



Deep Learning-based Bean Leaf Disease Classification: A Comparison of ResNet50 and VGG19

Yu Yan Xu¹, Hui Qing Chen², Qing Mei Lin²

10.18805/LRF-953

ABSTRACT

Background: Legumes, such as beans, are important in worldwide agriculture because of their nutritional value and soil-enriching qualities. However, bean crops are susceptible to diseases such as angular leaf spot and rust, which may reduce production and quality. Disease identification that is both effective and timely is essential to crop health and output. Traditional diagnostic approaches are often labor-intensive and susceptible to inaccuracy. Recent advances in deep learning (DL) provide interesting possibilities for automating disease categorization, possibly improving accuracy and efficiency.

Methods: This study evaluates and compares the performance of two deep learning architectures, ResNet50 and VGG19, for the classification of bean leaf diseases. The dataset, sourced from Kaggle, comprises 1295 images categorized into three classes: Angular Leaf Spot, Rust and Healthy. Both systems relied on pre-trained ImageNet weights, with adjustments customized to the classification objective. The models were trained for 25 epochs and their performance was assessed based on overall accuracy.

Result: The performance of the models is evaluated in terms of the confusion matrix, classification report and ROC(AUC) curves. The ResNet50 model achieved an overall accuracy of 93.75%, while the VGG19 model attained an accuracy of 91.41%. The findings indicate that ResNet50 performs better than VGG19 in terms of classification accuracy. This work demonstrates ResNet50's performance for bean leaf disease classification tasks, providing important insights for future research and practical applications in agricultural disease control.

Key words: Bean leaf diseases, Convolutional neural networks, Image processing, Legume crops, ResNet50, VGG19.

INTRODUCTION

Plant diseases are a major risk to global agriculture, hurting crop output and food security (Jena *et al.*, 2022). Early and precise identification of these disorders is necessary for effective care and prevention (FAO, 2023). Deep learning and computer vision breakthroughs in recent years have transformed the detection and classification of plant diseases. Bean plants are susceptible to a variety of diseases that can have a substantial influence on their health and production. Beans represent an important source of dietary protein and play a significant role in food security and agricultural sustainability in many regions of the world. Consequently, disease outbreaks in bean crops can lead to substantial economic losses and reduced agricultural productivity. Among them, Angular Leaf Spot and Rust are two major diseases that must be accurately identified for appropriate treatments. Angular Leaf Spot, caused by the bacterium *Pseudomonas syringae*, is distinguished by angular lesions on the leaves, which result in decreased photosynthetic ability and yield loss. Rust, produced by *Uromyces appendiculatus*, appears as pustules on the leaf surface, resulting in comparable decreases in plant health and production. The successful treatment of these disorders is strongly dependent on correct and prompt categorization.

Traditional detection procedures, such as eye examination and laboratory testing, are frequently laborious and time-consuming. Recent advances in deep learning provide interesting alternatives by using enormous

¹School of Quanzhou Huaguang Vocational College, Fujian, 362121, China.

²School of Fujian Agriculture and Forestry University, Fujian, 350002, China.

Corresponding Authors: Hui Qing Chen, Qing Mei Lin; School of Fujian Agriculture and Forestry University, Fujian, 350002, China. Email: 3136552527@qq.com; 3296593716@qq.com

How to cite this article: Xu, Y.Y., Chen, H.Q. and Lin, Q.M. (2026). Deep Learning-based Bean Leaf Disease Classification: A Comparison of ResNet50 and VGG19. *Legume Research*. 49(10): 1789-1797. doi: 10.18805/LRF-953.

Submitted: 01-04-2026 **Accepted:** 20-08-2026 **Online:** 01-09-2026

datasets of plant images to build models capable of autonomous disease diagnosis (Kim and AlZubi, 2024; Lugito *et al.*, 2022; Mohammad *et al.*, 2026). CNNs such as AlexNet, DenseNet, GoogleNet, ResNet50 and VGG19, have performed well in image classification challenges owing to their capacity to learn hierarchical features from raw image data. ResNet50, a deep residual network, is well-known for its unique feature of residual learning to solve the vanishing gradient issue in extremely deep networks. This design contains skip connections, which enable gradients to flow more efficiently during training, allowing the model to learn more complicated features with greater accuracy. VGG19, another notable CNN architecture, is known for its simplicity and depth, using a

succession of convolutional layers followed by max-pooling layers to capture fine-grained information.

