Nature Environment and Pollution Technology
An International Quarterly Scientific Journal
p-ISSN: 0972-6268
(Print copies up to 2016)
e-ISSN: 2395-3454
Vol. 24 | No. 2 | Article ID B4260 | 2025
Original Research Paper
https://doi.org/10.46488/NEPT.2025.v24i02.B4260
Open Access Journal

Using Deep Learning for Plant Disease Detection and Classification

G. N. Balaji†, G. Parthasarathy, A. K. P. Kovendan and Aakash Jha
School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, Tamil Nadu, India
†Corresponding author: G. N. Balaji; balaji.gnb@gmail.com

Abbreviation: Nat. Env. & Poll. Technol.
Website: www.neptjournal.com

Received: 15-08-2024
Revised: 11-10-2024
Accepted: 16-10-2024

Key Words:

InceptionNet
MobileNet
Plant disease classification
ResNet
ResNeXt

Citation for the Paper:

Balaji, G. N., Parthasarathy, G., Kovendan, A.K.P. and Jha, A., 2025. Using deep learning for plant disease detection and classification. *Nature Environment and Pollution Technology*, 24(2), p. B4260. https://doi.org/10.46488/NEPT.2025.v24i02.B4260

Note: From year 2025, the journal uses Article ID instead of page numbers in citation of the published articles.

ABSTRACT

In India's economy, farming is crucial, making early detection of plant diseases an important task. This helps in reducing crop damage and preventing the diseases from spreading further. Numerous plants, such as corn, tomatoes, and potatoes, display evident symptoms of diseases on their leaves. These noticeable patterns can be employed to accurately predict the diseases and facilitate prompt intervention to reduce their impact. The customary method involves farmers or plant pathologists visually inspecting plant leaves and identifying the specific disease. This project involves a deep learning model designed for classifying plant diseases, utilizing CNNs for their proficiency in image classification. The model, which utilizes architectures like MobileNet, InceptionNet, ResNet, and ResNeXt, delivers faster and more accurate predictions than traditional manual methods. Notably, ResNeXt, with its added dimension of cardinality that aids in learning more complex features, achieved the highest accuracy, reaching 98.2%.

INTRODUCTION

Agriculture serves as the bedrock of India's economy, playing a pivotal role in human sustenance. The scrutiny of plant health is imperative for the success of agriculture. In India, an estimated 15-25% of potential crop production is lost to pests and diseases, posing a significant challenge (Baranwal et al. 2019). Early disease detection is essential, and the manual monitoring of plants is a labor-intensive and time-consuming task that requires specialized expertise. This research focuses on efficient disease identification to enhance agricultural outcomes and secure food production.

Plant leaves are a primary indicator for identifying leaf infections, as they often display visible disease symptoms (Zhao et al. 2021). The most effective method for detecting plant infections involves recognizing various symptoms indicative of different diseases on leaves. Common manifestations of plant leaf diseases include Chlorosis, which causes leaf yellowing; Stem rust in wheat; powdery mildew, common Leaf rust in corn, Sclerotinia also known as white mold; Septoria brown spot, known as Leaf spot; and Birds-eye spot on berries caused by anthracnose.

Deep Learning, a branch of Machine Learning that traces its roots back to 1943 with the advent of threshold logic, represents a pivotal step towards developing computer models that emulate human biological processes. In this context, our operational model utilizes convolutional neural networks and transfer learning to accurately classify different plant leaf diseases. Renowned for their efficacy in image classification, CNNs-a specialized type of deep learning neural network—enhance this system (Harte 2020). This approach significantly outperforms traditional methods, offering both increased speed and accuracy in detecting disease signs on plant leaves.

The paper is structured as follows: Section II encompasses a Literature Review assessing existing research. Section III explains the Models used, detailing their design and function. Section IV, 'Experimental Detail', describes our experimental methodologies. In Section V, we present Results and Discussion, analyzing our findings with visual aids. The paper concludes with a succinct summary of our key insights in the Conclusion section.

LITERATURE REVIEW

Lakshmanarao et al. (2021) addressed the significant challenge of plant disease detection using CNNs and Kaggle's Plant Village dataset. They experiment with datasets for Potato, Pepper Bell, and Tomato plants, training and testing the CNN models. The results of these experiments are promising, demonstrating high accuracy in disease detection across all three datasets. However, the study highlights a gap in current research, particularly in the practical implementation of such models in real-world agricultural settings. Despite the high accuracy achieved, there is a need for further research to adapt these deep learning models for large-scale, field-level applications.

Elfatimi et al. (2022) proposed an automated classification system using MobileNet models. The methodology involves employing MobileNet with TensorFlow and testing various architectures and optimizers to determine the most effective configuration for disease classification. The research showcases the performance of the MobileNet model, which was trained on a publicly available dataset comprising 1,296 bean leaf images categorized into two diseased classes and one healthy class. The results indicate that the model achieves an accuracy of over 97% on the training dataset and over 92% on the test dataset. However, it's important to note that the study's scope is limited to the classification of bean leaf diseases. Consequently, the model's applicability to other scenarios remains untested. Its performance may remain high if other datasets exhibit similar distribution and sampling characteristics as the one used in this research.

