MJFAS
Malaysian Journal of Fundamental and Applied Sciences

RESEARCH ARTICLE

Generative AI and Computer Vision-based Hybrid Approach for Plant Diseases Classification

Huma Jamshed[a], Shahbaz Qammar[b], Syeda Tehreem Naqvi[a], Yusra Mansoor[a,*], Ahmad Hussain[c]

[a]Department of Computer Science, DHA Suffa University, Karachi 75500, Pakistan; [b]Marines Training Center PNS Qasim Karachi, Pakistan; [c]Faculty of Engineering and Applied Sciences, DHA Suffa University, Karachi 75500, Pakistan;

Abstract Plant Cotton leaf diseases pose a serious hazard to crop productivity. These diseases contribute to increased use of pesticide resulting in environmental degradation and unsustainable agricultural practices. The problem is particularly serious in precision agriculture, where timely and accurate analysis is crucial for effective interference. Deep learning (DL) models for image classification have shown strong potential for automated disease detection. However, its performance is often affected by inadequate and diversified training data sets, which have reduced generalization under practical field conditions. To address this challenge, this study proposes a generative Artificial Intelligence (AI) and computer vision based hybrid framework for cotton leaf disease classification with implications for sustainable precision agriculture. The proposed framework integrates StyleGAN3-ADA for synthetic image augmentation, ConvNeXt for disease classification, and a retrieval-augmented generation (RAG) module for contextual agricultural decision support. StyleGAN3-ADA was trained on class-specific training subsets, and the generated synthetic images were added only to the training set to increase visual diversity. Validation and testing were performed exclusively on real cotton leaf images to ensure realistic performance assessment. Experimental results demonstrate that the StyleGAN3-based augmentation process produced visually realistic samples with a Fréchet Inception Distance (FID) score of 26.47. The ConvNeXt classifier achieved a training accuracy of 99% and a testing accuracy of 96%, indicating strong convergence and robust classification performance across the six cotton leaf categories. Furthermore, the integration of the RAG module enables the system to provide context-aware, knowledge-driven recommendations for disease management. The proposed framework establishes practical potential for intelligent decision support in precision agriculture. By enabling early and accurate disease identification, the proposed framework can reduce unnecessary pesticide application, thereby minimizing environmental impact and promoting more sustainable agricultural practices.

Keywords: Cotton leaf disease classification, StyleGAN3-ADA, ConvNeXt, synthetic image augmentation, precision agriculture, retrieval-augmented generation.

***For correspondence:**
yusra.mansoor@dsu.edu.pk

Received: 09 Mar. 2026
Accepted: 28 Jul 2026

Introduction

Cotton is one of the most important cash crops in the world and serves as a basic raw material for the textile industry worldwide. With its extremely high economic value and comprehensive industrial value, many people regard cotton as white gold. Major countries engaged in cotton production, such as China, India, the United States, Brazil, and Pakistan, have displayed visible leading positions in both national economies and the global cotton supply chain. Cotton accounts for about 35% of the world's annual fibre consumption and thus is indispensable to sustain textile production and rural livelihoods [1] [2]. As shown in Table 1, China and India hold a dominant position, while Pakistan ranks constantly among leading producers, which indicates the appropriate degree of dependence on cotton-growing across both developed and developing economies.

Table 1. Countries and cotton crop production

Major Countries	Production in thousand metric tons
China	5879
India	5334
United State	3815
Brazil	2678
Pakistan	1306
Australia	1197
Turkey	827
Uzbekistan	577
Argentina	327
Mali	311

Cotton production underpins both the agricultural and industrial industries in many nations, including Pakistan [3]. The textile industry is one of the major industries that consumes cotton as raw material and contributes about 73% of the total export earnings of the country [4]. The contribution of agriculture to the total national GDP is around 10% and hence cotton is a part of the economy and employment. While cotton is becoming more economically significant, a variety of plant diseases that have a negative effect on yield and fiber quality continually threaten cotton production and quality, and often require more reliance on chemical treatments that can have a negative impact on soil health and surrounding ecosystems [5].

Cotton plants are susceptible to various bacterial, fungal, viral, and pest-related diseases such as bacterial blight, alternaria leaf spot, cotton leaf curl virus (CLCV), infestation by aphid pest species, armyworms, and powdery mildew diseases [6]. These diseases have adverse effects on photosynthesis, plant growth, and boll formation, resulting in extensive losses in yield quantity and quality. Of all the stress factors affecting the crop, the leaf diseases are very destructive since they interfere directly with the physiological activity of the crop during early developmental stages. The early identification and correct identification of cotton leaf diseases are essential for proper management of the crop. Plant diseases across the world result in losses ranging from 20% to 40% yield loss in crop quantity each year, posing a major threat to plant sustainability, and often encouraging excessive pesticide application that contributes to environmental pollution [7].

Farmers and specialists typically diagnose cotton disease through visual inspections, which are typically performed in the field. While assessments that rely on the expertise of an agricultural specialist can be very effective, they require significant labour and time, and are usually dependent on the experience of the person performing the assessment. Additionally, because of the similarities between the symptoms presented by different diseases in the early stages of development, there can be significant inconsistencies in visual assessments and misclassifications of diseases. This makes the scalability of visual assessments very low, and these types of assessments are especially difficult to perform in large agricultural regions, or when the agricultural producers are resource-constrained. Therefore, there is an increasing need for automated, accurate, and efficient disease identification systems that enable agricultural producers to make timely and informed management decisions while potentially reducing unnecessary chemical interventions [8].

Recent DL [9], computer vision [10], and generative AI advances have shown significant promise for the automation of plant disease detection from leaf images. Specifically, DL models have achieved state-of-the-art performance in various image classification applications due to their capability of automatically learning discriminative features without reliance on handcrafted descriptors. These models, especially CNNs [11] and transformer-inspired architectures, provide non-destructive, fast, and highly accurate solutions for crop disease diagnosis, thus being especially suitable for precision agriculture [12]. The performance of DL models is contingent on having an abundance of varied, representative, and accurately annotated training datasets. One of the primary challenges associated with DL in agriculture specifically for crop disease detection is the lack of good quality annotated datasets that are created based on true-to-life/on-field conditions [13]. Although public datasets may have a large number of available diseases, the number of ways each category has variations in disease symptoms, is represented differently by the orientation of each leaf, was recorded under different backgrounds and was recorded in different lighting including time of day, will often reduce how robust the model is to using that data and cause overfitting, and ultimately result in the model's inability to generalize when it is implemented in actual agricultural scenarios [14] [15]. These limitations will hinder these highly advanced and well-designed deep learning models from fully utilizing their representational capacity as well as limiting their reliability when finally executed in a real environment.