Singh *et al.* (2023) used transfer learning to classify bean leaf diseases using three pre-trained deep learning models: MobileNetV2, EfficientNetB6 and NasNet. They used 1295 images and used optimization techniques to assess performance. EfficientNetB6 achieved the highest accuracy at 91.74%. The study highlights the importance of optimizers in improving CNN model performance, suggesting that the most effective model can help farmers in disease-prone areas. Kumar *et al.* (2020) employed residual networks (ResNet34) to identify plant leaf disease. The work intended to improve precision agriculture *via* automated disease diagnosis, making use of the worldwide rise in digital cameras and advances in computer vision technology. The ResNet34 model obtained an impressive 99.40% accuracy on the test set using an accessible dataset of 15,200 crop leaf images. Kalaivani *et al.* (2024) investigated the application of machine learning and deep learning algorithms to identify leaf diseases in tomato and potato plants. Before implementing the Support Vector Machine (SVM) and Residual Network (ResNet) algorithms, considerable data augmentation and preparation are required. The performance study demonstrates that ResNet beats SVM, with an accuracy of 94% vs 88% for SVM. As a consequence, ResNet is selected for real-time usage in detecting plant leaf diseases and guiding the application of suitable chemical fertilizers.

Tavakoli *et al.* (2021) presented a method for classifying 12 common bean cultivars using CNN model. The method uses sophisticated loss functions, such as additive angular margin loss and large margin cosine loss, to enhance class discrimination and reduce inter-class variability. The method is evaluated across three levels: species classification, cultivar classification within the same species and cultivar classification across different species. The results show that backside images achieve maximum mean accuracies of 95.86% for species, 91.37% for cultivars within the same species and 86.87% for cultivars across different species. Serttaş and Deniz (2023) explored the use of deep learning for detecting diseases in bean leaves, an important aspect of smart agriculture. They applied CNNs and transfer learning to analyze images of healthy and diseased bean leaves. The study utilized a dataset of 1295 images to train the CNN model. The proposed method, leveraging the ResNet50 model, achieved a high accuracy of 98.33% in identifying disease status. This demonstrates the effectiveness of deep learning techniques in enhancing disease detection and advancing smart farming practices. Recent studies have explored the use of AI in managing large data sets and emphasized the effectiveness of various machine learning techniques (Min *et al.*, 2024).

Although previous studies have demonstrated the effectiveness of CNN-based approaches for plant disease recognition, the reported findings vary considerably with

respect to crop species, datasets, model architectures and evaluation procedures. Consequently, direct comparisons between widely used deep learning architectures for bean leaf disease classification remain underexplored. A comparative assessment of ResNet50 and VGG19 within a common experimental framework may therefore provide additional insights into their suitability for detecting bean leaf disease.

The present study aims to evaluate and compare the efficacy of ResNet50 and VGG19 in categorizing bean leaf diseases by using a dataset obtained from Kaggle. The study encompasses a range of performance parameters. The specific objectives of this study are: (i) to develop deep learning-based classification models using ResNet50 and VGG19 architectures, (ii) to compare their performance in identifying diseased and healthy leaves and (iii) to determine the most effective model for automated bean leaf disease diagnosis. It is hypothesized that the residual learning mechanism of ResNet50 will provide superior classification performance compared with VGG19 due to its enhanced ability to extract discriminative features from complex leaf images. This study aims to evaluate these measures to establish the model that exhibits the best classification performance and to get insights into their efficacy in differentiating between different disease phases.

