Enhancing Apple Leaf Disease Classification Using Transfer Learning and CLAHE-Based Image Preprocessing

Gireesh Kumar[1], Weiwei Jiang[2]
[1]Post Doctoral Researcher, Lincoln University College, Petaling Jaya, Selangor Darul Ehsan-47301, Malaysia; [2]School of Information and Communication Engineering Beijing University of Posts and Telecommunications
pdf.Gireeshkumar@lincoln.edu.my; jww@bupt.edu.cn

Abstract: The diseases of the apple leaf are among the factors affecting agricultural production; thus, there is a need to detect them efficiently and accurately. The study investigates how CLAHE image preprocessing affects the cateGorizing of appLe leAf disEases using transfer learning with MobileNetV2. Plantvil- lage data containing four disease types were used to develop the classifier using raw (unprocessed) images and CLAHE enhanced images. An individual CNN model was also trained as the baseline model. Validation accuracy of 99.31%, 98.74% and 97.08% were achieved using raw image MobileNetV2, CLAHE enhanced MobileNetV2 and CNN, respectively. It is clear that CLAHE processing improves the local contrast of leaf images, making them more visually appealing. However, it did not improve the classification accuracy of the models. It implies that the feature representations of MobileNetV2 pretrained networks are good enough to extract the various features without necessarily im- proving the contrast. The model trained using CLAHE enhanced images still recorded an impressive macro-averaged F1 score of 0.98%. These results show that transfer learning works well for classifying plant diseases and give us more information about how image preprocessing works in deep learning pipelines.
Keywords: Apple leaf disease, CLAHE, MobileNetV2, transfer learning, deep learning, image preprocessing, Plantvil- lage, CNN, plant pathology

---

Introduction

In world economy, agriculture plays an important role. It feeds billions of people and gives them jobs. Plant diseases are one of the major problems that modern farming has to deal with. They can cost a lot of money. Apple farming is especially prone to many types of fungal and bacterial infections that can greatly affect the mass as well as quality of the crop if they go undetected. Diseases like appLe scAb, cedar appLe rust and bLack rot have caused a lot of damage to apple orchards all over the world in the past [1]. Conventional disease identification methods depend significantly on visual checks conducted by trained agronomists, a procedure that is not only labor-intensive and costly but also susceptible to human error, particularly in extensive farming operations.

In the last ten years, cv and dl have made huge pace. This, in turn, has resulted in the appearance of promising approaches that would enable automation of plant disease detection with great efficiency and speed. Methods of deep learning such as CNNs prove extremely useful when it comes to classification of pictures due to their ability to learn organized feature presentation [2]. Training of deep neural networks from the start requires large volumes of data along with enormous computational power, which might not be always possible in agriculture research. Fortunately, there is an approach to solve this problem. Transfer learning enables scientists to apply models built on big datasets, for example, ImageNet, on tasks related to other domains with limited resources [3].

Besides picking up the appropriate architecture, one should pay special attention to the choice of images as it affects greatly the accuracy of the created algorithm. Photos of crops can be taken under various and uncontrollable lighting conditions which lead to insufficient level of contrasts and impossibility of distinguishing diseases. Many researchers use techniques such as CLAHE for pre-processing of images which increases local contrast and noise reduction. Applying them to plant disease images is a promising direction of research [4].

In this paper, there is a discussion about three comparative models – custom-built CNN baseline model, MobileNetV2 trained on unenhanced images, and MobileNetV2 trained on CLAHE-enhanced images. The main goal of this Research is to explore whether the use of CLAHE enhancement technique for pre-processing has a statistically significant effect on the efficiency of apple leaf disease classification. The Plant Village dataset [5] that is frequently used as a standard test dataset in plant disease classification studies is applied for the repro- ducibility and fair comparison of all models. Continuing to the next section, Section II will discuss the literature analysis on plant disease detection using dl and image preprocessing. The next section will deal with the data, preprocessing, and architecture model in Section III. ExperIment resuLts will be discussed in section IV, while conclusions will foLlow in section V.

