Rethinking GAN-Augmented Data: A Case Study on Leaf-GAN for Tomato Leaf Disease Classification

Prudhvi Sai Mandalapu
Department of Artificial Intelligence
Kyungpook National University,
Daegu, South Korea
Email: prudhvisaimandalapu12@gmail.com

Adeyinka P. Adedigba
Department of Artificial Intelligence
Kyungpook National University,
Daegu, South Korea
Email: yinkpeace@gmail.com

Rammohan Mallipeddi
Department of Artificial Intelligence
Kyungpook National University,
Daegu, South Korea
Email: mallipeddi.ram@gmail.com

***Abstract*—Deep learning models have been developed for automated plant disease classification. However, these models often suffer from limited data diversity, especially when training data does not reflect real environmental conditions. To address this issue, generative models have been proposed for data augmentation. Despite their potential, generative models face challenges such as mode collapse, which prevents them from guaranteeing sufficient diversity in the generated data. As a result, data augmentation using generative adversarial networks (GANs) can sometimes degrade algorithm performance. Recently, Leaf-GAN was proposed to generate diseased leaf images from healthy ones, to augment datasets and improve the performance of deep learning models. While prior work claimed performance gains, this study investigates the effect of Leaf-GAN–generated synthetic images on a real-world tomato leaf disease dataset. Our results show that when 10% synthetic data is added, the accuracy of all tested models decreases. Increasing the proportion of synthetic data to 25% leads to a continued decline in performance. These findings suggest that Leaf-GAN–generated data may introduce inconsistencies that hinder model generalization in the tomato disease classification task.**

Keywords—Deep Learning, LeafGan, Data Augmentation, Plant disease detection, Tomato leaf, Mode Collapse

I. INTRODUCTION

Plant disease detection plays a critical role in ensuring crop health and productivity, especially in large-scale agricultural systems where timely intervention can prevent severe yield losses. In recent years, deep learning [1] has emerged as a powerful tool in plant disease detection due to its ability to automatically learn hierarchical and discriminative features from large image datasets without the need for handcrafted feature engineering. Among deep learning approaches, convolutional neural networks (CNNs) and, more recently, transformer-based models have demonstrated impressive performance on benchmark datasets, enabling reliable detection of disease symptoms from leaf images [1]. These advancements hold significant promise for precision agriculture, as they can drastically reduce diagnostic time, enable large-scale automated monitoring, and facilitate early intervention—a critical factor in managing highly infectious crop diseases such as those affecting tomatoes [2].

Despite these advances, real-world deployment of plant disease classification models remains challenging. High performance on curated benchmark datasets does not always translate to real-world agricultural settings. One of the primary obstacles is the lack of large, diverse, and representative datasets. Collecting diseased leaf samples that capture the full range of environmental conditions, disease stages, and symptom variations is both labor-intensive and time-consuming. Moreover, certain diseases occur only under specific weather or seasonal conditions, making it difficult to collect sufficient training samples for all classes. This process often results in class imbalance, where some diseases are over-represented while others are severely under-represented. Such an imbalance and limited diversity can cause deep learning models to overfit to the training data, leading to poor generalization when deployed in new environments. Both the quantity and diversity of the data, encompassing various lighting conditions, leaf orientations, background clutter, and symptom expressions, are critical determinants of model performance.

To mitigate these challenges, researchers have explored various data augmentation techniques. Traditional augmentation strategies, such as image rotation, flipping, scaling, color jittering, and cropping, help increase dataset variability by simulating different viewpoints and lighting conditions. However, these methods only alter existing images and cannot introduce entirely new disease appearances. To address this limitation, more advanced generative approaches, particularly Generative Adversarial Networks (GANs) [3], have been proposed. GAN-based methods can synthesize novel, realistic-looking diseased leaf images, potentially expanding the diversity of training data beyond what is available from real-world collections.

