International Journal of Science and Research (IJSR)
ISSN: 2319-7064
SJIF (2022): 7.942

Advancing Plant Disease Detection: Harnessing Deep Learning and Machine Vision

Nidhi Singh

M. Tech – AI & ML, School of Computer Science and Engineering, Vellore Institute of Technology, Vellore – 632014, Tamil Nadu, India
Email: nidh*isingh8888888[at]gmail.com*

Abstract: ***This research explores the application of deep learning and machine vision in detecting plant leaf diseases in agricultural settings, specifically focusing on datasets from farm villages. By combining real farm village data with synthetic data generated by Generative Adversarial Networks (GANs), three advanced convolutional neural network (CNN) models VGG16, ResNet50, and InceptionNet V3 are utilized through transfer learning. Transfer learning enhances model performance by fine-tuning pre-trained networks. The study evaluates the models systematically using metrics such as accuracy, precision, recall, and F1 score. The findings demonstrate the effectiveness of the methodology, with ResNet50 achieving the highest performance at 83.23%. This research contributes to the advancement of precision agriculture, offering promising implications for sustainable farming practices and optimizing crop yields.***

Keywords: Plant leaf diseases, Convolutional neural networks, Transfer learning, Generative Adversarial Networks, Machine Vision

1. Introduction

In modern agriculture, timely and precise detection of plant diseases is essential for maintaining crop health and maximizing yields. Plant leaf diseases represent a significant threat to global food security, highlighting the need for innovative approaches to early detection and management. This research tackles these challenges by investigating the integration of deep learning and machine vision technologies for automating the identification of plant leaf diseases, particularly in the complex setting of agricultural communities, including farm villages. By employing advanced convolutional neural network (CNN) architectures such as VGG16, ResNet50, and InceptionNet V3 with transfer learning, the study aims to improve the accuracy and efficiency of disease identification. Traditional methods reliant on manual inspection by agronomists are subjective and limited in scalability, underscoring the importance of adopting deep learning and machine vision techniques for automated, rapid, and accurate diagnosis.

The study acknowledges the intricate nature of farm village environments, which involve various crop types, environmental elements, and farming techniques, necessitating an adaptable and robust approach. Transfer learning is utilized to customize pre-trained models to the specific characteristics of farm village datasets, mitigating the shortage of labeled data tailored to these environments. Furthermore, the research introduces the application of Generative Adversarial Networks (GANs) to produce synthetic data, supplementing the training dataset and improving the models' ability to generalize in real-world agricultural scenarios. Its objective is to advance the comprehension of plant leaf disease detection in rural communities by developing, testing, and evaluating deep learning models that contribute to precision agriculture. Through the assessment of ResNet50, InceptionNet V3, and VGG16 in the context of farm villages, leveraging transfer learning and GAN-generated data, the study aims to offer valuable insights into tackling the unique challenges associated with identifying agricultural diseases and promoting sustainable farming practices.

2. Literature Review

In Pranesh Kulkarni et al.'s investigation titled "Plant Disease Detection Using Image Processing and Machine Learning," although they achieved an impressive accuracy of 93%, it's essential to recognize limitations such as difficulties in extrapolating results to novel datasets and coping with diverse environmental conditions [1]. The effectiveness of the model could be affected by the evolving nature of plant diseases and the computational requirements in extensive agricultural operations. Moreover, depending solely on image-based detection might neglect other significant factors such as weather patterns or variations in soil composition.

Yan Guo and colleagues' study titled "Plant Disease Identification Based on Deep Learning Algorithm in Smart Farming," presents a mathematical model that utilizes deep learning to enhance the efficiency of plant disease detection in smart farming applications [2]. While achieving an accuracy of 83.57% and proving effective against specific diseases, Yan Guo et al.'s research highlights challenges associated with the iterative nature of the Chan–Vese algorithm. This suggests the need for potential enhancements to expedite identification processes in future research endeavors.

Andrew J. et al. research on "Deep Learning-Based Leaf Disease Detection in Crops Using Images for Agricultural Applications" underscores the pivotal importance of the agricultural sector. The study delves into the utilization of CNN-based pre-trained models for detecting leaf diseases in crops [3]. DenseNet-121 achieves an outstanding classification accuracy of 99.81%, showcasing the substantial potential of deep learning in improving the early diagnosis of plant diseases. Future research endeavors will focus on tackling challenges related to real-time data collection and the development of a multi-object deep learning model.

