Using Machine Learning to Identify Crop Diseases with ResNet-18

Rihan Rahman
Skyline High School, 228th Ave SE, Sammamish, WA, U.S.A.

Keywords: ML, Machine Learning, Agriculture, Machine Learning Applications, Machine Learning in Agriculture, ResNet-18, Deep Learning, Computer Vision.

Abstract: Plant diseases are a highly prevalent issue in agriculture, causing countless farmers annually to face career threatening damages such as diminished profits and crop yields and environmental damages. Consequently, it is imperative that these diseases are quickly detected and treated against. An increasingly effective solution is to train convolutional neural networks (CNNs) using deep learning (DL). DL has several effective applications in a variety of major fields such as healthcare and fraud detection and has a high potential to solve issues of global significance. This research's goal is to create a machine learning (ML) model with DL to identify plants' diseases using photos of infected leaves. Many farmers in rural areas struggle to treat blights due to limited access to technology and information regarding them. Therefore, an ML model which can automatically identify these diseases would be highly useful for these people. After sourcing a comprehensive dataset with images of 88 types of plants and diseases, I used it to train a CNN model using several data augmentation techniques. With the model architecture ResNet-18, while evaluating its performance with a validation dataset, the model achieved a loss of 4.541%. This value demonstrates ResNet-18's applicability to the task of identifying plant diseases and illustrates the potential for classification-based DL networks to support rural farmers and the field of agriculture. If a superior model is created to identify blights more accurately, it should be used to help the billions of farmers who would greatly benefit from such technology.

1 INTRODUCTION

1.1 Background and Significance

Machine Learning has a nearly infinite number of possible applications and can have large effects on areas such as in data mining (Maglogiannis et al., 2007). With its Deep Learning capability, ML has the potential to solve many complex problems of global concern. DL is created through many hidden layers, with iterations of transformations and abstractions (Vargas, 2017). Consequently, DL models can adapt to new knowledge extremely quickly, making them the ideal solution for a variety of scenarios.

For instance, Mohsen et al. used DL to identify brain tumours using scanned images of the brain, demonstrating the potential to save lives by detecting cancer in early stages with ML (Mohsen et al., 2018). In another instance, Lee and Yoo used DL to predict domestic market prices based on information about foreign markets (Lee & Yoo, 2019).

1.2 Research Question/Objective

Many famers lack access to technology and information to identify and effectively treat plant diseases. As a result, I believe that creating a convolutional neural network (CNN) to identify these diseases would greatly assist these farmers. My goal is to explore the question of if ML can be used to identify blights from images of infected plant leaves.

When several trees in my local community became infected, identifying and treating the diseases became a long and tedious process. As a result, I was inspired to use a CNN to automate this. Creating an ML model would not only be more efficient, but also make disease identification and treatment far more accessible to farmers with limited experience and resources, enabling new growers to grow a variety of crops without fearing potential diseases.

Rahman, R.
Using Machine Learning to Identify Crop Diseases with ResNet-18.
DOI: 10.5220/0013123500003911

In *Proceedings of the 18th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2025) - Volume 1*, pages 311-315
ISBN: 978-989-758-731-3; ISSN: 2184-4305

2 LITERATURE REVIEW

2.1 Overview of Deep Learning

DL is a type of ML which utilizes artificial neural networks to replicate the human brain, producing a more capable model which can be used to solve problems of greater complexity. DL is used to address many modern issues in a variety of fields for this advantage, such as medicine, economics, and agriculture. It functions utilizing "neurons", which oversee processing data and are critical to the functionality of a DL model (Nielsen, 2015). Using ML models to classify objects within images is called Object Classification, an extremely popular branch of DL.

Furthermore, DL networks use several "layers" to separate the processes required for object classification. For instance, CNNs uses convolutional layers to process image input and scale it down using techniques such as Max Pooling, aiding networks to process information faster by simplifying images. This helps models understand what exact features to search for in an image during classification. As a result, object detection and classification are some of DL's major strengths (Zilong, 2018).

2.2 Related Studies and Projects

Deep learning for classification is a rapidly growing field and is becoming highly prevalent in our society. One example of image-based ML classification being utilized is in cancer research. Lung cancer classification research uses DL networks to analyse lung scans and determine if they are cancerous (Asuntha & Srinivasan, 2020).

In Asuntha and Srinivasan's study, the researchers used factors such as gradient, texture, and shape, to differentiate between images of cancerous and healthy lungs. The capability to precisely analyse data while focusing on factors like these demonstrates the flexibility of DL classification to be effectively used for many different purposes.