Among these approaches, Leaf-GAN [4] stands out as a model specifically designed for plant disease image generation. Leaf-GAN employs an attention-guided mechanism to transform healthy leaf images into diseased ones while preserving the leaf's natural texture, venation patterns, and realistic disease spread characteristics. The intention is to aid in cases where disease samples are limited, producing synthetic data that captures fine-grained details of lesions and discolorations. While promising, GAN-based methods, including Leaf-GAN, are susceptible to well-known challenges, most notably mode collapse [5].

Mode collapse occurs when the GAN's generator produces a limited range of outputs, failing to capture the full diversity of the target data distribution. Instead of generating a wide variety of disease symptoms, lesion shapes, and severities, the generator may repeatedly produce highly similar outputs. In plant disease classification, this lack of variability undermines the core purpose of data augmentation, as it results in synthetic datasets that overrepresent certain visual features while underrepresenting others. Consequently, models trained on such data risk becoming biased toward a narrow subset of disease appearances, reducing their ability to generalize across real-world scenarios that contain greater variability.

To systematically evaluate the impact of GAN-generated data—and Leaf-GAN in particular—on tomato leaf disease classification, this study employs five state-of-the-art deep learning architectures: ConvNeXt [6], EfficientNetB3 [7], MobileViT [8], Vision Transformer (ViT) [9], and the Compact Convolutional Transformer (CCT) [10]. These models were chosen to represent diverse architectural paradigms, ranging from pure convolutional designs to hybrid CNN-transformer models and lightweight vision transformers. Each model was trained and evaluated on three versions of the dataset:

1. Original Dataset (Baseline) – real tomato leaf images only.
2. A1 Augmentation – original dataset plus 10% Leaf-GAN–generated synthetic images.
3. A2 Augmentation – original dataset plus 25% Leaf-GAN–generated synthetic images.

This experimental design allows us to measure not only whether GAN-based augmentation improves performance but also how different proportions of synthetic data affect accuracy and generalization.

In our experiments, applying Leaf-GAN [4] to tomato leaf disease classification revealed a notable limitation. While Leaf-GAN had previously reported promising results on cucumber leaf datasets, its performance did not generalize to tomatoes. The synthetic tomato leaf images frequently displayed repetitive lesion patterns and discolorations, suggesting a lack of diversity indicative of mode collapse [5]. This visual repetition led to augmented datasets that over-represented certain disease traits while omitting others, ultimately biasing model training.

The results were unexpected and counterintuitive: even when only 10% of the training data was replaced with Leaf-GAN–generated images, accuracy decreased across all four deep learning models. Increasing the synthetic data proportion to 25% further amplified the decline in performance. These findings challenge the common assumption that synthetic augmentation inherently benefits model training, and they underscore the necessity for crop-specific validation of generative augmentation methods [11].

Ultimately, our results suggest that the quality and diversity of synthetic data are as important—if not more important—than the quantity of generated samples. In the case of Leaf-GAN for tomato disease classification, the generated data introduced inconsistencies that reduced model robustness and generalization, highlighting the risks of uncritically adopting generative augmentation in agricultural AI pipelines.

II. RELATED WORK

Plant disease classification has been extensively investigated across a wide range of crop types [1], [2], [11], leveraging both laboratory-acquired datasets and real-world field imagery [12], [13]. Early research predominantly relied on controlled laboratory images characterized by uniform illumination, uncluttered backgrounds, and high-resolution captures [14]. Such conditions minimized environmental noise, enabling deep learning models—particularly convolutional and transformer-based architectures—to detect disease-specific symptom patterns with high precision and minimal confounding factors. While these laboratory datasets often yielded remarkably high classification accuracies [1], [15], subsequent studies have underscored the importance of assessing model performance under realistic field conditions [2], [13], where variability in lighting, occlusion, background complexity, leaf orientation, and physical damage is inevitable [11]. Models trained exclusively on lab-acquired imagery typically exhibit substantial performance degradation when deployed directly in the field [3], [12], highlighting the domain shift between controlled and in situ environments.