Khattak et al. (2021) tackled the crucial challenge of detecting diseases in citrus fruits and leaves using a Convolutional Neural Network (CNN) model. The proposed CNN model effectively categorizes prevalent citrus diseases, including black spot, canker, scab, greening, and Melanose, by integrating multiple layers to extract distinctive features. This system surpasses numerous existing state-of-the-art deep learning models when evaluated on the Citrus and PlantVillage datasets, achieving an impressive test accuracy of 94.55%. The methodology involves using 2293 images from these datasets, pre-processed and classified into different disease categories. The research highlights the need for improved parameter and layer selection in CNN models for more accurate classification, a gap that this study addresses by experimenting with various CNN model variants and parameters.

Zhou et al. (2021) this study addresses the pressing issue of tomato leaf disease identification amid the increasing concern for food security, especially highlighted during the COVID-19 pandemic. They present an innovative approach that leverages an RRDN, which combines the advantages of deep residual networks and dense networks. This hybrid model is designed to improve accuracy in calculations and information flow while simultaneously reducing the number of training parameters. Originally developed for image super-resolution tasks, the RRDN model has been adapted for classification tasks, demonstrating an impressive top-1 average identification accuracy of 95% on the Tomato test dataset sourced from the AI Challenger 2018 datasets. This exceptional performance represents a significant advancement over most existing state-of-the-art models for crop leaf identification, all while demanding less computational resources to achieve high performance.

Sunil et al. (2022) introduced an inventive method for detecting diseases in cardamom plants by employing the EfficientNetV2 model. This study specifically addresses two significant diseases that impact cardamom plants: Colletotrichum Blight and Phyllosticta Leaf Spot. The unique aspect of this research lies in its adoption of the EfficientNetV2 model, which is a state-of-the-art deep learning model renowned for its efficiency and exceptional performance in image classification tasks. The suggested method incorporates the utilization of U2-Net for removing unwanted backgrounds from the input images and selecting multiscale features, enhancing the model's accuracy in disease detection.

Amin et al. (2022) addressed the crucial challenge of classifying corn leaf diseases in this study. The research project introduces an end-to-end deep learning model that combines two pre-trained CNNs: EfficientNetB0 and DenseNet121. The model underwent training using a subset of the PlantVillage dataset, with a specific emphasis on corn plant diseases such as northern leaf blight, common rust, and gray leaf spot. Impressively, the proposed model achieved a remarkable classification accuracy of 98.56%, surpassing the performance of other pre-trained CNN models like ResNet152 and InceptionV3.

Hassan & Maji (2022) proposed a novel lightweight CNN model for plant disease identification, focusing on reducing the number of parameters required in deep learning models. This model incorporates depthwise separable convolution, Inception, and Residual connections, significantly cutting down computational complexity and the time required for training. Traditional methods for identifying diseases in plants, which rely heavily on visual inspections conducted by experts, are both time-consuming and expensive. The implementation of automatic identification techniques using smart devices, especially deep learning models like CNNs, presents a promising solution to this challenge.

Wang et al. (2020) introduced a new database dedicated to classifying plant diseases and pests. They highlight the significance of employing deep learning models in the field of agriculture. This database was created by using keyword-based retrieval to gather images of various plant diseases and pests, resulting in a comprehensive dataset for analysis. The research underscores the critical role of early detection and treatment of these issues in enhancing crop yields. To classify plant diseases and pests, the authors have explored a hierarchical multi-task learning approach that capitalizes on the connections between different plant species and pests. To address the limitations of existing databases, the authors construct their database, selecting images based on relevance ranking from search engines. The database contains images of three different gramineous plants: rice, wheat, and corn, with each type having five diseases and one pest. The classification algorithms introduced in the research include classical CNN structures like AlexNet, VGG, and ResNet. In addition, hierarchical tree classifiers and fine-grained Triplet network structures are used to enhance the accuracy of disease and pest classification.

Rithik et al. (2022) emphasized the challenge of limited dataset availability in implementing deep learning frameworks and proposed the use of transfer learning with the VGG-19 model to address this issue. The model, pre-trained on the ImageNet and CIFAR 100 datasets, achieved a high accuracy of 94% in disease classification. The study suggests that such AI-based systems can be extended to other crops and used to evaluate crop protection models.

Pherry et al. (2023) investigated the use of transfer learning in machine learning models like VGG16, VGG19, ResNet50, and InceptionV3 to classify rice plant diseases. The study, utilizing a Kaggle dataset, found that data augmentation did not significantly improve results, while regularization techniques did. VGG19 emerged as the most accurate model with 84.4% accuracy. The paper highlights the potential of image-based machine learning over traditional methods for early and precise detection of plant health, although it acknowledges the challenge of overfitting in the validation phase and suggests future research should focus on more distinguishable datasets for improved performance.

A framework for detecting plant leaf damage was proposed by Roshini Polly et al. using a combination of pixel-level and image-level classification models. The suggested approach makes use of CNN for precise disease classification, DeepLabV3+ for background removal and disease classification assistance, and YOLOv8 for identifying regions of interest from drone photos. Additionally, they increased the evaluation accuracy by employing pixel-level UNet semantic detection. This approach not only forecasted the disease's severity but also its treatment (Polly & Devi 2024.).

