ITM Web of Conferences **70**, 03017 (2025) https://doi.org/10.1051/itmconf/20257003017
DAI 2024

Exploring CycleGAN Technique for Improved Plant Disease Detection and Analysis

Luyi Ouyang

Olin Business School, Washington University in St. Louis, St. Louis, MO, 63l30, United States of America

Abstract. The threat of plant disease is a significant problem affecting the world, when untreated these diseases can affect food production. Diagnosis of these diseases in an un-delayed manner is very important, however, methods described in current use that only involve the use of sight are inefficient and are also subject to errors. This paper tackles the problem by using Cycle-Consistent General Adversarial Networks (CycleGAN) to create artificial images of diseased plant leaves. The advantage of this approach is that augmenting the training data with images that do not exist in the real world helps improve the performance of disease classifications. The research takes into consideration the apple leaves diseased images, is of various pathogens, and CycleGAN creates images to even it. The results indicate that CycleGAN is indeed able to generate artificial images for the less complicated sicknesses associated with a mere shift in color, with an achieved micro-average Area Under the Curve (AUC) of .98 and macro-average AUC of 0.94. On the contrary, this model has problems in striking a balance while dealing with more complex diseases that have problems that are underlying structural deformation. However, adding such images in training datasets increases the classification accuracy in total. Future work should involve making the model more robust to complex and rich visual details as well as employing more sophisticated models for better applicability in real farming settings.

1 Introduction

Plant diseases pose a greater problem to agriculture. One of the significant problems related to plant disease detection is the lack of sufficiently diverse and balanced datasets. This is especially true concerning the development of strong generalist models considering the variations in disease symptoms of crops and the limited number of case studies of rarer diseases. Generative models, and in particular Generative Adversarial Networks (GANs), solve this issue effectively by resourcing. GANs, proposed for the first time by Goodfellow et al. in 2014, have the particular power of creating realistic synthetic data that can be used as supplementary datasets in a wide range of tasks and lead to improved models' performance.

The Special Value of these and similar models resides in the fact that they can generate some uncommon and sophisticated forms of illness that may be present in the data generation approach but not sufficiently in the datasets themselves. To overcome the data imbalance issues, GANs help with producing synthetic imagery across a spectrum of disease forms which enhances the performance of the machine learning models over such conditions. For example, Wang and Cao's study showed that including pictures constructed by GANs in the used dataset led to 12% better detection accuracy, which proves the modeling efficiency for rare disease models [1]. This shows how important these generative models are in solving the issues of improving the precision and reliability of plant disease detection approaches.

In a related study, Bing employed GANs to tackle similar challenges, focusing on enhancing Convolutional neural network (CNN) performance through data augmentation [2]. Their approach resulted in a significant 15% improvement in model performance by enriching training datasets with high-quality, GAN-generated images. This method not only addressed the scarcity of diverse training examples but also reduced the need for extensive and costly data collection efforts in varied agricultural settings. Amreen Abbas further explored the utility of GANs in plant disease detection, particularly through advanced image augmentation techniques [3]. Her research contributed to a 14% improvement in detection accuracy, highlighting the potential of GANs to produce complex, varied images that significantly enhance the training data quality for machine learning models in agriculture. This progress demonstrates the practical benefits of synthetic data in environments where collecting diverse and comprehensive datasets is often challenging and resource-intensive.

In conclusion, the use of GANs for detecting plant diseases marks a milestone in AI adoption within the field of agriculture. By addressing class imbalance, variability in agricultural conditions, and scarcity of data. GAN may revolutionize diagnosis management control systems against crop ailments. Using the dataset from the plant pathology which contains high-resolution images of leaves categorized by various plant diseases, this study employs CycleGAN which was introduced by Zhu to explore innovative methods for enhancing crop disease detection [4].

This article uses CycleGAN-based methods to generate synthetic images of diseased plant leaves and studies their impact on improving classification performance. To develop scalable and generalizable models, the approach is designed to be integrated with diverse farming practices to bolster global food security and promote sustainable agriculture.