A critical bottleneck in advancing robust and generalizable plant disease detection systems is the limited availability of large-scale, high-quality field datasets [1], [2]. This scarcity has motivated the adoption of Generative Adversarial Network (GAN)–based augmentation techniques, which have demonstrated success in various agricultural applications, such as synthetic data generation for cotton disease detection [12], groundnut leaf stress recognition [11], and tomato plant disease classification [2], [13]. Several studies [1], [16] stress that domain alignment between synthetic and real-world data is crucial for improving model performance in deep learning–based plant disease detection. Despite promising results, the literature consistently emphasizes the necessity for crop-specific validation before integrating GAN-generated imagery into operational pipelines [5], [17], given differences in leaf morphology, disease manifestation, and environmental backgrounds.

In this context, the present study employs Leaf-GAN [4] to synthetically expand tomato leaf disease datasets using realistic field imagery. Leaf-GAN generates diverse, biologically plausible diseased leaf images embedded within natural scene backgrounds [3], [12], aiming to reduce the domain gap between laboratory-controlled and real-world scenarios. Designed explicitly for plant disease detection, Leaf-GAN builds upon CycleGAN [18] by introducing two key innovations: (1) Attention-guided translation for focusing disease synthesis exclusively on relevant leaf regions, and (2) a Label-Free Leaf Segmentation (LFLSeg) [13] to isolate leaves without manual annotations, preserving the authenticity of field backgrounds [3], [12]. These enhancements enable Leaf-GAN to produce highly realistic and contextually coherent diseased leaf images, outperforming conventional augmentation techniques, such as geometric transformations and color perturbations [14], [15], particularly in low-data and class-imbalanced settings [12].

Leaf-GAN's initial application to cucumber leaf disease classification [4] achieved substantial performance gains [1]. This success was partly attributed to the relatively uniform leaf shape, texture, and lesion manifestation in cucumber datasets [3], [13], which simplified the disease–healthy leaf transformation process. The generated synthetic data improved model generalization on unseen samples, positioning Leaf-GAN as a promising augmentation strategy for addressing dataset scarcity, inter-class imbalance, and limited environmental diversity in agricultural computer vision [11].

However, its applicability to crops with greater morphological and pathological variability, such as tomatoes, remains underexplored. Tomato leaf disease classification [2], [11], [19] presents additional challenges, including irregular leaf margins, heterogeneous venation patterns, visually complex field backgrounds, and overlapping lesions. These complexities raise an important question: *Can Leaf-GAN maintain diagnostic relevance and morphological fidelity when synthesizing disease symptoms across diverse leaf types and complex environments?*

Figure 1: Healthy and diseased tomato images.

A persistent challenge in GAN-based image synthesis is mode collapse [5], [17], [20], [21], wherein the generator produces limited and repetitive outputs that fail to capture the diversity of the target distribution. In the context of plant disease datasets, mode collapse may lead to overrepresentation of specific lesion shapes or colorations, biasing downstream classifiers and reducing their generalization capacity. While prior research has proposed strategies to mitigate mode collapse, such as architectural modifications, domain-specific loss functions, and progressive training stabilization [3], [4], [12], there has been no systematic evaluation of this phenomenon within Leaf-GAN for tomato disease detection.

To address this gap, the present study applies Leaf-GAN to a tomato leaf disease dataset and evaluates its augmentation efficacy across five state-of-the-art image classification architectures: ConvNeXt [6], EfficientNetB3 [7], MobileViT [8], Vision Transformer (ViT) [9], and Compact Convolutional Transformer (CCT) [10].

This evaluation provides a comprehensive analysis of Leaf-GAN's capacity to generate diverse, diagnostically relevant synthetic data for tomato leaf disease detection, its potential to reduce domain shift, and its resilience against mode collapse in high-variability agricultural domains.

III. The Proposed Evaluation Framework

This section outlines the overall methodology employed to evaluate the impact of Leaf-GAN–generated synthetic data on tomato leaf disease classification. The proposed workflow consists of four main components:
(A) Synthetic data generation using Leaf-GAN,
(B) Data preparation,
(C) Model training, and
(D) Performance evaluation.

