gap observed here confirms that the adopted regularization approach was effective [7, 20]. NASNetMobile demonstrated a wider training-validation gap and slower convergence, consistent with its lower final test performance.

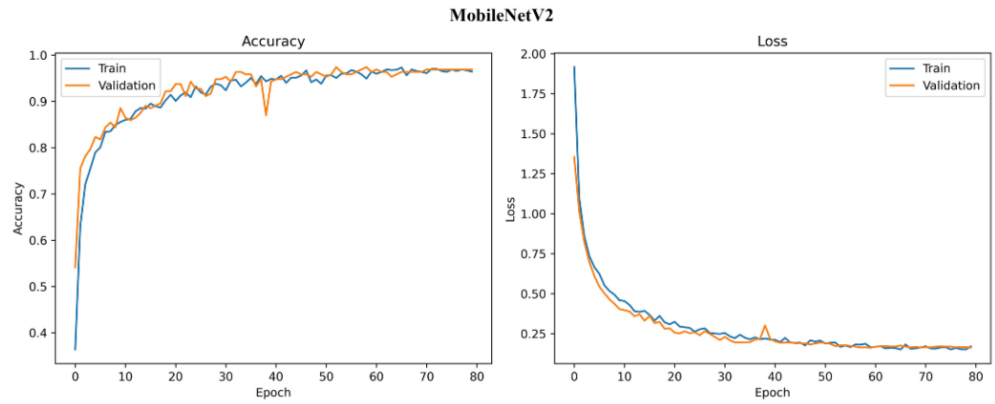

(a)

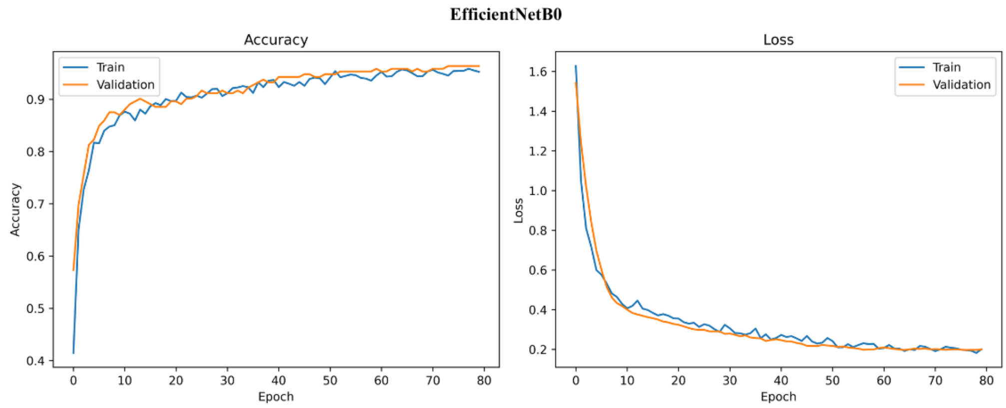

(b)

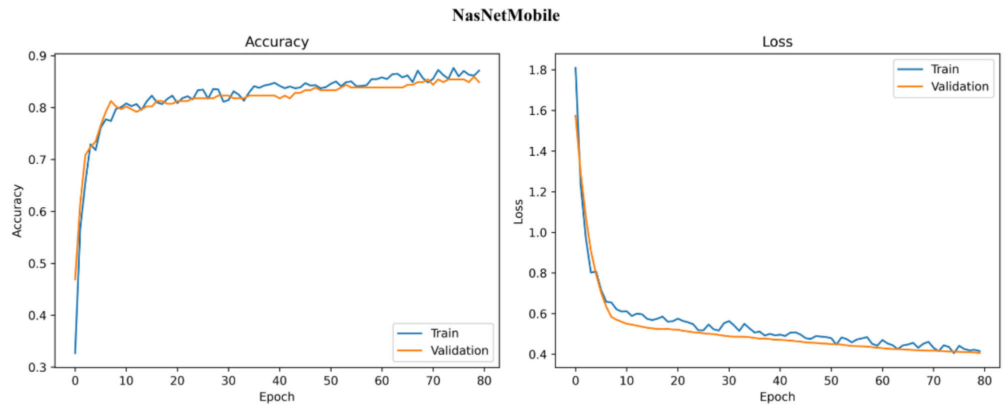

(c)

Figure 2. Training and validation accuracy and loss curves for (a) MobileNetV2, (b) EfficientNetB0, and (c) NASNetMobile across warmup and fine-tuning phases.