However, there are some drawbacks in the current research. Previous studies have shown high classification accuracy, but this is frequently achieved using well-annotated data that was gathered under relatively controlled conditions [16] [17]. In practical agricultural scenarios, the datasets for disease detection on cotton leaves are typically small, with limited visual diversity within classes, variations in symptom expression, and a lack of field conditions, which leads to poor robustness and generalizability of deep learning models [18] [19]. Moreover, the majority of the current disease detection models are mainly classification-based with no context of interpretation, treatment and decision making, which reduces their practical use in sustainable agriculture [20] [21].

Generative AI, specifically those using GAN techniques, is a viable solution to augment training data in image-related deep learning models which face challenges such as low visual diversity and limited variation [22]. The use of adversarial learning enables the creation of synthetic images that share the same statistics as real images but also have extra variation in terms of texture, structure, and appearance [23]. This can be helpful for agriculture deep learning applications where the data is sparse and algorithms are not able to generalize sufficiently because of the lack of diversity in the training data.

The absence of visual diversity and limited real-world variability in agricultural image datasets have been a major challenge, and generative AI, specifically GAN-based methods [24] have emerged. GAN-based frameworks can generate realistic synthetic images that resemble real images, but also introduce some variations at the appearance level, which can enhance the diversity of the training samples and boost the performance of deep learning models [25]. Previous research has demonstrated that synthetic images can enhance the accuracy of plant disease classification, particularly when the data sets are limited in scope and do not adequately represent the spectrum of disease symptoms and field conditions [26]. However, most of the previous research has focused on using GANs as a standalone pre-processing technique and failed to combine high fidelity synthetic augmentation with a strong classification model [27] and to explicitly connect these techniques to environmentally sustainable agricultural practices in a single cotton specific framework.

In addition to the correct disease classification, intelligent decision-support systems (DSS) are an important part of modern precision agriculture. In recent years, the feasibility of integrating machine learning (ML) and deep learning (DL) models with recommendation systems for disease management, fertilization, and pesticide application in agriculture has been explored [28]. DL classifiers have been used to create many web-based and mobile applications for providing real-time disease diagnosis and advisory services [29]. With the advent of large language models (LLMs) and Retrieval-Augmented Generation (RAG) frameworks, the scope of AI-driven decision-support systems in agriculture has expanded even further. Structured knowledge bases and generative models can be integrated into RAG-based systems to deliver almost instant, contextualized, explainable, and domain-specific answers [30]. However, most existing decision-support systems remain either loosely connected to image-based disease detection pipelines or dependent on conventional classifiers trained on limited datasets [31]. Consequently, little attention has been paid to integrated systems that combine synthetic image augmentation, accurate cotton disease classification, and intelligent decision support within a unified framework while promoting sustainable agricultural practices.

The objective of the research was to design a combined generative AI and DL model capable of classifying the diseases of cotton plants, and simultaneously promote sustainable agricultural practices through better management of these diseases. In the design of the proposed model, synthetic images created by StyleGAN3-ADA were used to enhance the variety in the data set for training, while ConvNeXt provided accurate classification results. To enhance this experience, a RAG module was added to the framework to allow access to disease information and management recommendations based on context; thus providing the basis for transforming a classification model into a real-world precision agriculture tool that provides informed and directed remedial actions that can be taken to alleviate the negative effects on the environment.

The proposed study is significant because it combines the synthetic data augmentation using StyleGAN3-ADA and the robust cotton leaf disease classification using ConvNeXt in a single framework. A hybrid training approach was used, where synthetic augmentation was only applied to the training set, and validation and testing were only done on real images. Furthermore, a Retrieval-Augmented Generation (RAG) module was added to enable the extension of disease classification into an intelligent decision support system that can generate agricultural recommendations based on the context. The proposed framework integrates synthetic data generation, disease identification and knowledge-based decision support in a single pipeline, while supporting sustainable agriculture through early disease detection and more targeted pesticide application.

This paper is organized as follows. In Section II, the methodology employed in this study is described, which consists of the dataset, the generation of synthetic images with StyleGAN3-ADA, the classification of the diseases with ConvNeXt, and the decision-support module based on RAG. The experimental results and comparative analysis are given in Section III. The findings, practical implications, limitations, and future research directions are discussed in Section IV. Finally, Section V concludes the paper.

Materials and Methods.

This research work presents an integrated approach to develop a generative AI and a DL framework for effectively classifying diseases found in the leaves of cotton plants and providing intelligent, data-based decisions. The methodology includes a series of steps where various parts of the methodology can be processed as stand-alone or through multiple processes completed in tandem to achieve the required results. The framework is made up of four major components or elements, including (1) preparing the Cotton dataset, (2) generating synthetic Cotton Plant Disease Data using StyleGAN3, (3) classifying the diseased cotton plant images using a ConvNeXt model and (4) providing intelligent, decision-based responses via, a RAG system for diseases affecting cotton. A visual representation of the project framework is provided in Figure 1.

Overall Framework

The proposed system utilized the openly accessible Cotton Leaf Disease Dataset found on Kaggle (Dhamur), which included RGB images of cotton leaves categorized into six different classes: Aphids, Army Worm, Bacterial Blight, Powdery Mildew, Target Spot, and Healthy Leaves. This data set had natural differences in background settings, illumination levels, posture of the leaves, and manifestation of symptoms, making it well-suited for practical applications involving image-based disease recognition. Images in the data set were preprocessed by input size and normalization prior to the training process. After labels were verified, the data set was split into training, validation, and test sets with a stratified split to ensure a good representation of each class. Since class imbalance was not the dominant issue in this particular data set, artificial data augmentation techniques were performed exclusively on the training group. Validation and testing data sets only contained genuine images to provide accurate assessment results.

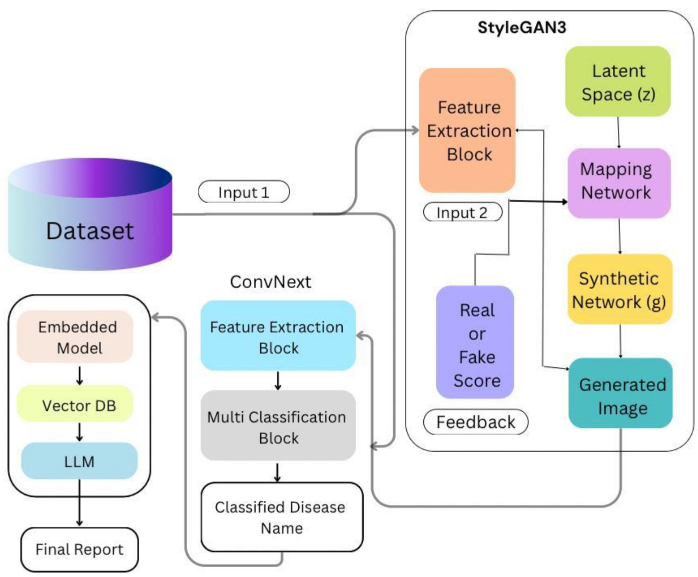

Figure 1. Proposed hybrid framework architecture integrating StyleGAN3 for data augmentation, ConvNeXt for disease classification, and RAG for decision support.

