EDUMATIC Jurnal Pendidikan Informatika

e-ISSN 2549-7472
Edumatic: Jurnal Pendidikan Informatika
Volume 10, Issue 2, **300-309.**
https://e-journal.hamzanwadi.ac.id/index.php/edumatic
DOI: 10.29408/edumatic.v10i2.34588

Two-Stage Transfer Learning with EfficientNetB0 for Four-Class Banana Ripeness Classification

Anhar [1,*], Farid Wajidi [1], Chairi Nur Insani [1]

[1] Universitas Sulawesi Barat, Indonesia
[*] Corresponding author: Anhar, Universitas Sulawesi Barat, Indonesia
✉ anhar@unsulbar.ac.id

Received: 15 April 2026 | Revised: 9 May 2026 | Accepted: 20 May 2026 | Published: 11 June 2026

Abstract
Manual visual inspection to assess banana ripeness is subjective and unable to meet the scale of industrial needs, while previous studies have used heavy CNN architectures and have not systematically explored the depth of fine-tuning. This study proposes a two-phase transfer learning framework using EfficientNetB0 on a pre-augmented dataset of 13.478 images across four ripeness classes: unripe, ripe, overripe, and rotten. Class imbalance is addressed through class weighting during training. In Phase 1 (Feature Extraction), all base layers are frozen, and the classification head is trained until it achieves a best validation accuracy of 98.58%. In Phase 2 (Fine-Tuning), the optimal unfrozen layer depth was determined through systematic ablation across five configurations (5, 10, 15, 20, and 25 layers), with the 25-layer configuration yielding the highest validation accuracy of 98.75%. Evaluation on 562 test images yielded an accuracy of 99.46% and a test loss of 0.0538, with an F1-score of 1.00 for the overripe and ripe classes, and 0.99 for the rotten and unripe classes. The ROC curve confirmed high discriminative capability with an AUC of 1.000 for the overripe and ripe classes, and 0.999 and 0.998 for the unripe and rotten classes. These results demonstrate that the combination of a two-phase strategy, depth ablation, and fine-tuning with class weights yields a robust classification system with potential for application in automated banana sorting using edge devices.

Keywords: bananas; efficientnetb0; ripeness classification; transfer learning

To cite this article: Anhar, A., Wajidi, F., & Nur Insani, C. (2026). Two-Stage Transfer Learning with EfficientNetB0 for Four-Class Banana Ripeness Classification. *Edumatic: Jurnal Pendidikan Informatika*, *10*(2), 300–309. https://doi.org/10.29408/edumatic.v10i2.34588

INTRODUCTION
Bananas (*Musa spp.*) represent an important horticultural commodity with high economic value and are widely consumed both in Indonesia and globally. Indonesia is among the largest banana producers in Southeast Asia, with production exceeding 9.33 million tons in 2023 (Voronoi, 2025). This large production volume requires efficient postharvest management, particularly in determining fruit ripeness because it directly affects quality, shelf life, and market value (Shezi et al., 2023). In practice, ripeness assessment still relies on manual visual inspection, which is subjective and difficult to maintain consistently in large-scale sorting environments (Rizzo et al., 2023). In addition, banana ripeness datasets often exhibit class imbalance, causing deep learning models to become biased toward dominant classes. Xu et al. (2022) explained that adaptive class-balanced weighting can help maintain balanced learning performance across minority and majority classes. Therefore, handling class imbalance becomes an important aspect in developing robust banana ripeness classification systems.

Advances in the field of artificial intelligence, particularly Convolutional Neural Networks (CNNs), have opened up significant opportunities for developing reliable visual classification systems (Philip, 2025). CNNs have proven capable of hierarchically extracting visual features from digital images and have been widely applied to fruit ripeness classification (Yang, 2023). Various CNN architectures, such as VGG16, MobileNet, and ResNet, have been applied to fruit ripeness classification with accuracy ranging from 90% to 96% (Mathew et al., 2023). However, these architectures have a fundamental weakness in terms of computational efficiency: VGG16 has 138 million parameters, while ResNet50 has 25 million parameters, making both impractical to implement on edge devices commonly used in small- to medium-scale sorting facilities. Although these studies report adequate accuracy, none have explored the depth of fine-tuning as a controlled experimental variable. Consequently, it remains unclear whether the reported results reflect optimal domain adaptation (Lee & Lee, 2023). The inconsistency in results across studies points to an unresolved issue: to what depth must the layers be pre-trained so that knowledge transfer from ImageNet can be optimally adapted to the domain of banana ripeness? This question becomes increasingly critical given the visual nature of banana ripeness a gradual color transition from green to yellow to speckled brown which requires domain-specific high-level feature representations.