A deep learning-based plant disease prediction system was proposed by Ali Hussein et al. Several deep learning architectures are combined in this ensemble technique. To improve the accuracy of plant leaf disease classification, DenseNet201, efficientNetB0, inceptionresnetV2, and efficientNetB3 have been introduced. To improve the effectiveness of deep learning models, a unique image processing method is put forth in this work. An innovative approach to the problem of plant disease classification is presented in this research. After a thorough evaluation of several image processing methods, the research presents a novel image processing algorithm (CLAHE with AMF). Then, to improve photos, a brand-new image preprocessing technique is used (Ali et al.).

MATERIALS AND METHODS

CNN models excel in recognizing and classifying objects in image datasets. However, despite their numerous advantages, CNNs do come with certain challenges. For instance, training CNNs can be time-consuming, and they often demand large datasets to perform effectively. Deep CNN models are essential for extracting detailed and minor features from images, but this requirement significantly complicates the training process. One effective way to tackle these challenges is through the use of transfer learning techniques. Transfer learning involves leveraging pre-trained networks, where the knowledge gained from training on one dataset can be applied to solve other related problems. In the next section, we explore the techniques used in this research to utilize transfer learning effectively for achieving our goals.

Transfer Learning

Training cutting-edge models typically demands considerable time, sometimes extending to days or weeks, even on advanced GPU systems. Creating a model from scratch is especially lengthy. For example, constructing a CNN model from the beginning with a typical plant disease dataset resulted in moderate accuracy after numerous training epochs. However, using a pre-trained CNN model via transfer learning achieved significantly higher accuracy in fewer epochs. Transfer learning encompasses different approaches, and their choice depends on the kind of pre-trained network model used for classification and the unique features of the dataset (Sagar & Jacob 2020).

ResNet

ResNet-50, a type of convolutional neural network, is notable for its depth, comprising 50 layers. The architecture of ResNet-50 is divided into five distinct stages, each containing convolution and identity blocks. These blocks are the fundamental components of the network, extensively utilized in a range of computer vision tasks. The key innovation introduced with ResNet is the concept of layer stacking, where convolution layers are piled on top of each other (Sharma & Singh 2021)

A significant feature of ResNet-50 is its utilization of skip connections. These connections enable the network to circumvent certain layers by directly transmitting the input to the output of a convolution layer. Skip connections are strategically positioned before activation functions to tackle the problem of vanishing gradients, which is prevalent in deep neural networks. This design helps mitigate the error accumulation that tends to occur in deeper models.

In ResNet-50, an input image is denoted as x, and the nonlinear layers that fit mappings are represented as $F(x)$. The residual mapping in this context is expressed as $H(x)$. Hence, the formula for the residual mapping in ResNet-50 is defined as:

$$H(x) = F(x) + x \quad \ldots(1)$$

Each identity block within ResNet-50 contains three convolutional layers. The network overall boasts over 23 million trainable parameters. For the skip connection to function effectively, the dimensions of the input x and the shortcut x need to align. If the dimensions differ, the shortcut x undergoes an additional convolution layer and batch normalization to ensure compatibility in dimensions (Yadav et al. 2020).

ResNeXt

ResNeXt is an advanced CNN architecture that extends the fundamental concepts of ResNet. It introduces "cardinality" as an additional dimension, alongside depth and width, in network architectures. The key innovation of ResNeXt lies in its block design. Each ResNeXt block comprises a set of transformations with the same topology, referred to as aggregated transformations.

The general formulation for a ResNeXt block can be expressed as:

$$Y = \sum_{i=1}^{C} T_i(x) \quad \ldots(2)$$

Where x is the input, Y is the output, C is the cardinality (number of parallel paths), and T_i represents the transformation function of the ith path.

This aggregated transformation approach enables the network to learn more complex features without a substantial increase in computational complexity compared to widening or deepening the network. ResNeXt models are more parameter-efficient, making them easier to train and scale. The architecture is particularly effective in image recognition tasks, where it leverages its cardinality feature to handle a wide range of image types and complexities.

InceptionNet

Images, rich in detail and varying in size, present a challenge in selecting appropriate filter sizes for feature extraction in CNNs. Smaller kernel sizes are typically used for extracting local features, while larger kernels are better suited for capturing global information. However, stacking too many convolution layers can lead to issues like overfitting and vanishing gradients. The Inception architecture addresses this by incorporating multiple kernel sizes within each block, allowing the network to become wider rather than deeper. A basic Inception module might employ filters of sizes 3x3, 1x1, and 5x5 in consecutive convolution stages. Following this, max-pooling is applied, and the results are concatenated and fed into the subsequent layer. The "stem" of the Inception layer initiates basic operations that come before the actual Inception module. Additionally, Inception V4 incorporates reduction blocks, which alter the height and width of the grid dimensions (Appe et al. 2023).

MobileNet

MobileNet is a category of efficient convolutional neural networks crafted specifically for mobile and embedded vision applications. Its core innovation lies in the use of depthwise separable convolutions, which significantly reduce the model size and computational complexity without substantially sacrificing performance. This design makes MobileNet highly suitable for environments with limited computational resources. Depthwise separable convolutions consist of two layers: a depthwise convolution, which implements a single filter for each input channel, and a pointwise convolution, which employs a 1x1 convolution to merge the outputs from the depthwise layer. This method creates a model that is both lightweight and rapid, rendering MobileNet suitable for real-time tasks such as object detection, classification, and facial recognition on mobile devices (Gupta et al. 2023).