Dataset Description and Preprocessing

The Cotton Leaf Disease Dataset (Dhamur) is publicly available through Kaggle, containing RGB images of cotton leaves classified into 6 classes: Aphids, Army Worm, Bacterial Blight, Powdery Mildew, Target Spot and Healthy Leaves. The dataset contains substantial variations in background, illumination, leaf orientation, and symptom severity. Therefore, as a result of this extensive variability, it can be considered to have usefulness for practical application of image-based disease identification. Image labels were verified for annotation accuracy, and the images were then resized to a standardized resolution, before being normalized and trained with the models. After labelling validation, images were stratified into three subsets using a stratified split which allowed for maintenance of class representation across all 6 classes. As the main problem was that rather than being a case of severely imbalanced classes, the class composition was an issue (meaning the images in 1 class would have very few differences within them), a synthetic augmentation method was applied only to increase visual variability for improving generalization. The validation and test samples were comprised entirely of real imagery, providing a more accurate representation of actual performance.

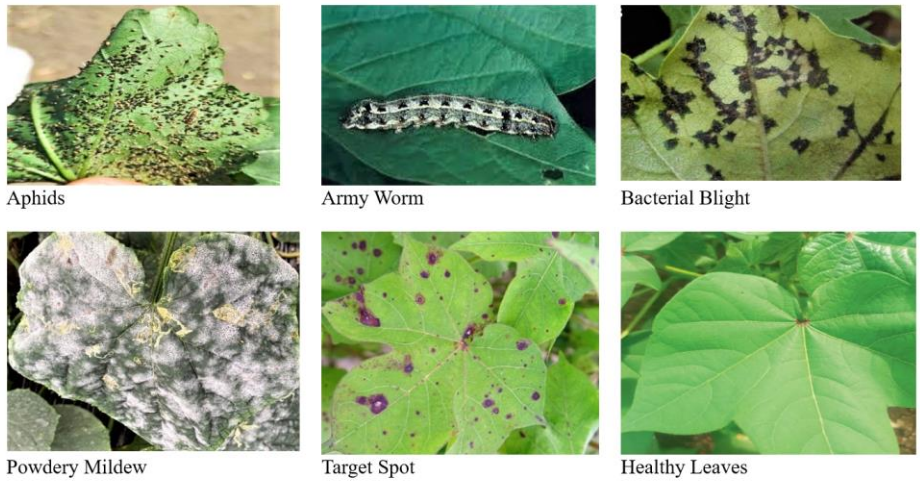

Figure 2. Visual samples of the cotton crop disease types from the utilized dataset.

StyleGAN3-Based Synthetic Data Generation

In order to achieve more diversity of training data and enhance the robustness of the classifiers, StyleGAN3-ADA was chosen for the generation of images to include in the framework. Instead of assuming that the training data is heavily skewed, the model was used to generate more realistic cotton leaf images within each class. This is particularly useful in agricultural imaging, where visually similar symptoms, limited variation within classes, and constrained sample diversity can reduce the generalization ability of deep learning models. StyleGAN3 was selected because of its ability to generate high-fidelity, alias-free images while preserving fine-grained texture and disease-relevant visual patterns. The StyleGAN3 architecture is composed of two adversarial neural networks: A Generator (G) and a Discriminator (D). The generator learns to synthesize realistic cotton leaf disease images by mapping a latent noise vector z, sampled from a normal distribution, into a high-dimensional image space. This process is mathematically expressed as $G(z)$ where $z \sim N$ (0, I). Through iterative adversarial training, the generator progressively improves its ability to produce synthetic images that closely resemble real samples from the dataset. The discriminator operates in opposition to the generator and is responsible for distinguishing between real images x from the dataset and synthetic images generated by $G(z)$. The discriminator output is defined as D(x) where the output represents the estimated probability that the input image x is real. During training, the discriminator is optimized to correctly classify real and generated images, while the generator is optimized to deceive the discriminator. The adversarial learning process is governed by the following objective function as shown in equation 1.

$$\min_{G} \max_{D} \; V(D, G) = \; E_{x \sim p_{\text{data}}(x)}[\log D\,(x)] + E_{z \sim p_z(z)}\left[\log\left(1 - D\big(G(z)\big)\right)\right] \quad (1)$$

The purpose of this objective is to prompt the generator to reduce the statistical divergence between the real and synthetic image distributions while still allowing the discriminator to perform the differentiation operation effectively. The main advantage of StyleGAN3 stems from the style-based modulation mechanism, which focuses on image quality as well as stability. Instead of inserting the latent vector into the generator, StyleGAN3 uses a mapping network M to convert the latent vector into an intermediate latent space as given in equation 2.

$$w = M(z) \quad (2)$$

The generator's convolutional layers receive their modulation through the style vector w which uses affine transformations A(w) to control their operation. The generator uses this modulation system to control different image attributes because it allows them to create intricate cotton leaf disease textures and morphological patterns. StyleGAN3 produces images that appear realistic and accurately preserve disease-related visual patterns because it removes aliasing artifacts and enhances spatial consistency. In this study, StyleGAN3-ADA was trained on class-specific training subsets from the cotton leaf disease dataset. Synthetic images were generated only from the training partition and were added to the corresponding classes during classifier training. This strategy enriched the visual diversity of training samples while preventing information leakage into validation and test data. Class-related structural and textural information present in original images was preserved in generated images and they were utilized as additional data rather than a substitute for actual images. The augmented dataset, consisting of both real and synthetic images, was then used to train the ConvNeXt classifier for stable multi-class cotton leaf disease classification.

