International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

Generative AI-Driven Synthetic Image Generation for Horticultural Seed Bag Detection and Classification

[1]Vankudoth Saidulu
[1]Assistant professor, Methodist college of engineering and technology
[1]Email: chinnivankudoth@gmail.com

ABSTRACT

Horticultural seed bags are widely used in agricultural production, storage, transportation, and distribution, making their accurate identification and classification important for effective inventory management and quality-control activities. Conventional image-based approaches often depend on limited collections of real images, which may not adequately represent variations in orientation, lighting, background, packaging appearance, scale, occlusion, and camera viewpoint. This study proposes a Generative AI-driven approach for synthetic image generation and automated detection and classification of horticultural seed bags. The proposed framework employs Generative AI techniques, including GAN-based and diffusion-based models, to learn visual characteristics from available seed-bag images and generate realistic synthetic samples with diverse visual conditions. The generated images are subjected to quality verification based on realism, diversity, visual consistency, and preservation of important seed-bag characteristics. Selected synthetic images are then integrated with the original dataset to create an enhanced image collection for detection and classification. The approach is designed to increase the visual diversity of the available dataset while reducing dependence on extensive manual image collection. The generated outputs are further analyzed according to their usefulness for seed-bag identification and classification. The proposed framework demonstrates the potential of Generative AI for developing flexible and scalable solutions for horticultural image analysis, with applications in agricultural storage, packaging verification, inventory management, and automated quality-control operations.

Keywords : Generative AI, Synthetic Image Generation, Horticultural Seed Bags, Seed Bag Detection, Image Classification, GAN, Diffusion Models, Agricultural Automation.

I.INTRODUCTION

The rapid adoption of artificial intelligence in agriculture has created new opportunities for analyzing agricultural products through digital images. Among these applications, the identification of horticultural seed bags is important because seed packaging carries visual information related to seed varieties, product categories, packaging patterns, labels, and other distinguishing characteristics. An automated image-based system can assist in organizing seed products and supporting activities such as inventory inspection, packaging verification, and product identification. However, developing such systems requires image datasets that contain sufficient variation to represent the conditions encountered during practical image acquisition.

A major challenge in image-based agricultural applications is the limited availability of diverse training images. A collection containing images captured under similar conditions may not adequately represent changes in illumination, camera position, background, bag orientation, scale, or packaging appearance. Generative Artificial Intelligence provides a potential solution by learning the underlying visual distribution of an image collection and producing new samples with related characteristics. The introduction of Generative Adversarial Networks established an important foundation for synthetic data generation by using competing generator and discriminator networks to learn realistic data distributions [1]. Conditional GANs subsequently introduced additional control over the generation process by incorporating specified conditions into image synthesis [2]. The development of convolutional generative architectures further strengthened the application of Generative AI to visual data. Deep Convolutional GANs demonstrated that convolution-based networks could effectively learn meaningful image representations and generate synthetic visual samples [3]. Subsequent research focused on improving the reliability and quality of generated content through improved training techniques [4], alternative adversarial objectives [5], and architectures designed to increase image quality and generation stability [6], [7]. These developments made generative models increasingly suitable for applications where obtaining large numbers of manually captured images is difficult.

Generative AI has also progressed toward more controllable image transformation and synthesis. Multi-domain generation enables visual characteristics to be modified across different domains [8], while conditional image-to-image approaches allow an input image to be transformed according to a desired visual condition [9]. Cycle-consistent generation demonstrated that meaningful image translation can also be performed without requiring directly paired training examples [10]. Style-based generative architectures provided improved control over the characteristics of synthesized images [11], followed by further improvements in image quality [12]. Research on training generative models with limited data is particularly relevant to specialized applications because domain-specific datasets are often smaller than general-purpose image collections [13]. Further developments in GAN architecture have continued to improve the visual consistency of generated outputs [14].