Experimental Details

In this section, we outline the methodology adopted for implementing the plant disease classification project, highlighting the key steps from data handling to model evaluation and prediction. Fig. 1

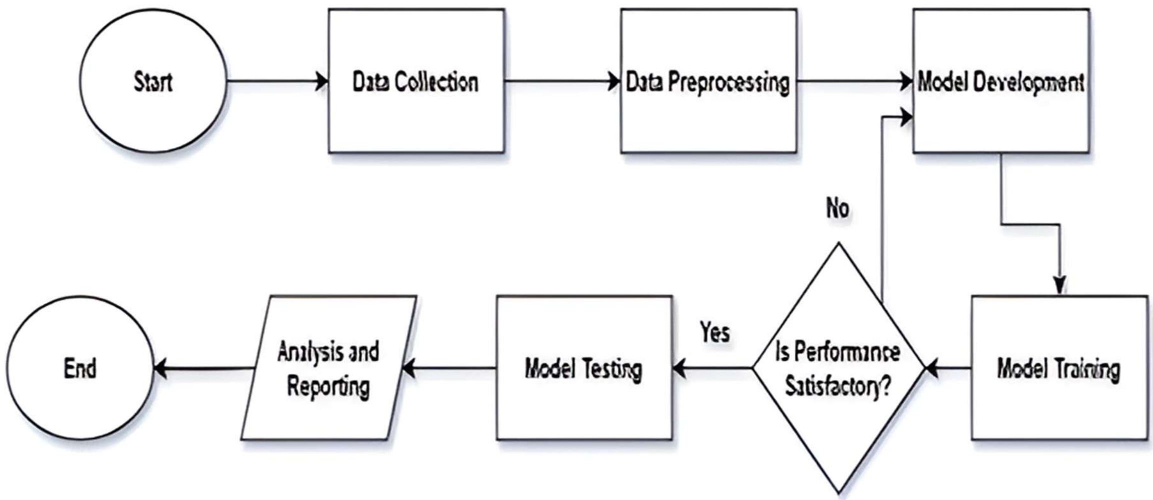

Fig. 1: Methodology flowchart.

demonstrates the methodology flowchart followed in this work.

Experimental Setup

In this study, all models were optimized for GPU use and conducted on Google's cloud platform using a Windows 11 OS system with an Intel Core i7-10750H CPU at 2.60GHz and a 16 GB Nvidia GeForce GTX GPU. This system featured a 64-bit OS with an x64-based processor. Programming was executed using the Python framework, version 3.12.0.

Description of Dataset

This study utilized data from the PlantVillage dataset modified and combined with the Chili dataset, which

Table 1: List of plants and diseases.

S. No.	Plant Name	Plant Disease Class
1.	Apple	Cedar Rust
		Scab
		Healthy
		Black Rot
2.	Cassava	Brown Streak
		Healthy
		Mosaic Disease
		Bacterial Blight
		Green Mottle
3.	Cherry	Healthy
		Powdery Mildew
4.	Chili	Healthy
		Leaf Curl
		Whitefly
		Yellowish
		Leaf Spot
5.	Corn	Leaf Spot
		Healthy
		Common Rust
		Leaf Blight
6.	Blueberry	Healthy
7.	Grape	Black Measles
		Black Rot
		Healthy
		Leaf Spot
8.	Orange	Citrus Greening
9.	Peach	Healthy
		Bacterial Spot
10.	Pepper	Healthy
		Bacterial Spot
11.	Raspberry	Healthy
12.	Potato	Late Blight
		Healthy
		Early Blight
13.	Strawberry	Healthy
		Leaf Scorch
14.	Squash	Powdery Mildew
15.	Tomato	Leaf Curl Virus
		Early Blight
		Bacterial Spot
		Mosaic Virus
		Leaf Mold
		Late Blight
		Target Spot
		Leaf Spot
		Spider Mite
		Healthy

Table cont....

includes over 50,000 images representing 47 classes across 15 plant species. To enhance performance and prevent overfitting, the images were resized and augmented using various techniques. The dataset was split into a training set containing 47,762 images and a validation set comprising 10,417 images, adhering to an 80:20 distribution ratio. Table 1 provides a detailed overview of the different plants and their respective diseases included in this study.

Preprocessing

Different preprocessing techniques, including image resizing, normalization, and data augmentation (Sudhakar & Priya 2023), were applied to guarantee uniform formatting and diversity of the images. This enhances the model's capability to generalize and minimizes the risk of overfitting during training. The pictures were resized to a uniform size of 256 x 256 pixels, as the training involved the use of various pre-trained neural network models, each with its requirements for input dimensions. For ResNet and ResNeXt, the input dimensions are 256 x 256 x 3, representing height, width, and channel depth. In contrast, InceptionNet requires an input size of 150 x 150 x 3, while MobileNet operates with an input size of 224 x 224 x 3.