MATERIALS AND METHODS

Data source and its preprocessing

Data preparation

The dataset was obtained from the Kaggle repository (Abode, 2022). A total of 1,295 images were available and organized into predefined training (1,034 images), validation (133 images) and testing (128 images) subsets. These images were arranged into three subsets: training, validation and testing. To properly handle these datasets, images are loaded from designated directories and preprocessed to a standard size of 224×224 pixels. The training set is used to train the model, the validation set helps in tuning hyperparameters and performance monitoring through training and the test set is used to assess the final model's performance. Images are tagged according to their directory structure and batch sizes are limited to 32 images per batch to maximize processing performance. Categorical label encoding was employed for the three leave classes. Fig 1 shows sample images from the training dataset.

Pixel rescaling

Before feature extraction, images were processed using the model-specific `preprocess_input()` function provided by TensorFlow/Keras for both ResNet50 and VGG19 architectures. This preprocessing ensures compatibility with the ImageNet pre-trained weights used in transfer learning. The pixel values of the images were rescaled to a range of 0 to 1.

Data augmentation

Data augmentation was implemented using TensorFlow image transformation layers. The augmentation parameters included random horizontal flipping, random rotation with a factor of 0.15, random zooming with a factor of 0.10 and random translation with horizontal and vertical factors of 0.10. The augmented images were generated dynamically during training, increasing image variability and enhancing the robustness of the deep learning models.

```
data_augmentation = tf.keras.Sequential([
    tf.keras.layers.RandomFlip("horizontal"),
    tf.keras.layers.RandomRotation(0.15),
    tf.keras.layers.RandomZoom(0.10),
    tf.keras.layers.RandomTranslation(0.10, 0.10)])
```

Model architecture of ResNet50

The model is based on the pre-trained ResNet50 architecture and is known for its deep residual learning capabilities and image classification performance (Fig 2). Transfer learning was employed using ImageNet pre-trained weights. The convolutional layers of ResNet50 were frozen (non-trainable) during training, resulting in 23,587,712 non-trainable parameters and 525,315 trainable parameters in the custom classification head.

The ResNet50 layers are configured to be non-trainable, which ensures that the pre-learned features do not change throughout the training process. A new set of layers is added to the ResNet50 basis to tailor it to the particular categorization challenge. This includes:

- A Global Average Pooling layer reduces the dimensionality of the feature maps and aggregates feature data.
- A dense layer with 256 units with ReLU activation provides nonlinearity and allows for complicated pattern learning.
- To minimize overfitting, a dropout layer with a rate of 0.5 drops units at random during training.
- A final Dense layer with three units and a softmax activation function corresponds to the classification task's three classes and returns probabilities for each.

Model compilation and training

The model is built utilizing the adam optimizer with a learning rate of 0.0001. The model was trained for 25 epochs using a batch size of 32. Validation accuracy was monitored during training and the best-performing model was automatically saved using the ModelCheckpoint callback. Random seeds were fixed (NumPy = 42, TensorFlow = 42 and Python random = 42) to improve reproducibility.




Fig 1: Leaves with healthy and diseased states.




Fig 2: Architecture of ResNet50.

Model architecture of VGG19

For the model architecture, a pre-trained VGG19 network is used (Fig 3). The VGG19 model was initialized with ImageNet pre-trained weights and used as a fixed feature extractor. All convolutional layers were frozen during training, resulting in 20,024,384 non-trainable parameters and 132,099 trainable parameters in the classification layers. The VGG19 model is modified by adding a global average pooling layer, followed by a dense layer of 256 units and a dropout layer to avoid overfitting. The final output layer consists of three dense units and a softmax activation function, which correspond to the classification task's three classes.

Model compilation and training

The model is built using the Adam optimizer with a learning rate of 0.0001 and the loss function is categorical cross-entropy due to the problem's multi-class nature. Training

was conducted for 25 epochs with a batch size of 32. The best-performing model based on validation accuracy was saved using the ModelCheckpoint callback for subsequent evaluation on the test dataset.