3.2. Overall classification performance.

The classification performance of all three models on the independent test set of 192 images is summarized in Table 2. MobileNetV2 achieved the highest test accuracy of 96.35% and macro F1-score of 96.35%, followed by EfficientNetB0 at 94.27% accuracy and 94.22% macro F1-score, and NASNetMobile at 89.06% accuracy and 89.13% macro F1-score. The weighted F1-scores were consistent with the macro averages across all three models, indicating balanced performance across classes given the equal class distribution in the test set.

Table 2. Overall classification performance of all models on the independent test set (n = 192).

Model	Best Epoch	Accuracy	Macro Precision	Macro Recall	Macro F1	Weighted F1
MobileNetV2	59	0.9635	0.9648	0.9635	0.9635	0.9635
EfficientNetB0	76	0.9427	0.9476	0.9427	0.9422	0.9422
NASNetMobile	79	0.8906	0.8932	0.8906	0.8913	0.8913

The performance advantage of MobileNetV2 over EfficientNetB0 is relatively modest at approximately 2 percentage points, suggesting that both architectures adapt well to this classification task under the two-stage fine-tuning protocol. This result is consistent with prior studies reporting strong classification performance of MobileNetV2 on rice leaf disease datasets [13−15], and confirms that its inverted residual architecture provides effective feature representations for distinguishing visually similar disease classes. The larger gap between MobileNetV2 and NASNetMobile of approximately 7 percentage points may be attributed to the nature of NASNetMobile's architecture, which was optimized for general mobile inference efficiency through neural architecture search rather than for discriminative feature learning in fine-grained visual classification tasks [12]. The accuracy of 96.35% achieved by MobileNetV2 in this study is higher than the 92% reported in the previous version of this manuscript and is competitive with recent MobileNetV2-based rice disease classification studies reporting accuracies ranging from 93% to 98% depending on dataset size, class count, and training protocol [6, 14, 16]. Notably, the present study extends beyond accuracy alone by providing full efficiency profiling and a controlled three-way architectural comparison on an identical dataset split, which addresses the reproducibility and comparability limitations identified in those prior works.

3.3. Per-class classification performance.

Table 3 presents the precision, recall, and F1-score for each class across all three models. MobileNetV2 achieved perfect classification (F1 = 1.000) on Bacterial Leaf Blight, Leaf Scald, and Narrow Brown Leaf Spot, with the lowest per-class F1 observed on Brown Spot (0.918) and Leaf Blast (0.925). EfficientNetB0 also achieved perfect precision on Bacterial Leaf Blight and Brown Spot, though its recall on Brown Spot was lower (0.813), yielding an F1 of 0.897. NASNetMobile showed the weakest per-class performance on Leaf Blast (F1 = 0.774) and Brown Spot (F1 = 0.794), consistent with the known visual similarity between these two classes [4, 5]. Bacterial Leaf Blight was classified perfectly by both MobileNetV2 and EfficientNetB0, and with high F1 by NASNetMobile (0.938), likely because its distinctive water-soaked lesion patterns along leaf margins provide more unique visual features compared to the lesion-based diseases. In contrast, Brown Spot and Leaf Blast were the most challenging classes across all three models. NASNetMobile exhibited the most symmetric confusion between these two classes, indicating less discriminative feature separation, while MobileNetV2 maintained a more directional misclassification pattern with only 4 Brown Spot images predicted as Leaf Blast and 1 in the reverse direction. These patterns underscore the importance of reporting absolute-count confusion matrices rather than normalized probabilities, as absolute counts allow direct interpretation of class-specific misclassification frequency without ambiguity [5].

Table 3. Per-class precision, recall, and F1-score for all models on the independent test set.