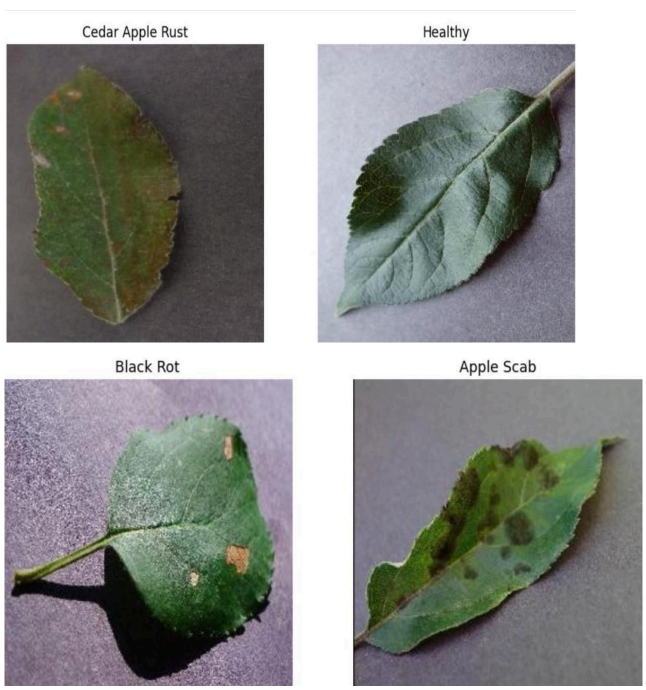

Fig 1. Enhanced leaf images

Related work

A. Observation of plant disease using DL

The use of dl to find plant diseases really took off after the groundbreaking work of Mohanty et al. [1], who showed that CNNs trained on the PlantVillage dataset could get more than 99% accuracy in a controlled lab setting. Their research established PlantVillage as the de facto standard for research on classifying plant diseases, and it led to a lot of follow-up work aimed at making it more accurate, generalizable, and efficient.

Further studies investigated various CNN architectures. For instance, Ferentinos [6] conducted a performance assessment of different deep learning algorithms, such as AlexNet, VG- GNet, GoogLeNet, and others, with the goal of identifying deep CNN architectures able to classify plant diseases and reported accuracy rates exceeding 99%. Additionally, Too et al.

[7] carried out a comparative analysis of DenseNet, Inception, ResNet, and VGG network architectures. According to them, DenseNet was determined to be the most accurate one while also being computationally efficient. Furthermore, it should be noted that Charan and Mohammed Abrar [14] published a study devoted to CNN architecture used for the classification of apple leaf disease by means of transfer learning models on the PlantVillage dataset. Specifically, VGG16 yielded the high- est test accuracy of 96%, beating InceptionV3. At the same time, it should be stressed that the custom CNN architecture is a necessary baseline in the analysis of other more complex models [15].

B. Transfer Learning in Agricultural Applications

Transfer learning has become the de facto methodology employed for the grouping of pLant diseases, especially when training data is sparse. The use of prepared weights in im- age classification helps models learn highly representative features even with fewer agricultural datasets. Saleem et al.

[8] provided a full survey of the usage of transfer Learning appRoaches in disease identification of plants and found out that models such as VGG16, InceptionV3, and MobileNet are superior compared to those which are manually built using convolutional neural networks from scratch.

One example of this kind is MobileNetV2 introduced by Sandler et al. [9]. The incorporation of reversed leftover blocks and linear projection layers ensures that MobileNetV2 achieves significant reductions in the no. of parameters relative to dense structures like VGG and ResNet. Thus, it is ideal for resource-limited agricultural contexts, including field-based smartphone applications.

C. Image Preprocessing Techniques in Deep Learning

Preprocessing of the input image is critical in influencing the quality of the features that may be extracted from the image using neural network models. In domains where the lighting of the scene is not optimal, contrast enhancement techniques may prove useful in improving visibility of pathogenic features. Histogram Equalization (HE) is one of the earliest and simplest contrast enhancement techniques; however, it is global and may over-amplify the effects of noise in homogeneous areas [10].

CLAHE is an improvement on conventional histogram equalization proposed by Zuiderveld [11] through segmenting the image into tiles and applying equalization on each tile independently using a clip limit to avoid over-amplifying noise. It has been widely applied to medical imaging including retinal fundus and dermoscopy and has proven effective in improving detection of pathological lesions [4].

Use of CLAHE for plant disease identification is relatively novel compared to its application to medical images. Ramcha- ran et al. [12] have applied CLAHE and other preprocessing techniques on cassava disease images captured using smart- phones. Nevertheless, results are inconsistent among different crops, datasets, and model types, thus the motivation behind this comparative analysis.