Evaluation metrics

Accuracy evaluates the classification model's overall performance by calculating the ratio of properly predicted instances (including true positives and true negatives) to total events. It gives an overview of the performance across all classes. The accuracy formula is as follows:

$$\text{Accuracy} = \frac{(\text{True positive} + \text{True negative})}{(\text{True positive} + \text{True negative} + \text{False positive} + \text{False negative})}$$

Precision, also known as positive predictive value, is the fraction of accurately predicted positive cases (true positives) among all instances predicted as positive. It is




Fig 3: Architecture of VGG19 model.




Fig 4: Accuracy and loss functions vs epochs for both ResNet50 and VGG19 models.

important when the consequences of false positives are significant. The precision formula is as follows:

$$\text{Precision} = \frac{\text{True positive}}{(\text{True positive} + \text{False positive})}$$

Recall, also known as sensitivity or true positive rate, is a measure of the model's ability to properly identify true positive cases. It is important when the effects of false negatives are important. The formula for recall is:

$$\text{Recall} = \frac{\text{True positive}}{(\text{True positive} + \text{False negative})}$$

The F1 score is a statistic that calculates the harmonic mean of accuracy and recall to get a single rating. It strikes a compromise between accuracy and recall, which is especially effective when working with unbalanced datasets when both metrics are relevant. The F1 score is calculated as follows:

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

A confusion matrix is a method for assessing the effectiveness of a classification model by presenting a thorough breakdown of its predictions. The matrix is $n \times n$, with n representing the number of classes in the collection. Each member of the matrix gives insight into the model's performance by displaying the number of predicted vs actual classifications.

		Predicted Class	
		A	B
Actual Class	A	TP	FN
	B	FP	TP

Receiver Operating Characteristic (ROC) curves are used to assess the performance of classification models by displaying the trade-off between true positive rate (sensitivity) and false positive rate (1-specificity) at different threshold values. The true positive rate is determined as the ratio of properly recognized positive instances to the total number of positive cases, while the false positive rate is the percentage of mistakenly identified positive instances

to the total number of negative cases. By showing these rates, ROC curves give a visual depiction of a model's ability to differentiate across classes. Each point on the ROC curve represents a distinct threshold value, demonstrating the model's performance at various degrees of sensitivity and specificity. The Area Under the Curve (AUC) is a commonly used measure of model performance, with a value of 0.5 indicating no discriminative power and a value of 1.0 representing perfect classification. ROC curves and AUC are especially helpful for evaluating model performance in situations with unbalanced datasets or when comparing several models.

$$\text{ROC Curve} = \frac{\text{True positive rate}}{\text{False positive rate}}$$

To assess the robustness and generalization capability of the developed models, a 5-fold cross-validation procedure was employed. The dataset was partitioned into five subsets of approximately equal size. In each iteration, four subsets were used for training and one subset was used for validation, ensuring that each subset served as the validation set once. Model performance was evaluated using accuracy and F1-score across all folds.

To determine whether the differences in classification performance between ResNet50 and VGG19 were statistically significant, a paired Student's t-test was performed on the fold-wise accuracy values obtained during cross-validation. Statistical significance was assessed at a confidence level of 95% ($p < 0.05$).

RESULTS AND DISCUSSION

The experimental work is carried out in Jupiter notebook in the Anaconda environment. Python packages like Keras, pandas, etc. are used to develop the model. After 25 epochs of training, the ResNet50 model obtained a 0.0178 loss and a 99.81% accuracy on the training set (Fig 4). This high accuracy implies that the model was successfully trained to categorize images of bean leaf diseases and healthy leaves with minimum error on the training data. For the validation set, the model had a loss of 0.1161 and an accuracy of 93.23%. While the validation accuracy is




Fig 5: Confusion matrices.

slightly lower than the training accuracy, it performs well in differentiating between diverse classes, such as leaf diseases and healthy leaves. The difference between training and validation metrics indicates that the model generalizes effectively, while there may be potential for improvement in managing unknown data or decreasing overfitting.