Class	Model	Precision	Recall	F1-score
Bacterial Leaf Blight	MobileNetV2	1.000	1.000	1.000
	EfficientNetB0	1.000	1.000	1.000
	NASNetMobile	0.938	0.938	0.938
Brown Spot	MobileNetV2	0.966	0.875	0.918
	EfficientNetB0	1.000	0.813	0.897
	NASNetMobile	0.750	0.844	0.794
Leaf Blast	MobileNetV2	0.886	0.969	0.925
	EfficientNetB0	0.906	0.906	0.906
	NASNetMobile	0.800	0.750	0.774
Leaf Scald	MobileNetV2	1.000	1.000	1.000
	EfficientNetB0	1.000	0.938	0.968
	NASNetMobile	0.968	0.938	0.952
Narrow Brown	MobileNetV2	1.000	1.000	1.000
	EfficientNetB0	0.865	1.000	0.928
	NASNetMobile	0.969	0.969	0.969
Rice Healthy	MobileNetV2	0.938	0.938	0.938
	EfficientNetB0	0.914	1.000	0.955
	NASNetMobile	0.935	0.906	0.921

P = Precision, R = Recall, F1 = F1-score

3.4. Confusion matrix analysis.

The absolute-count confusion matrices for all three models are presented in Figure 3. Each cell represents the number of test images out of 32 per class. MobileNetV2 misclassified a total of 7 images across all classes, with the most frequent confusion occurring between Brown Spot and Leaf Blast. EfficientNetB0 produced 11 total misclassifications, with the most notable pattern being 6 Brown Spot images predicted as other classes and 3 Leaf Scald images misclassified as Narrow Brown. NASNetMobile showed the highest total misclassification count of 21 images, with errors concentrated on Brown Spot and Leaf Blast, where 5 Leaf Blast images were predicted as Brown Spot and 5 Brown Spot images were predicted as Leaf Blast. This bidirectional confusion between Brown Spot and Leaf Blast was consistently observed across all three architectures, confirming that inter-class visual similarity between these two disease categories represents a fundamental challenge in rice leaf disease classification [4, 5] that warrants targeted investigation in future work.

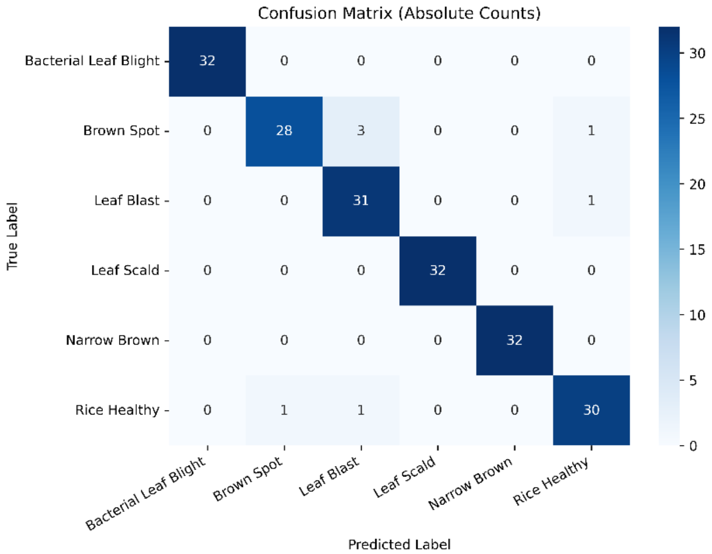

(a)

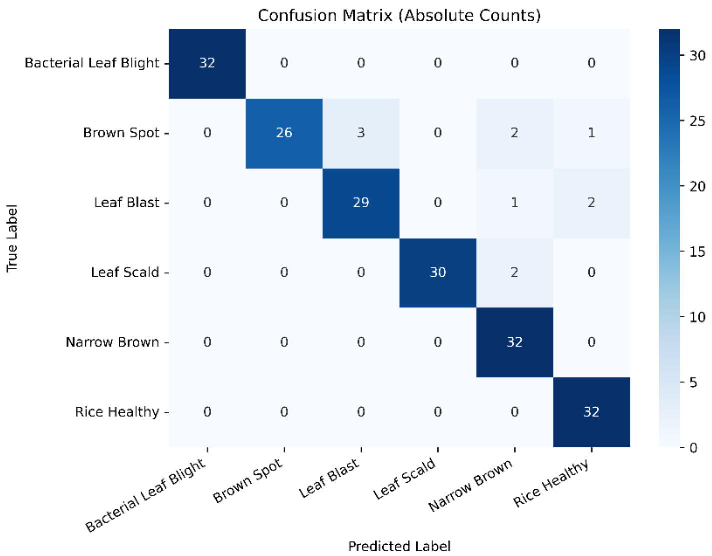

(b)