2 Method

2.1 Dataset and resource description

This investigation uses the Plant Pathology dataset available on Kaggle, which consists of high-quality photos of the leaves of apples and corresponding annotations for scab, rust, multiple disease, and healthy state samples [5]. These images are important in the training process for the models in the detection and classification of plant disease ensuring the models can perform well regardless of different agricultural practices employed. The variation in the datasets with regards to the various types of diseases patients are affected by provides a very strong base for the implementation of advanced image processing techniques that are essential in the creation of efficient algorithms aimed at disease detection. This dataset has also been useful in effectively training and testing independent generative models as conducted by Wang who used GANs to balance plant disease datasets and Bing and Hu who focused on GANs to enhance classifiers with limited data.

2.2 Methodology

This research uses cutting-edge techniques based on CycleGAN to improve plant disease detection by generating synthetic images for training machine learning models. First of all, CycleGAN helps in solving problems related to lack and imbalance of data by creating pictures of different plants filled with healthy leaf samples devoid of any disease. These images ensure rather more balanced and varied datasets which are important in building strong machine learning models. In the training stage, different learning rates and batch sizes are tuned among several other parameters that seek to improve the learning effectiveness of both CycleGAN. Indeed, the tuning processes are in themselves necessary as some studies have indicated for example by Karthik where attention-embedded CNN models were perfectly optimized for tomato leaf disease detection using CNN optimally [6].

Additionally, Fawakherji and Wang aimed at incorporating multispectral data for crop and weed segmentation and defect detection of lychee surfaces [7, 8]. Similarly, such studies also extend the use of GANs beyond plant disease detection convincing that these kinds of models can be applied in a wide range of agricultural practices from crop monitoring to precision farming. On the other hand, the entropy-based feature selection framework for grape leaf disease detection developed by Adeel also gives credence to the assertion that improvement in data quality regardless of whether it is featuring selection or synthetic augmentation entails enhancement in model accuracy [9]. This is also in conformity with the enhancements in this study through the use of CycleGAN. In addition, Jung talks about how deep learning models for the agriculture crop can be constructed especially with the help of superiorities that the models receive from the images created by GAN [10]. Their work gives further insight into the methodology used in this research work, stating the fact that deep learning models for the detection of plant diseases are elevated on the pedestal of augmentation and balancing of data sets.

Consistent with earlier work on GAN applications in agriculture, this study tests the efficacy of CycleGAN in plant disease detection using different measures. Loss plots which are employed to assess the stability of learning during training are used to present the performance of the model throughout training its patients. Receiver Operating Characteristic (ROC) curves are employed to evaluate the model's capacity to classify diseased and healthy leaves. The Area Under the Curve (AUC) describes the sensitivity and specificity of the model about several disease statuses, thus determining the accuracy of its classification task. With these methods, the investigation provides a conclusive evaluation of the advantages and shortcomings of the model providing the basis for the application of CycleGAN within more complex agricultural systems.

Achieving these aims involves the use of high-end generative models such as CycleGAN, which the researcher intends to apply to create advanced detection and classification systems for plant diseases that are broad, effective as well as cost-efficient and have the potential to greatly improve plant disease management and farming. Thanks to recent developments in synthetic image generation and optimally tweaking the model parameters, the models can solve the problems of plant diseases and their management in both optimal experimental designs as well as in real agricultural settings [11]. This approach shows that Artificial-Intelligence-GAN-based technologies can receive a widespread application in the fundamental changes in plant pathology aimed at a more effective, early identification and treatment of diseases, decreased losses of cultivated crops, and healthier farming technologies.