In Zhu et al.'s research, the researchers described many types of DL that can benefit the field of agriculture and educated readers about its functional benefits (Zhu et al., 2018).

3 METHODOLOGY

3.1 Dataset Description

This study uses the Kaggle dataset *Plant Disease Classification Merged Dataset* by Aline Dobrovsky, which contains pictures of whole infected and healthy leaves (Dobrovsky, 2023). This is a collection of 14 sub datasets and has 88 classes with over 17,000 images. Each subset contains images with varying backgrounds. None of the images were generated using pre-processing techniques such as flipping or zooming in (Dobrovsky, 2023). However, in this research, a reduced version is used to train the model due to technical limitations.

Figure 1: Sample labelled images from dataset. (Dobrovsky, 2023).

3.2 Model Selection and Architecture

For the task of identifying diseases from images of plant leaves, I chose a pre-trained ResNet-18 model due to its image classification capability. An example of this model's usage on a similar scenario is an ML study by Al-Falluji et al. (Al-Falluji et al., 2020). In this study, the researchers were able to use ResNet-18 to identify if a patient had COVID-19 by feeding X-ray images through the model. Finally, with ResNet-18, the researchers were able to classify patient's diseases with an accuracy of 96.73%.

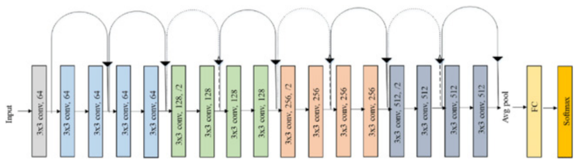

Figure 2: Diagram of ResNet-18 architecture (Ramzan et al., 2020).

ResNet-18 uses 18 layers to process data, supporting the complexity required for DL, with the activation function ReLU (Rectified Linear Unit). ReLU determines if a neuron should be activated and is critical for CNNs. By selectively activating only the necessary neurons, the model can run with much higher efficiency. ResNet-18 has about 11 million parameters which contribute to its precision and adaptability.

3.3 Data Pre-Processing and Augmentation

According to Wang and Perez. "Data augmentation has been shown to produce promising ways to increase the accuracy of classification tasks." (Wang & Perez, 2017). Therefore, I used pre-processing and data augmentation to artificially create more data while improving the model's accuracy and its ability to tolerate a greater variety of images. During pre-processing, I first conformed all the images to a format of (244x244) pixels, helping the network to quickly analyse each image without having to adapt to multiple resolutions.

The augmentation Horizontal Flip, which feeds the model a reflection of the original image, significantly improved the dataset's variety, helping to reduce the model's angular bias and preparing it to classify real photographs. Additionally, the Random Rotation augmentation addressed images which are not horizontally centered, improving the model's ability to analyse realistic photographs, irrespective of the angle. Another mechanism I used was feeding randomly cropped images to the network, which prepared the model to analyse non-perfect images, in which leaves are not fully captured. Finally, centering the images' leaves made the data more consistent and helped the model to adapt easier.

3.4 Model Training Process

Post-processing the data, I chose the optimizer ADAM (Adaptive Movement Estimation) to accelerate training speed, shorten training windows, and increase acquired accuracy. According to Liu et al., ADAM improved tested accuracy by roughly 3 percent (Liu et al., 2021). Also, with a learning rate value of 0.001, the model spends more time on each image, thus improving its growth. Additionally, using a batch size of 400 images and several epochs expedited the training process by effectively utilizing the machine's processing power.

4 EXPERIMENTAL RESULTS

4.1 Experimental Setup

The CNN was developed using Google Colab with a T4 GPU, which helped to increase the model training speed and assisted with running the code faster. Machine configuration as follows: Windows 11 OS, RTX 3080 GPU, and Ryzen 9 5900x CPU. For ML, I used the ResNet-18 model from Pytorch's ML libraries and Anaconda Python Environment tools. ResNet-18 is a CNN that is 18 layers deep and pretrained with more than 1 million images from the ImageNet Database.

I used a large dataset from Kaggle, which encompassed multiple sub datasets. Both RestNet and the dataset played a key role in learning rate and accuracy of the training. With Torchvision from Pytorch, I was able to augment the data, expanding the data set artificially. This helped to increase the volume and variety of dataset, helped to train the model with a varied set of images and increased accuracy on the non-training dataset. For translations, I used random horizontal flips, random rotations, and random crops.