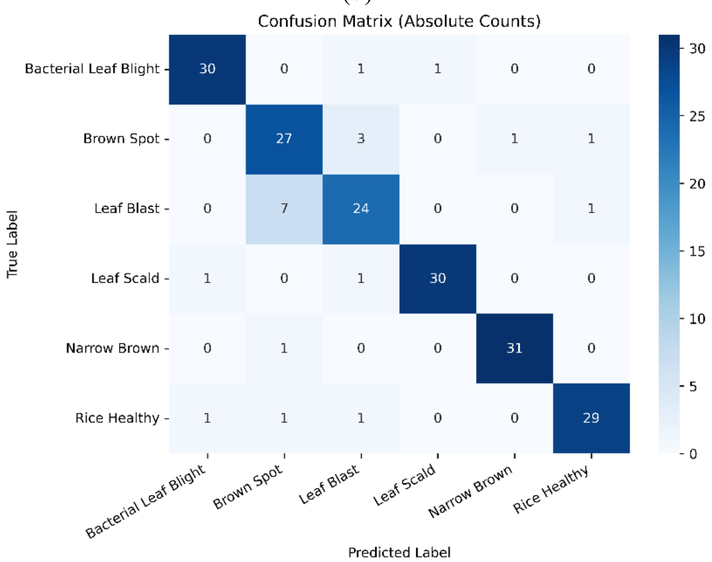

(c)

Figure 3. Absolute-count confusion matrices for (a) MobileNetV2, (b) EfficientNetB0, and (c) NASNetMobile on the independent test set.

3.5. Computational efficiency.

Table 4 presents the computational efficiency indicators for all three models. MobileNetV2 demonstrated the most favorable deployment profile among the three architectures, with a total parameter count of 2,624,710, a TensorFlow Lite model size of 2.75 MB, an inference latency of 3.22 ms per image, and FLOPs of 613,459,300, all of which were the lowest among the three models evaluated. EfficientNetB0 had a TFLite size of 4.69 MB and a comparable inference latency of 3.23 ms per image. NASNetMobile had the highest parameter count at 4,578,202, the largest TFLite size at 5.07 MB, and the slowest inference latency at 4.00 ms per image despite having the fewest trainable parameters among the three models (474,806), a consequence of its frozen backbone layers contributing a large share of non-trainable parameters.

Table 4. Computational efficiency comparison of all models.

Metric	MobileNetV2	EfficientNetB0	NASNetMobile
Total Parameters	2,624,710	4,416,297	4,578,202
Trainable Parameters	1,874,886	1,847,526	474,806
Model Size (.keras)	24.90 MB	31.77 MB	23.70 MB
TFLite Size	2.75 MB	4.69 MB	5.07 MB
Latency (ms/image)	3.22	3.23	4.00
FPS	310.48	309.61	249.92
FLOPs	613,459,300	801,518,267	1,146,375,664

MobileNetV2 achieved the highest classification accuracy while simultaneously maintaining the smallest TFLite model size and the lowest computational cost, representing the most favorable accuracy-efficiency trade-off among the three architectures. Its TFLite size of 2.75 MB is approximately 1.7 times smaller than EfficientNetB0 and 1.8 times smaller than NASNetMobile, making it considerably more suitable for storage-constrained deployment environments [9, 24]. It must be noted, however, that these efficiency measurements were obtained in a controlled GPU notebook environment rather than on physical mobile or edge hardware. On-device latency on ARM-based processors may differ from the values reported here, and empirical on-device validation is recommended before making conclusive deployment claims. Therefore, deployment suitability in this study is presented as indicative, representing potential for resource-constrained deployment pending on-device empirical validation [9, 15].

4. Conclusions

This study presented a systematic benchmark evaluation of three lightweight CNN architectures, MobileNetV2, EfficientNetB0, and NASNetMobile, for the classification of six rice leaf disease categories using a dataset of 1,920 images under a unified experimental framework. All three models were trained on an identical stratified split using a two-stage transfer learning protocol with a shared custom classification head comprising Global Average Pooling, Batch Normalization, two regularized dense layers with dropout, and a Softmax output layer. MobileNetV2 achieved the highest test accuracy of 96.35% and macro F1-score of 96.35%, with perfect per-class classification on Bacterial Leaf Blight, Leaf Scald, and Narrow Brown Leaf Spot. It also demonstrated the most favorable computational efficiency profile, with a TensorFlow Lite model size of 2.75 MB, an inference latency of 3.22 ms per image, and