Sharada P. Mohanty et al. research on deep learning-based plant disease detection attains a notable accuracy rate of 99.35% [4]. Nevertheless, the study faces limitations such as decreased accuracy when tested under diverse conditions and challenges in accurately classifying single leaves against homogeneous backgrounds. Ongoing research endeavors seek to mitigate these drawbacks to enable practical real-world applications in agriculture.

Riyao Chen and colleagues' study introduces CACPNET as a promising model for plant disease identification, demonstrating high accuracy. However, the study acknowledges limitations in the attention mechanism and highlights the trade-off between accuracy and the demands of real-time deployment [5]. Despite its limitations, CACPNET exhibits notable advantages, particularly its potential for lightweight deployment in precision agriculture.

Kowshik B and colleagues' literature survey on "Plant Disease Detection Using Deep Learning" emphasizes the importance of agriculture and highlights the potential of deep learning techniques to enhance accuracy in disease detection [6]. While emphasizing the positive effects on early disease detection, the review acknowledges the necessity for further research to address existing gaps in disease detection transparency. Additionally, it proposes future expansions to incorporate additional features such as pesticide price lists and market information. However, specific drawbacks or limitations are not explicitly outlined in the review.

In summary, these studies collectively contribute to the progression of plant disease detection through the utilization of deep learning and machine vision techniques. While demonstrating impressive accuracies and potential applications, each study acknowledges specific limitations. This underscores the ongoing necessity for research and enhancement in the crucial domain of precision agriculture.

3. Materials and methods

a) Data Acquisition

The study utilizes the PlantVillage dataset, which comprises 19,458 meticulously selected photos aimed at enhancing plant disease identification through computer vision and deep learning. To enhance dataset diversity, data augmentation techniques and a Generative Adversarial Network (GAN) model contribute an additional 9,973 augmented photos. In total, the dataset comprises 29,928 photos categorized into 20 different classifications representing various plant health and disease categories. This dataset offers a comprehensive depiction of plant conditions and covers a variety of crops, including blueberries, apples, cherries, corn, grapes, potatoes, peppers, and strawberries.

b) Data Preparation

In this phase, the dataset containing 29,928 leaf images is split into training and testing sets with a ratio of 65% for training and 35% for testing. The training set consists of 19,458 plant leaf images categorized into 20 classes, representing diverse plant species and health or disease conditions. This dataset is strategically divided, with 85% (16,545 images) allocated for training and 15% (2,911 images) for validation, ensuring robust model development and evaluation.The testing set comprises a subset of 9,973 images distributed across the same 20 classes, with each class representing a specific category of plant health or disease. The distribution of images per class is as follows:

Table I: Dataset Details

No.	Class Names	Images count
0	Blueberry___healthy	2002
1	Apple___Cedar_apple_rust	777
2	Apple___healthy	2146
3	Cherry_(including_sour)___healthy	1363
4	Cherry_(including_sour)___Powdery_mildew	1564
5	Corn_(maize)___Cercospora_leaf_spot Gray_leaf_spot	1015
6	Corn_(maize)___Common_rust_	1707
7	Corn_(maize)___healthy	1662
8	Corn_(maize)___Northern_Leaf_Blight	1485
9	Grape___Black_rot	1681
10	Grape___Esca_(Black_Measles)	1884
11	Grape___healthy	916
12	Grape___Leaf_blight_(Isariopsis_Leaf_Spot)	1577
13	Pepper,_bell___Bacterial_spot	1487
14	Pepper,_bell___healthy	1976
15	Potato___Early_blight	1476
16	Potato___healthy	552
17	Potato___Late_blight	1496
18	Strawberry___healthy	956
19	Strawberry___Leaf_scorch	1609

