Paper ID: 352

Disease Detection in Coffee Plants Using Computer Vision

DTN Senanayake[1], MWP Maduranga[2]

[1]*Department of Information Technology, Faculty of Computing, General Sir John Kotelawa Defence University, Sri Lanka*
[2]*Department of Computer Engineering, Faculty of Computing, General Sir John Kotelawa Defence University, Sri Lanka*
[#]37-it-0031@kdu.ac.lk

Abstract— *Coffee is one of the most widely consumed beverages worldwide and an essential crop for many economies However, several illnesses that might negatively affect coffee yield and quality can affect coffee plants. For crop losses to be kept to a minimum, early detection of these diseases is essential. This research suggests a technique that makes use of convolutional neural networks (CNN). The suggested method entails several steps. Gather a dataset of coffee plants first, including both healthy plants and unhealthy plants. After that, the dataset is pre-processed to improve the quality of the images. The pre-processed dataset is then used to create and train a CNN architecture. The CNN develops the ability to automatically recognize patterns and traits. Once trained, the CNN model can be used to identify diseases in coffee plants. This forecast can help farmers and agricultural professionals spot sick plants quickly and take appropriate action. Extensive tests and comparative analyses are carried out to assess the performance of the proposed method. The outcomes show how well the CNN-based method for detecting coffee plant diseases performs in terms of accuracy and dependability. The suggested approach provides a potentially effective response to the difficulties involved in manually identifying diseases. Our proposed model with CNN three-layer classifier with a 0.01 learning rate achieved an overall classification accuracy of 0.89% with the 28th iteration of the training process out of a total of 100 planned epochs. This research utilizes the capability of CNNs to construct automated systems for identifying agricultural diseases, ultimately assisting in sustainable coffee production, and securing the livelihoods of coffee producers.*

Keywords— *Convolutional Neural Network, coffee leaf disease detection, AI*

I. INTRODUCTION

Coffee is one of the most important traded products in the world, with one-third of the world's population drinking it as a beverage. (Yimeru, 2020). Arabians brought coffee to Sri Lanka in the year 1503. The Food and Agriculture Organization reports that 7,688 tons of coffee were produced in 2019 from a total area of 6,445. The nation ranked as the world's 35th largest producer in 2019. Arabica coffee, which thrives particularly well in wetlands and hill country, has increased during the last ten years. Nearly 80% of coffee production came from the small holding sector. Arabica, Robusta, and Liberica coffee are the main coffee varieties farmed in Sri Lanka. However, Liberica coffee is not thought to be crucial from a business standpoint.

Expert manual visual inspection has always been the primary method of disease diagnosis in coffee plants. This method can be laborious, subjective, and time-consuming, which could delay the detection of an illness or result in a false positive. Additionally, as coffee farming expands in scope, there is a rising need for automated and effective disease-detection techniques to assist farmers and guarantee the sustainability of coffee production. In the old days, Sri Lanka exported nearly 50 000 tons of coffee annually. Early diagnosis and prevention are the foundations of successful plant disease avoidance and control, which play critical roles in management and decision-making (Fang and Ramasamy, 2015). But due to a severe disease called coffee rust that occurred in 1870 coffee cultivation was destroyed. Coffee plants are susceptible to many diseases: coffee leaf rust (CLR) coffee leaf miner and Cercospora leaf spot. The deep learning technique clearly shows that the Convolutional Neural Network (CNN) is an excellent tool for addressing computer vision problems in a wide range of application domains. CNN does not necessitate manual feature extraction (Yamashita et al., 2018). Based on that high level of efficiency, convolutional neural networks are well suited for disease detection and diagnosis of plant diseases. (Ferentinos, K.P., 2018) To create a reliable and precise system for disease detection in coffee plants, this research intends to take advantage of CNNs' capabilities. A CNN model may be taught to identify between healthy and ill plants using visual patterns and characteristics by training it on a dataset of coffee plant photos that includes both healthy plants and different disease symptoms. Once trained, CNN may be used to detect the presence of disease quickly and accurately in new photos by automatically analysing them. In this project, we will develop and train a CNN model using a carefully selected dataset to investigate the effectiveness of CNNs in detecting diseases in coffee plants. Through extensive testing, we will assess the model's performance by monitoring important parameters like accuracy, precision, recall, and F1-score. This study suggests a mobile application that relies on deep learning to identify problems affecting coffee plants quickly and precisely. We have created an innovative technique that is capable of precisely recognizing a wide range of diseases in coffee plants by leveraging the power of machine learning. These are Coffee Leaf Rust (CLR), coffee leaf miner, Cercospora, and Phoma. As input, both infected and healthy coffee leaf images are used for training and constructing the model. The study uses CNN architecture for learning and coffee disease detection and after recommending some solutions to cure the diseases. If the user wants to use the system via phone, he can do it very easily. Agriculture instructors can identify the locations which affect the diseases most by using Google Maps. Smartphone-based AI apps might alert farmers and speed up disease diagnostics, perhaps preventing pest and disease outbreaks, and offering coffee farmers a holistic approach to disease management. The study contributes by revolutionizing the detection of coffee plant diseases, enabling early intervention, and reducing crop losses. Our efforts lead the way for environmentally friendly coffee production, ensuring that there will always be a supply of high-quality coffee for future generations.