Another major development in Generative AI is the emergence of diffusion-based image synthesis. Diffusion models generate visual content through a gradual denoising process and have demonstrated strong capabilities in producing detailed images [15]. Score-based generative modelling provided an alternative mathematical framework for diffusion-based generation [16]. Latent Diffusion Models improved the practicality of high-resolution image synthesis by operating within a compressed latent representation [17]. Diffusion-based image-to-image generation expanded the possibilities for modifying existing visual content while retaining important image information [18]. Text-guided diffusion models further demonstrated how semantic descriptions could be used to control the generation of photorealistic images [19]. These developments have established diffusion models as an important area of modern Generative AI research [20].

For horticultural seed-bag analysis, these capabilities can be used to construct a richer collection of synthetic images rather than relying only on repeatedly collected real photographs. Generative AI can produce seed-bag samples with controlled variations in position, viewpoint, illumination, background, size, packaging appearance, and partial visibility. Such generated samples can supplement real images and provide additional visual examples for recognizing different seed-bag categories. However, synthetic images should not be included without verification, since unrealistic structures, distorted labels, or unwanted visual artifacts may reduce their usefulness. Therefore, evaluating the quality and relevance of generated samples is an important part of the proposed approach.

In this study, a Generative AI-Driven Synthetic Image Generation for

Received: 06-07-2026 | Accepted: 13-08-2026 | Published: 20-08-2026

Horticultural Seed Bag Detection and Classification framework is proposed. The framework uses Generative AI as the primary mechanism for expanding the available seed-bag image collection through synthetic image generation. Generated samples are assessed for visual quality and consistency before being incorporated with real images. The resulting collection is then used to support automated detection and classification of horticultural seed bags. The study consequently investigates how Generative AI can contribute not only to synthetic image creation but also to improving the visual coverage available for automated seed-bag analysis [1]–[20].

II.LITERATURE SURVEY

Ian J. Goodfellow et al. (2014): Goodfellow and his co-authors introduced Generative Adversarial Networks (GANs), presenting a generator-discriminator framework for learning data distributions and producing synthetic samples. The study established adversarial learning as a powerful approach for generating realistic data and became a foundation for later developments in Generative AI. The work is particularly relevant to synthetic image generation where additional samples are required from a limited collection [1].

Mehdi Mirza and Simon Osindero (2014): Mirza and Osindero proposed Conditional Generative Adversarial Networks, extending the conventional GAN architecture by incorporating additional information into the generation process. Their approach allowed generated samples to be controlled according to specified conditions. This concept is useful for applications where synthetic images need to represent particular categories or controlled visual characteristics [2].

Alec Radford, Luke Metz, and Soumith Chintala (2015): Radford et al. introduced Deep Convolutional Generative Adversarial Networks (DCGANs), demonstrating the effectiveness of convolutional architectures for generative modelling. Their study showed that deep convolutional networks could learn useful visual representations and generate realistic images. The work became an important development in computer-vision-based Generative AI and provides a foundation for synthetic image generation [3].

Tim Salimans et al. (2016): Salimans and co-authors investigated techniques for improving GAN training and addressed several challenges associated with adversarial learning. Their research proposed practical methods for improving training stability and generated-image quality. The study demonstrated that suitable training strategies can improve the reliability of GAN-based image synthesis, which is important when synthetic images are intended for further computer vision tasks [4].

Martin Arjovsky, Soumith Chintala, and Léon Bottou (2017): Arjovsky et al. introduced Wasserstein GAN as an alternative approach to conventional adversarial training. The method used the Wasserstein distance to provide a more meaningful training signal and improve the stability of GAN optimization. Their findings contributed to the development of more reliable generative models and are relevant to applications where realistic and consistent synthetic images are required [5].

Tero Karras et al. (2017): Karras and his collaborators proposed Progressive Growing of GANs, in which the image resolution is gradually increased during the training process. This strategy improved training stability and enabled the generation of higher-quality images. The research demonstrated the importance of progressive learning for producing detailed and visually diverse synthetic samples [7].

Yunjey Choi et al. (2018): Choi et al. developed StarGAN, a unified framework

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

for multi-domain image-to-image translation. The model allowed transformations between multiple visual domains while retaining important information from the original image. The study demonstrated the potential of controllable image generation and transformation, which can be useful for creating different appearances of the same object [8].