FLOPs of 613,459,300, all lower than those of EfficientNetB0 and NASNetMobile. These findings confirm that MobileNetV2 provides a superior balance of classification accuracy and computational efficiency for rice leaf disease identification compared to the two baseline architectures under the same experimental conditions. Overfitting was effectively controlled through the combination of dropout regularization, L2 weight decay, EarlyStopping, and ReduceLROnPlateau, as evidenced by stable training-validation convergence and strong generalization on the independent test set. Brown Spot and Leaf Blast remained the most challenging class pair across all three architectures, consistent with their visual similarity. Future work should investigate attention mechanisms or feature disentanglement strategies specifically targeting the discrimination of lesion-based disease classes. Additionally, on-device validation on physical mobile or edge hardware is recommended to empirically substantiate the deployment efficiency claims beyond the proxy metrics reported in this study.

Acknowledgments

The authors gratefully acknowledge Universitas Syiah Kuala (USK), particularly the Department of Electrical Engineering, Faculty of Engineering, for providing an enabling academic environment, research facilities, and institutional support that contributed to the completion of this study. The authors also extend their appreciation to colleagues and laboratory staff who provided technical and administrative assistance during the research process.

Author Contribution

Melinda contributed to the conceptualization, methodological supervision, and overall direction of the study. Yunidar contributed to validation, technical review, and refinement of the manuscript. Rahmat Maulana conducted the experiments, implemented the model, analyzed the results, and prepared the initial manuscript draft. Nurlida Basir and Elizar supported data collection/curation and contributed to data interpretation and manuscript review. Muhammad Saifullah Nur and Muhammad Irhamsyah supported data preparation, software implementation, and result interpretation. All authors reviewed and approved the final manuscript.

Competing Interest

The authors declare that they have no financial, personal, or professional conflicts of interest that could influence or appear to influence the outcomes or interpretation of this research.

References

[1] Kim, Y.S.; Choi, S.H.; Park, J.H. (2025). Changes in the occurrence patterns of rice fungal diseases due to climate change. *Research in Plant Disease*, *31*, 17–33. https://doi.org/10.5423/RPD.2025.31.1.17.

[2] Shen, M.; Cai, C.; Song, L.; Qiu, J.; Ma, C.; Wang, D.; et al. (2023). Elevated CO_2 and temperature under future climate change increase severity of rice sheath blight. *Frontiers in Plant Science*, *14*, 1115614. https://doi.org/10.3389/fpls.2023.1115614.