ConvNeXt-Based Disease Classification

For the classification stage of the suggested framework, a mixed dataset (real cotton leaf images and synthetic images generated by StyleGAN3) was used to train the ConvNeXt model, while the validation and test steps were carried out only on real images. ConvNeXt model is an up-to-date CNN architecture which takes advantage of vision transformer models and retains well-known efficiency of CNNs, thus is the most suitable CNN for the agricultural image classification under problematic field conditions.

i. Patch Embedding and Feature Representation

Given an input cotton leaf image $I \in \mathbb{R}^{H \times W \times C}$, is first transformed to patches by applying a convolutional layer which has a user-defined stride (s) to the image. The convolution operation converts the image into patches and then applies a feature map of (c) to the patch. Therefore, the total number of patches created (N) is defined through equation 3.

$$N = \frac{H \times W}{s^2} \quad (3)$$

Each patch is then represented as a feature embedding, forming the initial input to the ConvNeXt backbone.

ii. ConvNeXt Block Architecture

The system processes embedded patches through hierarchical ConvNeXt blocks which enable it to capture local texture details and global contextual information. The transformation within a ConvNeXt block is expressed in equation 4 and 5::

$$\mathrm{Y} = \mathrm{X} + \mathrm{DW\ Conv}\big(\mathrm{LN}(\mathrm{x})\big) \quad (4)$$

$$\mathrm{X}' = \mathrm{Y} + \mathrm{MLP}\big(\mathrm{LN}(\mathrm{Y})\big) \quad (5)$$

where:

- DWConv denotes a depth-wise convolution with a large kernel size, enabling an expanded receptive field,
- LN represents layer normalization for training stability,
- MLP consists of two linear projections separated by a non-linear activation function, typically GELU.

This architecture allows ConvNeXt to effectively model fine-grained disease spots as well as broader leaf-level patterns, which are critical for distinguishing visually similar cotton diseases.

iii. Global Feature Aggregation and Classification

In the final stage of the network, the feature maps are combined using a global average pooling layer to produce a compact representation of features, as shown in equation 6. The feature vector f is then fed into a fully connected classification layer that projects the feature vector into the respective classes of diseases. The final output is produced using a softmax activation function.

$$\mathrm{f} = \mathrm{GAP}(\mathrm{X}') \quad (6)$$

iv. Loss Function and Model Optimization

For training the ConvNeXt classifier, a categorical cross-entropy loss function is used as expressed in equation 7.

$$\mathrm{L\ ce} = -\sum_{\mathrm{i}=1}^{\mathrm{c}} \mathrm{y_i} \log(\hat{y}\,\mathrm{i}) \quad (7)$$

Here C is the number of classes for the diseases, yi is the actual label, and $\hat{y}$ i is the predicted probability for class i. The loss function is designed to help the model learn very discriminative features and also perform correct classification for all classes of diseases.

v. Training with Augmented Data

To improve generalization and enrich training diversity, the ConvNeXt model was trained on a hybrid dataset consisting of real cotton leaf images and synthetic samples generated by StyleGAN3-ADA. The addition of high-quality synthetic images led to higher intra-class variance and better feature learning for the 6 disease classes. The experimental results showed that this hybrid training method can achieve high accuracy and stable convergence.

Integration of Retrieval-Augmented Generation (RAG)

To extend the proposed framework beyond image-based disease classification, a RAG system is incorporated to offer disease-specific decision support. Although the DL classifier ConvNeXt is capable of effectively identifying cotton leaf diseases, it lacks the ability to offer contextual explanations, management plans, and recommendations necessary for farmers and agricultural experts. The RAG system is designed to overcome this challenge by integrating retrieval processes with generative language modeling. In the proposed method for gathering information based on a disease classification and a predicted disease class, a ConvNeXt classier's output is used as a query into the retrieval aspect of a RAG system. The agricultural knowledge base contains disease descriptions, symptoms, causes of disease, and methods of management that was used as the retrieval source. The retrieval of relevant text was done using semantic similarity-based matching. This retrieved text was used by a Large Language Model (LLM) to generate a response. Formally, when the classifier provides a predicted disease for a cotton leaf, the retriever for the RAG system retrieved a set of documents {d1, d2, ……, dk } from the knowledge base as described by Equation 8

$$\{d_i\}_{i=1}^{k} = \mathrm{Retriever}(c) \quad (8)$$

These retrieved documents are concatenated with the original query and provided as contextual input to the generative model as shown in equation 9:

$$\hat{r} = \mathrm{LLM}(c, \{d_i\}_{i=1}^{k}) \quad (9)$$

Here $\hat{r}$ denotes the generated response. The generated output establishes its authenticity through verified domain knowledge which the system uses to create its response. The response contains visual symptoms and potential causes and recommended management practices which are specific to the disease.

The proposed system achieves greater explanatory capability and practical utility through RAG integration, which transforms the framework into a complete intelligent decision-support system. The system enables users to make better decisions through accurate disease classification and contextual textual assistance which supports its use in actual precision agriculture applications.

Results and Performance Evaluation

Evaluation Metrics

The proposed framework was evaluated using quantitative measures for both generative performance and disease classification. The Fréchet Inception Distance (FID) was used to assess the visual similarity between real and synthetic images generated by StyleGAN3-ADA. Lower FID values indicate that the generated samples are more closely aligned with the distribution of real cotton leaf images.

For classification evaluation, the ConvNeXt model was assessed using accuracy, categorical cross-entropy loss, precision, recall, and F1-score. The categorical cross-entropy loss defined in Section 3.4 was used to monitor learning behavior and convergence during training, whereas precision, recall, and F1-score were derived from the confusion matrix to evaluate class-wise classification performance across the six cotton leaf categories.

Moreover, these measures allow the most accurate measurement of the performance at the whole-system as well as per-class level which would be very important for their usage in agriculture.

Performance of StyleGAN3-Based Data Augmentation

The performance of the StyleGAN3-based augmentation module was evaluated in terms of image quality, stability of training and similarity between real and generated images quantitatively. Class-specific training subset from the cotton leaf disease dataset were used to train StyleGAN3-ADA for generating synthetic images to be added for augmentation. Some of the generated samples can be found in Figure 3. The produced images are able to preserve the main characteristics of leaves in terms of structure and texture, and also contain the symptom related appearance of the leaves.

