Advancements in precision agriculture: Edge Impulse-based convolutional neural networks for robust apple disease diagnosis

Yijie Liao

The Department of Computer Science, Hong Kong Baptist University, Hong Kong Special Administrative Region, 999077, China

20250576@life.hkbu.edu.hk

Abstract. Apple cultivation is essential to the global economy, but it is vulnerable to diseases that can have severe impacts on crop yields and cause economic losses. Accurate and timely detection of these diseases is crucial for effective prevention and management. This study addresses the limitations of traditional machine learning methods, such as manual feature extraction and high computational costs, by employing convolutional neural networks (CNNs) and transfer learning using the MobileNetV2 architecture for the detection and categorization of four common apple ailments: Apple Scab, Black Rot, Cedar Apple Rust, and Healthy apples. The model produced excellent accuracy of 96.7% on the testing dataset while working with an enriched dataset of 12,000 photos in various dimensions. The model's performance on the validation set was 89.7% with a loss of 0.60. The transfer learning model demonstrates superior on-device performance, with faster inferencing time and lower RAM usage, making it suitable for real-world applications in apple orchards. In conclusion, this study successfully demonstrates the potential of using convolutional neural networks and transfer learning with the MobileNetV2 architecture for apple disease detection. The model's high accuracy and performance make it a promising tool for assisting apple growers in monitoring and managing apple diseases, ultimately leading to more efficient and sustainable apple production.

Keywords: apple disease detection, CNN, transfer learning, MobileNetV2 architecture.

1. Introduction

Apple production is important to the world economy and is the main source of income for numerous farmers throughout the world. However, there are a number of diseases that can affect the production of apples and cause major financial losses. Therefore, accurate and timely detection of apple diseases is essential for effective disease management and prevention. And the conventional method of manual plant monitoring for disease identification is a time-consuming and cumbersome task, requiring expertise in plant diseases [1]. A more effective and trustworthy substitute to manual inspection is now being offered by computer vision and machine learning approaches, which have emerged as promising instruments for automatic categorization and identification of plant diseases. Thus, automating plant disease detection using machine learning methods like convolutional neural networks can greatly benefit large-scale apple field monitoring, enabling early disease recognition and reducing potential losses for farmers [1].

DOI: 10.54254/2755-2721/18/20230983

Numerous studies have used methods based on deep learning and machine learning to identify plant diseases, with convolutional neural networks (CNNs) exhibiting exceptional performance in image classification tasks. Studies have employed various pre-trained networks, such as ResNet, VGG, and Inception, and fine-tuned them on specific plant disease datasets. However, the current research in apple disease classification mainly focuses on limited datasets, with limited attention being given to the potential benefits of data augmentation techniques that can enhance the model's generalization capabilities.

Support vector machines (SVM) and K-means clustering, in particular, have made great strides in recent years toward automating the identification of plant diseases [2]. But these conventional machine learning techniques have some drawbacks, including reduced effectiveness and suitability for uniform-background plant images taken in ideal lab settings [2]. Furthermore, CNNs have gained popularity for object detection and plant disease identification due to their ability to extract features automatically and directly from input images [2]. Despite the growing interest in CNNs for crop disease detection, difficulties remain in implementing real-time identification of apple leaf ailments because they may involve multiple diseases on the exact same the leaf, which is varying disease spot sizes, and interference from environmental factors like shadows, illumination, and soil [2]. Thus, there remains a need to develop more effective and robust methods for apple disease detection that can overcome these challenges and provide real-time, accurate results in diverse field conditions.

Singh et al. underscored the crucial significance of early detection of diseases in apple leaves for better crop yield and recovery chances [3]. The early diagnosis of apple illnesses has been studied in the past using traditional machine learning techniques like SVM and KNN approaches. These techniques, however, necessitate engineering in the fields of feature extraction and organization, as well as discriminative information. In a variety of image-related challenges, CNN have become a well-known method for pattern identification.

In order to classify and identify apple diseases using a comprehensive dataset, this study uses deep learning techniques, notably CNNs, to overcome these restrictions. The research focuses on four common apple diseases, namely Apple Scab, Black Rot, Cedar Apple Rust, and Healthy apples. An augmented dataset obtained from Kaggle was leveraged, which consists of 3000 images per class, with 1, 536 augmented images and 1, 464 original images. The data augmentation techniques used include zooming, horizontal flipping, vertical flipping, rotating, brightness changing, and shearing.