EfficientNet uses a compound scaling strategy that simultaneously balances the network's depth, width, and image resolution (Lin et al., 2023). As the base model of the EfficientNet family, EfficientNetB0 achieves competitive classification accuracy with just 5.3 million parameters a 26-fold reduction compared to VGG16. This compound scaling strategy enables EfficientNetB0 to generate efficient feature representations: the early layers capture edges and surface textures, the middle layers extract color patterns and spatial distributions, while the final layers produce high-level semantic representations that reflect ripeness. This parameter efficiency makes it particularly well-suited for banana sorting systems operating under computational resource constraints, including the need for fast inference on edge devices such as the Raspberry Pi or NVIDIA Jetson Nano (Maylianti et al., 2025; Sugianti et al., 2022).

Transfer learning further enhances this advantage by leveraging weights trained on ImageNet. This allows the network to be initialized with general visual features such as edges, textures, and color gradients that can be directly reused for fruit surface analysis (Baker & Zengeler, 2022). Nastitie & Handayani (2025) demonstrates the superiority of EfficientNetB0 over MobileNetV2 for banana variety classification, achieving 80% accuracy on a dataset of 2.248 images, though without a systematic two-phase strategy. Similarly, Manik et al. (2025) reported that EfficientNetB0 outperformed VGG16 and ResNet50 for the classification of horticultural commodities. Unlike previous studies, this study systematically evaluates the effect of fine-tuning depth on EfficientNetB0 using two-phase transfer learning for the classification of four levels of banana ripeness, while taking computational efficiency into account for edge deployment.

This study highlights three important gaps in previous research. First, many existing studies classify banana ripeness into only two or three categories, making it difficult to capture the gradual visual transitions that occur between ripeness stages. Second, prior research has largely depended on heavy CNN architectures, which often require substantial computational resources and are less suitable for deployment on low-resource devices. Third, limited attention has been given to the systematic evaluation of two-phase transfer learning strategies with different fine-tuning depths for four-class banana ripeness classification.

This study aims to develop and evaluate a four-class banana ripeness classification system using EfficientNet with a two-phase transfer learning strategy on a large-scale dataset. The study investigates the effect of fine-tuning depth on model adaptation and generalization performance while identifying the optimal balance between preserving pretrained representations and adapting domain-specific semantic features. The contributions of this study include the development of a structured two-phase transfer learning framework, systematic ablation experiments on multiple fine-tuning depths (5, 10, 15, 20, and 25 layers), and empirical evidence showing that lightweight CNN architectures can achieve highly competitive performance for agricultural image classification when transfer learning optimization is conducted systematically. The framework was evaluated on a dataset of 13.478 banana images representing unripe, ripe, overripe, and rotten classes, making it relevant for real-time postharvest sorting and edge-based smart agriculture applications.

METHOD

This study employs a quantitative experimental design using a two-phase transfer learning approach on a large-scale dataset of banana ripeness images. All experiments were conducted using Google Colaboratory, accelerated by an NVIDIA Tesla T4 GPU, using the TensorFlow 2.20 and Keras frameworks, and Python 3.12. The overall research workflow is shown in Figure 1.