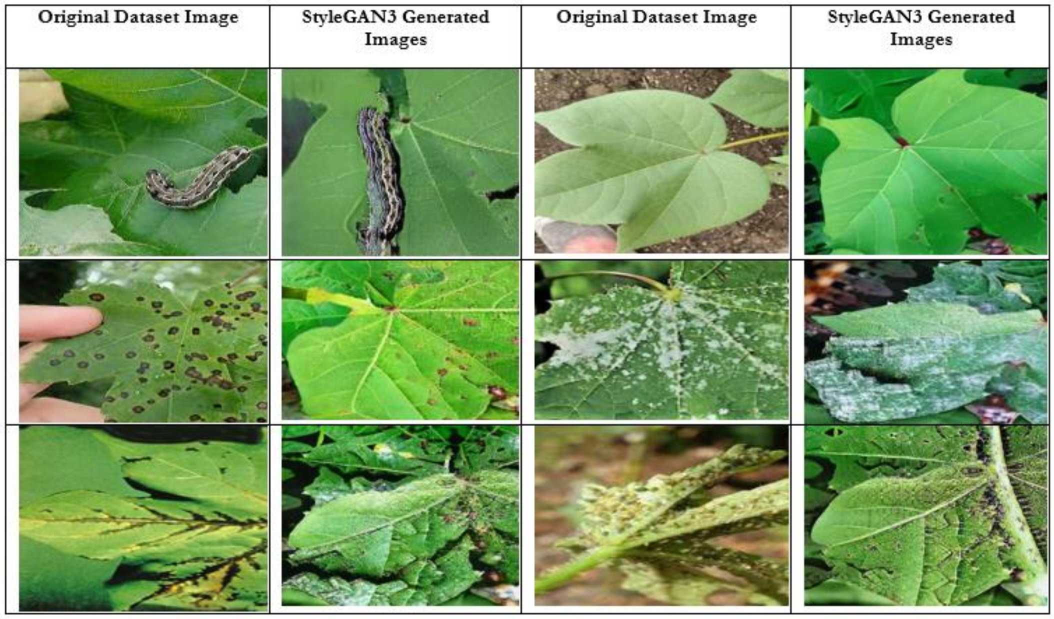

Figure 3. Qualitative comparison between original dataset images (columns 1 and 3) and synthetic images generated by StyleGAN3 (columns 2 and 4).

In addition, the Fréchet Inception Distance (FID) was also computed to find the distance between the real images and synthetic images. It turns out the FID score is 26.47 which shows the generated images have reasonable similarity and learned data distribution well.

While a lower FID value is ideal, the recorded value of 26.47 is reasonable considering that Agricultural images often have large within class variation and it was trained on this domain. Figure 4 below, shows the behavior of FID during training. Additionally, as observed in the plots of generator and discriminator loss (Figures 5 and 6) throughout the training it suggests a stable adversarial learning. There is no clear mode collapse.

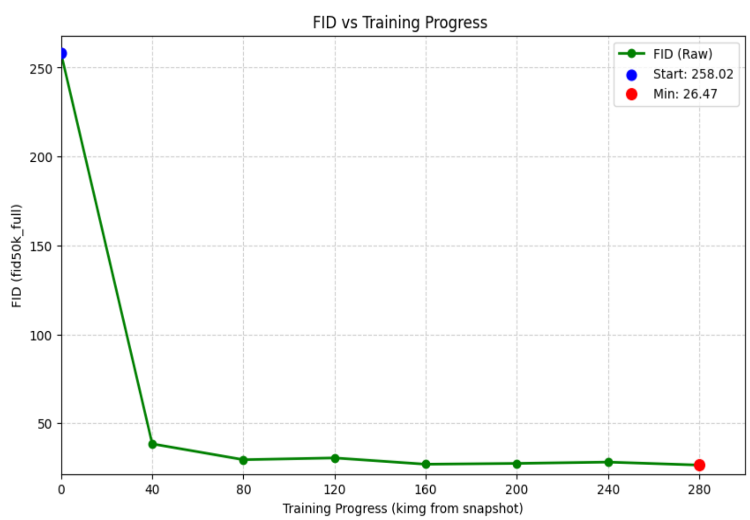

Figure 4. Quantitative evaluation of image quality via Fréchet Inception Distance (FID) over training iterations.

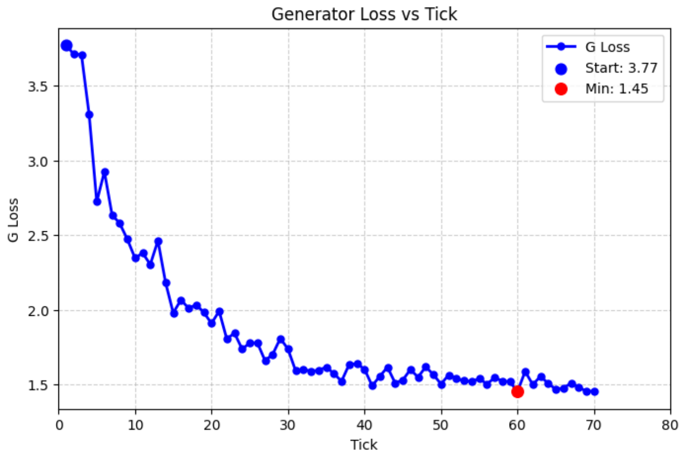

Figure 5: Training stability analysis depicting Generator Loss (G Loss) vs. training ticks.

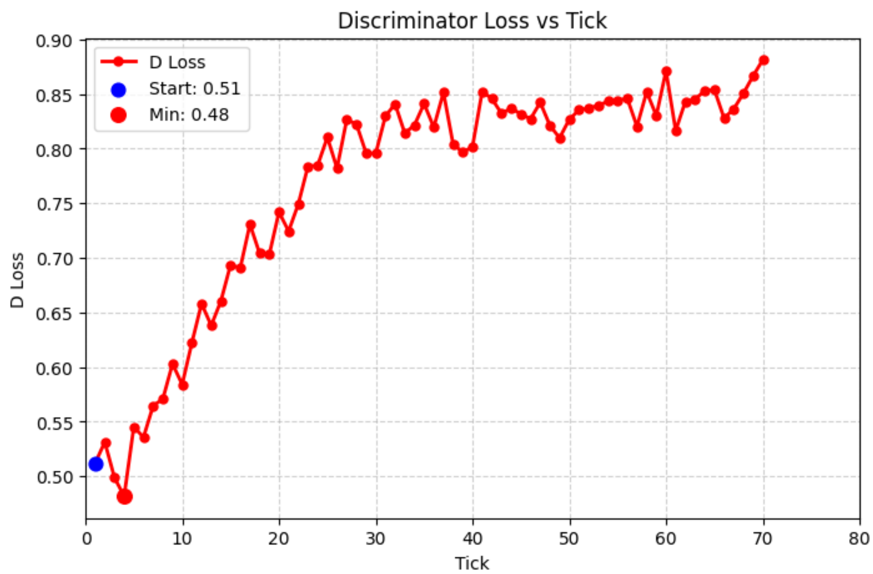

Figure 6. Training stability analysis depicting Discriminator Loss (D Loss) vs. training ticks. iterations.

ConvNeXt Classification Results

Trained the ConvNeXt classifier using a mix of real cotton leaf images and synthetic samples made with StyleGAN3. For validation and testing real images were used. The point was to see if using synthetic data during training actually helped the model handle real-world scenarios better.

Figure 7 displays the training accuracy and validation accuracy for the ConvNeXt model. The training loss and validation loss curves for ConvNeXt can be observed in Figure 8. ConvNeXt was able to train to 99% training accuracy and test accuracy was 96% at the end of training. The training and validation curves were closely matched which implies that ConvNeXt was not overfitting significantly, even with augmentation data.