[3] Singh, P.K.; Ray, S.; Thakur, S.; Rathour, R.; Sharma, V.; Sharma, T.R. (2018). Co-evolutionary interactions between host resistance and pathogen avirulence genes in rice-Magnaporthe oryzae pathosystem. *Fungal Genetics and Biology*, *115*, 9–19. https://doi.org/10.1016/j.fgb.2018.04.005.
[4] Kusuma, B.; Hermanto, T.I.; Lestari, C.D. (2025). Classification of disease types in rice plants using convolutional neural network algorithm. *Jurnal Informatika dan Komputer*, *9*, 40. https://doi.org/10.26798/jiko.v9i1.1395.
[5] Akter, S.; Sumon, R.I.; Ali, H.; Kim, H.C. (2024). Utilizing convolutional neural networks for the effective classification of rice leaf diseases through a deep learning approach. *Electronics*, *13*, 4095. https://doi.org/10.3390/electronics13204095.
[6] Ayyappan, A.B.; Gobinath, T.; Kumar, M.; Sivaramakrishnan, A. (2025). Rice plant disease detection using convolutional neural networks. *Discover Artificial Intelligence*, *5*, 277. https://doi.org/10.1007/S44163-025-00277-X.
[7] Dutta, M.; Saikia, M.J.; Bora, K.; Barman, D.; Sarma, K.K.; Paul, R.; et al. (2024). Rice leaf disease classification: a comparative approach using CNN, CAAR-U-Net, and MobileNetV2. *Technologies*, *12*, 214. https://doi.org/10.3390/technologies12110214.
[8] Shankar, S.J.; Palanivel, S.; Venkateswarlu, S.C. (2025). Automating paddy crop disease classification with deep learning models. *Journal of Computer Science*, *21*, 1772–1784. https://doi.org/10.3844/jcssp.2025.1772.1784.
[9] Putra, O.V.; Mustaqim, M.Z.; Muriatmoko, D. (2023). Transfer learning for classification of rice diseases and pests using MobileNetV2. *Techno.Com*, *22*, 562–575. https://doi.org/10.33633/tc.v22i3.8516.
[10] Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; Chen, L.C. (2018). MobileNetV2: inverted residuals and linear bottlenecks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4510–4520. https://doi.org/10.1109/CVPR.2018.00474.
[11] Tan, M.; Le, Q.V. (2019). EfficientNet: rethinking model scaling for convolutional neural networks. *Proceedings of the International Conference on Machine Learning*, *97*, 6105–6114.
[12] Zoph, B.; Vasudevan, V.; Shlens, J.; Le, Q.V. (2018). Learning transferable architectures for scalable image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8697–8710. https://doi.org/10.1109/CVPR.2018.00907.
[13] Zhang, Z.; Gu, Y.; Hong, Q. (2021). Rice disease identification system using lightweight MobileNetV2. *Proceedings of the International Conference on Image, Signal and Information Processing*, 29–36. https://doi.org/10.12792/icisip2021.007.
[14] Tiwari, R.; Vora, N. (2024). Enhancing paddy leaf disease classification using CNN and MobileNetV2. *Journal of Soft Computing Paradigm*, *6*, 324–340. https://doi.org/10.36548/jscp.2024.3.008.
[15] Pai, P.; Amutha, S.; Basthikodi, M.; Pai, B.H.; Pai, R.M.; Nayak, S.R. (2025). A twin CNN-based framework for optimized rice leaf disease classification with feature fusion. *Journal of Big Data*, *12*, 89. https://doi.org/10.1186/s40537-025-01148-z.
[16] Xu, Y.; Li, D.; Li, C.; Yuan, Z.; Dai, Z. (2025). LiSA-MobileNetV2: an extremely lightweight deep learning model for accurate rice disease classification. *Frontiers in Plant Science*, *16*, 1–14. https://doi.org/10.3389/fpls.2025.1619365.
[17] Munda, K.K.; Patil, N. (2025). A deep learning framework for plant disease detection. *Lecture Notes in Networks and Systems*, *1239*, 353–366. https://doi.org/10.1007/978-981-96-1188-1_26.
[18] Tanwar, V.; Lamba, S.; Sharma, B. (2023). Deep learning-based hybrid model for severity prediction of leaf smut sugarcane infection. *Proceedings of the 2023 International Conference on Artificial Intelligence and Smart Energy*, 1004–1009. https://doi.org/10.1109/ICAIS56108.2023.10073880.

[19] Jayasooriya, G.R.I.L.; Arachchi, S.M. (2022). Diagnosis of bacterial leaf blight, brown spots, and leaf smut rice plant diseases using Light GBM. *International Journal of Computer Applications*, *183*, 53–58. https://doi.org/10.5120/ijca2022921895.

[20] Pal, C.; Karmakar, S.; Mukherjee, I.; Chakrabarti, P.P. (2025). A lightweight and explainable CNN model for empowering plant disease diagnosis. *Scientific Reports*, *15*, 30720. https://doi.org/10.1038/s41598-025-94083-1.

[21] Lu, J.; Liu, X.; Ma, X.; Tong, J.; Peng, J. (2023). Improved MobileNetV2 crop disease identification model for intelligent agriculture. *PeerJ Computer Science*, *9*, e1595. https://doi.org/10.7717/peerj-cs.1595.

[22] Rice Leaf - v2 2023-06-19. *Roboflow Universe*. (accessed on 12 May 2026) Available online: https://universe.roboflow.com/utopialab/rice-leaf-wfax3/dataset/2.

[23] Hartono, A.C.; Muslikh, A.R. (2025). The application of MobileNetV2 transfer learning for image classification of fruit types. *Journal of Information Systems and Application Development*, *3*, 102–110. https://doi.org/10.26905/jisad.v3i2.16187

[24] Pranav Dhole, D.S.; Nandwanshi, I. (2025). Crop disease classification for edge devices: a quantized MobileNetV2 approach. *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering*, *13*, 39–45. https://doi.org/10.17148/ijireeice.2025.131006.