c) Image preprocessing and Data Augmentation

In the model training process, a batch size of 64 images per iteration is specified to efficiently process and optimize the deep learning model. Two distinct image data generators are defined: one for training and validation (train_generator) and another for testing (test_generator). The training and validation data generator incorporates various transformations, including a rotation range of 90 degrees, varied brightness between 0.1 and 0.7, horizontal and vertical shifts within a range of 0.5, as well as horizontal and vertical flips. A validation split of 15% is employed to allocate a portion of the training data for validation purposes. The VGG16 preprocessing function is applied to enhance data compatibility. Using flow_from_directory, batches of training and validation data are created from the original and augmented datasets, specified by the directories train_data_dir and test_data_dir. The class_subset, containing the list of class names, is retrieved from the original dataset. For testing data, the test_generator employs the VGG16 preprocessing function, pulls batches from the augmented dataset directory, and resizes images to (299, 299). With a batch size of 1, images are processed independently to ensure consistent evaluation without data scrambling, while reproducibility is maintained through the use of a seed. These generators collectively contribute to the model's ability to generalize effectively and recognize diverse patterns, thereby enhancing its performance across different datasets.

4. Experimental Procedure

The development of a deep learning model for image classification involves a carefully orchestrated series of steps

International Journal of Science and Research (IJSR)
ISSN: 2319-7064
SJIF (2022): 7.942

to ensure robust performance and interpretability. Beginning with loading and partitioning the dataset into test, validation, and training sets, diversity is introduced through data augmentation techniques applied specifically to the training set. Efficient batch processing is facilitated by a data loader, optimizing the training process by providing the model with batches of augmented data. Following this, essential image preparation steps such as pixel value normalization and scaling are performed to standardize input data, ensuring consistency and aiding in model convergence.

Utilizing transfer learning, pre-trained models like VGG16, ResNet50, or InceptionV3 serve as the foundation of the architecture, with the feature extraction phase initiated by removing the classification head of the pre-trained model. The model is then fine-tuned for task-specific classification by adding a custom classification head. Following this, the entire model is compiled, specifying the optimizer, loss function, and metrics for training. The training strategy involves initially freezing convolutional layers, gradually unfreezing them, and monitoring performance on the validation set. Techniques such as early stopping and learning rate schedules are implemented to optimize the training process and prevent overfitting. Following model training, fine-tuning strategies are explored to develop a versatile model suitable for deployment in inference or storage. Detailed insights into hyperparameter selections, such as learning rates and batch sizes, are crucial for understanding the model's sensitivity. Additionally, the integration of regularization techniques, such as dropout rates or weight decay, not only helps mitigate overfitting but also enhances model interpretability.

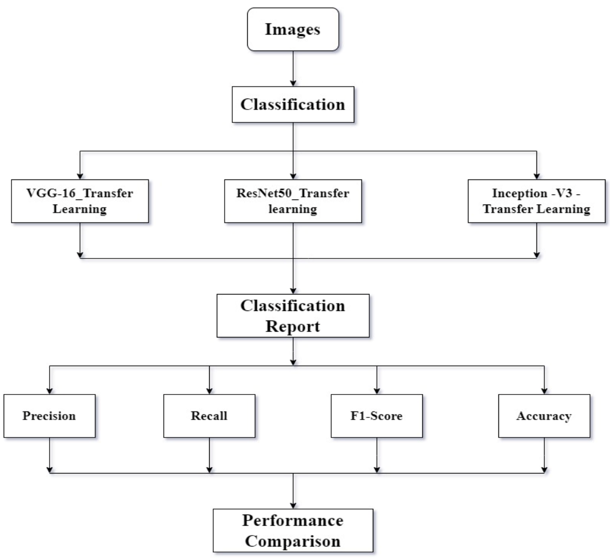

Figure I: Proposed Architecture

The systematic approach outlined ensures the adaptability of image classification models to diverse datasets and task requirements, promoting transparency and reproducibility, which are essential for guiding practitioners, particularly in the domain of plant disease detection. Transparency in model configuration not only enhances the reproducibility of the study but also provides valuable guidance for practitioners seeking to implement or extend similar approaches in this field. With a focus on adaptability, this systematic methodology addresses the nuanced challenges posed by different datasets and specific classification tasks in image analysis. It serves as a reliable reference for practitioners in the complex field of plant disease detection, facilitating the effective implementation or extension of similar approaches.