Model Training

In this research, the datasets were randomly split into training and validation sets at a ratio of 80% and 20%, respectively. The training and validation datasets were exclusively utilized for training and fitting the model, while the test set was employed to assess the model's predictive performance on samples it had not previously encountered. The transfer learning approach presents the benefit of faster learning compared to models developed from the ground up. This method also allows for selective training of the model's final layers while keeping earlier layers fixed, enhancing the accuracy of classifications. In our configuration, we first standardized hyperparameters for different pre-trained models, including ResNet, ResNeXt, MobileNet, and InceptionNet (Amudha & Brindha 2022).

For each model, distinct learning rates and optimization functions were applied, as detailed in the accompanying Table 2. The training duration for each model was set at 10 epochs. During our experimentation, we observed that the output graphs started showing convergence within these 10 epochs, which helped in mitigating issues related to overfitting and model degradation.

Model Evaluation

The evaluation phase was designed to rigorously assess the models using two primary performance metrics: accuracy metric and validation loss.

Accuracy Metric: Accuracy is the principal metric, measuring the rate of correct predictions against the total predictions made. It's a direct indicator of a model's performance, with a higher accuracy reflecting a model's effectiveness in making correct classifications. The accuracy is determined by employing the following formula:

$$Acc = \frac{Number\ of\ Correct\ Predictions}{Total\ Number\ of\ Predictions} \times 100\% \quad ...(3)$$

Validation Loss: Complementing accuracy and validation loss provided insight into the model's error rate on unseen data. It served as a measure of how well the model generalized beyond the training samples. Validation Loss is often computed using a loss function, such as categorical cross-entropy for classification tasks, defined as:

$$L_{valid} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{i=1}^{M}((y_{ij}\log(pij)) \quad ...(4)$$

This formula for calculating accuracy takes into account various components:

N represents the total number of samples in the validation set, M stands for the total number of classes involved in the classification task, y_{ij} is a binary indicator signifying whether the class label j is the correct classification for observation i and p_{ij} Represents the predicted probability produced by the model, indicating the likelihood that observation i belongs to class j. A lower validation loss pointed to better generalizability and indicated that the model was learning the data patterns rather than memorizing specific examples.

RESULTS AND DISCUSSION

Fig. 2a illustrates the accuracy of the InceptionNet model across different epochs throughout both the training and

Table 2: Hyperparameter specifications.

Model	Batch Size	Input Shape	Epoch	Max Learning Rate	Optimization Function	Total Parameters (M)
ResNet	32	3×256×256	5	$1e^{-2}$	Adam	6
ResNext	32	3×256×256	10	$1e^{-3}$	SGD	23
InceptionNet	32	3×150×150	10	$1e^{-4}$	RMSprop	21
MobileNet	32	3×224×224	10	$1e^{-3}$	Adam	4

validation phases. The training accuracy shows a steady increase as the number of epochs progresses, while the validation accuracy also improves, peaking at the best validation accuracy of 90.23%. This peak indicates the model's highest performance on unseen data. As the epochs continue, both accuracy lines demonstrate learning and improvement in the model's predictions. The graph in Fig. 2b represents model loss over epochs for both the training and validation phases. The loss serves as a metric to assess the model's performance, where lower values indicate better model performance. Initially, the training loss decreases sharply and then levels off, while the validation loss decreases more gradually and shows some fluctuation, suggesting areas where the model's generalization could potentially be enhanced.

The model is effectively learning the patterns in the training data, which is why the loss decreases rapidly in the initial epochs. The validation loss typically decreases more

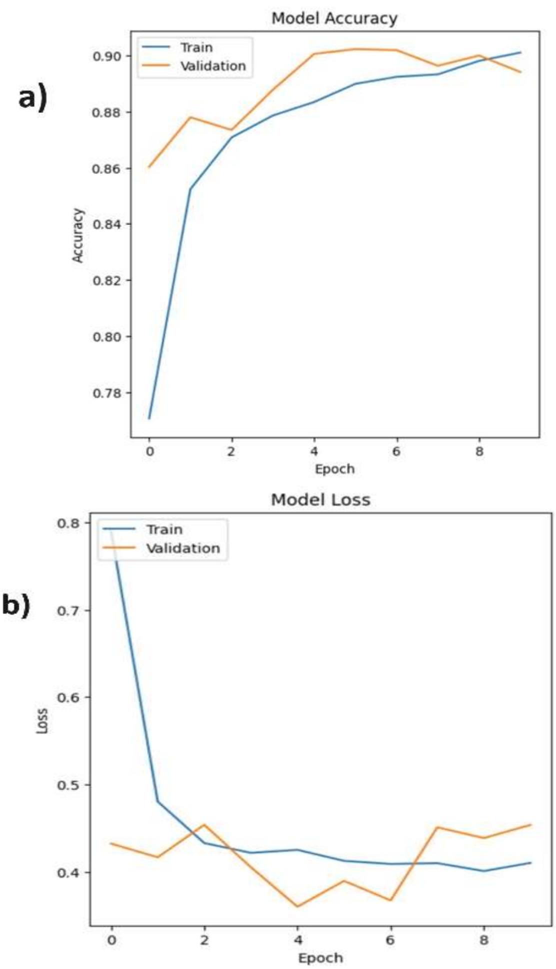

Fig. 2: Evaluation of the InceptionNet model showing (a) Training and validation accuracy, (b) Training and validation loss.