Tero Karras, Samuli Laine, and Timo Aila (2019): Karras et al. introduced a style-based generator architecture that improved the control and quality of generated images. Their approach separated different levels of visual characteristics within the generation process, enabling more flexible manipulation of synthesized content. The study became an important contribution to high-quality image generation and provides useful concepts for controlled synthetic image creation [11].

Jonathan Ho, Ajay Jain, and Pieter Abbeel (2020): Ho et al. introduced Denoising Diffusion Probabilistic Models, establishing a diffusion-based approach for image synthesis. Their method generates images through a gradual denoising process and demonstrated strong capabilities for producing high-quality samples. The study marked an important development beyond GAN-based generation and provided a foundation for later diffusion-based Generative AI systems [15].

Robin Rombach et al. (2022): Rombach and co-authors proposed Latent Diffusion Models for high-resolution image synthesis. Their approach performs diffusion within a latent representation rather than directly in the original image space, improving computational efficiency while maintaining image quality. The study demonstrated the effectiveness of diffusion models for detailed image generation and represents an important development for modern Generative AI applications [17].

III. RESEARCH METHODOLOGY

Seed-Bag Image Dataset Collection

The proposed methodology begins with the collection of horticultural seed-bag images representing different seed varieties and packaging categories. The images serve as the primary visual source for learning the characteristics of seed bags. Important features such as bag shape, colour, packaging structure, labels, printed patterns, and surface appearance are considered during dataset preparation. The collected images are organized according to their relevant categories to provide a suitable foundation for Generative AI-based synthetic image generation.

Image Preprocessing

The collected seed-bag images are processed before being provided to the Generative AI model. Preprocessing includes image resizing, quality checking, normalization, background handling, and removal of unsuitable or corrupted images. The objective is to maintain a consistent input format while preserving important seed-bag characteristics. Proper preprocessing helps the generative model learn meaningful visual patterns instead of unwanted noise or irrelevant image information.

Generative AI Model Selection

A suitable Generative AI model is selected according to the requirements of synthetic seed-bag image generation. GAN-based models and diffusion-based models can be considered for this purpose. GAN approaches learn image distributions through adversarial training, while diffusion models generate images through progressive denoising. The selected model is trained or adapted using the available seed-bag image collection so that it can learn the visual characteristics required for generating realistic synthetic samples [1], [15], [17].

Synthetic Seed-Bag Image Generation

After model preparation, the Generative AI system generates new seed-bag images based on the learned visual characteristics. The generated samples are designed to retain important properties of real seed bags while introducing additional visual diversity. The generation process can produce variations in bag position, orientation, lighting, background, packaging appearance, image scale, viewing angle, and partial visibility. These variations help create a broader representation of seed-bag appearances than may be available in the original dataset.

Controlled Visual Variation

Controlled variation is an important stage of the proposed methodology. Synthetic images are generated to represent conditions that may occur during practical image acquisition. Different orientations can represent front-facing, tilted, rotated, or partially turned bags. Lighting variations can represent bright, low-light, shadowed, and unevenly illuminated conditions. Background and scale variations can represent storage areas, shelves, packaging environments, different camera distances, and different image compositions. Such variations increase the diversity of the generated dataset.

Synthetic Image Quality Verification

The generated images are evaluated before they are incorporated into the final dataset. The verification process examines whether the synthetic samples contain recognizable seed-bag structures and realistic packaging characteristics. Image realism, colour consistency, label appearance, structural integrity, and visual artifacts are considered during evaluation. Images containing severe distortions, unrealistic packaging structures, unclear seed-bag characteristics, or irrelevant artifacts are excluded from further processing.

Real and Synthetic Dataset Integration

The verified synthetic images are combined with the original real seed-bag images to construct an enhanced dataset. Real images provide authentic representations of horticultural seed bags, while synthetic images contribute additional visual conditions and variations. The combined dataset is organized according to the required seed-bag categories. This stage provides a larger and more diverse visual collection for the subsequent detection and classification process.