3 Result

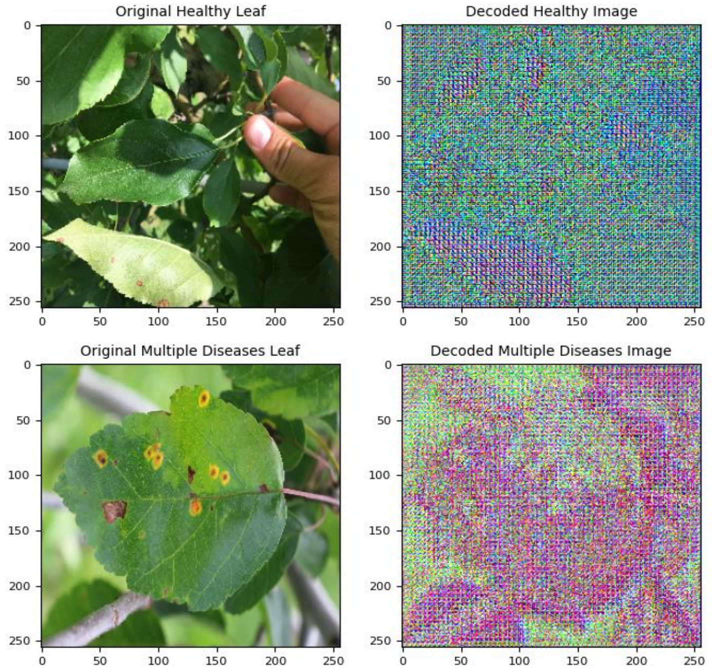

Fig. 1 Original and decoded leaf images using CycleGAN model: The top row shows the original healthy leaf image (left) and its corresponding decoded healthy image (right). The bottom row displays the original leaf affected by multiple diseases (left) and its decoded diseased image (right) (Photo/Picture credit: Original).

The CycleGAN model can generate images of healthy and diseased leaves as shown in Fig. 1. These images are accurately reconstructed down to finer details including the texture and vein patterns within the healthy leaves although the reconstructed well images do contain more information about the latent space which cannot be clearly discerned and therefore is seen as pixelated. In the case of the diseased leaves, the model does well to illustrate major attributes of the disease such as yellowing and browning patches, though more emphasis was laid on color and lesions than the delicate body structures The initial observation after differentiating the images of healthy leaves and the diseased leaves has been the gross visual manifestation of the variations between them which in this case depicts the healthy leaves as smooth and less damaged in comparison to diseased leaves. This shows the model's capability to correctly classify the images of healthy and diseased trees and even create synthetic diseased images that could be used in a dataset where there may be a disproportional number of diseased and healthy images.

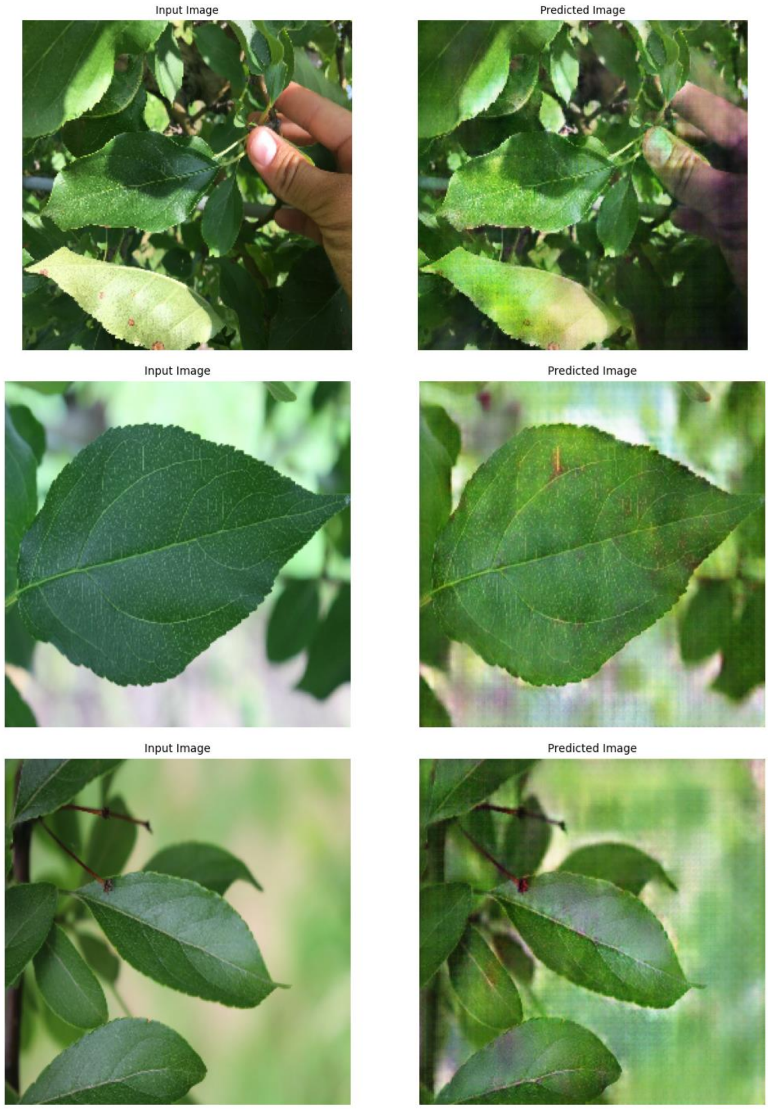

Fig. 2 Comparison of original and predicted leaf images: the figure presents three sets of input images) and corresponding predicted images generated by the CycleGAN model (Photo/Picture credit: Original).