II. RELATED WORKS

This proposal's problem area has been the subject of research that has been done by Burhanudin Syamsuri, and Gede Putra Kusuma. This study uses data from PlantVillage. Use CNN's deep learning network to identify plant disease signs. These models included InceptionV3, a model renowned for its accuracy on personal computers (PCs), MobileNet, and Mobile Nasnet (MNasNet), two models designed specifically for mobile devices. With 32 filters, MobileNet contains a full convolution layer. With the top 5 accuracies of 92%, this model. The major goal of MNAS (Mobile Neural Architecture Automated Search), which has minimal latency, is to strike a fair balance between accuracy and latency. With a 389 ms latency, this achieves top-5 accuracy of 92.55%. InceptionV3 had the highest accuracy (95.79%), MNasNet came in second (94.87%), and MobileNet came in third (92.83%) while MobileNet had the lowest time delay (394.70 ms, 430.20 ms, and 2236.10 ms, respectively).

In (Pranathi et al. 2018) To ensure that agricultural losses are kept to a minimum, it is crucial to monitor the crop's growth. Tomato leaf disease comes in several forms. This study uses the LeNet convolutional neural network model to identify and classify diseases in tomato leaves. 18160 photos are included in the dataset, which has 10 different classes. Neural network models use automatic feature extraction to classify the input image into the relevant disease classifications. The suggested system's average accuracy ranges from 94 to 95 percent, proving the neural network approach's feasibility even under difficult conditions. (Shita and Shashi 2021) Authors proposed automatic identification of major coffee plant leave diseases. These are brown eye spot (BES), coffee berry disease (CBD), coffee leaf rust (CLR), and coffee wilt disease (CWD). Using deep learning features paper suggested a method for identifying Ethiopian coffee leaf diseases. The paper discusses Gaussian filtering, median filtering, and the hybrid of the two filtering approaches which are used to eliminate the noise from the photos. The hybrid of the two filtering techniques produced better results. Use CNN for feature extraction and KMeans clustering for segmentation as well. Finally, compared the CNN-Softmax classifier to the CNN-SVM classifier. SVM outperforms the softmax classifier in terms of performance and computation efficiency, according to the experimental results. It succeeded in achieving an overall classification accuracy of 96.5% using the recommended model and SVM classifier.

Research by Mengistu et al (2016) compared different classifiers such as artificial neural network (ANN), k-Nearest Neighbors (KNN), Naïve and a hybrid of self- organizing map (SOM) and Radial basis function (RBF). And use five different segmentation techniques to extract meaningful information and understand the structure of dataset. Such as Otsu, FCM, K-means, Gaussian distribution, and the combinations of K-means and Gaussian distribution. Authors conduct different experiments to find the best one which performs best. Overall, the combined segmentation technique outperformed Otsu, FCM, K-means, and Gaussian distribution, and the combined classifiers of RBF (Radial basis function) and SOM (Self-organizing map) together with a combination of k-means and Gaussian distribution performed 92.10% better. The overall target is to identify suitable segmentation techniques to detect coffee diseases in CLR, CBD, and CWD.