A. Synthetic Images Generation Using Leaf-GAN

As described in the preceding section, Leaf-GAN [4] is an attention-guided image-to-image translation framework specifically designed for synthesizing plant disease images. It generates diseased leaf images from healthy ones while preserving background realism through its integrated Label-Free Leaf Segmentation (LFLSeg) module.

In this study, the Leaf-GAN generator was trained using a subset of healthy tomato leaf images. The training was performed for 40 epochs with the Adam optimizer, employing a fixed learning rate of 0.001 and instance normalization in both the generator and discriminator networks. Before training the GAN, input images were augmented via random horizontal flips and small-angle rotations to increase variability in the generator's input space and improve robustness.

Once trained, the generator was applied to selected healthy leaf images to produce synthetic diseased samples. The proportion of synthetic images integrated into the training set was systematically varied to create three dataset configurations:

A0 – Original dataset (no synthetic augmentation)

A1 – Original dataset + 10% synthetic Leaf-GAN images

A2 – Original dataset + 25% synthetic Leaf-GAN images

The synthetic images were only added to the training set, while the validation and test sets contained exclusively real images. This design prevents synthetic bias in evaluation and reflects a real-world deployment scenario, where augmentation is applied solely during model training.

B. Data Preparation

The tomato leaf disease dataset used in this study was sourced from a publicly available repository, containing high-resolution field images that capture both healthy leaves and those affected by common tomato diseases. As shown in Figure 1, the dataset comprises 11 classes: Bacterial Spot, Early Blight, Late Blight, Leaf Mold, Septoria Leaf Spot, Target Spot, Yellow Leaf Curl Virus, Spider Mites, Mosaic Virus, Powdery Mildew, and Healthy.

The images exhibit substantial variability in lighting conditions, leaf orientation, occlusion, and background clutter, closely reflecting real-world agricultural environments. Class distributions range from 1,000 to 2,100 images per class, ensuring a balanced yet diverse dataset suitable for evaluating both baseline performance and the effect of synthetic augmentation on model generalization.

All images were resized to 224×224 pixels and normalized using standard preprocessing techniques. The dataset was split into training, validation, and test subsets with a 70:20:10 ratio. The same test set was used across all experiments to ensure consistency in comparative analysis.

C. Deep Learning Models

To assess the effect of synthetic augmentation across different architectural paradigms, five state-of-the-art image classification models were selected:

1. **ConvNeXt** [6] – A modernized CNN architecture inspired by Vision Transformers, incorporating design elements such as large kernel sizes and layer normalization to balance accuracy and efficiency.
2. **EfficientNetB3** [7] – A convolutional network that optimally scales depth, width, and resolution, achieving high accuracy with fewer parameters through compound scaling.
3. **MobileViT** [8] – A lightweight hybrid architecture that fuses convolutional feature extraction with transformer blocks, optimized for mobile and edge deployment.
4. **Vision Transformer (ViT)** [9] – A pure transformer-based model that processes images as sequences of non-overlapping patches (tokens), enabling global context modeling without convolution layers.
5. **Compact Convolutional Transformer (CCT)** [10] – A transformer variant incorporating convolutional token embedding, which enhances local feature modeling while retaining global attention capabilities.

These architectures span purely convolutional, purely transformer-based, and hybrid CNN–transformer approaches, enabling a comprehensive evaluation of how different inductive biases respond to GAN-based augmentation.

D. Training and Evaluation Protocol

All models were trained on an NVIDIA RTX 4090 GPU for 40 epochs with a batch size of 32. The Adam optimizer was used with a learning rate of 0.001, and regularization techniques include dropout and early stopping based on validation loss to mitigate overfitting. Each model was trained independently under the three dataset configurations (A0, A1, and A2), ensuring a fair comparison of baseline performance versus synthetic augmentation.