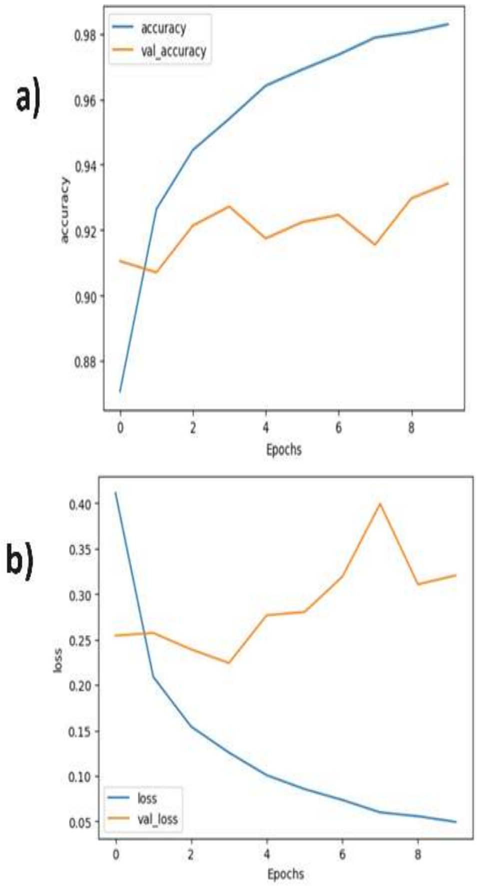

Fig. 3: Evaluation of the MobileNet model showing (a) the training and validation accuracy, (b) the training and validation loss.

slowly because the model is trying to generalize to unseen data. Validation loss shows how well the model performs on data it hasn't seen before, which is why the decrease is more gradual. Increasing the amount of training data or improving its quality can help the model generalize better.

The graph in Fig. 3a illustrates the MobileNet model's accuracy, with the training accuracy showing a consistent upward trend and the validation accuracy also improving, reaching a peak of 93.42%. The graph in Fig. 3b displays the model's loss, with training loss rapidly decreasing and leveling off, while validation loss declines with greater fluctuations, indicating some inconsistency in the model's performance on unseen data.

The accuracy graph for the ResNet model in Fig. 4a demonstrates a positive trend in learning performance across epochs. After an initial steep climb, the model's accuracy stabilizes, showcasing the effectiveness of the training process, and notably achieves a peak validation accuracy of 96.8%. The elevated validation accuracy suggests that the model possesses robustness and an aptitude for generalizing

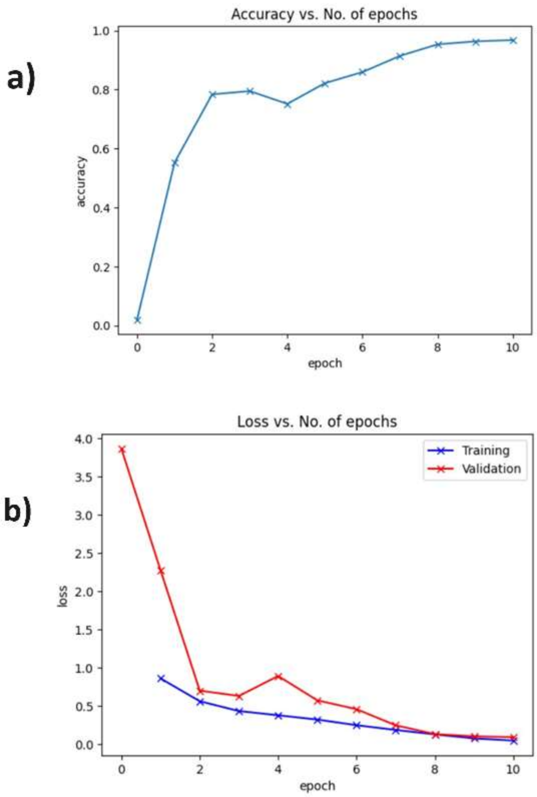

Fig. 4: Evaluation of the ResNet model showing (a) the training and validation accuracy, (b) the training and validation loss.

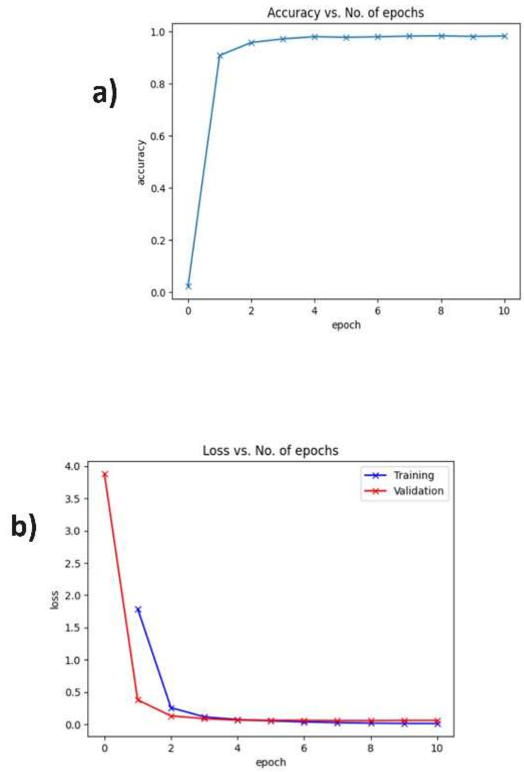

Fig. 5: Evaluation of the ResNeXt model showing (a) Validation accuracy vs Epoch (b) Training and validation loss.

effectively to new, unseen data. The loss graph in Fig. 4b complements this by showing a rapid decrease in training loss, which is a sign that the model effectively minimizes errors during training. The validation loss follows a downward trend with some fluctuations, which is common in model training dynamics. These fluctuations are minor, and the overall trajectory indicates good convergence without signs of overfitting, affirming the model's strong predictive capability throughout its training epochs.