After 25 epochs of training, the VGG19 model obtained a training loss of 0.3173 and an accuracy of 89.17% on the training set. This suggests that the model performed well in categorizing images of bean leaf diseases and healthy leaves, but with a larger error rate than the ResNet50 model. On the validation set, the VGG19 model had a 0.2118 loss and an accuracy of 94.74%. This greater validation accuracy compared to training accuracy shows that the VGG19 model generalizes well to new data and distinguishes well across classes. The higher validation performance over training accuracy suggests that the model is more resilient and less susceptible to overfitting than the ResNet50 model.

The ResNet50 and VGG19 models' performance was assessed using confusion matrices, which offer information on their classification accuracy for bean leaf diseases such as Angular Leaf Spot, Rust and Healthy leaves (Fig 5). The ResNet50 model exhibited good accuracy, with 41 true positives for Angular Leaf Spot, 38 for Rust and 41 for Healthy, proving its ability to discriminate between these classes with few misclassifications. It misclassified just two Angular Leaf Spots as Rust and one Healthy leaf as Rust, demonstrating its effectiveness in class identification and differentiation.

In comparison, the VGG19 model accurately recognized 40 Angular Leaf Spot, 37 Rust and 40 Healthy pictures. However, it revealed a little more misunderstanding, with three Angular Leaf Spot photos mislabeled as Rust and two Healthy images mistakenly labeled as Angular Leaf Spot. Compared to ResNet50, the VGG19 model has a greater rate of misclassification, notably between Rust and other classes.

Error analysis revealed that most misclassifications occurred between angular leaf spot and rust classes, indicating that these diseases share similar visual symptoms in certain leaf images. A small number of healthy leaves were also incorrectly classified as diseased leaves. These errors may be attributed to

variations in lighting conditions, image quality, disease severity and symptom overlap. Despite these challenges, ResNet50 exhibited fewer misclassifications than VGG19, suggesting a greater ability to extract discriminative disease-related features.

The performance of the ResNet50 and VGG19 models was evaluated using four essential classification metrics (Table 1 and Fig 6). The findings show that ResNet50 performs well in comparison to VGG19 in most tests. ResNet50 achieved 89.13% accuracy and 95.35% recall for Angular Leaf Spot, resulting in an F1-score of 92.13%, whereas VGG19 achieved slightly lower precision (86.96%) and recall (93.02%), with an F1-score of 89.89%. This demonstrates that ResNet50 is slightly more accurate in differentiating Angular Leaf Spot, despite VGG19 having a high F1-score in this class. ResNet50 performed well in the Rust classification, with an accuracy of 92.68%, recall of 88.37% and an F1-score of 90.48%. VGG19 showed comparable accuracy (92.50%) but lesser recall (86.05%), yielding an F1-score of 89.16%. This shows that, although both models perform well, ResNet50 marginally exceeds VGG19 in properly finding Rust instances.

ResNet50 attained a good precision of 100.00% and recall of 97.62% for healthy leaves, with an F1-score of 98.80%. In comparison, VGG19 achieved a precision of 95.24% and a recall of 95.24%, yielding an F1-score of 95.24%. This reveals that ResNet50 is better at recognizing




Fig 6: Comparison of classification performance metrics for ResNet50 and VGG19.

Table 1: Classification report.

	ResNet50				VGG19			
	Precision	Recall	f1-Score	Support	Precision	Recall	f1-Score	Support
0	89.13	95.35	92.13	43	86.96	93.02	89.89	43
1	92.68	88.37	90.48	43	92.50	86.05	89.16	43
2	100.00	97.62	98.80	42	95.24	95.24	95.24	42
Accuracy	93.75	93.75	93.75	128	91.41	91.41	91.41	128
Macro avg	93.94	93.78	93.80	128	91.56	91.44	91.43	128
Weighted avg	93.89	93.75	93.76	128	91.54	91.41	91.40	128

healthy leaves with greater accuracy and a better mix of precision and recall.