The CrossEntropyLoss function from Pytorch helped to calculate the loss during training, which was then used to update the model's neural parameters to decrease the loss and improve its accuracy. Adjusting in response to the loss value calculated by the loss function was critical to the model's growth and helped to improve the model's performance. The loss function showed similar effectiveness in the validation loop as well.

Pytorch CrossEntropyLoss Function:

$$H(p, q) = -\sum_{x \epsilon X} p(x) \log q(x) \tag{1}$$

Where x is the number of classes, p(x) is the actual label for class x, and q(x) is the predicted probability for class x.

4.2 Performance Evaluation and Analysis

The CrossEntropyLoss loss function from Pytorch, also known as Softmax or Log Loss, was used as the main evaluation metric in all accuracy calculations. After running the validation and test loops, the loss function returned a loss value of 4.541%. This result matched my expectations given the limitations of my experiments; however, it demonstrates the viability of using ML to identify plant diseases. This loss value indicates that the model is working well, although there is room for improvement. According to Muhammad et al., relevant state-of-the-art classification models can achieve accuracies in the range of 96-97.5% (Muhammad et al., 2023). While my model's accuracy was below this, its development demonstrates and supports the spread of ML and DL technology to be applied to agriculture. On the test set, my model showed that its accuracy increased with the number of epochs until the epoch count reached

15. After reaching 15 epochs, the accuracy on the test set started to decrease because of the model becoming too over-specified for the training dataset. This matches the findings of Komatsuzaki's research, stating that too many epochs lead to models becoming not generalizable and losing accuracy due to being over-specified (Komatsuzaki, 2019).

4.3 Discussion of Findings

The model's loss value of 4.541% shows that there is a large area for growth in terms of overall accuracy. However, it is reasonable given I was only able to use a fraction of the overall training data specified earlier due to technical issues. However, the fact that the model still scored a loss of 4.541% despite this demonstrates that my training methodology was correct and can be utilized to create a more accurate model. Given better testing conditions and software, the technical issues which prevented the use of the entire training dataset to train the model could have been avoided. In this case, the model would be much more accurate and applicable to helping farmers identify diseases. However, it is clear that a strength of the model is its quickness to learn. Given such a small amount of training data, it still managed to achieve a decent accuracy, showcasing the effectiveness of ResNet-18 to train classification models.

When testing the data, the constant growth in response to more epochs demonstrated the benefits of iterations. I initially believed that the model would become over-specified at 11 epochs. However, the consistent growth until 15 epochs impressed me and demonstrated ResNet-18's strength in being able to grow consistently without becoming over-specified. I believe with 15 epochs of a large dataset, a model using ResNet-18 could become very accurate and possibly rival state-of-the-art models. Therefore, if this experiment were to be replicated, I would recommend 15 epochs to be used to train the ResNet-18 model. This would lead to the greatest accuracy with the lowest loss and produce the most effective model.

5 CONCLUSIONS

5.1 Summary of Findings

In summary, the model I created using ResNet-18 and *Plant Disease Classification Merged Dataset* performed well and identified plant diseases accurately. This research demonstrates the applicability of ML in addressing global issues such as plant diseases in agriculture and aligns with the research of Liu et al. (Liu et al., 2021) and Huang (Huang, 2007). There is a great opportunity for ML to be used to support farmers worldwide, and with improved designs, data, and computing power, this goal is highly achievable in the near future.

5.2 Applications, Limitations, and Future Works

With its loss value of 4.541%, this model can be improved by using the entirety of the original dataset. Additionally, the technical limitations can be addressed with a better design environment and Compute power. With similar methods on a larger scale, the model could further be tuned for wider range of diseases and can be integrated with a tool to identify plant diseases. A model capable of detecting plant diseases accurately could drastically improve crop yield and help to feed the millions of people lacking access to food. I suggest my research be used and fine-tuned to aid the millions of farmers who lack information about plant diseases worldwide by creating a physical tool implementing ML technology.

Such innovation could immensely improve farmers' crop yields, potentially providing billions of people across the world access to a consistent food source. Similarly, further developed classification models could be used for a wider range of purposes beyond the field of agriculture. Classification as a tool can be used in numerous ways to help billions of people worldwide. Using DL classification effectively could be the key solution to a multitude of significant global issues faced by billions of people across the world.