The accuracy graph for the ResNeXt model in Fig. 5a demonstrates excellent learning efficiency, with a swift rise to high accuracy levels after the initial epoch and sustaining a commendably high accuracy thereafter. Notably, the model attained a peak validation accuracy of 98.2%, reflecting its superior capability in generalizing from the training data to unseen data. In parallel, the loss graph in Fig. 5b presents a favorable decrease in both training and validation loss, suggesting that the model is accurately capturing the underlying patterns without overfitting. The convergence of loss values at low levels towards the end of the training indicates that the model's performance is both stable and highly optimized.

The graphs in Fig. 6 depict the learning rate schedules for the ResNet and ResNeXt models across various batches of training. Both show a learning rate that initially ramps up, peaks at an intermediate point, and then decays. This pattern follows a cyclical learning rate policy, where the learning rate is increased to potentially help the model escape local minima and then decreased to allow for finer convergence to a minimum. The shape of the curve is smooth

Table 3: Showing comparative performance analysis of the four models used.

Model	Accuracy at the last epoch	Loss at last epoch	Train Loss Trend	Valid Loss Trend	Accuracy Stability
ResNet	0.9680	0.0925	Decreasing	Decreasing	High
ResNeXt	0.9818	0.0607	Decreasing	Decreasing	High
InceptionNet	0.8941	0.4536	Decreasing to stable	Fluctuating	Moderate
MobileNet	0.9342	0.3205	Decreasing	Fluctuate	Moderate

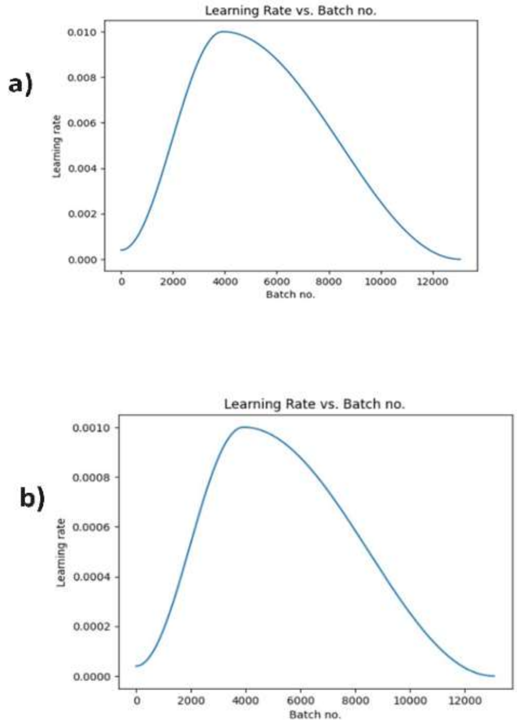

Fig. 6: Graphs showing the learning rate schedule for (a) the ResNet model and (b) the ResNeXt model.

and symmetrical for both models, suggesting a well-tuned learning rate strategy that could contribute to effective training and optimization of the neural networks.

Early detection of crop diseases plays a crucial role in achieving high agricultural yields. To ensure optimal productivity, it's vital to incorporate advanced technologies in the early identification of plant diseases. Deep learning models have been recognized for their effectiveness in image classification through a review of relevant literature. Specifically, models built upon transfer learning are well-known for their efficiency in simplifying the training process and diminishing the necessity for extensive datasets. In this study, our primary goal was to evaluate the performance of four pre-trained models—ResNet, ResNeXt, InceptionNet, and MobileNet—to identify the most effective model for the classification of various plant diseases. We evaluated these models using metrics such as accuracy and validation loss. Notably, the ResNeXt model demonstrates superior results, likely due to its more complex architecture that allows for better feature representation and generalization, especially in handling diverse plant disease image data. The validation accuracy for each model was computed, and a visual representation of these accuracies is presented in Table 2. Our analysis showed that the ResNeXt model (Fig. 5) surpassed the other models (Figs. 2-4) in performance, achieving the highest validation accuracy of 0.982. From Table 3, we can see that the accuracy stability is high, indicating consistent performance across epochs, and the loss at the final epoch is also the lowest among the three models, which is generally a good indicator of model performance.