The figure obtained from the CycleGAN (Fig. 2) used to equalize leaves' diseased and healthy surfaces is further illustrated and proved by the visual comparison of real leaf images with predicted leaf images. The original images of leaves contain healthy ones while the predicted images by the CycleGAN model contain abnormal disease symptoms which were not presented. This means that the model can imitate the symptoms of diseases thereby fully accomplishing the task of mapping from normal leaf condition to diseased leaf condition.

 https://doi.org/10.1051/itmconf/20257003017

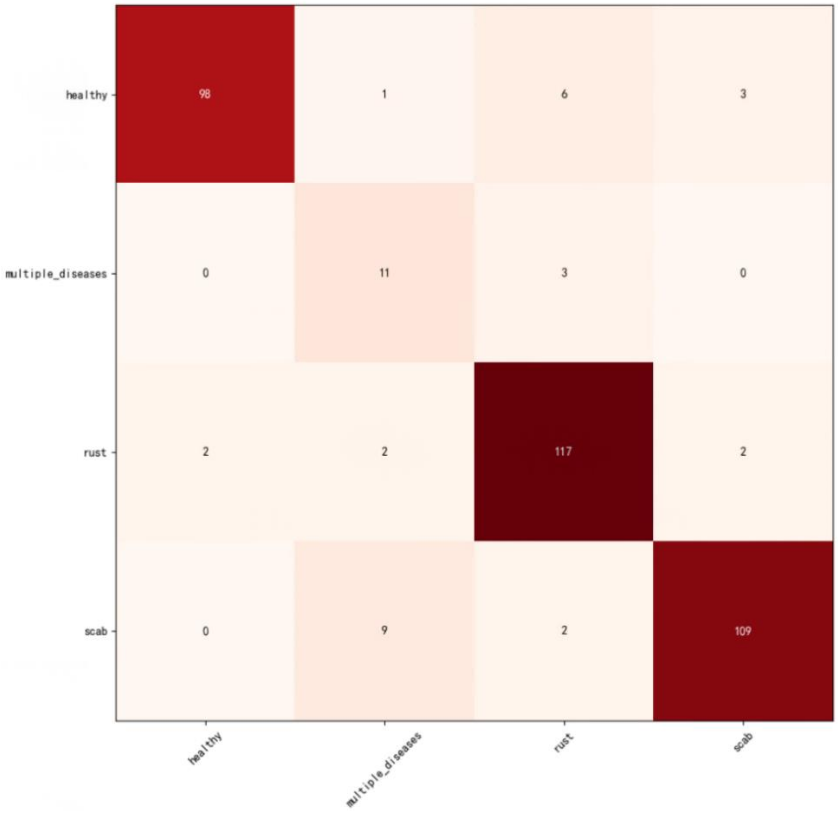

Fig. 3 Confusion matrix of the CycleGAN model for multi-class plant disease classification (Photo/Picture credit: Original).

As for the model evaluation, the CycleGAN model proved beneficial in producing imitation images in the enhancement of plant disease detection and this was verified both visually and quantitatively. A confusion matrix (Fig. 3) was used to analyze the classification accuracy of the model concerning the four disease classes: healthy, rust, scab, and multiple diseases. The matrix of confusion affords other useful parameters concerning the merits and demerits of the model. Much success in this task was achieved in the classification of healthy foliage with a perfect 91 of 107 instances classified correctly but a few of these correct identifications were in leaves with rust and scab diseases which caused some errors in classification with 7 and 6 instances respectively. Rust and scab-infected leaves classification were generally superior with 112 and 118 out of 129 and 124 correctly classified instances respectively. The problem with this model was the classification of samples with multiple diseases, where 5 samples were classified correctly and others were misclassified to rust and scab. Thus, the above analysis, along with the ROC curve analysis, points to the fact that the model handles the least complex disease cases very well but when it comes to handling the complicated cases like a leaf with multiple diseases, some work still needs to be done.