The commitment to transparency in model configuration enhances the reliability of results and promotes a standardized approach for future research and applications in image classification, especially within the context of agricultural and plant health monitoring.

a) *The ResNet50 Transfer Learning model*

The model architecture employs the ResNet50 convolutional neural network, which has been pre-trained on ImageNet, for plant disease detection. The top layers of ResNet50 are excluded, repurposing the network as a feature extractor. A Global Average Pooling layer is utilized to condense features, followed by a Dense layer with Softmax activation for multi-class classification. The base weights of ResNet50 remain non-trainable, leveraging the prior knowledge gained from ImageNet to enhance performance. Stochastic Gradient Descent optimizes the model with carefully adjusted parameters.

For multi-class classification, categorical crossentropy serves as the loss function, along with relevant evaluation metrics. Training is performed iteratively using generators, and critical callbacks are implemented to ensure optimal model weights and early termination if necessary. The resulting preserved model provides a robust solution for accurate image categorization.

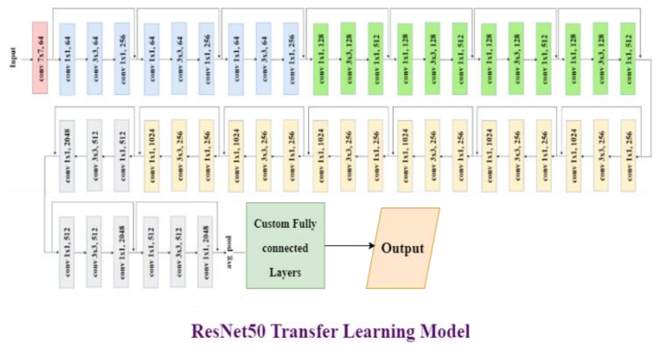

Figure II: The ResNet50 Transfer Learning Model

b) *The VggNet16 Transfer Learning model*

The model architecture for plant disease classification is constructed based on the VGG16 convolutional neural network, leveraging transfer learning with pre-trained weights from ImageNet. The adapted model incorporates a custom classification head and offers fine-tuning options for selective training.The classification head comprises two densely connected layers, a dropout layer, and a softmax output layer.The model is compiled with categorical crossentropy loss, an optimizer (initially 'rmsprop'), and the accuracy metric. Early stopping and learning rate adjustments are implemented to enhance training convergence.

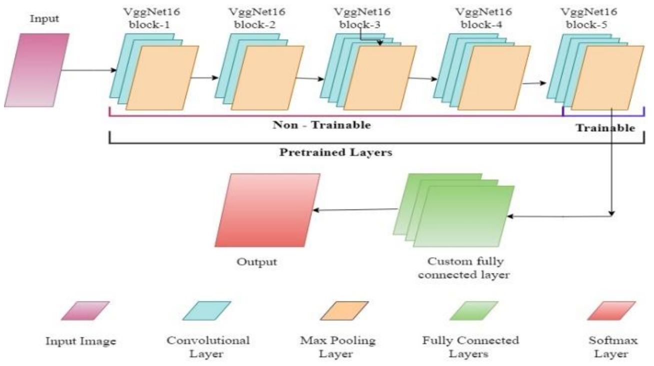

Figure III: The VggNat16 Transfer Learning Model

c) *The InceptionNet V3 Transfer Learning model*

Leveraging the InceptionV3 architecture with pre-trained weights from ImageNet, a plant disease identification model is developed. Customizable fine-tuning, which specifies trainable layers, aims to strike a balance between adapting to plant disease datasets and leveraging pre-trained knowledge. The model comprises global average pooling, a densely connected layer, and a dropout layer to address overfitting concerns.The output layer utilizes softmax activation to classify outputs into probability distributions across plant disease classes. Assembled with the Adam optimizer, categorical crossentropy loss, and accuracy metric, the model is configured for training.