This study focuses on the application of convolutional neural networks (CNNs), specifically using transfer learning with the MobileNetV2 architecture, to classify and recognize apple diseases. By leveraging the power of transfer learning, this study aims to improve the model’s performance in recognizing apple diseases while reducing the training time and computational resources required.

In summary, by examining MobileNetV2's migration learning potential in the categorization and diagnosis of apple diseases, this study aids in apple planting and production. The findings will help enhance the accuracy and robustness of apple disease detection systems, ultimately benefiting apple growers and the agricultural industry as a whole.

2. Methods

2.1. Dataset preparation

The dataset utilized in this study was sourced from Kaggle and consisted of 12,000 apple leaf images, categorized into four types of leaf illnesses, namely Apple Scab, Black Rot, Cedar Apple Rust, and Healthy [4]. Each category contained 3,000 images, with 1,536 augmented images and 1,464 original images. The augmentation techniques employed included zooming, horizontal flipping, vertical flipping, rotating, brightness changing, and shearing. The dataset consisted of RGB images with varying dimensions. Some sample images from each class are presented in Figure 1.

Figure 1. The sample images on the collected dataset.

2.2. Edge impulse-based apple disease recognition model

2.2.1. Introduction of edge-impulse. Edge Impulse is a cloud-based machine learning operations (MLOps) platform aimed at creating embedded and edge ML systems [5]. It provides users with a range of ready-made neural network models that can be directly used to train. Edge Impulse facilitates the direct conversion of trained models into code that is executable on microcontroller units (MCUs), which renders its use notably straightforward. Edge Impulse provides the implementation of extremely optimized ML on technology ranging from MCU to CPU and specialized AI accelerators, making it possible to employ ML on embedded devices for detectors, sound, and image recognition directly. This study used the Edge Impulse platform to develop a CNN recognition model for Apple disease. The platform offers an accessible and efficient way to develop ML models for embedded systems, addressing data collection, preprocessing, development, deployment, and monitoring challenges.

2.2.2. Introduction of CNN. The CNN is a prominent deep machine learning model in computer vision [6-8]. CNN is made up of fully connected layers, multi-layer convolutional layers, and pooling layers, which were all modeled after the biological visual brain [9]. The convolution layer, which is constructed by swiping various convolution kernels on the input image and carrying out particular operations, is in charge of extracting the local features of the image. Finally, the computational complexity and the number of parameters is reduced, and the spatial information is preserved. Pooling layers are a nonlinear form of downsampling intended to decrease the spatial dimensionality and the number of parameters, which lowers the risk of overfitting. After these layers, a fully connected layer maps the extracted features to a specific output space, such as for classification or regression tasks. Additionally, all of the activations in the preceding layer are connected to the neurons in the fully connected layer. Finally, the backpropagation algorithm trains the CNN to minimize the loss function.

2.2.3. Transfer learning with MobileNetV2. Transfer learning is a technique that involves the fine-tuning of a pre-existing neural network to accommodate a new task. This approach capitalizes on the insights acquired from the initial task to enhance the performance of the network on the new task. The present investigation utilized transfer learning methodology in conjunction with the MobileNetV2 architecture [10], a well-established approach for image classification tasks. The MobileNetV2 is a CNN architecture that has been specifically developed for mobile and edge devices. It is known for its ability to provide a balance between computational efficiency and classification performance, making it a lightweight and efficient option for these types of devices.

2.3. Implementation details

The MobileNetV2 model was used for the transfer learning strategy in this study, and it had a depth multiplier of 0.35 and an input size of 96×96. The topmost layer of the MobileNetV2 network had a dropout rate of 0.1 and was made up of 16 neurons. The transfer learning model underwent 20 iterations of training with a learning rate of 0.0005 for each iteration. Figure 2 provides the architecture of the model used in this study.