The primary evaluation metric was overall classification accuracy on the held-out real-image test set. To provide a more granular performance assessment, per-class precision and recall were also computed, enabling evaluation of each model's ability to accurately identify specific tomato disease classes and distinguish between visually similar conditions.

To further investigate the effect of synthetic data augmentation, confusion matrices were generated for every model–dataset configuration pair. These matrices visualize class-wise misclassification patterns, allowing for the identification of specific disease classes that were prone to confusion and assessing whether the inclusion of Leaf-GAN–generated images reduced such confusion or enhanced feature separability.

This multi-level evaluation strategy—combining global accuracy, class-specific metrics, and confusion matrix analysis—ensures a comprehensive understanding of model behavior and quantifies the contribution of synthetic augmentation to classification robustness in real-world deployment scenarios.

TABLE 1: MODEL ACCURACY COMPARISON

Model	A0	A1	A2
ViT	0.9642	0.9815	0.9857
CCT	0.8967	0.8825	0.8713
MobileViT	0.9757	0.9792	0.9861
EfficientNetB3	0.9881	0.9892	0.9888
ConvNeXt	0.9884	0.9803	0.9842

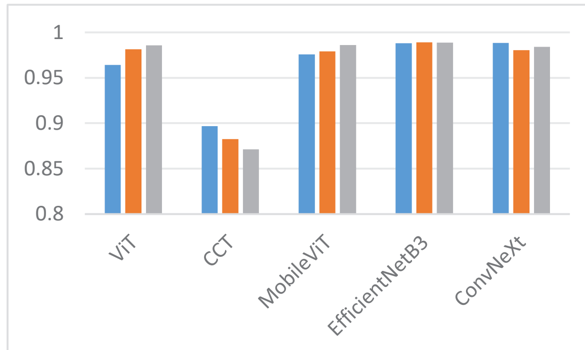

Figure 2: Model accuracy comparison across the three configurations: A0 (blue), A1 (orange), and A2 (gray).

IV. EXPERIMENTAL RESULTS

All models were trained using the Pytorch framework with identical hyperparameters in a GPU-enabled environment to ensure consistency and reproducibility. Basic data augmentation techniques, such as horizontal flipping and random rotation, were applied uniformly across all experiments to maintain diversity.

The classification accuracy for all five models across the three dataset configurations (A0, A1, and A2) is presented in Table 1 and visualized in the corresponding bar-graph shown in Figure 2. On the original dataset (A0), ConvNeXt and EfficientNetB3 achieved the highest accuracy (98.8%), followed closely by MobileViT, ViT, and CCT.

When 10% synthetic Leaf-GAN data was added (A1), most models experienced a slight decline in accuracy, with the exception of MobileViT and ViT, which recorded marginal improvements. Increasing the synthetic data proportion to 25% (A2) led to further decreases for CCT and ConvNeXt, while MobileViT and ViT maintained upward trends. These results indicate that while transformer-based or hybrid architectures may adapt better to moderate synthetic augmentation, heavy reliance on Leaf-GAN-generated data can degrade performance in CNN-dominant models.

To further investigate these trends, confusion matrix analysis was performed, with results illustrated in Figure 3 using the ViT model, which displayed the most notable class-level variation across A0, A1, and A2. On the original dataset (A0), ViT performed well in differentiating most disease classes. However, with synthetic augmentation in A1 and A2, confusion increased, particularly between visually similar diseases such as Leaf Mold and Septoria leaf spot. This suggests that the synthetic images, while visually realistic, may have lacked sufficient intra-class variability or introduced inconsistencies that reduced separability between similar classes.