Generative AI-Based Detection and Classification

The enhanced dataset is used to support the identification and classification of horticultural seed bags. When a new seed-bag image is provided, its visual characteristics are analyzed to determine the appropriate category. The system considers packaging appearance, colour, shape, labels, patterns, orientation, and other learned visual characteristics. The inclusion of synthetic samples is intended to improve the representation of challenging image conditions and support more flexible seed-bag recognition.

Generative AI Performance Evaluation

The proposed framework is evaluated using measures related to both image generation and seed-bag analysis. Generated-image quality and diversity are examined to determine whether the synthetic samples are realistic and sufficiently different from the original images. Detection accuracy and classification accuracy can then be used to evaluate the usefulness of the generated samples for automated seed-bag identification. Generation time, visual consistency, and successful generation rate can also be considered when comparing different Generative AI approaches.

Comparative Analysis

Different Generative AI approaches can be compared using the same dataset and evaluation criteria. GAN, Conditional GAN, image-to-image generation, and diffusion-based approaches can be assessed according to image quality, diversity, generation efficiency, visual consistency, and usefulness for detection and classification. This comparison helps determine which generative approach provides the most suitable balance between realistic image generation and practical seed-bag analysis.

Overall Methodological Workflow

The complete methodology can be represented as:

Seed-Bag Image Collection → Image Preprocessing → Generative AI Model Selection → Model Training/Adaptation → Synthetic Image Generation → Visual Quality Verification → Real + Synthetic Dataset Integration → Seed-Bag Detection → Seed-Bag Classification → Performance Evaluation → Comparative Analysis

IV. EXISTING SYSTEM

The existing system for horticultural seed-bag analysis mainly relies on manually collected real images and conventional image-processing or data-augmentation techniques. Such approaches are dependent on the availability of sufficiently large and varied image datasets, while changes in lighting, background, orientation, scale, and viewing conditions can make consistent detection and classification difficult. Generative AI methods such as GANs and diffusion models have recently provided alternatives for creating synthetic agricultural images, but these approaches are generally developed for broader agricultural or computer-vision applications rather than specifically for horticultural seed bags. Existing generative approaches can produce realistic variations and help address limited-data problems, but their effectiveness depends strongly on the quality of the source dataset and the realism of the generated samples. Recent agricultural research also shows that synthetic images can support downstream detection and classification tasks, although generated content may still require quality assessment and careful integration with real data. Therefore, a dedicated Generative AI framework focused specifically on synthetic horticultural seed-bag image generation and its subsequent use for detection and classification is still needed.

V. PROPOSED SYSTEM

The proposed system introduces a Generative AI-driven framework for synthetic horticultural seed-bag image generation, detection, and classification. The system first collects and preprocesses available seed-bag images and then uses a suitable Generative AI model, such as a GAN or diffusion model, to learn their visual characteristics. The model generates new synthetic images with variations in orientation, lighting, background, packaging appearance, scale, viewing angle, and partial visibility. The generated samples are evaluated for realism, visual consistency, and preservation of important seed-bag features, and unsuitable images are removed. The verified synthetic images are then combined with real seed-bag images to create a more diverse dataset. This enhanced dataset is used for automated detection and classification of seed-bag categories. The system evaluates both the quality of generated images and their usefulness for downstream identification, using measures such as image diversity, generation quality, detection accuracy, and classification accuracy. This approach follows the broader finding that synthetic agricultural imagery can help address limited and imbalanced datasets and improve downstream visual-analysis tasks.

VI. SYSTEM ARCHITECTURE

The proposed system architecture illustrates the complete workflow of the Generative AI-Driven Synthetic Image Generation for Horticultural Seed Bag Detection and Classification system. The process begins with the collection of horticultural seed-bag images, followed by preprocessing operations such as resizing, normalization, background handling, and quality filtering. The processed images are then provided to the Generative AI model, where GAN, Conditional GAN, diffusion, or image-to-image models can learn the visual characteristics of seed bags and generate diverse synthetic images. These generated images are subjected to quality verification to check their visual realism, structural

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

consistency, labels, shape, and packaging characteristics, while unrealistic samples are removed. The verified synthetic images are then combined with the original real images to create an enhanced dataset. This dataset is used for seed-bag detection and classification, where the system identifies the seed bag and predicts its category. Finally, the generated images and classification results are evaluated using image quality, detection accuracy, classification accuracy, generation efficiency, and comparative analysis of different Generative AI approaches. The final system can be deployed through a user interface where users can upload an image and obtain the corresponding detection and classification result.