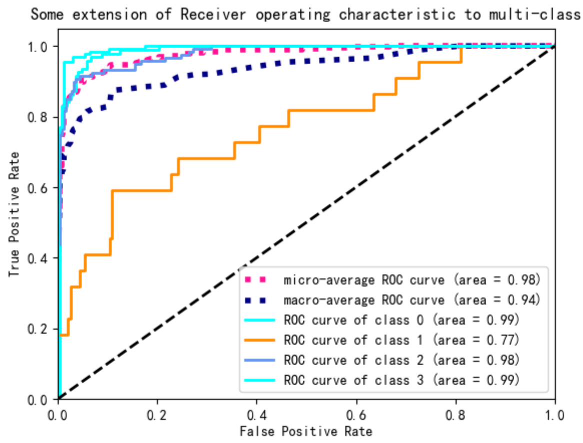

Fig. 4 Multi-class ROC curve for plant disease classification using CycleGAN (Photo/Picture credit: Original).

To better assess the validity of the model, it was evaluated further with the ROC analysis that was performed (Fig. 4), which assesses the discriminatory power of the model about various disease classes. The micro-average ROC curve gave a remarkable AUC of 0.98, which also demonstrates a good overall classification of the target attributes. The macro-average AUC was 0.94 signifying consistency in the model across various disease classes.

Class 0 (Healthy Leaves): The classification of the healthy leaves was impressive recording an AUC score of 0.99 implying perfect discrimination.

Class 1 (Scab-Infected Leaves): In the case of scab-infected leaves, however, this particular model had its lowest AUC of 0.77. This indicates difficulties in recognizing more complex structural perturbations due to this disease.

Class 2 (Rust-Infected Leaves): In the case of rust-infected leaf detection model yielded an AUC of 0.98, which serves to support the management of diseases with gross signs of disease in the visual aspects.

Class 3 (Multiple Diseases): For leaves with multiple diseases, the model also performed strongly, with an AUC of 0.99.

The current analysis also supports the previous findings in particular the model's capability to generate synthetic images of less complicated diseases such as rust or multiple diseases although scab is a more advanced disease characterized by complex structural variation which the model tends to have problems with. These results illustrate that CycleGAN is for sure able to tackle the issues of dataset imbalance in the field of plant disease detection in Goa but optimization for other more complex feature diseases is required for better performance in classification.

For further assessment of the performance of the model, the training and validation loss curves were plotted and analyzed. Both curves display some level of instability in the training process, oscillating within a span of 50 epochs

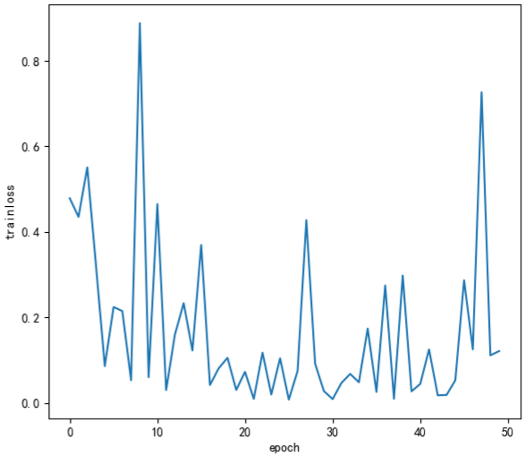

Fig. 5 Training loss curves of the CycleGAN model (Photo/Picture credit: Original).

The graph representing training loss (Fig. 5) does not remain constant and has significant variation, especially at the beginning epoch and the later epochs indicating that there was no stability during the training phase. The spikes in the first 10 epochs could be viewed as indications that the model was failing to learn properly at the early stages and these troubles continued throughout the training process as well sometimes with peaks even at the latter stages, for instance, epoch 35 and epoch 47. Despite the general trend of loss decreasing, suggesting learning by the model, there are sudden spikes that may indicate overfitting and or difficulty in learning the data. It is recommended that if hyperparameters like learning rate are changed, or regularization or learning rate scheduling are added, this will prevent the late phase fluctuation and enhance the convergence of the model.

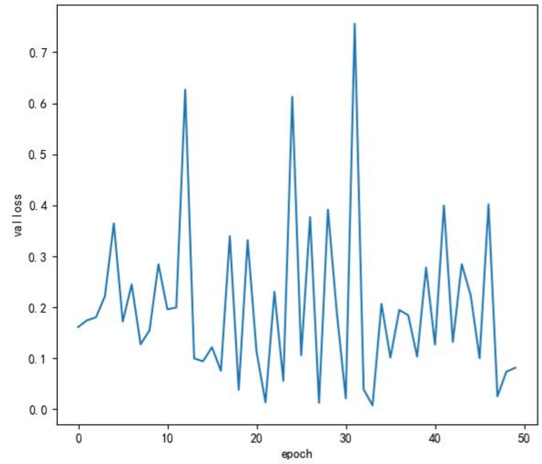

Fig. 6 Validation loss curves of the CycleGAN model (Photo/Picture credit: Original).