Overall, ResNet50 performed well than VGG19 (91.41%) with an accuracy of 93.75%. ResNet50 had greater accuracy and better macro and weighted averages (precision: 93.94%, recall: 93.78%, F1-score: 93.80%), indicating better overall performance. Although competitive, VGG19 has slightly lower average metrics (precision: 91.56%, recall: 91.44%, F1-score: 91.43%) and total accuracy. This comparison research demonstrates ResNet50's stronger classification skills, notably in recognizing Angular Leaf Spot and Healthy leaves. The good performance of ResNet50 may be attributed to its residual learning architecture, which facilitates efficient gradient propagation and enables the extraction of more representative features from complex leaf images.

The performance of the ResNet50 and VGG19 models to distinguish between classes of bean leaf disease was examined by analyzing their ROC curves (Fig 7). AUC values of 0.95 for Angular Leaf Spot, 0.92 for Rust and 0.99 for Healthy leaves were shown by the ROC curve for ResNet50. ResNet50 differentiates effectively between classes. It performed well in the classification of healthy leaves. Comparatively, VGG19 showed a slightly lower ROC curve performance, with AUC values of 0.96 for healthy leaves, 0.91 for rust and 0.93 for angular leaf spots. Although VGG19 also exhibits good performance, its AUC values are slightly lower than ResNet50's, especially when it comes to the categorization of Rust and Angular Leaf Spot. This implies

that, in comparison to ResNet50, VGG19 may have a greater false positive rate or lower sensitivity.

The ability of the developed models was further evaluated using 5-fold cross-validation. ResNet50 achieved fold accuracies of 93.75%, 93.75%, 95.09%, 95.54% and 94.62%, resulting in a mean accuracy of $94.55 \pm 0.71\%$ and a mean F1-score of $94.54 \pm 0.72\%$. In comparison, VGG19 achieved fold accuracies of 87.95%, 85.27%, 91.52%, 89.73% and 91.03%, yielding a mean accuracy of $89.10 \pm 2.28\%$ and a mean F1-score of $88.80 \pm 2.52\%$. The lower standard deviation observed for ResNet50 indicates more stable performance across different data partitions compared with VGG19.

A paired Student's t-test was conducted using the fold-wise accuracies of the two models. The analysis yielded a t-statistic of 6.014 and a p-value of 0.0039. Since the p-value was less than 0.05, the difference in classification performance between ResNet50 and VGG19 was statistically significant. These findings suggest that ResNet50 consistently provided superior classification performance for bean leaf disease identification under the experimental conditions considered in this study.

Table 2 presents a summary of the performance of several models for the categorization of bean leaf disease, comparing our findings with those of earlier studies. The comparison with other well-known models in the area shows the effectiveness of the ResNet50 and VGG19 models used in this study. Elfatimi *et al.* (2022) explored the use of MobileNetV2 model for classifying bean leaf diseases.




Fig 7: ROC(AUC) curves.

Table 2: Comparison of outputs of presented models with existing literature.

Author(s)	Model	Dataset and source	No. of images	Accuracy
Abed <i>et al.</i> (2021)	Densenet121, ResNet34, ResNet50, VGG-16, VGG-19	Extensive bean leaf dataset	1295	Binary: CAR 98.31%, Multi-class: CAR 91.01%
Elfatimi <i>et al.</i> (2022)	MobileNetV2	Kaggle	1296	<92%
Sahu <i>et al.</i> (2021)	GoogleNet, VGG16	Bean crop leaf images	1296	GoogleNet: 95.31%
Kumar <i>et al.</i> (2022)	MobileNetV2	Kaggle	1295	93.54%
This work	ResNet50; VGG19	Bean crop leaf images, Kaggle	1295	ResNet50:93.75%; VGG19:91.41%