In summary, there are some drawbacks to the current disease detection methods for coffee plants that may affect their efficacy and dependability. The model's accuracy could be poor. because of the algorithm and the dataset. As a result, these models lack the ability to produce results. Some systems concentrate on reducing the number of diseases. They might not include major diseases that can affect coffee plants.

III. METHODOLOGY

A. Dataset Gathering

In this step, gather a dataset comprising pictures of both healthy and unhealthy coffee plants. The CNN model is trained and tested using this dataset as its basis. While 376 photos were used to validate the model, 1749 images were used to train the CNN models. For model testing, the remaining 375 photos were kept. For training, validation, and testing, the selection was made while keeping in mind the 70%, 15%, and 15% image percentages, respectively. Data preprocessing: Make sure the collected dataset is in a format that will be useful for training. First applied denoising to the input image. To convert the denoised image to grayscale, use binary Thresholding. The binary image's contours are located by the image segmentation function, which chooses the greatest contour as the leaf region. Moreover, it normalizes the segmented image by scaling the pixel values between 0 and 1 and converting it to float 32. Finally, crop the normalized image to only show the leaf area. The clipped image is then resized to the required dimensions.

B. CNN Training:

Utilize the pre-processed dataset to train the CNN model. Adding four CNN layers with various learning rates, train the system. by including max pooling, batch normalizing, dense layers, and various filter sizes. This model employs early stopping as a technique to reduce overfitting and enhance model generalization. The maximum training accuracy of 0.89% was reported after 100 training epochs.

C. Model validation:

Utilize a different validation dataset for evaluating the effectiveness of the trained CNN model. This process evaluates the model's generalizability to new data. Metrics including accuracy, precision, recall, and F1 score are generated to assess how well the model performs in precisely identifying illnesses in coffee plants. The network was fed with training and validation datasets to start the training process. The validation dataset is used at the end of each epoch to validate the learning, and the network automatically modifies the weights after convolving the images before moving on to the next iteration. Initially, all the images in the training dataset were input into the network in batches of size 224 (224 images).

D. Hyperparameter Tuning:

To further improve the performance of the CNN model, adjust its hyperparameters. Hyperparameters. The ideal value is determined by comparing different learning rates (0.0001, 0.001, and 0.01) to get faster convergence and greater generalization. 64 is the batch size. This model employs a variety of convolutional layer sizes and filter sizes, including (32,64,128,256). Common activation functions are used in dense and convolutional layers. In output layers, the SoftMax function is used. Dropout rates (0.3, 0.2, and 0.1) increase the generalizability of the model and assist avoid overfitting.

Develop an accurate and efficient system for detecting coffee plant diseases by using this methodology, a properly selected dataset, pre-processing the data, training a CNN model, assessing its performance, and adjusting its hyperparameters. By providing a useful tool to monitor and manage the health of their crops, this strategy gives coffee farmers the power to intervene early and protect the quality and production of their farms.

Figure.1: Cerscospora

Figure.2: Healthy

Figure.3: Leaf rust

Figure.4: Miner

Figure.5: Phoma

IV. RESULT AND DISCUSSION

During the training phase, the model is fed with the images in the training dataset multiple times. Each complete pass of the training dataset through the model's layers is known as an epoch. For each image, the model's weights are adjusted based on the difference between the predicted output and the actual output. This adjustment is typically done using an Adam optimization algorithm. The model's weights are updated iteratively for each batch of images (a subset of the training dataset) until all batches have been processed, completing one epoch. The training phase continues for a predefined number of epochs, allowing the model to learn and refine its weights over time.

At the end of each epoch, the model evaluates its performance on the validation dataset. The images in the

validation dataset are fed through the model, and the model generates predictions for each image. These predicted outputs are compared to the ground truth labels for the validation dataset, and evaluation metrics such as accuracy, loss, precision, recall, etc., are calculated. The validation metrics provide an indication of how well the model is generalizing and whether it is overfitting or underfitting the training data. The validation results help in monitoring the model's progress, identifying the best-performing epoch, and potentially making decisions like early stopping or adjusting hyperparameters. This iterative training-validation process continues until the predefined number of epochs is completed or certain convergence criteria are met. The goal is to find a balance between minimizing the loss on the training set and achieving good performance on the validation set. The outcomes of the training technique generate two graphs for each model. Fig.6 shows the training and validation accuracy plot and Fig.7 denotes the loss graph, on the other hand, displayed the training loss plot vs the model's validation loss plot.