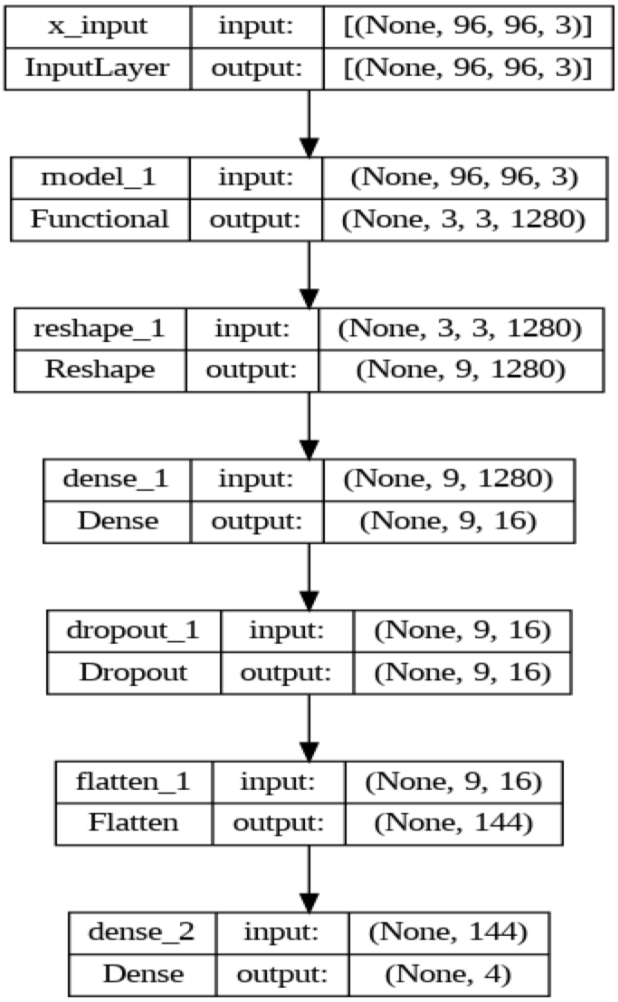

Figure 2. The architecture of the model.

3. Result and discussion

The results of this study demonstrate that the transfer learning approach using MobileNetV2 outperforms the classifier-only model in apple disease classification. The transfer learning model

DOI: 10.54254/2755-2721/18/20230983

achieved an accuracy of 89.7% on the validation set, accompanied by a loss of 0.60. On the other hand, the classifier model achieved an accuracy of 91.8% along with a loss of 0.27. However, the transfer learning model demonstrated superior performance on the testing dataset, with an impressive accuracy of 96.7%, compared to the classifier model's 82.5% accuracy. A comprehensive presentation of the overall performances of each model is provided in Table 1 and Table 2.

Table 1. Classifier model validation and testing performance.

Classifier	Performance
Accuracy (Validation)	91.8%
Loss (Validation)	0.27
Accuracy (Testing)	89.25%
Inferencing Time	226 ms
Peak RAM Usage	364.0K
Flash Usage	66.5K

Table 2. Transfer learning model validation and testing performance.

Classifier	Performance
Accuracy (Validation)	89.7%
Loss (Validation)	0.60
Accuracy (Testing)	96.7%
Inferencing Time	110ms
Peak RAM Usage	346.9K
Flash Usage	579.8K

The confusion matrix for the validation set of both models reveals that the transfer learning model has a more balanced performance across all disease classes, with F1 scores ranging from 0.85 to 0.92. In contrast, the classifier model exhibits a higher disparity in F1 scores, ranging from 0.87 to 0.97. This disparity in F1 scores is indicative of the fact that the transfer learning model has a more consistent performance in identifying various apple diseases and is less prone to produce misclassifications. A detailed depiction of the two confusion matrices related to the validation set can be found in Table 3 and Table 4.

Table 3. Classifier model confusion matrix (validation set).

	Apple Scab	Black Rot	Cedar Apple Rust	Health
Apple Scab	92.1%	4.5%	1.1%	2.3%
Black Rot	5.6%	91.9%	0%	2.5%
Cedar Apple Rust	4.0%	1.0%	94.4%	0.6%
Health	7.5%	3.8%	0%	88.7%
F1 Score	0.87	0.92	0.97	0.91

Table 4. Transfer learning model confusion matrix (validation set).

	Apple Scab	Black Rot	Cedar Apple Rust	Health
Apple Scab	91.9%	0.9%	3.4%	3.8%
Black Rot	7.5%	86.2%	1.0%	5.4%
Cedar Apple Rust	9.4%	0.8%	87.1%	2.7%
Health	4.8%	0.4%	0.4%	94.3%
F1 Score	0.85	0.92	0.91	0.91

The on-device performance of the transfer learning model is also noteworthy. The inferencing time of the transfer learning model is recorded at 110 ms, indicating a significantly faster speed than the classifier model's 226 ms. Additionally, the transfer learning model consumes less RAM (346.9K) compared to the classifier model (364.0K) but requires more flash memory (579.8K) compared to the classifier model (66.5K). This enhanced on-device performance is crucial for real-time applications, such as in-field disease detection in apple orchards.

