

# Enhancing Cocoa Pod Disease Classification via Transfer Learning and Ensemble Methods: Toward Robust Predictive Modeling

Devina Anduyan

*College of Computing and Information Sciences (CCIS)  
Caraga State University (CSU)  
Butuan City, Philippines  
devina.anduyan@carsu.edu.ph*

Nyza Cabillo

*College of Computing and Information Sciences (CCIS)  
Caraga State University (CSU)  
Butuan City, Philippines  
nyza.cabillo@carsu.edu.ph*

Navy Gultiano

*College of Computing and Information Sciences (CCIS)  
Caraga State University (CSU)  
Butuan City, Philippines  
navy.gultiano@carsu.edu.ph*

Mark Phil Pacot

*College of Computing and Information Sciences (CCIS)  
Caraga State University (CSU)  
Butuan City, Philippines  
mbpacot@carsu.edu.ph*

**Abstract**—This study presents an ensemble-based approach for cocoa pod disease classification by integrating transfer learning with three ensemble learning strategies: Bagging, Boosting, and Stacking. Pre-trained convolutional neural networks, including VGG16, VGG19, ResNet50, ResNet101, InceptionV3, and Xception, were fine-tuned and employed as base learners to detect three disease categories: Black Pod Rot, Pod Borer, and Healthy. A balanced dataset of 6,000 cocoa pod images was curated and augmented to ensure robustness against variations in lighting, orientation, and disease severity. The performance of each ensemble method was evaluated using accuracy, precision, recall, and F1-score. Experimental results show that Bagging consistently achieved superior classification performance with a test accuracy of 100%, outperforming Boosting (97%) and Stacking (92%). The findings confirm that combining transfer learning with ensemble techniques improves model generalization and reliability, making it a promising direction for precision agriculture and automated crop disease management.

**Index Terms**—Cocoa pod disease classification, Transfer learning, Ensemble learning, Convolutional neural networks (CNNs), Deep learning

## I. INTRODUCTION

Cocoa (*Theobroma cacao*), a high-value tropical crop essential to global chocolate production, is severely threatened by diseases such as Black Pod Rot and Pod Borer. In the Philippines, it is recognized under Republic Act 7900 as a priority crop, yet disease outbreaks continue to impact yield and quality. Accurate and timely detection of these diseases is critical for effective management.

Recent developments in computer vision and machine learning have shown potential in automating plant disease diagnosis. In particular, **transfer learning** and **ensemble methods** have demonstrated improved performance in classification tasks. Ensemble learning combines the predictions of multiple base classifiers to reduce variance, enhance stability, and improve accuracy [1]–[3]. However, its effectiveness relies on the diversity and quality of the

base models used. Despite its promise, ensemble learning has not been widely explored for cocoa pod disease classification. Prior studies often utilized single classifiers or traditional machine learning techniques, which may fall short in addressing the complexity of disease patterns and visual variability [4], [5]. This study proposes an ensemble-based framework that integrates **transfer learning models**—such as VGG, ResNet, and others—as weak learners within **bagging**, **boosting**, and **stacking** schemes. By fine-tuning these pre-trained CNNs and systematically combining their predictions, the research aims to identify the optimal ensemble configuration for maximizing classification accuracy and reliability in cocoa pod disease detection. To address the limitations of single-model classifiers and advance automated disease detection in cacao cultivation, this research leverages ensemble learning techniques in conjunction with weak learners derived from fine-tuned transfer learning models. The framework aims to improve the accuracy and robustness of cocoa pod disease classification by evaluating the performance of different ensemble strategies and determining the most effective method for maximizing predictive performance. The remainder of this paper is organized as follows: Section 2 presents the materials and methods used in this study, including details on geometric correction technique and cloud removal algorithm. Section 3 discusses the experimental results and comparative performance against existing approaches. Finally, Section 4 concludes the study with a summary of findings and potential directions for future research.