Here, Figure 9 demonstrates the classification accuracy across the six cotton leaf classes in terms of confusion matrix. It can be clearly noticed that almost all the classes were predicted correctly and there was no considerable confusion between the closely resemble disease classes. The model is able to identify the subtle visual differences between disease classes, thus, practical application can be assumed. Furthermore, the class-wise precision, recall and F1-score are also provided in Table 2 to indicate that the classifier performs equally well on the six classes. The equally good precision and recall indicates a good classifier that does not produce much false positives or false negatives for the detection of the diseases.

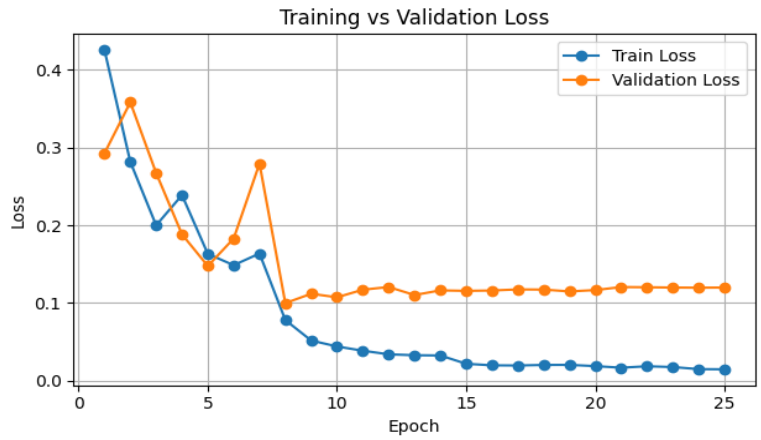

Figure 7. Training and testing loss curves over 25 epochs

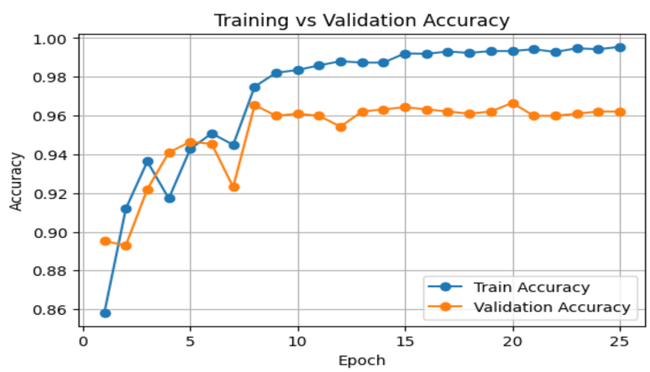

Figure 8. Training and testing accuracy curves demonstrating strong

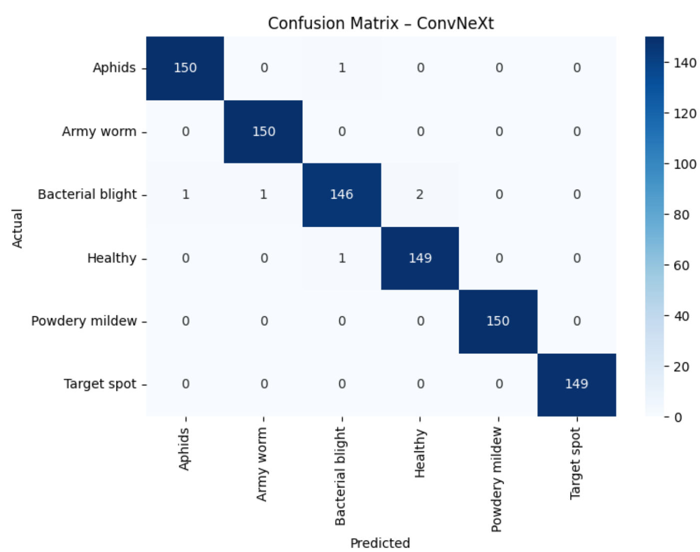

Figure 9. Confusion matrix illustrating the classification performance across the six cotton leaf categories.

Table 2: Class-wise performance of the ConvNeXt classifier

Class	Precision	Recall	F1 Score
Aphids	0.9934	0.9934	0.9934
Army Worm	1.000	0.9933	0.9967
Bacterial Blight	0.9551	0.9933	0.9739
Powdery Mildew	1.0000	1.0000	1.0000
Target Spot	1.0000	0.9799	0.9898
Healthy leaf	0.9932	0.9800	0.9866

Comparative Analysis

To further assess the effectiveness of the proposed framework, its classification performance was compared with existing studies on cotton leaf disease detection reported in the literature. The comparison is made based on classification accuracy. Table 3 provides a comparison of recent DL techniques used for cotton and plant leaf disease identification through their classification results. The majority of current research studies concentrate on enhancing classification and detection performance by utilizing different GAN designs and DL techniques. The methods developed by these researchers provide solutions for specific tasks through their image-level prediction capabilities but generally lack integrated decision-support capabilities and environmental considerations. The proposed framework combines StyleGAN3-based data augmentation, ConvNeXt-based classification, and a RAG-enabled decision-support interface to help users identify diseases while also receiving contextual agricultural recommendations.

The proposed StyleGAN3–ConvNeXt framework achieved a classification accuracy of 96%, which is competitive with existing methods and demonstrates strong performance relative to several previously reported techniques. In addition to classification accuracy, the designed framework integrated data augmentation, disease identification, and the provision of context-aware advice into a single pipeline making it more suitable for practical precision agriculture applications.

It is important to note that direct comparison across studies is influenced by factors such as size of dataset, amount of disease classes or experimental configurations. Nevertheless, the proposed method is shown to work well with natural agricultural images and restricted within-class variations while it enables accurate decision-making which potentially helps agriculture to be sustainable.