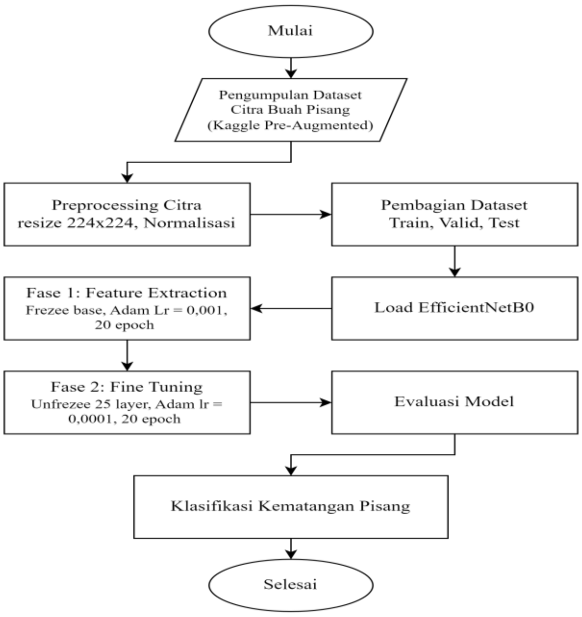

Figure 1. Banana Ripeness Classification Flowchart

As shown in Figure 1, this stage is designed to produce an optimal banana ripeness classification model using the EfficientNetB0 architecture. The dataset used was obtained from the Kaggle platform and consists of 13.478 banana images that have undergone a pre-augmentation process including horizontal/vertical flipping, 90° rotation (clockwise, counterclockwise, and inverted), cropping with a zoom range of 0% to 20%, rotation between -15° and +15°, as well as variations in hue, saturation, brightness, and exposure between -10% and +10%, and blurring up to 1px. It should be noted that this dataset pre-augmentation was performed by the dataset provider prior to subset division; consequently, the augmentation variations are evenly distributed across the training, validation, and test partitions. This condition is acknowledged as a methodological limitation, as ideal augmentation should only be applied to the training data after the data splitting process to prevent potential information leakage.

The distribution in Table 1 shows a moderate class imbalance, particularly in the unripe class, which has relatively fewer samples compared to the rotten and ripe classes. To address this bias, class weights were applied using the formula: $w_i = n_{total} / (n_{class} \times n_i)$, resulting in weights: unripe (1.55), overripe (1.25), ripe (0.84), and rotten (0.73). All images were resized to 224×224 pixels to match the standard input dimensions of EfficientNetB0, then normalized using the `preprocess_input` function from `tf.keras.applications.efficientnet`. Mixed-precision training (`mixed_float16`) was enabled to optimize GPU memory efficiency, with a batch size of 32 (Das et al., 2022).

Table 1. Dataset distribution

Subset	Overripe	Ripe	Rotten	Unripe
Train	2.349	3.522	4.020	1.902
Validation	229	339	388	167
Test	113	154	185	110
Total	2.691	4.015	4.593	2.179

The classification model was built using EfficientNetB0 as the base model with ImageNet weights (include_top=False), to which a custom classification head was added consisting of a GlobalAveragePooling2D layer, BatchNormalization, a Dense layer (256, ReLU), Dropout (0.5), and a Dense layer (4, softmax). Training was conducted in two sequential phases using the Adam optimizer with identical callback configurations: EarlyStopping (patience=5, monitor=val_accuracy), ReduceLROnPlateau (factor=0.3, patience=3, monitor=val_loss), and ModelCheckpoint. Phase 1 (Feature Extraction) was performed by freezing all base layers of EfficientNetB0 and training only the classification head for a maximum of 20 epochs with an initial learning rate of 1e-3. Phase 2 (Fine-Tuning) involved unfreezing the top layers of the base model and training them together with the classification head for a maximum of 20 epochs with a learning rate of 1e-4 (Arastu et al., 2024).

RESULTS AND DISCUSSION

Results

The experimental results are presented in the form of a training history by epoch. Visualization is performed using accuracy and loss curves on the validation data. The evaluation also includes a confusion matrix to assess the model's classification capabilities. Additionally, a classification report is used to evaluate performance metrics in detail. The ROC curve is displayed to provide a comprehensive overview of the model's performance across each maturity category. Table 2 shows the training history for Phase 1, illustrating the progression of accuracy and loss during the feature extraction process.