## II. RELATED WORK

Ensemble learning has gained prominence in improving classification accuracy by combining multiple base models to reduce variance and bias [1], [3]. The three main types—bagging, boosting, and stacking—operate either in

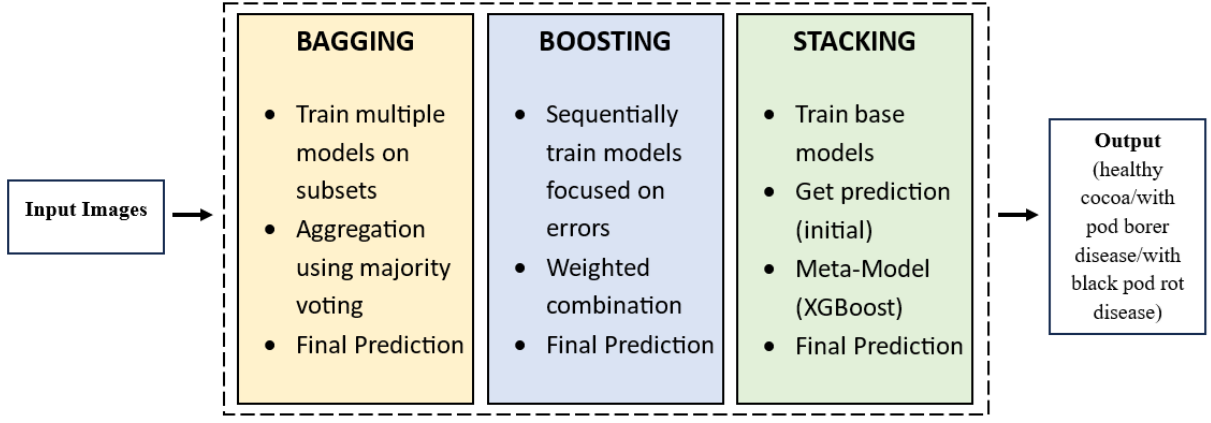


Fig. 1: The conceptual framework of the proposed solution.

parallel or sequentially. Bagging methods, such as Random Forest and Extra Trees, aggregate predictions from models trained on bootstrap samples, while boosting methods like AdaBoost, Gradient Boosting, and XGBoost sequentially refine model errors. Stacking leverages meta-learners to combine outputs from diverse base models [6], [7]. In agriculture, ensemble methods have been successfully applied to plant disease classification. [8] proposed an ensemble classifier combining SqueezeNet, VGG, InceptionV3, and DeepLoc for fruit disease detection, achieving up to 99% accuracy. [9] classified tomato leaf diseases using Random Forest, SVM, and MLP, attaining over 92% accuracy. Similarly, stacking-based ensemble classifiers have proven effective in diagnosing fruit tree diseases, with [10] reporting a test accuracy of 97.34% on a 10,000-image dataset. [11] achieved 100% accuracy in cherry classification using maximum voting over deep models. [12] applied ensemble deep learning (VGG16, VGG19, Xception) to grape leaf disease classification, achieving 99.82% accuracy. [13] utilized ensemble CNNs (EfficientNet, InceptionV3, MobileNetV2, VGG19) with bagging and weighted averaging for pear disease detection, demonstrating superior performance under real-field conditions. [8] employed hard and soft voting strategies using five deep models for plant disease classification, with hard voting nearing 100% accuracy. [14] implemented an optimized ensemble to detect wheat diseases, achieving classification accuracy up to 98.25%. While ensemble learning is well-explored across various crops, its application to cocoa pod disease remains limited. This study builds on the demonstrated success of ensemble frameworks in agriculture by combining transfer learning with ensemble strategies to improve the accuracy and reliability of cocoa pod disease classification.

### III. MATERIALS AND METHODS

The focus of this study is to determine the most effective ensemble learning method for cocoa pod disease classification by combining fine-tuned transfer learning models. Considering variations in pod appearance, lighting conditions, and disease severity, the study evaluates ensemble techniques such as bagging, boosting, and stacking.

Classification performance is measured using standard quantitative metrics to assess accuracy, robustness, and practical effectiveness in real-world agricultural settings.

#### A. Dataset

The dataset used in this study, titled *Cacao Disease*, was manually obtained from the Kaggle platform. It comprises approximately 4,300 images of cocoa pods with a resolution of  $1080 \times 1080$  pixels, categorized into three classes: *Black Pod Rot*, *Healthy*, and *Pod Borer*.

To ensure consistency and improve model performance, the images were pre-processed by resizing them to a standard resolution of  $240 \times 240$  pixels. In addition, data augmentation techniques such as random rotation and horizontal flipping were applied to increase dataset diversity and enhance the model's ability to generalize across different orientations and visual perspectives of cocoa pods.