The graph for validation loss (Fig. 6) shows wild ups and downs over the 50 epochs along with deep spikes at different epochs with more intense spikes at about ten epochs such as at epochs 10, 25, and 30 where the loss exceeds 0.6. However, even though the general pattern depicts a loss reduction, there are still ongoing fluctuations that depict further potential instability in the validation phase where the model might be having difficulties generalizing throughout the varied dataset. The more abnormal the peaks and quick losses are, the greater the opportunities for overfitting or there are mismatches in training and validation data.

Improving hyperparameters tuning, regularization or early stopping can be better practices to have validations that are more even and steady thereby increasing generalization.

That individual losses decrease is a good sign as this means the model can learn and reduce errors satisfactorily; there is however a large rise in the training as well as validation losses raising alarm on issues of convergence especially in the more complicated situations like scab infected leaves where previously an area under the curve of 0.77 was achieved. Such instability could be attributed to the model's inability to capture diseases that have non-linear geometries as evidenced by the ROC plot. Despite these variations, the increasing trend in loss means that there is some degree of improvement over time, which enhances the model being learned. These conclusions also correspond with earlier findings that the model could handle simple diseases like rust and multiple diseases but further improvement is still required to handle scabs and other visual complexity.

However, in the case of generating synthetic images for certain diseases, the CycleGAN model performed adequately in some aspects but had some difficulties with more complex features. The easier visual manifestations of the diseases like a uniform color change were generated fairly well by the model as compared to more complex structural disease features. While this could be so, the model often has a problem with comprehending the axial and textural parts because of the pixel transform skill which is difficult. Also, the dataset might have been imbalanced and biased toward less complicated disease features, thus leading to an overrepresentation of synthetic generation. The model, therefore, produced results that focused on visual features such as color that are more prominent, instead of concentrating on subtle features that are equally or even more important. That is to say, while CycleGAN has its uses, more sophisticated tools may be required when dealing with image-related illnesses that have more intricate features.

4 Discussion

From this research, it has been shown clearly that the CycleGAN model is a good method for creating synthetic images, especially in situations with simple plant diseases like rust. The images created helped in the increase in classification performance, with a micro-average AUC of 0.98 and macro-average AUC of 0.94. This illustration is also supported in this study which shows that the model effectively learned visual characteristics of diseases such as rust, and multiple simpler diseases such as bruchid beetles infested turnips. This supports findings from previous studies, stating the advantage of GANs in generating synthetic data necessary for training models in situations where data is sparse.

While the limitation of the CycleGAN model in terms of performance has been demonstrated in the results, other limitations also exist. As for scab or any other disease with more complicated shifts in the structure of the tissues, the results from the model were much worse. As such, this can be seen in the avoidance score of 0.77 for those infected with scab leaves. A sit is evident from the analysis that the lower AUC means that the model cannot be able to integrate the complex structural changes in the scab with a mixture of different colors and fine textures. The oscillation present in the training and validation loss graphs also confirms further the existence of certain disease complexities that the model finds challenging to handle. The variation within the loss curves, particularly during the initial 50 epochs, also suggests some convergence concerns. This instability might be caused by the model's shortcomings in the learning and representation of certain diseases of non-linear structure such as those of scabbed leaves. Alternately, the instability could be caused by the imbalance of the dataset, whereby the model is often presented with simpler diseases, which it masters well hence acquiring those representations instead of the complex or less common ones. It is not difficult to find that CycleGAN is not without its shortcomings as it has difficulty managing changes that are usually geometric in nature or delicate in detail [12].

Also, it is important to control at the same time such errors as approximation, that arise due to the model not learning the best mapping fully, and the estimation error, which is caused by the small set of training data, since it is critical to limiting overall risk and attaining good quality translations [13].