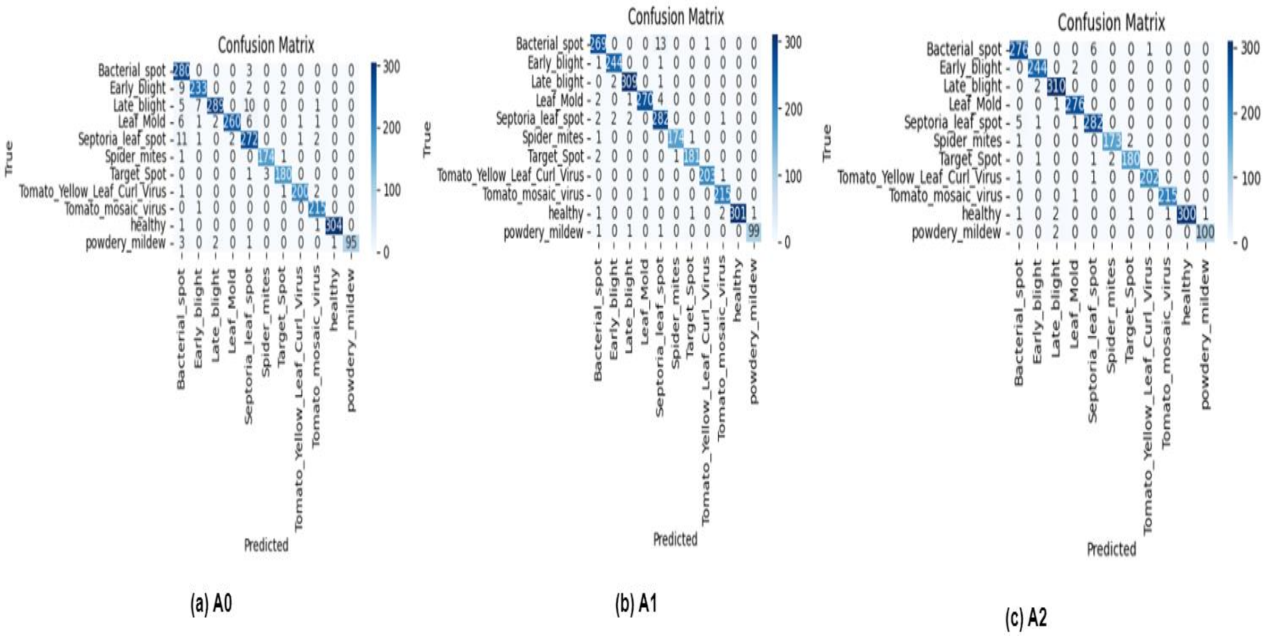

Figure 3: Confusion matrices for the ViT model trained on (a) A0-original data, (b) A1- with 10% Leaf-GAN data, and (c) A2- with 25% Leaf-GAN data. Misclassification increases with synthetic data, demonstrating model sensitivity to augmentation.

TABLE 2: MODEL PRECISION COMPARISON

Model	A0	A1	A2
ViT	0.9658	0.9818	0.9858
CCT	0.9016	0.8860	0.8753
MobileViT	0.9761	0.9794	0.9863
EfficientNetB3	0.9881	0.9892	0.9889
ConvNeXt	0.9885	0.9809	0.9846

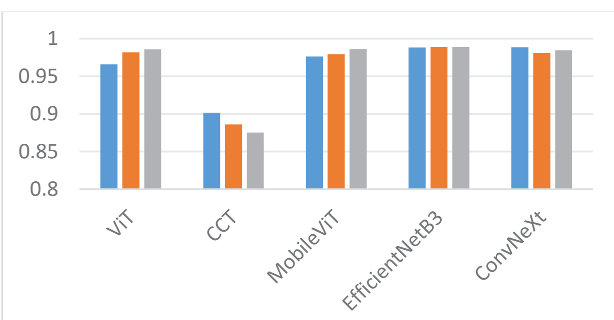

Figure 4: Precision Scores of Models Under Varying Synthetic Data Proportions: A0 (blue), A1 (orange), and A2 (gray).