Following preprocessing and augmentation, the dataset was balanced to contain 2,000 images per class, all stored in .jpg format. Representative samples from each class are shown in Figures 2, 3, and 4.



Fig. 2: Sample cocoa pod images affected by Pod Borer disease.

#### B. Base Models

This study employs established convolutional neural network (CNN) architectures—VGG16 [15], VGG19 [15], InceptionV3 [16], Xception [17], ResNet50 [18], and



Fig. 3: Sample of healthy cocoa pod images.



Fig. 4: Sample cocoa pod images affected by Black Pod Rot disease.

ResNet101 [18]—as base learners within the ensemble framework. These models were selected due to their proven performance in image classification tasks and their capacity to learn complex features and visual patterns. The use of these pre-trained networks aims to enhance the accuracy and robustness of the cocoa pod disease classification system by leveraging their transfer learning capabilities.

### C. Ensemble Learning Techniques

In this study, ensemble learning techniques—bagging, boosting, and stacking—are implemented to improve the performance of transfer learning models for cocoa pod disease classification. These methods aim to enhance prediction accuracy and robustness by leveraging multiple base learners and aggregating their outputs through different ensemble strategies.

1) *Bagging*: Bagging (Bootstrap Aggregating) is a parallel ensemble technique where multiple base learners are trained independently on different subsets of the dataset generated through bootstrapping. This reduces variance and mitigates overfitting. For each new sample, predictions from all learners are aggregated using majority voting.

Given a dataset  $D = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ ,  $m$  bootstrap samples  $D_i$  are generated by random sampling with replacement. Each base model  $M_i$  is trained on a corresponding  $D_i$ . For a test sample  $x$ , the final predicted class  $\hat{y}$  is obtained via:

$$\hat{y} = \arg \max_{c_j \in C} \sum_{i=1}^m \mathbb{I}[M_i(x) = c_j] \quad (1)$$

where  $\mathbb{I}[\cdot]$  is the indicator function and  $C$  is the set of possible classes.

2) *Boosting*: Boosting is a sequential ensemble method where each base learner is trained to correct the errors of its predecessor. The models are trained iteratively, and their predictions are combined using a weighted majority scheme. Boosting focuses on difficult samples by assigning them higher weights in subsequent rounds.

The boosting process follows:

- 1) Initialize weights  $w_i = \frac{1}{n}$  for all training instances.
- 2) For  $t = 1$  to  $T$  (number of rounds):
  - a) Train base learner  $M_t$  on weighted data.
  - b) Compute weighted error  $\varepsilon_t$ :

$$\varepsilon_t = \sum_{i=1}^n w_i \cdot \mathbb{I}[M_t(x_i) \neq y_i] \quad (2)$$

- c) Compute model weight:

$$\alpha_t = \frac{1}{2} \ln \left( \frac{1 - \varepsilon_t}{\varepsilon_t} \right) \quad (3)$$

- d) Update sample weights:

$$w_i \leftarrow w_i \cdot \exp(-\alpha_t y_i M_t(x_i)) \quad (4)$$

- e) Normalize weights so that  $\sum w_i = 1$ .

The final prediction is computed as:

$$\hat{y} = \text{sign} \left( \sum_{t=1}^T \alpha_t M_t(x) \right) \quad (5)$$

3) *Stacking*: Stacking is a layered ensemble technique that combines predictions from multiple base models using a meta-model. Base learners are trained on the original training data, and their outputs are used as input features for a higher-level model that learns how to best combine them.

Let  $M_1, M_2, \dots, M_k$  be base models. For each training instance  $x_i$ , meta-features are generated as  $Z_i = [M_1(x_i), M_2(x_i), \dots, M_k(x_i)]$ . A meta-model  $M_{\text{meta}}$  is then trained using the feature matrix  $Z = \{Z_1, Z_2, \dots, Z_n\}$  and original labels  $Y = \{y_1, y_2, \dots, y_n\}$ . The final prediction for a new sample  $x$  is:

$$\hat{y} = M_{\text{meta}}([M_1(x), M_2(x), \dots, M_k(x)]) \quad (6)$$

Stacking enables the system to leverage the strengths of different base models and reduce generalization error.