The MobileNetV2 model achieved a classification accuracy of approximately 92% using a dataset of 1296 images. This result illustrates MobileNetV2's ability to handle bean leaf disease classification tasks well, but significantly lower than the accuracy attained by the presented ResNet50 model. Kumar *et al.* (2022) also employed MobileNetV2 to classify bean leaf diseases on a dataset of 1295 images. Their study reported a slightly higher accuracy of 93.54% for MobileNetV2 compared to Elfatimi *et al.* (2022) indicating a slight performance improvement. Abed *et al.* (2021) presented a comprehensive deep-learning framework for bean leaf disease identification using multiple models, including Densenet121, ResNet34, ResNet50, VGG-16 and VGG-19. The study used a dataset of 1295 images and evaluated the models' performance in binary and multi-class classification. Densenet121 achieved the highest accuracy with a Classification Accuracy Rate (CAR) of 98.31% for binary classification and 91.01% for multi-class classification. Among the models tested, ResNet50 achieved notable performance with a CAR of 91.01% in the multi-class classification scenario. Sahu *et al.* (2021) investigated GoogleNet and VGG16 for classifying bean leaf diseases. With a dataset of 1296 images, GoogleNet performed better than VGG16, achieving an accuracy of 95.31%. This indicates that GoogleNet was particularly effective for this classification task, surpassing VGG16 in accuracy.

In this study, two models ResNet50 and VGG19 were evaluated on a dataset of 1295 bean leaf images. The ResNet50 model achieved an accuracy of 93.75%, while VGG19 reached 91.41%. These results indicate that ResNet50 provides slightly superior performance compared to VGG19, aligning closely with the accuracy of MobileNetV2 reported by Kumar *et al.* (2022) and slightly exceeding the performance of VGG19. Overall, our results demonstrate that ResNet50 outperforms VGG19 and is competitive with other top-performing models like Densenet121 and GoogleNet in the literature. The findings highlight ResNet50's effectiveness in bean leaf disease classification, particularly in comparison to other DL models.

The findings of this study are consistent with previous research reporting the effectiveness of residual learning-based architectures for plant disease classification. Although some studies reported slightly higher accuracies using alternative architectures such as GoogleNet and DenseNet121, differences in datasets, preprocessing methods, training strategies and evaluation protocols limit direct comparison. Nevertheless, the performance achieved by ResNet50 demonstrates its competitiveness with established deep learning models reported in the literature.

From a practical perspective, the developed ResNet50 model could support automated bean disease diagnosis in precision agriculture systems. Integration of such models into mobile applications, drone-based monitoring systems, or smart farming platforms may assist farmers in early disease detection and timely intervention. Early identification of bean leaf diseases can reduce crop losses,

improve disease management practices and contribute to sustainable agricultural production.

CONCLUSION

This study can help to enhance farming processes by improving the accuracy of disease detection in bean crops, which might lead to improved disease control and crop health. The comparative investigation of ResNet50 and VGG19 for bean leaf disease classification demonstrates significant differences in the performance characteristics of the two models. ResNet50 demonstrated higher classification accuracy compared to VGG19. The limitations of the available dataset and differences in bean efficiency among conditions are among the drawbacks of the research. To fully comprehend the long-term impacts of legumes, future research should concentrate on growing longitudinal studies and creating more extensive datasets that span a range of meteorological and geographical situations.

ACKNOWLEDGEMENT

Funding details

This research received no external funding.

Author's contributions

All authors contributed toward data analysis, drafting and revising the paper and agreed to be responsible for all aspects of this work.

Disclaimers

The views and conclusions expressed in this article are solely those of the authors and do not necessarily represent the views of their affiliated institutions. The authors are responsible for the accuracy and completeness of the information provided but do not accept any liability for any direct or indirect losses resulting from the use of this content.

Data availability

The data analysed/generated in the present study will be made available from corresponding authors upon reasonable request.

Declarations

Authors declare that all works are original and this manuscript has not been published in any other journal.

Conflict of interest

The authors declare that they have no conflict of interest.

REFERENCES

- Abode, G. (2022). Bean Disease Dataset [Dataset]. Kaggle. <https://www.kaggle.com/datasets/therealoiase/bean-disease-dataset> (accessed August 1, 2024).
- Abed, S.H., Al-Waisy, A.S., Mohammed, H.J. and Al-Fahdawi, S. (2021). A modern deep learning framework in robot vision for automated bean leaves diseases detection. *International Journal of Intelligent Robotics and Applications*. **5**(2): 235-251. <https://doi.org/10.1007/s41315-021-00174-3>.