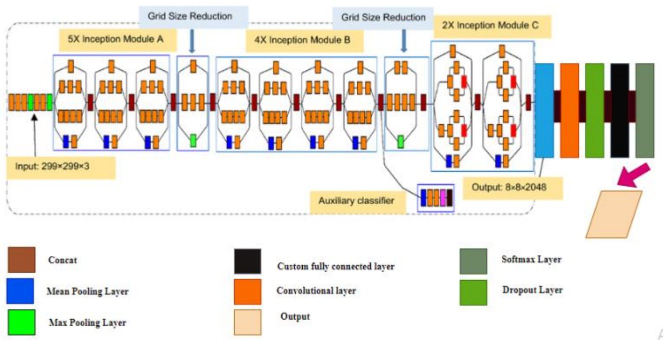

Figure IV: The InceptionNet V3 Transfer Learning Model

5. Results and Conclusion

The experimental results reveal distinct performance characteristics among the evaluated pre-trained models utilizing transfer learning for image classification. ResNet50 emerged as the most impressive performer, achieving the highest accuracy of 83.23%. This finding aligns with the general trend observed in the literature, which underscores the efficacy of deeper architectures, particularly those incorporating residual learning, for enhancing feature representation and classification accuracy. Furthermore, the accuracy of the VGG16 model notably improved to 79.27% after fine-tuning. This outcome underscores the flexibility of VGG16 when task-specific modifications are applied, corroborating previous studies advocating for fine-tuning as an effective approach for enhancing trained models.

InceptionV3 demonstrated competitive performance both with and without fine-tuning. With an accuracy of 76.43%, the untuned InceptionV3 showed a trade-off between computing efficiency and accuracy. The refined InceptionV3 further indicated its adaptability by achieving an accuracy of 74.82%, striking a balance between model complexity and performance.

These findings provide practical guidance for practitioners on selecting a model based on the requirements of a given task and available computational resources. ResNet50 is well-suited for scenarios where high accuracy is paramount, while VGG16 offers flexibility for task-specific optimization due to its fine-tuning capabilities. In terms of efficiency, InceptionV3 emerges as a reasonable choice due to its well-balanced performance.

Overall, this comprehensive research contributes valuable insights to the broader field of pre-trained models, enabling practitioners to make informed decisions tailored to their specific image classification applications.The comparison study on transfer learning-based pre-trained models for plant disease detection—ResNet50, VGG16, and InceptionV3—revealed distinct performance characteristics. ResNet50 emerged as the top performer with an accuracy of 83.23%, showcasing the effectiveness of deep residual learning. VGG16 demonstrated its flexibility through fine-tuning, resulting in a notable accuracy boost to 79.27%. InceptionV3, both with and without fine-tuning, demonstrated a balance between computing efficiency and accuracy, achieving accuracies of 76.43% and 74.82%, respectively.

These findings provide valuable guidance to practitioners in selecting models based on task demands and available computational resources. The research contributes significant insights to the field, informing decisions for practical image classification applications. Additionally, it suggests future directions for enhancing model flexibility, fine-tuning strategies, and generalization skills.

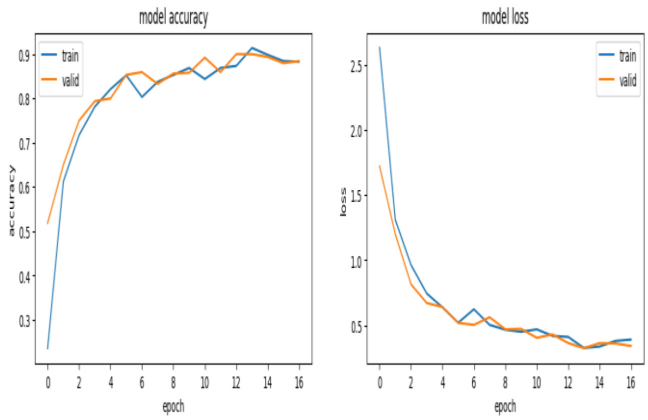

Figure V: Accuracy and loss plot diagram of ResNet50

Table 2: Accuracy and loss result of ResNet50 model

Parameter name	Accuracy (%)	Loss
Training	88.28	0.3898
Validation	88.44	0.3409
Test	83.23	--------