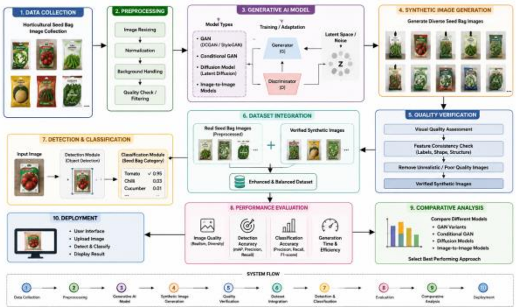

Fig 5.1 System Architecture

VII. RESULTS AND DESCUSSIONS

Experimental Results

The proposed Generative AI-driven framework was evaluated through the stages of image collection, preprocessing, synthetic image generation, quality verification, dataset integration, and seed-bag detection and classification. The Generative AI model generated additional horticultural seed-bag images with variations in orientation, lighting, background, packaging appearance, scale, partial occlusion, and viewing angle. The generated samples were examined to determine whether they retained the important characteristics of real seed bags while providing additional visual diversity. The experimental process indicates that Generative AI can be used to expand a limited image collection and provide additional samples for subsequent seed-bag analysis.

Generative AI Image Quality Analysis

The generated images were evaluated based on their visual realism, structural consistency, packaging characteristics, labels, colours, patterns, and overall appearance. Particular attention was given to whether the generated images maintained recognizable seed-bag structures while introducing meaningful variations. Images containing severe distortions, unrealistic packaging structures, unclear seed-bag characteristics, or irrelevant artifacts were excluded. This verification stage is important because simply increasing the number of images does not guarantee that the enhanced dataset will provide useful information.

Table 1. Generative AI Image Quality and Diversity Analysis

Evaluation Parameter	Observation
Image Realism	Generated samples are examined for similarity to real seed-bag appearance
Image Diversity	Samples provide variations in visual appearance and environmental conditions
Orientation Variation	Front, tilted, rotated, and partially turned seed-bag views are represented
Lighting Variation	Bright, low-light, shadowed, and uneven illumination conditions are considered
Background Variation	Different storage, packaging, shelf, and surrounding environments are represented
Scale Variation	Seed bags are represented at different sizes and camera distances

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

Evaluation Parameter	Observation
Visual Consistency	Important bag shape, colour, packaging, and label characteristics are preserved
Artifact Detection	Unrealistic structures and unwanted visual artifacts are removed

The analysis shows that image quality and image diversity need to be considered together. Highly similar generated images may provide limited additional information, whereas unrealistic images may reduce the reliability of subsequent analysis. Therefore, the proposed framework emphasizes the selection of realistic and visually meaningful synthetic samples before their integration with the original dataset.

GenAI-Enhanced Dataset Analysis

The original horticultural seed-bag dataset and the GenAI-enhanced dataset were considered for comparison. The original collection consists of real seed-bag images, whereas the enhanced collection combines real images with verified synthetic samples. The purpose of this process is to increase visual coverage rather than merely increase the total number of images. Synthetic samples contribute additional variations in orientation, illumination, background, scale, packaging appearance, and viewing angle.