Key Contribution

While numerous studies have independently explored trans- fer learning and image preprocessing for plant disease clas- sification, few have systematically compared the effect of CLAHE preprocessing against raw images within the same transfer learning framework for apple leaf diseases specifically. The present work attempts to bridge this knowledge gap by conducting an experiment on the PlantVillage data set through a systematic comparative study based on MobileNetV2, as well as a tailored CNN.

Method, Experiments and Results

Methodology

A. Dataset

The PlantVillage dataset [5], accessed via Kaggle [13], was used in this study. The dataset consists commentated pictures of heaLthy and diSeased pLant Leaves over many plant kind. For this study, only apple related classes were selected, having four categories:

- AppLe SCab (CLass 0)
- AppLe BlAck ROt (CLass 1)
- CedAr AppLe RuSt (CLass 2)
- AppLe HeaLthy (CLass 3)

The dataset was divided in a ratio of 80/20 where one set was named training and other set validation. The training set contained 243 batches of pictures, and the validation set contains 196 samples. All imAges were resized to 224 × 224 pixeLs to foLLow with MobileNetV2's inPut necessity.

B. Image Preprocessing: CLAHE

To assess the influence of contrast enhancement, a sec- ond parallel data set was constructed through application of CLAHE on each of the images before training the models. This was done by first transform ing the images into LAB colorspace from RGB colorspace and then applying CLAHE to the luminance channel. Contrast enhancement should be per formed using the luminance channel only, since this preserves the chrominance channels. The clip limit and tile grid size used were 2.0 and 8 ×8, respectively [4]. The preprocessing pipeline for CLAHE-enhanced images can be summarized as:

1) Convert RGB image to LAB coLor spAce.
2) AppLy CLaHE to the l (luminance) channeL with cLip Limit = 2.0.
3) Reconstruct the image by merging the enhanced L channeL with the original A aNd B channeLs.
4) ConVert bAck to rgb coLor spAce.
5) Resize to 224 × 224 pixeLs and organize pixeL vaLues to [0, 1].

Raw images underwent only resizing and normalization, without any contrast enhancement, to serve as a control condition.

C. Model Architectures

1) Baseline CNN: A custom CNN was used as a base- line model for comparison purposes. The model architecture consists of 3 convoLutional blocks wiTh Conv2D, BaTch normalization, relu activation, and MaxPooling operations. Next, the output of the featur maps gets leveled and passed through 2 dense laYers, which use the leave technique, beFore reaching the Softmax classification layer. The model is trained from scratch on images only.

2) MobileNetV2 with Transfer Learning: Two different MobileNetV2 models were trained with transfer learning. One Mobilenetv2 model was coached with raw images (MobileNet Raw), and the other was trained with CLAHE preprocessed images (MobileNet CLAHE). For the pre-trained imagenet, the initial weights were used to start the network weights. The first training phase consisted of keeping the base architecture, MobileNetV2, fixed while adding a classifier on top of it. It consisted of:

- Global Average Pooling layer
- dense layer (128 units, relu activation)
- dropout layer (rate = 0.3)
- output dense layer (4 units, softmax activation)

All networks have been built with the help of the ADam improve, learn rate of 1×10^{-4}, grouping cross-entropy lOss funCtion, and acCuracy as the main evaLuation criterion.

D. Training Configuration

All three models were trained for 5 epochs on an iden- tical hardware environment. The model.fit() function was used with the training generator and validation data as inputs. No additional data augmentation beyond CLAHE preprocessing was applied, so as to isolate the result of preprocessing on modeL perFormance. Table I outlines the training configuration.

Results

A. Training Performance

1) Baseline CNN: The custom CNN achieved 96.19% accuracy for training and 97.08% for validation in the final epoch. In this case, although this proves the fact that even a less complex model can learn relevant information from the dAta, the modeL showed slow convergence, achieving a training accuracy of 75.63% in its first epoch before showing steady progress. This aligns with the theoretical assumption that learning from scratch necessitates additional epochs to learn representative features.
2) *MobileNetV2 on Raw Images:* In terms of best results, MobileNetV2 trained on raw data performed exceptionally well throughout the study. This was evident from an accuracy for validation that increased from 98.45% in its first epoch, all the way up to **99.31%** after five epochs. In terms of the loss of validation, it too decreased significantly from 0.0443 down to 0.0197. Epoch-by-epoch results are presented in Table III.
3) *MobileNetV2 on CLAHE-Enhanced Images:* The Mo- bileNetV2 model trained on CLAHE-enhanced images also performed strongly, achieving a peak validation accuracy of **98.74%**. However, the training trajectory was notably less smooth compared to the raw model. Validation accuracy actually decreased between epochs 3 and 4 (from 98.34% to 98.11%) before recovering to 98.74% in the final epoch, sug- gesting that CLAHE preprocessing introduced some variability in the feature distributions that the model needed additional epochs to reconcile. Epoch-by-epoch results are shown in table IV.