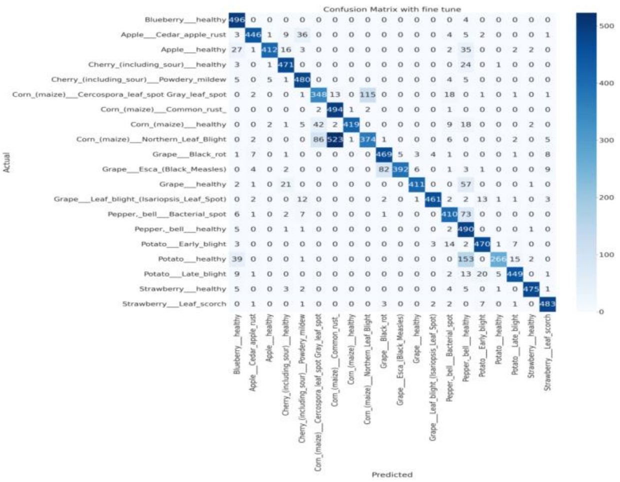

Figure VI: Confusion Matrix of ResNet50

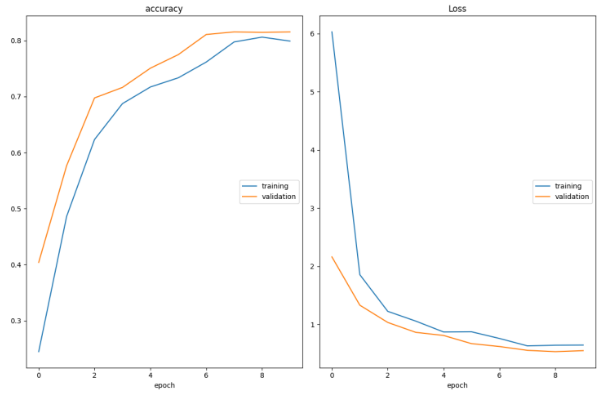

Figure VII: VggNet16 Training and Validation Accuracy and Loss over Epochs

Table 3: Accuracy and loss result of VggNet16 model

Parameter name	Accuracy (%)	Loss
Training	79.92	64.67
Validation	81.56	55.14
Test	79.27	-

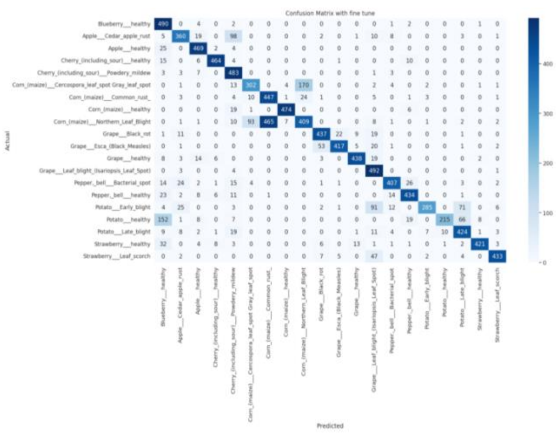

Figure VIII: VggNet16 Confusion Matrix

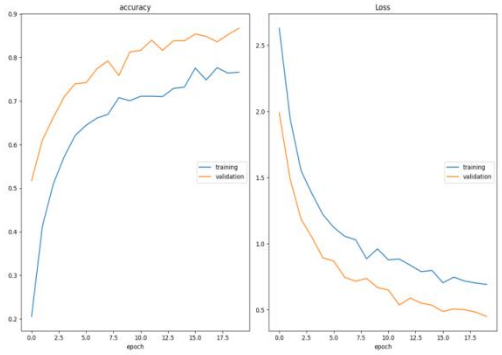

Figure IX: InceptionNet V3 Training and Validation Accuracy and Loss over Epochs

Table 4: Accuracy and loss result of InceptionNet V3 model

Parameter name	Accuracy (%)	Loss
Training	73.12	0.7790
Validation	88.13	0.3089
Test	76.43	---------

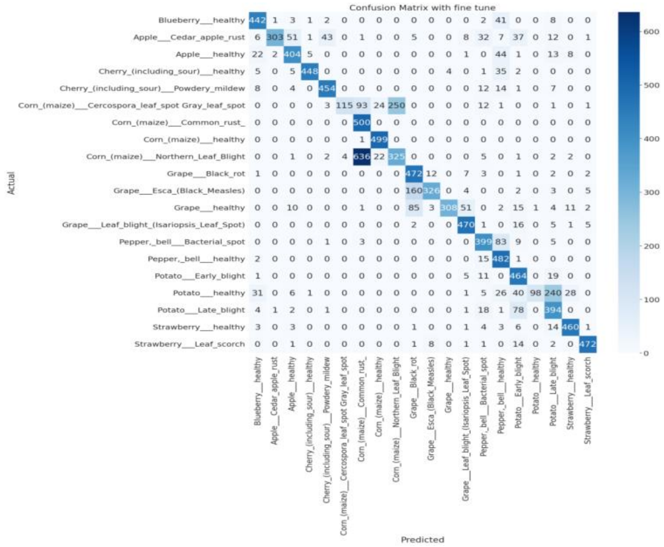

Figure X: InceptionNet V3 Confusion Matrix

6. Limitation and Future Scope

The systematic approach outlined for developing adaptable image classification models in the realm of plant disease detection provides a comprehensive framework. However, it's essential to acknowledge certain limitations and outline future directions for continued improvement.