### D. Evaluation Metrics

After fine-tuning and validating the pre-trained models, their performance is assessed on the test dataset using standard classification metrics: accuracy, precision, recall, and F1-score. These metrics provide a comprehensive evaluation of the model's predictive capabilities across all cocoa pod disease classes.

- **Accuracy** measures the proportion of correctly predicted instances over the total number of predictions. It provides an overall indication of model performance:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (7)$$

where:

- $TP$  = True Positives
- $TN$  = True Negatives
- $FP$  = False Positives
- $FN$  = False Negatives

- **Precision** indicates the correctness of positive predictions, i.e., the proportion of predicted positive cases that are truly positive:

$$\text{Precision} = \frac{TP}{TP + FP} \quad (8)$$

- **Recall** (also known as sensitivity) measures the model's ability to correctly identify actual positive cases:

$$\text{Recall} = \frac{TP}{TP + FN} \quad (9)$$

- **F1-Score** provides a harmonic mean of precision and recall, balancing both metrics into a single value:

$$\text{F1-Score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (10)$$

These metrics are especially relevant in multi-class classification tasks such as cocoa pod disease detection, where both over- and under-predictions of disease types can significantly affect model reliability.

#### IV. EXPERIMENTS

##### A. Cocoa pod disease classification results

Table 1 and the corresponding confusion matrix in Figure 5, demonstrate the high performance of the bagging ensemble with majority voting on the cocoa pod disease classification task. The model achieved an overall accuracy of 100%, with only four misclassifications out of 900 test samples: two healthy misclassified as black pod rot, and one pod borer each misclassified as black pod rot and healthy. Despite these minor errors, the precision, recall, and F1-scores remained at or near 1.00 across all classes, indicating excellent robustness and class balance.

TABLE I: Classification Report - Bagging (Majority Voting) Ensemble

	precision	recall	f1-score	support
Black Pod Rot	0.99	1.00	1.00	300
Healthy	1.00	0.99	0.99	300
Pod Borer	1.00	0.99	1.00	300
Accuracy			1.00	900
Macro Avg	1.00	1.00	1.00	900
Weighted Avg	1.00	1.00	1.00	900




Fig. 5: Confusion Matrix - Bagging (Majority Voting) Ensemble.

Table 2 and the associated confusion matrix in Figure 6, summarize the performance of the Boosting ensemble method. The model achieved an overall accuracy of 97%, with class-level F1-scores of 0.99 for Black Pod Rot, 0.96 for Healthy, and 0.97 for Pod Borer. The confusion matrix reveals minor misclassifications, including 4 Healthy samples incorrectly predicted as Black Pod Rot and 16 as Pod Borer. While Pod Borer was classified perfectly, Healthy class exhibited the most confusion, slightly reducing recall. Nonetheless, macro and weighted averages of 0.97 across all metrics demonstrate the model's strong and balanced performance across classes.

TABLE II: Classification Report - Boosting Ensemble

	precision	recall	f1-score	support
Black Pod Rot	0.99	0.99	0.99	300
Healthy	0.99	0.93	0.96	300
Pod Borer	0.94	1.00	0.97	300
Accuracy			0.97	900
Macro Avg	0.97	0.97	0.97	900
Weighted Avg	0.97	0.97	0.97	900




Fig. 6: Confusion Matrix - Boosting Ensemble.

Table 3 and the corresponding confusion matrix in Figure 7, present the evaluation results of the stacking ensemble model. Despite achieving a perfect confusion matrix with all 900 test samples correctly classified, the class-wise F1-scores reveal some inconsistencies. While Black Pod Rot was predicted with high consistency (F1 = 0.93), Healthy had a lower precision (0.83), and Pod Borer showed a lower recall (0.85). These discrepancies



suggest internal class imbalance or meta-feature misalignment despite perfect final predictions. The model achieved a macro and weighted average of 0.92 across all metrics, indicating good but not superior generalization compared to bagging or boosting.

TABLE III: Classification Report - Stacking Ensemble

	precision	recall	f1-score	support
Black Pod Rot	0.96	0.90	0.93	300
Healthy	0.83	0.99	0.90	300
Pod Borer	0.99	0.85	0.92	300
Accuracy			0.92	900
Macro Avg	0.93	0.92	0.92	900
Weighted Avg	0.93	0.92	0.92	900




Fig. 7: Confusion Matrix - Stacking Ensemble.