CONCLUSIONS

In this study, our main focus was on conducting a comparative analysis of different transfer learning models. The goal was to determine the most efficient architecture for accurately classifying a total of 47 distinct plant disease categories. Our comparative results revealed that ResNeXt surpassed the performance of ResNet, InceptionNet, and MobileNet in terms of accuracy. Therefore, ResNeXt is preferable for the task of identifying plant diseases, especially when integrating new diseases into the model, as it shows a lower complexity in training. The model we proposed attained a noteworthy classification accuracy of 98.2%. In future research, we aim to address the challenges of creating a deep learning model that can identify plant diseases from clusters of leaves, moving beyond the analysis of individual leaves. In addition, future research will focus on how meteorological conditions such as wind, rain, or dust affect leaf appearance while forecasting plant diseases. Another issue that needs to be focused on is the possibility of images containing extraneous background elements, which lowers the classification accuracy.

REFERENCES

Ali, A.H., Youssef, A., Abdelal, M. and Raja, M.A., 2024. An ensemble of deep learning architectures for accurate plant disease classification. *Ecological Informatics*, 81, pp.1–10.

Amin, H., Darwish, A., Hassanien, A.E. and Soliman, M., 2022. End-to-end deep learning model for corn leaf disease classification. *IEEE Access*, 10, pp.31103–31115.

Amudha, M. and Brindha, K., 2022. Multi techniques for agricultural image disease classification and detection: A review. *Nature Environment & Pollution Technology*, 21, pp.100–110.

Appe, S.N., Arulselvi, G. and Balaji, G.N., 2023. Detection and classification of dense tomato fruits by integrating coordinate attention mechanism with the YOLO model. In: *Handbook of Research on Deep Learning Techniques for Cloud-Based Industrial IoT*. IGI Global, pp.278–289.

Baranwal, S., Khandelwal, S. and Arora, A., 2019. Deep learning convolutional neural network for apple leaves disease detection. *SSRN Electronic Journal*, pp.1–8.

Elfatimi, E., Eryigit, R. and Elfatimi, L., 2022. Beans leaf disease classification using MobileNet models. *IEEE Access*, 10, pp.9471–9482.

Gupta, A., Yadav, D., Raj, A. and Pathak, A., 2023. Real-time object detection using SSD MobileNet model of machine learning. *International Journal of Engineering and Computer Science*, 12, pp.25729–25734.

Harte, E., 2020. Plant disease detection using CNN. *ResearchGate*, 8(1), pp.45–52.

Hassan, S.M. and Maji, A.K., 2022. Plant disease identification using a novel convolutional neural network. *IEEE Access*, 10, pp.5390–5401.

Khattak, A., Khan, A., Rehman, M., Ahmed, F. and Ali, S., 2021. Automatic detection of citrus fruit and leaves diseases using deep neural network model. *IEEE Access*, 9, pp.112942–112954.

Lakshmanarao, M.R.B. and Kiran, T.S.R., 2021. Plant disease prediction and classification using deep learning ConvNets. In: *Proceedings of the International Conference on Artificial Intelligence and Machine Vision (AIMV)*. Gandhinagar, India: IEEE, pp.1–6.

Pherry, F., Kristanto, G.J. and Kurniadi, F.I., 2022. Rice plants disease classification using transfer learning. In: *Proceedings of the 4th International Conference on Cybernetics and Intelligent Systems (ICORIS)*. Prapat, Indonesia: IEEE, pp.1–4.

Polly, R. and Devi, E.A., 2024. Semantic segmentation for plant leaf disease classification and damage detection: A deep learning approach. *Smart Agricultural Technology*, 9, pp.1–12.

Rithik, J.P., Devi, S.R.S., Brendan, M.A. and R., S., 2023. Plant disease classification using deep learning approach (VGG19). In: *Proceedings of the 9th International Conference on Advanced Computing and Communication Systems (ICACCS)*. Coimbatore, India: IEEE, pp.1715–1718.

Sagar, A. and Jacob, D., 2020. On using transfer learning for plant disease detection. *ResearchGate*, 5(2), pp.12–18.

Sharma, V. and Singh, N., 2021. Deep convolutional neural network with ResNet-50 learning algorithm for copy-move forgery detection. In: *Proceedings of the 7th International Conference on Signal Processing and Communication (ICSC)*. Noida, India: IEEE, pp.146–150.

Sudhakar, M. and Priya, R.M., 2023. Computer vision based machine learning and deep learning approaches for identification of nutrient deficiency in crops: A survey. *Nature Environment & Pollution Technology*, 22(3), pp.145–152.

Sunil, C.K., Jaidhar, C.D. and Patil, N., 2022. Cardamom plant disease detection approach using EfficientNetV2. *IEEE Access*, 10, pp.789–804.

Wang, Q., He, G., Li, F. and Zhang, H., 2020. A novel database for plant diseases and pest classification. In: *Proceedings of the IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC)*. Macau, China: IEEE, pp.1–5.

Yadav, D., Jalal, A.S., Garlapati, D., Hossain, K., Goyal, A. and Pant, G., 2020. Deep learning-based ResNeXt model in phycological studies for future. *Algal Research*, 50, pp.1–10.

Zhao, Y., Li, J., Xu, F. and Kim, D., 2021. Plant disease detection using generated leaves based on DoubleGAN. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18(6), pp.1–1.

Zhou, C., Zhou, S., Xing, J. and Song, J., 2021. Tomato leaf disease identification by restructured deep residual dense network. *IEEE Access*, 9, pp.28822–28831.