5.3 Limitations and Future Work

Although the model functioned decently well, it can be improved with superior methods and technology. For instance, as previously mentioned, the technical issue preventing the use of the entire dataset could be mitigated with advanced software and technology. Additionally, the capability for experimentation was shortened due to limited compute units on Google Colab which inhibited my ability to further refine the model. Furthermore, with more diverse data, the model could be improved to include a wider range of plants rather than just 88 total varieties. Implementing these methods could develop a model far more practically applicable for the task of aiding rural farmers. If a ML framework is created specifically to identify plant diseases, this would also drastically improve the accuracy of the model, as it would look for more specific details in the leaves to determine diseases. Or if multiple other model architectures, such as GoogLeNet and EfficientNet, were used to determine which model performed the best, this could also contribute to the development of a more accurate model.

I recommend my research be studied, utilized, and implemented, to aid the millions of farmers affected annually by plant diseases. In the future, with technological advancements, this research could be a stepping stone in the path toward the globalization of ML technologies in agriculture.

REFERENCES

Al-Falluji, R. A., Katheeth, Z. D., & Alathari, B. (2020). Automatic detection of COVID-19 using chest X-ray images and modified ResNet18-based Convolution Neural Networks. Retrieved from https://www.techscience.com/cmc/v66n2/40667

Asuntha, A., & Srinivasan, A. (2020). Deep learning for lung cancer detection and classification - multimedia tools and applications. Retrieved from https://link.springer.com/article/10.1007/s11042-019-08394-3

Dobrovsky, A. (2023). Plant Disease Classification merged dataset. Retrieved from https://www.kaggle.com/datasets/alinedobrovsky/plant-disease-classification-merged-dataset

Huang, K.-Y. (2007). Application of artificial neural network for detecting phalaenopsis seedling diseases using color and texture features. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0168169907000385

Komatsuzaki, A. (2019). One epoch is all you need. Retrieved from https://arxiv.org/abs/1906.06669

Lee, S. I., & Yoo, S. J. (2019). Multimodal Deep Learning for Finance: Integrating and forecasting international stock markets - The Journal of Supercomputing. Retrieved from https://link.springer.com/article/10.1007/s11227-019-03101-3

Liu, Z., Shen, Z., Li, S., Helwegen, K., Huang, D., & Cheng, K.-T. (2021). How do adam and training strategies help BNNS optimization. Retrieved from https://proceedings.mlr.press/v139/liu21t.html

Maglogiannis, I. G. (2007). *Emerging artificial intelligence applications in computer engineering: Real word AI systems with applications in eHealth, HCI, information retrieval and Pervasive Technologies*. Amsterdam: IOS Press.

Mohsen, H., El-Dahshan, E.-S. A., El-Horbaty, E.-S. M., & Salem , A.-B. M. (2017). Classification using deep learning neural networks for brain tumors. Retrieved from https://www.sciencedirect.com/science/article/pii/S2314728817300636

Nielsen, M. (2015). Retrieved from https://www.ise.ncsu.edu/fuzzy-neural/wp-content/uploads/sites/9/2022/08/neuralnetworksanddeeplearning.pdf

Ramzan, F., Khan, M. U. G., Rehmat, A., Iqbal, S., Saba, T., Rehman, A., & Mehmood, Z. (2019). A deep learning approach for automated diagnosis and Multi-class classification of alzheimer's disease stages using resting-state fmri and residual neural networks - Journal of Medical Systems. Retrieved from https://link.springer.com/article/10.1007/s10916-019-1475-2

Shoaib, M., Shah, B., EI-Sappagh, S., Ali, A., Ullah, A., Alenezi, F., … Ali, F. (2023). An advanced deep learning models-based Plant Disease Detection: A review of recent research. Retrieved from https://www.frontiersin.org/journals/plant-science/articles/10.3389/fpls.2023.1158933/full

Vargas, R., Mosavi, A., & Rulz, R. (2017). (PDF) Deep learning: A review. Retrieved from https://www.researchgate.net/publication/328285467_Deep_Learning_A_Review

Wang, J., & Perez, L. (2017). The Effectiveness of Data Augmentation in Image Classification using Deep Learning. Retrieved from http://vision.stanford.edu/teaching/cs231n/reports/2017/pdfs/300.pdf

Zhu, N., Liu, X., Liu, Z., & Hu, K. (2018). (PDF) Deep Learning for smart agriculture: Concepts, tools, applications, and opportunities. Retrieved from https://www.researchgate.net/publication/327016437_Deep_learning_for_smart_agriculture_Concepts_tools_applications_and_opportunities

Zilong, H., Jinshan, T., Ziming, W., Kai, Z., Ling, Z., & Qingling, S. (2018). Deep learning for image-based cancer detection and diagnosis − A survey. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0031320318301845