Table 2. Phase 1 Training history (feature extraction)

Epoch	Accuracy	Loss	Val Acc	Val Loss	LR
1	0.9119	0.2759	0.9564	0.1416	1e-3
2	0.9592	0.1244	0.9786	0.1026	1e-3
3	0.9675	0.0951	0.9858	0.0923	1e-3
7	0.9827	0.0468	0.9840	0.0875	3e-4
14	0.9913	0.0237	0.9840	0.0817	9e-5

The results in Table 2 show that the model experienced a very rapid improvement in performance during the early epochs (1–3), with validation accuracy peaking at 98.58% in the third epoch. After that, although training accuracy continued to rise to over 99%, validation accuracy tended to stagnate and even declined slightly, indicating mild overfitting that justifies the need for Phase 2. Table 3 presents a comparison of validation accuracy for each configuration of the number of layers unfrozen in Phase 2, which was used to determine the optimal depth of fine-tuning.

Table 3. Comparison of val accuracy across different numbers of fine-tuning layers

Number of Unfreeze Layers	Val Accuracy (%)	Val Loss	Description
5	98.75	0.0705	High Performance
10	98.75	0.0801	High Performance
15	98.39	0.0721	Decline in Performance
20	98.66	0.0746	Consistent Performance
25	98.75	0.0567	Best Performance

According to Table 3, the 5, 10, and 25-layer configurations each achieved a validation accuracy of 98.75%. The 15-layer configuration yields the lowest performance (98.39%), while the 20-layer configuration falls in the middle (98.66%). Although the accuracy of some configurations is the same, the 25-layer configuration has the lowest validation loss of 0.0567, indicating better model stability and generalization ability. The lower validation loss in the 25-layer configuration indicates that the model has a more stable level of prediction confidence compared to the other configurations. These results suggest that increasing the number of fine-tuning layers does not always improve performance linearly; however, the 25-layer configuration provides the best balance between feature adaptation and preservation of the pretrained representation.

Table 4. Phase 2 training history (fine-tuning)

Epoch	Accuracy	Loss	Val Acc	Val Loss	LR
1	0.9899	0.0302	0.9840	0.0760	1e-4
2	0.9938	0.0207	0.9849	0.0722	1e-4
3	0.9947	0.0185	0.9858	0.0998	1e-4
4	0.9959	0.0132	0.9866	0.0688	3e-5
5	0.9952	0.0158	0.9875	0.0714	3e-5

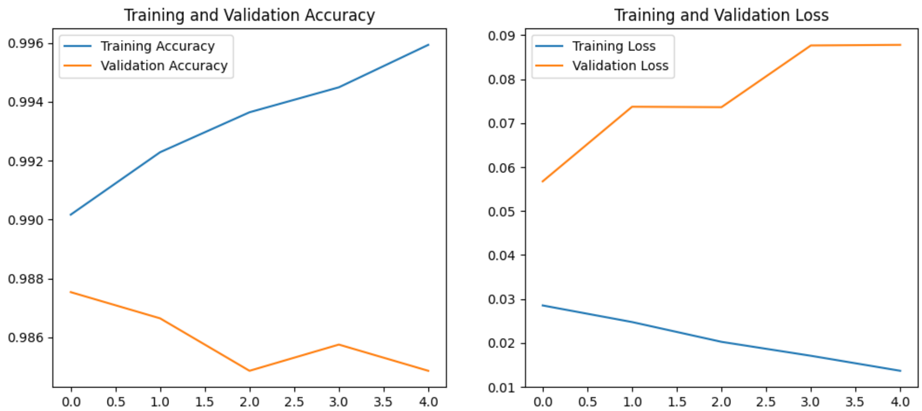

Figure 2. Accuracy and loss validations plot

The results in Table 4 show that the model in the fine-tuning phase demonstrated high performance from the start (accuracy >98%). Performance improved gradually, peaking at the 5th epoch with a validation accuracy of 98.75%. Reducing the learning rate from 1e-4 to 3e-5 proved to help improve validation performance, although fluctuations in the validation loss indicate the potential for mild overfitting.