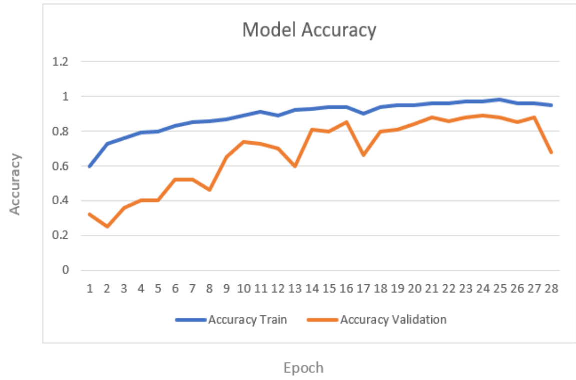

Figure.6. Training and validation Accuracy

The training and validation accuracy graph shows the change in loss values during the training process for the CNN model with three layers and a learning rate of 0.01. The initial epoch's training accuracy is 0.6009; it increases with each succeeding epoch. In the 28th epoch, it reaches 0.9560. The validation accuracy grows, however not as smoothly as the training accuracy. It starts at 0.3218 in the first epoch. In the 28th epoch, it gets to 0.8830. The model overfits the training data when there is a large initial difference between the training and validation accuracy. The validation accuracy increases as the training goes on, indicating that the model is generalizing to new data more effectively. Overall, both the training and validation accuracy of this model exhibits an upward trend. Further hyperparameter adjustment may be helpful to enhance generalization, according to validation accuracy.

Above training and validation loss graph shows the change in loss values during the training process for the CNN model with three layers and a learning rate of 0.01. The loss values are rather high at the start of training and continuously drop over the epochs, showing that the model is learning new information and performing better. The loss values significantly decrease in the initial epochs, indicating that the model is swiftly picking up on the patterns in the training data. The val_loss numbers, which stand for validation loss, give an estimate of how effectively the model generalizes to new data. The validation loss lowers in the early epochs, showing good generalization of the model. The validation loss, in contrast to the training loss, improves less at the sixth epoch. This may indicate that the model is having trouble generalizing adequately to new cases due to overfitting the training set of data.

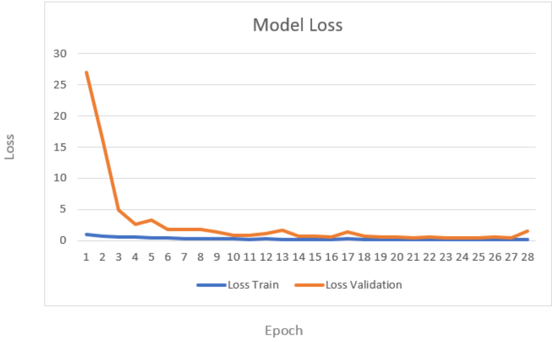

Figure.7. Training and validation loss

Table 1. Classification Report - Train Set:

	precision	recall	f1-score	support
Architecture 1	96%	0.99	0.97	352
Architecture 2	100%	1.00	1.00	360
Architecture 3	100%	1.00	1.00	341
Architecture 4	100%	0.96	0.98	345
Architecture 5	98%	0.99	0.99	351

In this classification report for each class, Architecture 1: With a precision of 0.96, the model correctly predicted 96% of the Cercospora samples. Recall is 0.99, indicating that 99% of the real Cercospora samples were accurately identified by the model. The precision and recall measures are balanced according to the F1-score of 0.97.

Architecture 2: The model obtained an F1-score of 1 and flawless precision, recall, and performance. This shows that all instances of the Healthy class were correctly categorised by the model. Architecture 3: The model accurately identified every instance of the Leaf rust class, achieving perfect precision, recall, and an F1-score of 1.