Table 3: Comparative summary of recent DL based approaches for cotton or plant leaf disease classification

Reference	GAN Model	Classification / Detection Model	Performance	Dataset	Web Interface / RAG
Hybrid Cotton Leaf Disease Detection using DCGAN and CNNs	DCGAN	InceptionV3, VGG16, ResNet50, MobileNetV2	Accuracy: 99%	Cotton leaf disease dataset	No
Faster R-CNN with Augmentation for Cotton Disease Detection	CycleGAN + traditional augmentation	Faster R-CNN	mAP: 90.2%	Cotton disease dataset	No
CSA-GAN for IoT-Enabled Cotton Disease Detection	CSA-GAN (conditional GAN)	CSA-GAN (end-to-end)	Not specified	Cotton plant images	No

Lightweight Cotton Disease Detection using StyleGAN 2-ADA	StyleGAN2-ADA	RT-DETR-DFSA	Accuracy: 87.14%	Custom dataset (4000 images, China)	No
Cotton Fusarium Wilt Diagnosis using GANs	DCGAN	Custom CNN	Test Accuracy: 95%	Field images (China)	No
Multi-CNN Cotton Disease Detection with StyleGAN	Customized StyleGAN	MobileNetV2 + VGG16 + StackNet + LSTM + SAM	Avg. Accuracy: 97%	PlantVillage + Roboflow + external dataset	No
Automated Lesion Detection in Cotton Leaves	GAN (unspecified)	VGG16 + InceptionV3 + ResNet50 (ensemble)	Accuracy: 95%	Public cotton leaf dataset	No
Proposed Method	StyleGAN3	ConvNeXt	Training: 99%, Validation: 96%	Cotton leaf disease dataset	Yes (Web Interface + RAG)

Discussion

The findings suggest that StyleGAN3-based synthetic augmentation can improve cotton leaf disease classification by enriching visual diversity within the training data and supporting better model generalization. Rather than functioning only as a class-balancing approach, the generative module contributed by producing realistic additional samples that exposed the classifier to a wider range of disease-related visual patterns. This was particularly important for distinguishing visually similar symptoms under practical agricultural imaging conditions.

The results further indicate that synthetic data augmentation can be an effective strategy in scenarios where collecting large-scale, diverse agricultural datasets is challenging, thereby improving the scalability of AI-based solutions in precision agriculture.

The ConvNeXt classifier showed good performance in feature extraction for the classification of visually similar diseases. The architecture's modern convolutional design allowed the system to effectively capture subtle texture and color variations that are associated with different cotton leaf conditions. The findings show that the use of GAN-generated augmentation along with an advanced convolutional architecture enhanced the generalization performance and facilitated more reliable disease classification. The high classification accuracy and class-wise performance indicate that the proposed model is robust in diverse visual conditions, crucial for its deployment in practical agricultural settings.

Furthermore, the proposed system enhances disease detection by incorporating a Retrieval-Augmented Generation (RAG) module which offers context-aware agricultural suggestions depending on the foreseen disease category. It makes the system more useful to the practitioner by adding the disease recognition capability with knowledge-based support. The retrieval-based mechanism also contributes to guaranteeing the coherence of the generated answers with the verified agricultural knowledge, which increases the confidence of the recommendations for the farmers. It is an important step in the development of predictive models that can be translated into practical, decision-making actions for smart agriculture systems.

In addition, the modular design of the framework supports scalability and improves its applicability to broader smart agriculture settings. Each module can be adapted to support different crop types and disease datasets, thus raising the applicability of the system for various forms of smart agriculture.

However, the proposed framework has certain limitations, including the computational cost associated with training GAN models and the dependency of the RAG module on the quality and coverage of the underlying knowledge base.

Furthermore, the use of a single dataset may limit the generalizability of the model across diverse geographical and environmental conditions, which should be addressed in future work through multi-source and field-based data collection.

Nonetheless, the proposed hybrid framework has proved that an effective approach to the management of diseases in agriculture can be achieved via the hybridization of generative augmentation, convolutional neural networks, together with retrieval-based decision support.

From an environmental perspective, the framework has the potential to support sustainable agriculture by enabling early and accurate disease detection, which can reduce unnecessary pesticide application, minimize environmental contamination, and promote more efficient use of agricultural resources.

Conclusion

This study presented a hybrid generative AI and deep learning framework for cotton leaf disease classification and intelligent decision support in precision agriculture. In this research, StyleGAN3-ADA was used to generate synthetic training data. ConvNeXt was used for disease classification. A RAG module was incorporated to provide context-aware farming recommendations that effectively recognize the specific situation. In the experiments, it can be seen that using synthesized data improve the generalization of the model, increasing variety of sample of each class. The classifier ConvNeXt, has achieved good result on classification of six kinds of cotton leaves disease with 96% of test accuracy. This shows that the classifier can classify into the disease type correctly. In addition to that RAG enabled the methodology to depart from the conventional classification methods and provided with decision support in managing plant diseases. Thus, this study creates a more useful and effective approach for designing classifiers as it will have a knowledge-based decision-making process.

The proposed system has great potential for application in practical use, particularly in situation where rapid disease diagnosis and guidance are needed in agriculture. There are several limitations that need to be addressed in future work for instance dependency on only one data set, lack of proper statistical evaluation, and computations involved in generative models. Future research will include data acquisition from multiple sources and from the field. Evaluation will be made statistically using cross validation and the performance of guidance module will be tested from user's perspective.

This system also benefits sustainable farming since it enables early discovery of plant diseases thereby reducing unnecessary pesticide use and less pollutant in the environment as well as farmers makes the best of resources. The development of intelligent agricultural systems which can scale and protect the environment receives support through the combination of generative AI and advanced convolutional architectures and knowledge-based decision support systems.

Conflicts of Interest

The author(s) declare(s) that there is no conflict of interest regarding the publication of this paper.

Acknowledgment

This research was funded by DHA Suffa University, Karachi, Pakistan. The authors wish to express their gratitude to the DHA Suffa University for its technical and financial supports.

References

[1] Bakthavatchalam, K., Karthik, B., Thiruvengadam, V., Muthal, S., Jose, D., Kotecha, K., & Varadarajan, V. (2022). IoT framework for measurement and precision agriculture: predicting the crop using machine learning algorithms. Technologies, 10(1), 13.

[2] Jamshed, H., Waheed, U., Mansoor, Y., & Hussain, A. (2025). Unsupervised Learning for Crop Suitability Clustering Based on Soil Nutrients. Malaysian Journal of Fundamental and Applied Sciences, 21(6), 2945-2957.

[3] Das, P., Ray, S., Dasgupta, M., & Govil, M. C. (2023, October). Fusion of IoT and Blockchain (BIoT): A Novel Architecture for Smart Agriculture. In International Conference on Internet of Things and Connected Technologies (pp. 1-11). Singapore: Springer Nature Singapore.

[4] Azfar, S., Nadeem, A., Ahsan, K., Mehmood, A., Ali, A., Al Qahtany, S. S., & Almoamari, H. (2025). Automated System for Detecting, Identifying, and Preventing Cotton Leaf and Boll Diseases Using Deep Learning. International Journal of Advances in Soft Computing & Its Applications, 17(1).

[5] Noor, D., Jamshed, H., Memon, N., & Iqbal, S. (2025). Artificial Intelligence-Based Approach in Precision Agriculture for Crop Yield Prediction. International Journal of Information Systems and Computer Technologies, 5(1), 65-79.

[6] Khursheed, A., Farooq, A., Shafique, H., Ali, S., & Gilanie, G. (2026). DEEP LEARNING-BASED COTTON PEST CLASSIFICATION USING COMPUTER VISION. Spectrum of Engineering Sciences, 4(6), 3869-3897.