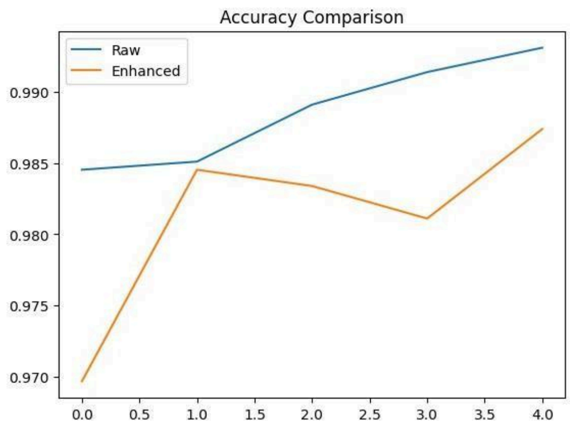

Fig. 2. Printed output confirming MobileNetV2 Raw (0.9931) and CLAHE (0.9874) peak validation accuracies

B. Model Comparison

A direct comparison of the peak validation accuracies achieved by each model is presented in Table V. The MobileNetV2 model trained on raw images outperformed all other models, followed closely by the CLAHE-enhanced variant, with the custom CNN achieving the lowest accuracy.

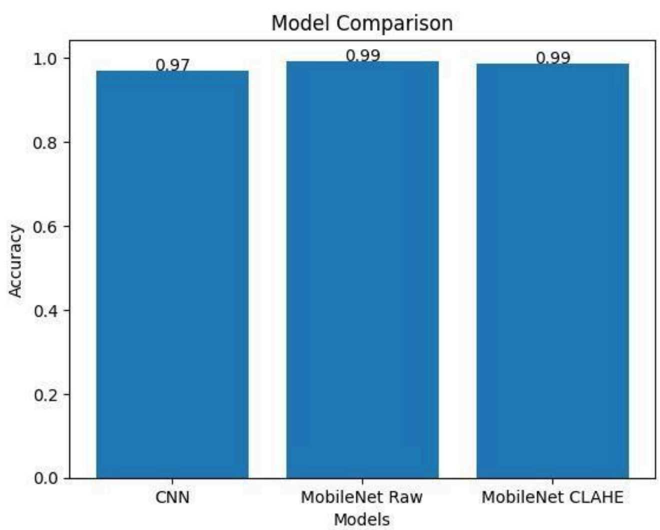

Fig. 3. Bar chart comparing peak validation accuracy of CNN (0.97), MobileNetV2 Raw (0.99), and MobileNetV2 CLAHE (0.99).

C. Classification Report

A comprehensive class-wise classification report was produced for the MobileNetV2 classifier trained on the CLAHE images tested using the validation dataset consisting of 196 in- stances. Results are presented in Table VI. The model achieves high levels of consistent performance in all four categories, with an overall average recall, precision, and F1 score of 0.98. None of the four categories appears to be significantly trailing behind others. The lowest recall is observed for Category 0 (Apple Scab), which has a recall value of 0.94.

Discussions

However, the obtained results demonstrate a quite unex- pected conclusion – the usage of CLAHE preprocessing, although visually enhancing the images by improving their contrast, did not lead to better classification accuracy in comparison with training the model on raw images. This is due to several reasons.

First, the base MobileNetV2 model, having pretrained on the Imagenet images, already contains a lot of valuable visual features in its weights, having been trained on a variety of images with different lighting and contrast levels. Therefore, the pretrained model is already capable of handling moderate differences in the contrast of the image. The application of CLAHE changes the distribution of pixel values in the image in a way that might interfere with the existing feature space. Another reason for the absence of improvements is that the PlantVillage images in this case were taken under quite con- trolled lab conditions, and, therefore, they lacked the problems with contrast. The effects of CLAHE could have been much more noticeable on the images acquired by smartphones in real conditions.