5 Conclusion

This study proves that the CycleGAN model can create artificial images of diseased plant leaves, which allows for improving the classification performance, especially in combat against imbalanced datasets. The model demonstrates the highest efficiency in image synthesis for less complicated diseases such as rust and a combination of diseases which led to high accuracy of classification. The micro-average AUC of 0.98 and the macro-average AUC of 0.94 indicate good discrimination in ROC analysis and thus support the general performance of the model in different disease classes. Almost perfect differentiation was observed between healthy leaves and rust-infected leaves while the model was successful on scab disease which had a lower AUC of 0.77 for more complex diseases. The confusion matrix additionally stresses that the model is well able to discriminate between healthy leaves and rust-infected ones but there are problems with leaves with other diseases and those that are scab-infected, where most of the errors occurred. The training and validation loss curves however show some fluctuations in the course of the training implying that some optimization of model hyperparameters including learning rate or batch size among others is required for optimal and consistent convergence.

In conclusion, CycleGAN architecture performs well in generating synthetic images for plant disease diagnosis, but further optimizations are necessary for its application to more sophisticated disease categories. These positive outcomes show in addition that the application of GANs to modify the dataset-related imbalance is workable and suggests that these methods could be helpful in other areas such as agricultural disease diagnosis. Thus, it is important to broaden this method and apply it to other crops and diseases where such an issue also prevails.

There is a need for better models that can deal with fine and gross deformation of the structure but not lose the important features of the image. Additionally, the validation of images using visual examination presents a major factor where the study is based entirely, which is quite useful although it is subjective. Subsequent studies should emphasize bringing in more quantitative approaches to address the concerns with the images produced. In evaluating the results, the use of distances like Fréchet Inception Distance (FID) may be practical in eliminating the bias in taking images. In this regard, further studies should concentrate on more sophisticated GAN architectures, such ‘as’ or ‘such and’ BigGAN networks to better diagnose various plant diseases. The model may also benefit from the use of additional sources of information, for instance, future work will also focus on the analysis of subtle features using infrared or multispectral imaging, which cannot be seen with standard RGB images There is also the hope that more reliable models for disease detection will be developed that are more applicable in various.

References

1. C. Wang, & L. Cao. GACN: Generative adversarial classified network for balancing plant disease dataset and plant disease recognition. Sensors, 23(15) (2023).
2. L. Bing, & G. Hu. Improving image-based plant disease classification with generative adversarial network under limited training set. Frontiers in Plant Science (2020).

3. A. Amreen, J. Sweta, G. Mahesh, & V. Swetha. Tomato plant disease detection using transfer learning with C-GAN synthetic images. Computers and Electronics in Agriculture (2021).
4. J. Y. Zhu, T. Park, P. Isola, A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint (2017).
5. C. Kaeser-Chen, & S. Dane. (2020). *Plant Pathology 2020 - FGVC7*. Kaggle. https://kaggle.com/competitions/plant-pathology-2020-fgvc7
6. R. Karthik, M. Hariharan, S. Anand, P. Mathikshara, A. Johnson, & R. Menaka. Attention embedded residual CNN for disease detection in tomato leaves. Applied Soft Computing (2020).
7. M. Fawakherji, V. Suriani, D. Nardi, & D. D. Bloisi. Shape and style GAN-based multispectral data augmentation for crop/weed segmentation in precision farming. Crop Protection (2020).
8. C. Wang, & Z. Xiao. Lychee surface defect detection based on deep convolutional neural networks with GAN-based data augmentation. Agronomy, 11(8) (2021).
9. A. Adeel, M. A. Khan, T. Akram, A. Sharif, M. Yasmin, T. Saba, & K. Javed. Entropy-controlled deep features selection framework for grape leaf diseases recognition. Expert Systems with Applications (2021).
10. M. Jung, J. S. Song, A. Y. Shin, B. Choi, S. Go, S. Y. Kwon, J. Park, S. Park, & Y. M. Kim. Construction of deep learning-based disease detection model in plants. Scientific Reports (2023).
11. C. Wang, & L. Cao. GACN: Generative adversarial classified network for balancing plant disease dataset and plant disease recognition. Sensors, 23(15) (2023).
12. L. G. Divyanth, D. S. Guru, P. Soni, R. Machavaram, M. Nadimi, & J. Paliwal. Image-to-image translation-based data augmentation for improving crop/weed classification models for precision agriculture applications. Algorithms, 15(11) (2022).
13. L. Sun, D. Shen, & H. Feng. Theoretical insights into CycleGAN: Analyzing approximation and estimation errors in unpaired data generation. arXiv preprint (2024).