Limitations include potential biases within the training data, which may hinder the model's ability to generalize effectively to unseen environmental conditions or emerging plant diseases. Additionally, the computational demands of training deep learning models, especially those with intricate architectures, may pose accessibility challenges for researchers with limited computational resources. Furthermore, the interpretability of deep learning models remains a challenge due to their inherent complexity.

Looking ahead, future directions involve exploring multi-modal fusion techniques to incorporate data from enhancing the diverse sources such as spectral and textual information, model's robustness across varied agricultural scenarios. Refining transfer learning strategies, particularly through domain adaptation tailored to agricultural contexts, is crucial for improving model performance.

The prospect of real-time deployment in agricultural settings, while considering computational constraints, is a critical area for practical implementation. Future work should focus on integrating explainable AI techniques to enhance model interpretability and conducting longitudinal studies to assess model performance over time in dynamic agricultural environments.

Addressing privacy and ethical considerations, collaborating closely with domain experts, benchmarking across different crops for broader applicability, and promoting open-source initiatives are also essential for advancing the field. Additionally, the integration of edge computing capabilities holds promise for enabling on-device processing, reducing dependency on centralized infrastructure, and facilitating real-time decision-making in the field.

Collectively, these efforts aim to refine and extend deep learning models in plant disease detection, addressing technical challenges while ensuring ethical and responsible deployment in agricultural settings.

Acknowledgment

With immense pleasure and a deep sense of gratitude, I wish to express my sincere thanks to my supervisor **RAJKUMAR S**, Associate Professor Sr., School of Computer science and Engineering (SCOPE), Vellore Institute of Technology (VIT), Vellore without his/her motivation and continuous encouragement, this research would not have been successfully completed.

Conflict of Interest

The author states that there is no conflict of interest.

References

[1] N, Banupriya & chowdry, Rajaneni & Yogeshwari, & V,Varsha,. ‘PLANT DISEASE DETECTION USING IMAGE PROCESSING AND MACHINE LEARNING ALGORITHM’, 10.37896/jxu14.7/012,2022.

[2] Yan Guo, Jin Zhang, Chengxin Yin, Xiaonan Hu, Yu Zou, Zhipeng Xue, Wei Wang. "Plant Disease Identification Based on Deep Learning Algorithm in Smart Farming", Discrete Dynamics in Nature and Society, vol. 2020, Article ID 2479172, 11 pages,. https://doi.org/10.1155/2020/2479172

[3] J., Andrew, Jennifer Eunice, Daniela Elena Popescu, M. Kalpana Chowdary, and Jude Hemanth."Deep Learning-Based Leaf Disease Detection in Crops Using Images for Agricultural Applications" *Agronomy* 12, no. 10: 2395. 2020,https://doi.org/10.3390/agronomy12102395.

[4] Mohanty, S. P., Hughes, D. P., & Salathé, M. “Using Deep Learning for Image-Based Plant Disease Detection”. *Frontiers in plant science*, *7*, 1419. https://doi.org/10.3389/fpls.2016.01419,2016.

[5] Chen, R., Qi, H., Liang, Y., & Yang, M. ”Identification of plant leaf diseases by deep learning based on channel attention and channel pruning”. *Frontiers in plant science*, *13*, 1023515. https://doi.org/10.3389/fpls.2022.1023515,2022.

[6] Jung, M., Song, J. S., Shin, A. Y., Choi, B., Go, S., Kwon, S. Y., Park, J., Park, S. G., & Kim, Y. M. “Construction of deep learning-based disease detection model in plants”. *Scientific reports*, *13*(1), 7331. https://doi.org/10.1038/s41598-023-34549-2,2023.

[7] Chohan, Murk & Khan, Adil & Chohan, Rozina & Katper, Saif & Mahar, Muhammad. “Plant Disease Detection using Deep Learning”. International Journal of Recent Technology and Engineering. 9. 909-914. 10.35940/ijrte.A2139.059120,2020.