Architecture 4: The model's precision was 1, which indicates that all Miner samples that were predicted were accurate. Recall is 0.96, which indicates that 96% of the actual Miner samples were properly identified by the model. With an F1-score of 0.98, the overall performance is strong. Architecture 5: With a precision of 0.98, the model successfully predicted 98% of the Phoma samples. Recall is 0.99, indicating that 99% of the actual Phoma samples were accurately identified by the model. With an F1-score of 0.99, the overall performance is strong. The model's total classification accuracy on the training set was 0.99, which indicates that 99% of the occurrences were correctly classified.

Table 2. Classification Report - Test Set:

	precision	recall	f1-score	support
Architecture 1	0.74	0.95	0.83	74
Architecture 2	0.99	1.00	0.99	66
Architecture 3	0.99	0.96	0.97	89
Architecture 4	0.92	0.73	0.81	84
Architecture 5	0.85	0.85	0.85	62

In this classification report for each class, Architecture 1: With a precision of 0.74, the model correctly predicted 74% of the Cercospora samples. Recall is 0.95, which means that 95% of the real Cercospora samples were accurately recognized by the model. The precision and recall are balanced according to the F1-score of 0.83.

Architecture 2: With a precision of 0.99, the model successfully predicted 99% of the Healthy samples. Recall is 1, indicating that the model accurately identified all real Healthy samples. With an F1-score of 0.99, the overall performance is strong.

Architecture 3: With a precision of 0.99, the model correctly predicted 99% of the samples of Leaf rust. The recall is 0.96, indicating that 96% of the actual Leaf rust samples were accurately identified by the model. With an F1-score of 0.97, the overall performance is strong.

Architecture 4: 92% of the projected Miner samples were accurate according to the model, which had a precision of

0.92. Recall is 0.73, indicating that 73% of the actual Miner samples were properly identified by the model. The F1- score of 0.81 indicates a balanced assessment of recall and precision. Architecture 5: Precision is 0.85 (The model correctly identifies 85%

of the instances predicted as Phoma). Recall is 0.85 indicate that model captures 85% of the true instances of Phoma). F1-score is 0.85. A balanced F1-score, indicating similar precision and recall and 62 of instances in this class in the test set. With an overall accuracy of 0.89 on the test set, the model correctly identified 89% of the cases.

A set of quantitative criteria, including accuracy, precision, recall, and F1-score, was utilized to assess the performance of the suggested model. Figure 3,4 summarizes the findings. They display the highest values of the quantitative metrics obtained up to the epoch number. Model for a maximum of 100 epochs. The validation loss is monitored, though, and training is interrupted early if the loss does not decrease after a specific number of epochs (patience=5). Therefore, if early stopping is initiated, the real number of epochs may be less than 100.With results, where the authors deal with plant diseases of multiple crops, the model thus offers an easy and efficient solution to the problem of plant disease identification. The model provides outcomes that are comparable to those of conventional state-of-the-art techniques with fewer resource restrictions and less data.