[7] Yadav, A., Sharma, V., Vij, A., Singh, J., & Chaudhary, A. (2025, July). Machine Learning-Based Image Analysis for Cotton Crop Health: A Comparative Study. In 2025 2nd International Conference On Multidisciplinary Research and Innovations in Engineering (MRIE) (pp. 85-91). IEEE.

[8] Kaneva, T., Valova, I., Gabrovska-Evstatieva, K., & Evstatiev, B. (2025). A data-driven approach for generating synthetic load profiles with gans. Applied Sciences, 15(14), 7835.

[9] Min, X., Ye, Y., Xiong, S., & Chen, X. (2025). Computer vision meets generative models in agriculture: Technological advances, challenges and opportunities. Applied Sciences, 15(14), 7663.

[10] Qadri, S. A. A., Huang, N. F., Wani, T. M., & Bhat, S. A. (2024). Advances and challenges in computer vision for image-based plant disease detection: a comprehensive survey of machine and deep learning approaches. IEEE Transactions on Automation Science and Engineering, 22, 2639-2670.

[11] Zualkernan, I., Abuhani, D. A., Hussain, M. H., Khan, J., & ElMohandes, M. (2023). Machine learning for precision agriculture using imagery from unmanned aerial vehicles (uavs): A survey. Drones, 7(6), 382.

[12] Umar, M., Altaf, S., Ahmad, S., Mahmoud, H., Mohamed, A. S. N., & Ayub, R. (2024). Precision agriculture through deep learning: Tomato plant multiple diseases recognition with cnn and improved yolov7. Ieee Access, 12, 49167-49183.

[13] Yang, Z. Y., Xia, W. K., Chu, H. Q., Su, W. H., Wang, R. F., & Wang, H. (2025). A comprehensive review of deep learning applications in cotton industry: From field monitoring to smart processing. Plants, 14(10), 1481.

[14] Logeshwaran, J., Srivastava, D., Kumar, K. S., Rex, M. J., Al-Rasheed, A., Getahun, M., & Soufiene, B. O. (2024). Improving crop production using an agro-deep learning framework in precision agriculture. BMC bioinformatics, 25(1), 341.

[15] Ruby, E. D., Amirthayogam, G., Sasi, G., Chitra, T., Choubey, A., & Gopalakrishnan, S. (2024). Advanced image processing techniques for automated detection of healthy and infected leaves in agricultural systems. Mesopotamian Journal of Computer Science, 4(1), 44-52.

[16] Faisal, H. M., Aqib, M., Rehman, S. U., Mahmood, K., Obregon, S. A., Iglesias, R. C., & Ashraf, I. (2025). Detection of cotton crops diseases using customized deep learning model. Scientific Reports, 15(1), 10766.

[17] Gani, R., Isty, M. N., Rashid, M. R. A., Ahmed, J., Islam, T., Hasan, M., ... & Reza, A. W. (2025). Attention guided convolutional neural network with explainable AI for papaya leaf disease detection in edge and drone agricultural systems. Scientific Reports, 15(1), 41489.

[18] Ankatwar, G., & Dhawale, C. (2025). Multispectral Imaging and CNN Architectures for Cotton Leaf Disease Classification: A Comprehensive Review. International Journal For Multidisciplinary Research, 7(3).

[19] Shafik, W., Tufail, A., Liyanage De Silva, C., & Awg Haji Mohd Apong, R. A. (2025). A novel hybrid inception-xception convolutional neural network for efficient plant disease classification and detection. Scientific Reports, 15(1), 3936.

[20] Adão, T., Chojka, A., Pascoal, D., Silva, N., Morais, R., & Peres, E. (2025). Synthetic Data-Driven Methods to Accelerate the Deployment of Deep Learning Models: A Case Study on Pest and Disease Detection in Precision Viticulture. Computers, 14(8), 327.

[21] Panchal, A. V., Patel, S. C., Bagyalakshmi, K., Kumar, P., Khan, I. R., & Soni, M. (2023). Image-based plant diseases detection using deep learning. Materials Today: Proceedings, 80, 3500–3506.

[22] Rahman, Z. U., Asaari, M. S. M., Ibrahim, H., Abidin, I. S. Z., & Ishak, M. K. (2024). Generative adversarial networks (GANs) for image augmentation in farming: A review. IEEE Access, 12, 179912-179943.

[23] Arshad, M. A., Lee, H., Lee, H., Kang, M., Kim, Y., & Bang, H. (2026). Structure-Consistent Contrastive Learning for Unpaired Image Translation With Gradient-Domain Constraints. IEEE Access.

[24] Klair, Y. S., Agrawal, K., & Kumar, A. (2024, March). Impact of generative AI in diagnosing diseases in agriculture. In 2024 2nd International Conference on Disruptive Technologies (ICDT) (pp. 870–875). IEEE.

[25] Aghabalayev, S., Abdalla, A. N., Hai, T., Liu, Y., & Wang, A. (2026). StyleGAN3-T: an alias-free generative framework for synthetic plant disease image augmentation and recognition. Scientific Reports, 16(1), 21901.

[26] Ali, T., Zakir, R., Ayaz, M., Murtaza, M., Hijji, M., & Hadi Aggoune, E. M. (2025). Cotton crop disease detection and classification using statistical prediction model in deep learning approach. Multimedia Tools and Applications, 84(41), 49503-49525.

[27] Song, L., & Chen, Z. (2025, July). Deep learning with multi-source data augmentation for accurate cotton disease recognition. In 2025 21st International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD) (pp. 183–188). IEEE.

[28] Usha Ruby, A., Chandran, J. G. C., Chaithanya, B. N., Jain, T. J. S., & Patil, R. (2024). Wheat leaf disease classification using modified ResNet50 convolutional neural network model. Multimedia Tools and Applications, 83(23), 62875–62893.

[29] Islam, M. M., Talukder, M. A., Sarker, M. R. A., Uddin, M. A., Akhter, A., Sharmin, S., ... & Debnath, S. K. (2023). A deep learning model for cotton disease prediction using fine-tuning with smart web application in agriculture. Intelligent Systems with Applications, 20, 200278.

[30] Aldossary, M., Alharbi, H. A., & Ul Hassan, C. A. (2024). Internet of Things (IoT)-enabled machine learning models for efficient monitoring of smart agriculture. IEEE Access, 12, 75718–75734.

[31] Dolatabadian, A., Neik, T. X., Danilevicz, M. F., Upadhyaya, S. R., Batley, J., & Edwards, D. (2025). Image-based crop disease detection using machine learning. Plant Pathology, 74(1), 18–38.