#### B. Model performance across ensemble methods

Figure 8 illustrates the training and validation accuracy of six pre-trained CNN models (VGG16, VGG19, ResNet50, ResNet101, InceptionV3, and Xception) under three ensemble learning strategies: *Bagging*, *Boosting*, and *Stacking*. Across all architectures, Bagging consistently achieved near-perfect validation accuracy ( $\geq 0.99$ ), indicating strong generalization and model stability.

Boosting also performed well, although minor drops were observed in VGG19 and ResNet50, where validation accuracy dipped to 0.98 and 0.97, respectively—suggesting slight overfitting or sensitivity to sample weighting. Stacking maintained competitive performance with validation accuracies closely matching Bagging, particularly for deeper architectures like ResNet101 and Xception.

Overall, Bagging delivered the most consistent performance across all models, with minimal variance between training and validation scores. Boosting showed excellent training accuracy but slight drops in generalization. Stacking provided robust results, though with more variability depending on the base architecture. These trends highlight the superior stability of Bagging for cocoa pod disease classification tasks, especially when paired with deep pre-trained models.

Figure 9 compares the overall test set accuracy of the three ensemble learning methods applied to cocoa pod

disease classification. Bagging achieved a perfect test accuracy of 1.00, clearly outperforming both Boosting (0.97) and Stacking (0.92). These results confirm Bagging’s superior generalization ability, as previously observed in the validation phase (see Figure 8), where it consistently yielded high and stable performance across all backbone models.

Boosting maintained strong performance but exhibited slight degradation, likely due to its sensitivity to noisy samples or class imbalance, which was reflected in both validation and test stages. Stacking, despite showing reasonable validation scores, had the lowest test accuracy, suggesting it may have overfit to the meta-features or lacked sufficient diversity among base learners.

Together, these findings underscore Bagging as the most robust and reliable ensemble method for cocoa pod disease classification in this study.

#### V. CONCLUSION

This study investigated the effectiveness of combining transfer learning and ensemble learning techniques for cocoa pod disease classification. By fine-tuning pre-trained CNN models—VGG16, VGG19, ResNet50, ResNet101, InceptionV3, and Xception—and integrating them into Bagging, Boosting, and Stacking frameworks, we evaluated their performance on a balanced and augmented dataset containing three disease classes. Among the ensemble methods, Bagging demonstrated the highest classification accuracy (100%), followed by Boosting (97%) and Stacking (92%). The results highlight Bagging’s superior generalization capabilities and consistent performance across different architectures. Boosting offered competitive results but showed sensitivity to class imbalance, while Stacking introduced variability likely due to the meta-model’s reliance on base learner outputs.

Overall, the proposed ensemble-based approach—particularly when leveraging Bagging—proves to be an effective and robust solution for automated cocoa pod disease detection. Future work may explore lightweight deployment for real-time UAV-based monitoring and integration of additional visual features to improve disease differentiation under real-field conditions.

#### DECLARATIONS

*Conflicts of interest/Competing interests:* No conflict of interest being declared.

#### REFERENCES

- [1] T. G. Dietterich, “Ensemble methods in machine learning,” in *International workshop on multiple classifier systems*. Springer, 2000, pp. 1–15.
- [2] G. Brown and L. I. Kuncheva, ““good” and “bad” diversity in majority vote ensembles,” in *International workshop on multiple classifier systems*. Springer, 2010, pp. 124–133.
- [3] L. Rokach, *Pattern Classification Using Ensemble Methods*. Singapore: World Scientific, 2010, vol. 75.
- [4] S. Zhang, S. M. H. Bamakan, Q. Qu, and S. Li, “Learning for personalized medicine: a comprehensive review from a deep learning perspective,” *IEEE reviews in biomedical engineering*, vol. 12, pp. 194–208, 2018.