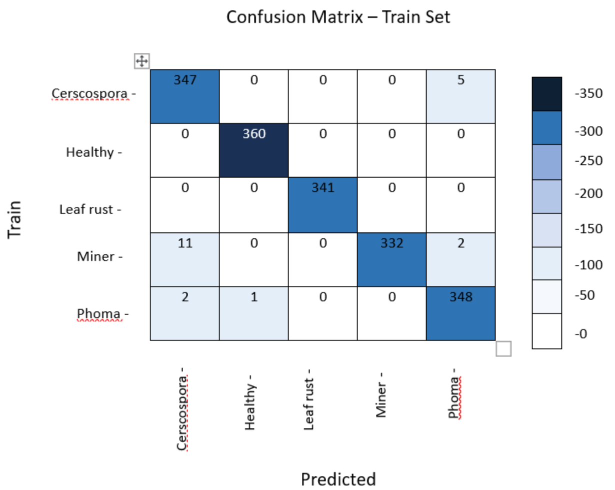

Figure.8. Confusion Matrix – Train Set

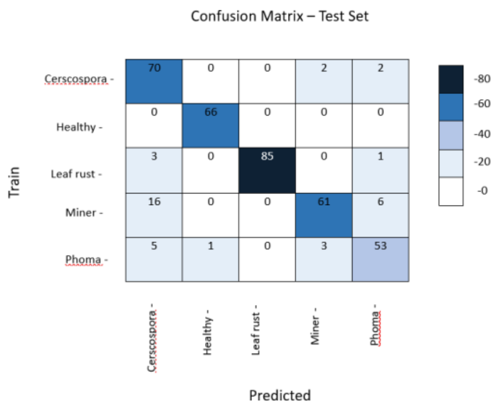

Figure.9. Confusion Matrix – Test Set

V. ONCLUSION

In farming, crop protection is a challenging endeavor. This depends on knowledge of the harmful weeds, diseases, and pests that affect crops. To categorize coffee leaf disorders into 5 different groups, a convolutional neural network with a limited number of layers was utilized as the architecture. A freely available dataset with 58,555 photos collected in genuine agricultural settings was used to train the models. The information mentions both harmful and advantageous plants. The Network's classification accuracy after training on images of the natural world was 0.89%. This displays CNN's ability to derive important components from nature that are required for the classification of plant diseases. So, the above-mentioned model can be made used of as a decision tool to help and support farmers in identifying the diseases that can be found in the coffee plant. The system only supports Android smartphone users. The mobile phone camera's quality has an impact otherwise, the generated model struggles to perform at its peak. These are the drawbacks of the suggested system. Only four different forms of coffee plant diseases can be identified by this research, which is still based on the existing dataset. Therefore, expanding the dataset and creating this model will allow farmers to diagnose more diseases accurately in the future.

VI. REFERENCES

Syamsuri, B. and Kusuma, G.P., 2019. Plant disease classification using Lite pretrained deep convolutional neural network on Android mobile device. *Int. J. Innov. Technol. Explor. Eng*, *9*(2), pp.2796-2804.

Tm, P., Pranathi, A., SaiAshritha, K., Chittaragi, N.B. and Koolagudi, S.G., 2018, August. Tomato leaf disease detection using convolutional neural networks. In *2018 eleventh international conference on contemporary computing (IC3)* (pp. 1-5). IEEE.

X. "ETHIOPIAN COFFEE LEAF DISEASES IDENTIFICATION USING DEEP LEARNING

FEATURES", International Journal of Emerging Technologies and Innovative Research (www.jetir.org | UGC and ISSN Approved), *ISSN:2349-5162, Vol.8, Issue 10, page no. ppb498-b508, October-2021, Available at: http://www.jetir.org/papers/JETIR2110156.pdf.*

Mengistu, A.D., Mengistu, S.G. and Alemayehu, D.M., 2016. Image analysis for

Ethiopian coffee plant diseases identification. *IJBB*, *10*(1).

Maduranga, M.W.P. and Nandasena, D., 2022. Mobile- based skin disease diagnosis system using convolutional neural networks (CNN). *IJ Image Graphics Signal Process*, *3*, pp.47-57 Esgario, J.G., de Castro, P.B., Tassis, L.M. and Krohling, R.A., 2022. An app to assist farmers in the identification of diseases and pests of coffee leaves using deep learning. *Information Processing in Agriculture*, *9*(1), pp.38-47.

Shrestha, G., Das, M. and Dey, N., 2020, October. Plant disease detection using CNN. In *2020 IEEE applied signal processing conference (ASPCON)* (pp. 109-113). IEEE.

Shrimali, S., 2021. Plantifyai: a novel convolutional neural network-based mobile application for efficient crop disease detection and treatment. *Procedia Computer Science*, *191*, pp.469-474.

Ferentinos, K.P., 2018. Deep learning models for plant disease detection and diagnosis. *Computers and electronics in agriculture*, *145*, pp.311-318.

Valdoria, J.C., Caballeo, A.R., Fernandez, B.I.D. and Condino, J.M.M., 2019, October. iDahon: An android based terrestrial plant disease detection mobile application through digital image processing using deep learning neural network algorithm. In *2019 4th International Conference on Information Technology (InCIT)* (pp. 94- 98). IEEE.

Ferentinos, K.P., 2018. Deep learning models for plant disease detection and diagnosis. *Computers and electronics in agriculture*, *145*, pp.311-318.

Yimeru Abnet Kifle “*Automatic Coffee Disease and Pest Damage Identification*.” Addis Ababa University (2020)

Fang, Y. and Ramasamy, R.P., 2015. Current and prospective methods for plant disease detection. *Biosensors*, *5*(3), pp.537-561.