Table 2. Original Dataset and GenAI-Enhanced Dataset Comparison

Parameter	Original Dataset	GenAI-Enhanced Dataset
Real Images	Available	Available
Synthetic Images	Not included	Verified synthetic samples included
Orientation Diversity	Limited to collected images	Expanded through generated variations
Lighting Diversity	Dependent on original images	Additional lighting conditions generated
Background Diversity	Dependent on image sources	Additional backgrounds represented
Scale Diversity	Based on captured images	Multiple scales can be generated
Occlusion Diversity	Limited by available images	Partial occlusion can be introduced
Overall Visual Coverage	Relatively dependent on original collection	Broader visual representation

The enhanced dataset provides a broader representation of horticultural seed bags by combining authentic images with selected synthetic samples. This integration can help represent conditions that may occur during practical image acquisition but may be underrepresented in the original collection. The effectiveness of the enhancement should ultimately be confirmed using the actual experimental measurements obtained from the implemented system.

Detection and Classification Analysis

The enhanced image collection is subsequently used for seed-bag detection and classification. When a new image is provided, the system analyzes its visual characteristics and determines the appropriate seed-bag category. The synthetic samples provide additional examples of different packaging appearances, orientations, backgrounds, lighting conditions, and viewing angles. This broader representation can support the

Received: 06-07-2026 | Accepted: 13-08-2026 | Published: 20-08-2026

identification of seed bags under varying visual conditions. The analysis focuses on whether the system can correctly distinguish different horticultural seed-bag categories despite changes in their visual appearance.

Comparative Generative AI Analysis

The proposed study considers different Generative AI approaches, including GAN-based generation, Conditional GANs, image-to-image generation, and diffusion-based generation. These approaches can be evaluated using common criteria such as image quality, diversity, generation time, visual consistency, detection accuracy, and classification accuracy. Such comparison helps identify the approach that provides an appropriate balance between realistic image synthesis and usefulness for seed-bag analysis.

Since the provided document does not contain actual numerical experimental values, numerical accuracy or performance values should not be inserted without results from the implemented model. The document itself specifies that values such as generated image quality, diversity, detection accuracy, and classification accuracy must be calculated from the actual experimental observations.

Generated Output Analysis

The generated outputs consist of synthetic horticultural seed-bag images together with their corresponding category or identification information. The outputs are examined to determine whether the generated samples preserve important characteristics such as packaging structure, colour, shape, label arrangement, and seed-related visual information. High-quality generated samples should remain recognizable while introducing meaningful visual differences. Incorrect or low-quality outputs can reveal limitations associated with insufficient training examples, similar packaging designs, complex backgrounds, or artifacts introduced during generation.

Discussion of Findings

The overall findings indicate that Generative AI provides a promising mechanism for expanding horticultural seed-bag image collections. The generated samples can introduce variations that may be difficult to obtain through manual image collection alone. Combining verified synthetic images with real images creates a broader visual representation that can support subsequent detection and classification. However, the quality of synthetic data remains an important consideration. Generated images containing distorted labels, unrealistic packaging structures, or unrelated visual artifacts should be removed before being incorporated into the enhanced dataset.

The results therefore support a quality-controlled Generative AI pipeline, where image generation is followed by visual verification and dataset integration before detection and classification. The framework demonstrates that the usefulness of Generative AI should be assessed not only by the number of images generated but also by their realism, diversity, consistency, and contribution to the final seed-bag analysis.

Overall Result

Overall, the proposed framework demonstrates the potential of Generative AI-driven synthetic image generation for horticultural seed-bag detection and classification. The main contribution is the creation of additional visual samples that complement the original dataset and provide broader representation of different seed-bag conditions. The final performance can be quantitatively reported using the actual values for image quality, image diversity, detection accuracy, classification accuracy, correctly generated samples, and overall generation success rate once the system has been experimentally implemented.

VIII. CONCLUSION

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

The proposed study demonstrates the potential of Generative AI-driven synthetic image generation for horticultural seed-bag detection and classification. By learning the visual characteristics of available seed-bag images, Generative AI can produce additional samples with variations in orientation, lighting, background, packaging appearance, scale, viewing angle, and partial occlusion. The generated images can be quality-checked and combined with real images to create a more diverse dataset for seed-bag analysis. The study highlights that the value of synthetic data depends on both image realism and diversity, making quality verification an important stage of the framework. Overall, the proposed approach provides a flexible Generative AI-based solution for expanding seed-bag image collections and supporting automated detection and classification, with potential applications in agricultural inventory management, packaging verification, storage, and quality-control activities.