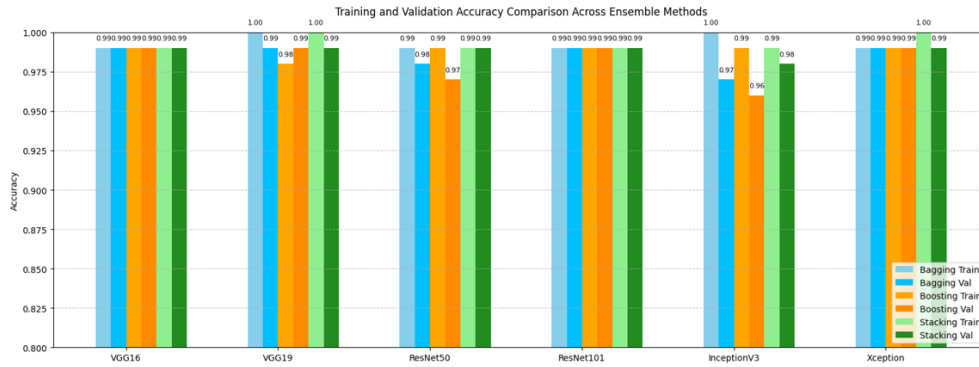


Fig. 8: Training and validation accuracy of pre-trained CNN models (VGG16, VGG19, ResNet50, ResNet101, InceptionV3, and Xception) using Bagging, Boosting, and Stacking ensemble methods.

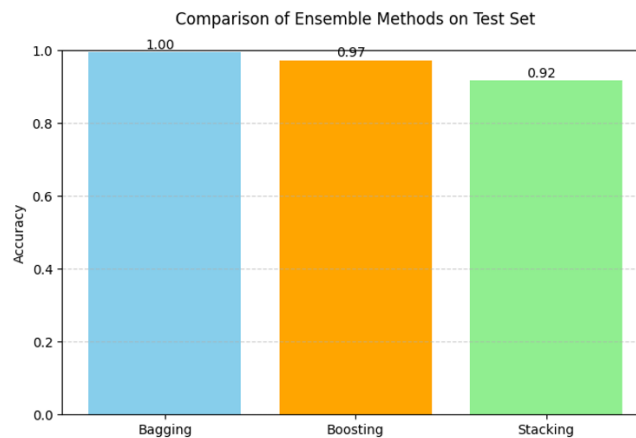


Fig. 9: Comparison of test set accuracy for Bagging, Boosting, and Stacking ensemble methods applied to cocoa pod disease classification.

- [5] P. S. Thakur, P. Khanna, T. Sheorey, and A. Ojha, "Trends in vision-based machine learning techniques for plant disease identification: A systematic review," *Expert Systems with Applications*, vol. 208, p. 118117, 2022.
- [6] C. Lersteau, T. T. Nguyen, T. T. Le, H. N. Nguyen, and W. Shen, "Solving the problem of stacking goods: mathematical model, heuristics and a case study in container stacking in ports," *IEEE Access*, vol. 9, pp. 25 330–25 343, 2021.
- [7] I. D. Mienye and Y. Sun, "A survey of ensemble learning: Concepts, algorithms, applications, and prospects," *Ieee Access*, vol. 10, pp. 99 129–99 149, 2022.
- [8] A. Ennoui, N. O. Sihamman, M. A. Sabri, and A. Aarab, "Analysis and classification of plant diseases based on deep learning," in *International Conference on Digital Technologies and Applications*. Springer, 2021, pp. 123–133.
- [9] M. Govardhan and M. Veena, "Diagnosis of tomato plant diseases using random forest," in *2019 global conference for advancement in technology (GCAT)*. IEEE, 2019, pp. 1–5.
- [10] H. Li, Y. Jin, J. Zhong, and R. Zhao, "A fruit tree disease diagnosis model based on stacking ensemble learning," *Complexity*, vol. 2021, no. 1, p. 6868592, 2021.
- [11] K. Kayaalp, "A deep ensemble learning method for cherry classification," *European Food Research and Technology*, vol. 250, no. 5, pp. 1513–1528, 2024.
- [12] A. Nader, M. H. Khafagy, and S. A. Hussien, "Grape leaves diseases classification using ensemble learning and transfer learning," *International Journal of Advanced Computer Science and Applications*, vol. 3, no. 7, pp. 563–571, 2022.
- [13] G. Fenu and F. M. Mallocci, "Classification of pear leaf diseases based on ensemble convolutional neural networks," *AgriEngineering*, vol. 5, no. 1, pp. 141–152, 2023.
- [14] A. K. Dewangan, S. Kumar, and T. B. Chandra, "Leaf-rust and nitrogen deficient wheat plant disease classification using combined features and optimized ensemble learning," *Research Journal of Pharmacy and Technology*, vol. 15, no. 6, pp. 2531–2538, 2022.
- [15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [16] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [17] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1251–1258.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.