Figure 2 shows a graph of the accuracy curve and training versus validation loss during Phase 2. The graph indicates that training accuracy consistently increased from approximately 99.0% to nearly 99.6%, while validation accuracy tended to plateau in the range of 98.5%–98.7%. On the other hand, training loss continues to decrease, but validation loss actually increases a pattern indicating the onset of overfitting. This confirms that the early stopping and ReduceLROnPlateau mechanisms play a crucial role in preventing further degradation.

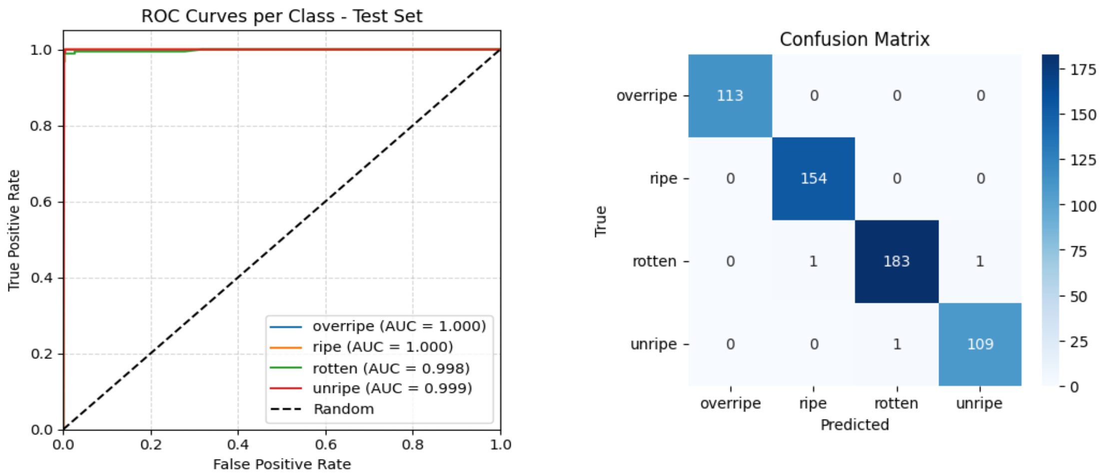

Figure 3. ROC-AUC plot and confusion matrix of test results

Figure 3 shows the ROC-AUC curve and confusion matrix from the evaluation results on the test data. The ROC curve indicates that all classes exhibit very high classification performance, with AUC values approaching 1.0. The overripe and ripe classes achieve a perfect AUC (1.000), while the unripe (0.999) and rotten (0.998) classes are also nearly perfect. A curve approaching the top-left corner indicates the model's excellent ability to distinguish between each class. Final evaluation on 562 test images yielded an accuracy of 99.46% and a test loss of 0.0538. The confusion matrix shows that the model is capable of classifying the majority of the data with high accuracy. The overripe (113) and ripe (154) classes were predicted perfectly with no errors. In the rotten class, there were 2 minor errors (1 data point misclassified as ripe and 1 as unripe), while the unripe class had only 1 error classified as rotten.

Table 5. Classification report on test data

Class	Precision	Recall	F1-Score	Support
overripe	1.00	1.00	1.00	113
ripe	0.99	1.00	1.00	154
rotten	0.99	0.99	0.99	185
unripe	0.99	0.99	0.99	110
accuracy			0.99	562
macro avg	0.99	1.00	0.99	562
weighted avg	0.99	0.99	0.99	562

The results in Table 5 show excellent model performance across all classes. The overripe class achieved a perfect score (precision, recall, and F1-score = 1.00), while the ripe, rotten, and unripe classes also achieved very high scores (0.99). The overall accuracy of 0.99 indicates that nearly all data points were classified correctly. The average macro and weighted scores in the 0.99 range indicate that the model's performance is consistent and balanced across all classes.