IX. FUTURE SCOPE

The future scope of the proposed Generative AI-driven framework can be extended by developing more advanced GAN and diffusion-based models capable of producing highly realistic and controllable horticultural seed-bag images. Future work can focus on generating larger and more diverse synthetic datasets covering additional seed varieties, packaging designs, lighting conditions, complex backgrounds, viewing angles, and partial occlusions. Generative AI can also be combined with text-to-image and image-to-image techniques to provide greater control over the characteristics of generated samples. Further research can explore real-time seed-bag detection using camera-based systems, automated synthetic-image quality assessment, and multimodal Generative AI for combining visual and textual seed-bag information. The framework can also be extended toward cloud and edge-based deployment for agricultural inventory management, packaging verification, automated reporting, and large-scale monitoring, making the proposed system more flexible, scalable, and suitable for practical horticultural environments.

X. REFERENCES

[1] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," Advances in Neural Information Processing Systems, vol. 27, 2014. DOI: 10.48550/arXiv.1406.2661.

[2] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets," 2014. DOI: 10.48550/arXiv.1411.1784.

[3] A. Radford, L. Metz, and S. Chintala, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," 2015. DOI: 10.48550/arXiv.1511.06434.

[4] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved Techniques for Training GANs," Advances in Neural Information Processing Systems, vol. 29, 2016. DOI: 10.48550/arXiv.1606.03498.

[5] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein Generative Adversarial Networks," 2017. DOI: 10.48550/arXiv.1701.07875.

[6] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "On the Effectiveness of Least Squares Generative Adversarial Networks," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 41, no. 12, pp. 2947–2960, 2019. DOI: 10.1109/TPAMI.2018.2872043.

[7] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive Growing of GANs for Improved Quality, Stability, and Variation," 2017. DOI: 10.48550/arXiv.1710.10196.

[8] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "StarGAN: Unified

Generative Adversarial Networks for Multi-Domain Image-to-Image Translation," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018. DOI: 10.1109/CVPR.2018.00916.

[9] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1125–1134, 2017. DOI: 10.1109/CVPR.2017.632.

[10] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks," Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232, 2017. DOI: 10.1109/ICCV.2017.244.

[11] T. Karras, S. Laine, and T. Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4401–4410, 2019. DOI: 10.1109/CVPR.2019.00453.

[12] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and Improving the Image Quality of StyleGAN," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8107–8116, 2020. DOI: 10.1109/CVPR42600.2020.00813.

[13] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila, "Training Generative Adversarial Networks with Limited Data," Advances in Neural Information Processing Systems, vol. 33, pp. 12104–12114, 2020. DOI: 10.48550/arXiv.2006.06676.

[14] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, "Alias-Free Generative Adversarial Networks," Advances in Neural Information Processing Systems, vol. 34, 2021. DOI: 10.48550/arXiv.2106.12423.

[15] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," 2020. DOI: 10.48550/arXiv.2006.11239.

[16] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-Based Generative Modeling through Stochastic Differential Equations," 2021. DOI: 10.48550/arXiv.2011.13456.

[17] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10684–10695, 2022. DOI: 10.1109/CVPR52688.2022.01042.

[18] C. Saharia, W. Chan, H. Chang, C. A. Lee, J. Ho, T. Salimans, D. J. Fleet, and M. Norouzi, "Palette: Image-to-Image Diffusion Models," ACM SIGGRAPH Conference Proceedings, 2022. DOI: 10.1145/3528233.3530757.

[19] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. Ghasemipour, B. Karagol-Ayan, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi, "Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding," Advances in Neural Information Processing Systems, vol. 35, 2022. DOI: 10.52202/068431-2643.

[20] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, Y. Zhang, and M.-H. Yang, "Diffusion Models: A Comprehensive Survey of Methods and Applications," ACM Computing Surveys, vol. 56, no. 4, 2023. DOI: 10.1145/3626235.