The Precision values for each model are summarized in Table 2 and depicted in the bar graph in Figure 4. Similar to the accuracy results, EfficientNetB3 and ConvNeXt displayed consistently high precision across all datasets, though both saw small reductions in A1 and A2. MobileViT and ViT achieved notable precision improvements as the proportion of synthetic data increased, suggesting these models were able to leverage additional training variation despite potential distributional inconsistencies. In contrast, CCT's precision declined steadily from A0 to A2, reflecting reduced ability to maintain class-specific boundaries in the presence of GAN-generated samples. These trends further highlight architectural differences in how models handle synthetic augmentation.

Table 3 and the corresponding recall bar graph in Figure 5 show that EfficientNetB3 and ConvNeXt retained high recall across all dataset configurations, with only minor fluctuations. MobileViT and ViT benefited from synthetic augmentation, registering incremental recall gains from A0 to A2. In contrast, CCT exhibited a continuous decline in recall, indicating that synthetic images negatively impacted their sensitivity to positive class detection. These results imply that while some architectures adapt to the increased variability provided by GAN augmentation, others, particularly those relying heavily on localized convolutional features, may experience a drop in detection sensitivity.

TABLE 3: MODEL RECALL COMPARISON

Model	A0	A1	A2
ViT	0.9642	0.9815	0.9857
CCT	0.8967	0.8825	0.8713
MobileViT	0.9757	0.9792	0.9861
EfficientNetB3	0.9881	0.9892	0.9888
ConvNeXt	0.9884	0.9803	0.9842

Although visually realistic, these synthetic images did not contribute positively to the tomato disease classification task. Notably, EfficientNetB3 and ConvNeXt demonstrated greater resilience, likely attributable to their enhanced local context learning capabilities, which may mitigate some of the limitations introduced by GAN-generated textures.

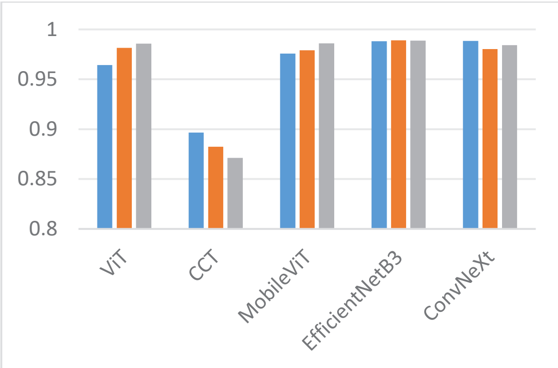

Figure 5: Recall Performance Across the three configurations: A0 (blue), A1 (orange), and A2 (gray).

Overall, the findings demonstrate that the addition of synthetic data not only failed to enhance model performance but also consistently decreased or remained the same. Even a modest augmentation of 10% led to a measurable reduction in accuracy, and further increasing the proportion of synthetic data (A2) resulted in a further decline in performance. These results contradict the expected benefit of GAN-based augmentation and highlight the need for task-specific validation before adopting synthetic data in practical diagnostic systems.

V. CONCLUSION

This study evaluated the effect of Leaf-GAN-generated synthetic images on tomato leaf disease classification using multiple deep learning models. Experimental results revealed that augmenting the original dataset with just 10% synthetic data (A1) consistently reduced classification accuracy, and using 25% synthetic data (A2) led to continued performance degradation. Although Leaf-GAN is capable of generating visually realistic images, the generated samples may lack critical diagnostic features or introduce distributional inconsistencies that negatively affect the models' learning and generalization capabilities. These findings emphasize the importance of task-specific validation when applying generative augmentation techniques, as their effectiveness can vary significantly across domains. Cautious integration and careful tuning are essential before adopting such methods in real-world plant disease diagnostic systems.