Discussion

The findings demonstrate that the proposed two-phase transfer learning framework effectively improves convergence stability and classification performance in four-class banana ripeness recognition. The model achieved a test accuracy of 99.46% with consistently high precision, recall, and F1-score values across all ripeness categories, indicating strong generalization capability despite moderate class imbalance. The experimental results also reveal that the depth of fine-tuning significantly influences model adaptation performance. Among the evaluated configurations, unfreezing the top 25 layers produced the most stable validation performance and the lowest validation loss, suggesting that selective adaptation of deeper semantic layers plays an important role in distinguishing visually adjacent ripeness stages. In addition, the extremely low number of misclassifications indicates that the model successfully learned the ordinal visual progression of banana maturation rather than relying solely on superficial color differences.

The results reinforce the feature transferability theory in deep transfer learning, which explains that lower CNN layers primarily learn generic visual features, whereas deeper layers capture increasingly domain-specific semantic representations (Parrillo et al., 2026). In the feature extraction phase, freezing the pretrained backbone enabled the model to retain generalized knowledge acquired from ImageNet, including fundamental visual characteristics such as edges, textures, and color gradients that remain highly relevant for agricultural image analysis (Baker & Zengeler, 2022). The fine-tuning phase subsequently allowed deeper layers to adapt more effectively to ripeness-related visual characteristics, including gradual skin color transformation and necrotic spotting patterns associated with banana maturation (Shezi et al., 2023). The strong performance achieved by the 25-layer configuration suggests that optimal transfer learning performance is obtained when generalized pretrained representations are preserved while higher-level semantic features are selectively adapted to the target domain.

The proposed framework demonstrates superior classification performance while maintaining computational efficiency. Hanifah & Hermawan (2023) reported 88% accuracy using a conventional CNN architecture, while Mathew et al. (2023) achieved 90–96% accuracy using heavier hybrid architectures such as ResNet50 and VGG16. In addition, Nastitie & Handayani (2025) obtained 80% accuracy using EfficientNetB0 without implementing a structured two-phase transfer learning strategy. The higher performance achieved in this study indicates that optimization of transfer learning strategy contributes substantially to classification effectiveness beyond architectural selection alone. The findings further suggest that systematic optimization of fine-tuning depth and adaptive representation learning plays a more important role than merely increasing model complexity or parameter size.

Several findings also reveal important characteristics of the model behavior that require scientific interpretation. The non-linear ablation pattern shows that increasing the number of unfrozen layers does not consistently improve classification accuracy. The 15-layer configuration produced lower performance compared with the 5-, 10-, and 25-layer configurations, indicating that partial adaptation of intermediate representations may disrupt previously stable pretrained features. This phenomenon is closely related to catastrophic interference during transfer learning adaptation, where excessive modification of mid-level representations can temporarily reduce representation stability (Sehra et al., 2024). Performance recovery in the 25-layer configuration suggests that deeper semantic adaptation allows the network to reconstruct more discriminative ripeness-specific representations, thereby restoring classification stability.

This study demonstrates that fine-tuning depth is a decisive factor in determining the effectiveness of transfer learning for agricultural image classification. While previous studies on fruit ripeness recognition primarily emphasized architecture comparison, the present study systematically investigates how the depth of layer adaptation influences semantic feature learning and model generalization. The findings provide empirical evidence that lightweight CNN architectures such as EfficientNet are capable of achieving highly competitive performance when adaptive representation learning is optimized in a structured manner. More importantly, the results reveal that optimal transfer learning performance emerges when deeper semantic layers are selectively refined while lower-level generalized visual representations remain preserved. This observation strengthens current understanding of adaptive representation learning by explaining how the interaction between pretrained knowledge retention and domain-specific adaptation affects classification stability and discriminative capability. The study therefore contributes to the development of efficient and resource-aware deep learning strategies for precision agriculture and agricultural computer vision applications.

The findings indicate strong potential for implementation in real-time agricultural automation systems, particularly in postharvest banana sorting environments that require rapid, scalable, and cost-efficient operation. The lightweight architecture of EfficientNet, which contains substantially fewer parameters than VGG16 and ResNet50, enables deployment on low-resource edge devices with lower memory consumption and reduced inference latency (Lin et al., 2023). The consistently balanced classification performance across all ripeness categories also indicates the capability of the proposed framework to reduce subjectivity and inconsistency commonly associated with manual visual inspection. In a broader agricultural context, the proposed system may support improvements in postharvest quality control, reduce economic losses caused by inaccurate ripeness assessment, and strengthen the adoption of smart agriculture technologies in environments with limited computational infrastructure.