Comparing results with previous studies, it can be clearly observing that the transfer learning approach using MobileNetV2 outperforms the classifier model. The high accuracy and fast inferencing time of the transfer learning model make it a promising solution for real-world applications, such as on-site disease detection in apple orchards. A thorough presentation of the final testing outcomes can be found in Table 5 and Table 6.

Table 5. Classifier model confusion matrix (testing set).

	Apple Scab	Black Rot	Cedar Apple Rust	Health	Uncertain
Apple Scab	91.0%	1.5%	0.2%	2.2%	5.1%
Black Rot	5.1%	87.9%	0%	1.0%	6.0%
Cedar Apple Rust	2.2%	0.7%	89.4%	0.8%	6.9%
Health	5.8%	0.7%	0%	88.7%	4.8%
F1 Score	0.89	0.92	0.94	0.92	NA

Table 6. Transfer learning model confusion matrix (testing set).

	Apple Scab	Black Rot	Cedar Apple Rust	Health	Uncertain
Apple Scab	95.1%	0.8%	1.02%	1.7%	1.4%
Black Rot	1.5%	97.2%	0%	0.7%	0.7%

DOI: 10.54254/2755-2721/18/20230983

Table 6. (continued).

Cedar Apple Rust	1.7%	0.2%	96.8%	0%	1.3%
Health	1.7%	0.2%	0%	97.7%	0.5%
F1 Score	0.95	0.98	0.98	0.98	NA

Furthermore, it is essential to consider the context of this study, which focused on a specific dataset and did not include a comparison of various transfer learning architectures. Future research could address these limitations by exploring other transfer learning models and investigating the potential of combining multiple pre-trained CNN architectures to further improve the performance of apple disease detection systems.

4. Conclusion

In conclusion, this study demonstrates the effectiveness of transfer learning with MobileNetV2 in apple disease classification and recognition. The transfer learning model outperforms the classifier-only model, achieving higher accuracy and more balanced performance across various apple disease classes. Additionally, the transfer learning model displays superior on-device performance, characterized by a faster inferencing time and lower RAM usage, rendering it suitable for practical applications in the real world.

This research contributes to the development of more accurate and efficient apple disease detection systems, which can greatly benefit apple growers and the agricultural industry. However, this study's limitations include reliance on a specific dataset and no comparison with other transfer learning architectures. Future research should explore different transfer learning models and potentially combine multiple pre-trained CNN architectures for better apple disease detection. This would lead to more robust tools, benefiting agriculture and promoting sustainable apple production practices.

References

[1] Baranwal S Khandelwal S and Arora A 2019 Deep learning convolutional neural network for apple leaves disease detection In Proceedings of international conference on sustainable computing in science technology and management (SUSCOM) Amity University Rajasthan Jaipur-India

[2] Jiang P Chen Y Liu B He D and Liang C 2019 Real-Time Detection of Apple Leaf Diseases Using Deep Learning Approach Based on Improved Convolutional Neural Networks IEEE Access 7 59069–59080

[3] Singh S Gupta I Gupta S Koundal D Sultan Aljahdali Mahajan S and Pandit A K 2022 Deep Learning Based Automated Detection of Diseases from Apple Leaf Images Computers Materials & Continua 71(1) 1849

[4] Rm1000 2023 Augmented Apple Disease Detection Dataset 12000 Apple Leaf Images with Apple Scab Black Rot Cedar Apple Rust or Healthy Kaggle https://www.kaggle.com/datasets/rm1000/augmented-apple-disease-detection-dataset

[5] Hymel S et al 2022 Edge Impulse: An MLOps Platform for Tiny Machine Learning arXiv preprint arXiv:2212.03332

[6] Arena P, Basile A, Bucolo M, et al. Image processing for medical diagnosis using CNN[J]. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 2003, 497(1): 174-178.

[7] Yu Q Yang Y Lin Z et al 2020 Improved denoising autoencoder for maritime image denoising and semantic segmentation of USV China Communications 17(3): 46-57

[8] Yang R Singh S K Tavakkoli M et al 2020 CNN-LSTM deep learning architecture for computer vision-based modal frequency detection Mechanical Systems and signal processing 144: 106885
[9] O'Shea K and Nash R 2015 An introduction to convolutional neural networks arXiv preprint arXiv:1511.08458.
[10] Manor E and Greenberg S 2022 Custom Hardware Inference Accelerator for TensorFlow Lite for Microcontrollers in IEEE Access vol 10 pp 73484-73493 2022 doi: 10 1109/ACCESS 2022 3189776