Finally, as the model was trained for just five epochs, it probably did not have enough time to adjust itself to new input data. It is quite likely that the increase in training iterations and leaRning rAte schedules will enhance the perfor of the model.

However, the high accuracy of the CLAHE model (98.74%, macro F1 = 0.98) proves that transfer learning based on Mo- bileNetV2 works well regardless of preprocessing operations and, in particular, that CLAHE cannot negatively affect the model performance.

Conclusions

In conclusion, this paper suggested a comparative analysis between three DL models for apple leaf disease classification on the PlantVillage dataset: a custom-designed CNN baseline, a MobileNetV2 model coached on the og images, and a MobileNetV2 model coached on CLAHE-enhanced images. The main results from this study include the following. First, transfer learning with MobileNetV2 achieves con siderably better accuracy when compared with training from scratch, as evidenced by MobileNetV2's 99.31% vs. the custom-designed CNN's 97.08%. Second, the hypothesis about the positive effect of CLAHE preprocessing was rejected, with the CLAHE-enhanced version yielding a slightly inferior peak validation accuracy (98.74% vs. 99.31%) and demonstrating that MobileNetV2 pretrained feature extraction layers can effectively cope with images even without additional contrast enhancement. Third, although CLAHE preprocessing did not help increase validation accuracy, CLAHE enhanced Mo bileNetV2 still managed to achieve outstanding performance for all four apple leaf diseases, with the overall balanced macro F1 score being 0.98. There are some important conclusions about designing plant disease detection systems. While CLAHE preprocessing is a highly valuable technique, especially useful in processing difficult-to-analyze real-world field images taken under chal- lenging lighting conditions, practitioners must assess its utility for their specific datasets and model architectures.

Further research directions include the following: (1) evalu- ating CLAHE enhancement effects on real-world field images taken under different environmental circumstances; (2) exper- imenting with fine-tuning the whole MobileNetV2 network; (3) testing more sophisticated augmentation techniques such as MixUp and CutMix; and (4) extending the model for use in different plants.

References

[1] S. P. Mohanty, D. P. Hughes, and M. Salathe´, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7,
p. 1419, 2016.

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[3] S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Transac- tions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345– 1359, 2010.

[4] A. M. Reza, "Realization of the contrast limited adaptive histogram equalization (CLAHE) for real-time image enhancement," Journal of VLSI Signal Processing Systems for Signal, Image and Video Technology, vol. 38, no. 1, pp. 35–44, 2004.

[5] D. P. Hughes and M. Salathe´, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv preprint arXiv:1511.08060, 2015.

[6] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311– 318, 2018.

[7] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," Computers and Electronics in Agriculture, vol. 161, pp. 272–279, 2019.

[8] G. Saleem, M. Akhtar, N. Ahmed, and W. S. Qureshi, "Automated analysis of visual leaf morphology for plant classification using deep convolutional neural networks," Ecological Informatics, vol. 51, pp. 18– 27, 2019.

[9] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 2018, pp. 4510–4520.

[10] R. C. Gonzalez and R. E. Woods, Digital Image Processing, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2008.

[11] K. Zuiderveld, "Contrast limited adaptive histogram equalization," in Graphics Gems IV, P. S. Heckbert, Ed. San Diego, CA, USA: Academic Press, 1994, pp. 474–485.

[12] A. Ramcharan, K. Baranowski, P. McCloskey, B. Ahmed, J. Legg, and D.
P. Hughes, "Deep learning for image-based cassava disease detection,"
Frontiers in Plant Science, vol. 8, p. 1852, 2017.

[13] T. Sharma, "Plant Village Dataset (Updated)," Kaggle, 2023. [Online]. Available: https://www.kaggle.com/datasets/tushar5harma/ plant-village-dataset-updated

[14] M. S. Charan and B. V. D. Mohammed Abrar, "Apple leaf diseases clas- sification using CNN with transfer learning," International Journal for Research in Applied Science & Engineering Technology (IJRASET), vol. 10, no. 6, pp. 1905–1912, Jun. 2022, doi: 10.22214/ijraset.2022.44176.

[15] "Deep learning models for classifying apple leaf diseases: A review," International Journal for Research in Applied Science & Engineering Technology (IJRASET), 2024. [Online]. Available: https://www.ijraset.com/research-paper/ deep-learning-models-for-classifying-apple-leaf-diseases