Several methodological constraints should be acknowledged to ensure a balanced interpretation of the findings. The dataset originated from a pre-augmented source prior to subset partitioning, which introduces the possibility of information leakage and may partially affect performance estimation. In addition, the dataset was collected under relatively controlled visual conditions and therefore may not fully represent real agricultural environments characterized by illumination variability, occlusion, heterogeneous backgrounds, and cultivar diversity, as also highlighted by Lee & Lee (2023). The absence of k-fold cross-validation further limits the statistical robustness and external generalizability of the proposed framework. Future studies should therefore incorporate real-field datasets, explainability analysis using Grad-CAM, cross-domain validation, and more rigorous validation protocols to strengthen the interpretability, robustness, and external validity of the model, particularly for real-world agricultural deployment scenarios (Lee & Lee, 2023).

CONCLUSION

This study demonstrates that a two-phase transfer learning strategy using EfficientNet is highly effective for four-class banana ripeness classification while remaining computationally efficient for edge-based agricultural applications. Systematic evaluation of fine-tuning depth shows that unfreezing the top 25 layers provides the best balance between preserving pretrained representations and adapting domain-specific semantic features, resulting in 99.46% test accuracy with consistently high precision, recall, and F1-score across all ripeness classes. The findings confirm that gradual fine-tuning, combined with class-weighting mechanisms, enables the model to capture subtle visual transitions between ripeness stages and improves generalization performance under class imbalance conditions. Beyond achieving high classification accuracy, this study also highlights the importance of fine-tuning depth optimization as a critical factor in transfer learning effectiveness for agricultural image classification. Furthermore, the lightweight architecture of EfficientNetB0 strengthens the feasibility of real-time deployment on low-resource edge devices for automated postharvest sorting systems. Nevertheless, limitations related to pre-augmented data and controlled imaging conditions remain, indicating the need for future studies involving real-world datasets, diverse environmental conditions, explainability analysis, and cross-validation strategies to further enhance the robustness and generalizability of the proposed framework

REFERENCES

Alshaari, H., & Alqahtani, S. (2025). Deep-EFNet : An Optimized EfficientNetB0 Architecture With Dual Regularization for Scalable Multi-Class Brain Tumor Classification in MRI. *IEEE Access*, *13*, 85682–85697. https://doi.org/10.1109/ACCESS.2025.3567919

Arastu, M. T. R., Muskan, T., Deepak, G., Sinha, K., & Kumari, K. (2024). Transformative Breast Cancer Diagnosis using CNNs with Optimized ReduceLROnPlateau and Early Stopping Enhancements. *International Journal of Computational Intelligence Systems*, *1*. https://doi.org/10.1007/s44196-023-00397-1

Baker, N. A., & Zengeler, N. (2022). A Transfer Learning Evaluation of Deep Neural Networks for Image Classification. *Machine Learning and Knowledge Extraction*, *4*, 22–41. https://doi.org/10.3390/make4010002

Das, S., Mullick, S. S., & Zelinka, I. (2022). On Supervised Class-Imbalanced Learning: An Updated Perspective and Some Key Challenges. *IEEE Transactions on Artificial Intelligence*, *3*(06), 973–993. https://doi.org/10.1109/TAI.2022.3160658

Hanifah, A. I., & Hermawan, A. (2023). Klasifikasi Kematangan Pisang Menggunakan Metode Convolutional Neural Network. *Komputika : Jurnal Sistem Komputer*, *12*(2), 49–56. https://doi.org/10.34010/komputika.v12i2.9999

Lee, J., & Lee, G. (2023). Unsupervised domain adaptation based on the predictive uncertainty of models. *Neurocomputing*, *520*, 183–193. https://doi.org/10.1016/j.neucom.2022.11.070

Lin, C., Yang, P., Wang, Q., Qiu, Z., Lv, W., & Wang, Z. (2023). Efficient and accurate compound scaling for convolutional neural networks. *Neural Networks*, *167*, 787–797. https://doi.org/10.1016/j.neunet.2023.08.053