- Elfatimi, E., Eryigit, R. and Elfatimi, L. (2022). Beans leaf diseases classification using MobileNet models. *IEEE Access*. **10**: 9471-9482. <https://doi.org/10.1109/ACCESS.2022.3142817>.
- Food and Agriculture Organization (FAO). (2023). FAOSTAT Database. https://openknowledge.fao.org/server/api/core/bits_treams/28cfd24e-81a9-4ebc-b2b5-4095fe5b1dab/content/cc8166en.html [Accessed on 01/08/2024].
- Jena, J., Maitra, S., Hossain, A., Pramanick, B., Gitari, H.I., Praharaj, S. et al. (2022). Role of legumes in cropping system for soil ecosystem improvement. *Ecosystem Services: Types, Management and Benefits*. Nova Science Publishers, Inc, 415.
- Kalaivani, S., Tharini, C., Viswa, T.M.S., Sara, K.Z.F. and Abinaya, S.T. (2024). RESNET-based classification for leaf disease detection. *Journal of the Institution of Engineers (India) Series B*. <https://doi.org/10.1007/s40031-024-01062-7>.
- Kim, S.Y. and AlZubi, A.A. (2024). Blockchain and artificial intelligence for ensuring the authenticity of organic legume products in supply chains. *Legume Research*. **47(7)**: 1144-1150. doi: 10.18805/LRF-786.
- Kumar, R., Dileep, B., B., Sairam, C. and Hemanth, K. (2022). Incorporating MobileNet models into the classification of bean leaf diseases. *Ijfans International Journal of Food and Nutritional Sciences*. **11(12)**: 9596-9597.
- Kumar, V., Arora, H. and Sisodia, J. (2020). ResNet-based Approach for Detection and Classification of Plant Leaf Diseases. In: 2020 International Conference on Electronics and Sustainable Communication Systems (ICESC). IEEE. (pp. 495-502). <https://doi.org/10.1109/ICESC48915.2020.9155585>
- Lugito, N.P.H., Djuwita, R., Adisasmita, A. and Simadibrata, M. (2022). Blood pressure lowering effect of *Lactobacillus*-containing probiotic. *International Journal of Probiotics and Prebiotics*. **17(1)**: 1-13. <https://doi.org/10.37290/ijpp2641-7197>.
- Min, P.K., Mito, K. and Kim, T.H. (2024). The evolving landscape of artificial intelligence applications in animal health. *Indian Journal of Animal Research*. **58(10)**: 1793-1798. doi: 10.18805/IJAR.BF-1742.
- Mohammad, A.A.S., Mohammad, S.I., Vasudevan, A. and Heindric, A. (2026). Intelligent decision support system for crop management using data analytics. *Journal of Experimental Biology and Agricultural Sciences*. **14(3)**: 9-14. <https://jebas.org/ojs/index.php/jebas/article/view/3796/633>.
- Sahu, P., Chug, A., Singh, A.P., Singh, D. and Singh, R.P. (2021). Deep learning models for beans crop diseases: Classification and visualization techniques. *International Journal of Modern Agriculture*. **10(1)**: 796.
- Serttaş, S. and Deniz, E. (2023). Disease detection in bean leaves using deep learning. *Communications Faculty of Sciences University of Ankara Series A2-A3 Physical Sciences and Engineering*. **65(2)**: 115-129. <https://doi.org/10.33769/aupse.1247233>.
- Singh, V., Chug, A. and Singh, A.P. (2023). Classification of beans leaf diseases using fine tuned CNN Model. *Procedia Computer Science*. **218**: 348-356. <https://doi.org/10.1016/j.procs.2023.01.017>.
- Tavakoli, H., Alirezazadeh, P., Hedayatipour, A., Nasib, A.B. and Landwehr, N. (2021). Leaf image-based classification of some common bean cultivars using discriminative convolutional neural networks. *Computers and Electronics in Agriculture*. **181**: 105935. <https://doi.org/10.1016/j.compag.2020.105935>.