ACKNOWLEDGEMENT

This research was supported by Kyungpook National University Research Fund, 2025

REFERENCES

[1] Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7. https://doi.org/10.3389/fpls.2016.01419

[2] A. Motwani and S. Khan, "Lightweight Recurrent and Convolutional Deep Learning Models for Tomato Disease Detection," Frontiers in Sustainability, Available: https://doi.org/10.3389/frsus.2024.1383182

[3] Zhao, Y., & Zhou, Y. (2024). A data augmentation method based on GAN for plant disease recognition. Lecture Notes in Electrical Engineering, 186-193. https://doi.org/10.1007/978-981-97-6937-7_22

[4] Cap, Q. H., Uga, H., Kagiwada, S., & Iyatomi, H. (2022). LeafGAN: An effective data augmentation method for practical plant disease diagnosis. IEEE Transactions on Automation Science and Engineering, 19(2), 1258-1267. https://doi.org/10.1109/tase.2020.3041499

[5] F. L. Barsha and W. Eberle, "An in-depth review and analysis of mode collapse in generative adversarial networks," Machine Learning, vol. 114, p. 141, 2025.

[6] Feng, J., Tan, H., Li, W., & Xie, M. (2022, November). Conv2NeXt: reconsidering Conv NeXt network design for image recognition. In *2022 international conference on computers and artificial intelligence technologies (CAIT)* (pp. 53-60). IEEE.

[7] Tan, M., & Le, Q. (2021, July). Efficientnetv2: Smaller models and faster training. In International conference on machine learning (pp. 10096-10106). PMLR.

[8] Mehta, S., & Rastegari, M. (2021). Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. *arXiv preprint arXiv:2110.02178*.

[9] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

[10] Hassani, A., Walton, S., Shah, N., Abuduweili, A., Li, J., & Shi, H. (2021). Escaping the big data paradigm with compact transformers. *arXiv preprint arXiv:2104.05704*.

[11] J, P., & Dinakaran, K. (2025). Enhancing groundnut leaf disease and stress detection using GAN-based deep learning models. 2025 International Conference on Computing and Communication Technologies (ICCCT), 1-4. https://doi.org/10.1109/iccct63501.2025.11019972

[12] Zhang Z, Ma L, Wei C, Yang M, Qin S, Lv X and Zhang Z (2023) Cotton Fusarium wilt diagnosis based on generative adversarial networks in small samples. Front. Plant Sci. 14:1290774. doi: 10.3389/fpls.2023.1290774

[13] Shoaib, M., Hussain, T., Shah, B., Ullah, I., Shah, S. M., Ali, F., & Park, S. H. (2022). Deep learning-based segmentation and classification of leaf images for detection of tomato plant disease. Frontiers in Plant Science, 13. https://doi.org/10.3389/fpls.2022.1031748

[14] Min, B., Kim, T., Shin, D., & Shin, D. (2023). Data augmentation method for plant leaf disease recognition. Applied Sciences, 13(3), 1465. https://doi.org/10.3390/app13031465

[15] Khare, O., Mane, S., Kulkarni, H., & Barve, N. (2024). LeafNST: An improved data augmentation method for classification of plant disease using object-based neural style transfer. Discover Artificial Intelligence, 4(1). https://doi.org/10.1007/s44163-024-00150-3

[16] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Comput. Electron. Agricult., vol. 145, pp. 311–318, Feb. 2018.

[17] H.-T. Thanh-Tung and T. Tran, "Catastrophic forgetting and mode collapse in GANs," arXiv preprint arXiv:1807.04015, Mar. 21, 2020.

[18] Zhu, J., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. 2017 IEEE International Conference on Computer Vision (ICCV), 2242-2251. https://doi.org/10.1109/iccv.2017.244

[19] Alam, T. S., Jowthi, C. B., & Pathak, A. (2024). Comparing pre-trained models for efficient leaf disease detection: A study on custom CNN. Journal of Electrical Systems and Information Technology, 11(1). https://doi.org/10.1186/s43067-024-00137-1

[20] P. Chong, L. Ruff, M. Kloft, and A. Binder, "Simple and Effective Prevention of Mode Collapse in Deep One-Class Classification," arXiv preprint arXiv:2001.08873, Jan. 19, 2021.

[21] M. Topal, "What Is Mode Collapse in GANs?" Medium, Jul. 23, 2023. [Online]. Available: https://medium.com/@miraytopal/what-is-mode-collapse-in-gans-d3428a7bd9b8.