Manik, F. Y., Efendi, S., Tarigan, J. T., & Lydia, M. S. (2025). Benchmarking Deep Learning Models for Visual Classification and Segmentation of Horticultural Commodities. *International Journal of Advanced Computer Science and Applications*, *16*(10), 440–448. https://doi.org/10.14569/IJACSA.2025.0161046

Mathew, M. B., Surya Manjunathan, G., Gokul, B., Mohana Ganesh, K., Sajidha, S. A., & Nisha, V. M. (2023). Banana Ripeness Identification and Classification Using Hybrid Models with RESNET-50, VGG-16 and Machine Learning Techniques BT - Machine Intelligence Techniques for Data Analysis and Signal Processing. *Machine Intelligence Techniques for Data Analysis and Signal Processing*, 259–273. https://doi.org/10.1007/978-981-99-0085-5_22

Maylianti, N. P., Wijayakusuma, I. G. N. L., & Arta Wiguna, I. P. C. (2025). Comparison of EfficientNet-B0 and ResNet-50 for Detecting Diseases in Cocoa Fruit. *Journal of Applied Informatics and Computing*, *9*(1), 115–120. https://doi.org/10.30871/jaic.v9i1.8868

Nastitie, C., & Handayani, I. (2025). Perbandingan metode mobilenetv2 dan efficientnetb0 untuk klasifikasi jenis pisang comparison of mobilenetv2 and efficientnetb0 methods for banana type classification. *Journal of Information Technology and Computer Science*, *8*(2023), 1990–1999. https://doi.org/10.31539/0240cx09

Parrillo, M., Laura, L., & Manna, A. (2026). Transfer Learning Strategies for Comic Character Recognition in Low-Data Regimes: A Comparative Study. *Future Internet*, *18*(4), 192.

https://doi.org/10.3390/fi18040192
Philip, A. R. (2025). Fine-Tuning Depth Analysis: Identifying the Sweet Spot for Maximum Accuracy in CNNs. *International Journal of Scientific Research in Computer Science and Engineering*, *13*(3), 48–63. https://doi.org/10.26438/ijsrcse.v13i3.703

Rizzo, M., Marcuzzo, M., Zangari, A., Gasparetto, A., & Albarelli, A. (2023). Fruit ripeness classification: A survey. *Artificial Intelligence in Agriculture*, *7*, 44–57. https://doi.org/10.1016/j.aiia.2023.02.004

Sehra, U., Dutta, R. R., & Pandit, M. K. (2024). Evaluating Neural Network Parameter Obliteration in Catastrophic Forgetting. *2024 IEEE Region 10 Symposium (TENSYMP)*, 1–6. https://doi.org/10.1109/TENSYMP61132.2024.10752250

Shezi, S., Magwaza, L. S., & Ncama, K. (2023). Importance of maturity indexing in postharvest management of selected subtropical and temperate fruits. In *Postharvest management of fresh produce* (pp. 59-79). Academic Press. https://doi.org/10.1016/B978-0-323-91132-0.00007-1

Sugianti, C., Imaizumi, T., Thammawong, M., & Nakano, K. (2022). Recent Postharvest Technologies in the Banana Supply Chain. *Reviews in Agricultural Science*, *10*, 123–137. https://doi.org/10.7831/ras.10.0_123

Voronoi. (2025). *Banana Production by Country 2023*. 2025; Voronoi App. https://www.voronoiapp.com/agriculture/Banana-Production-by-Country-2023-5557

Xu, C., Zheng, Y., Zhang, Y., Sun, C., Li, G., & Zhu, Z. (2022). Adaptive Class-Balanced Loss Based on Re-Weighting. *6th Asian Conference on Artificial Intelligence Technology (ACAIT)*, 1–8. https://doi.org/10.1109/ACAIT56212.2022.10137858

Yang, Y. (2023). Fruit Image Classification Using Convolution Neural Networks. *Highlights in Science, Engineering and Technology*, *34*, 110–119. https://doi.org/10.54097/hset.v34i.5430