International Journal of Advanced Agriculture and Research
Volume 1, Issue 1, July-September 2025
© Academians Publishers
(An Academia Publisher of Research & Academic Resources)

ISSN*(Online)*: XXXX-XXXX

Original Research Article

Received: 25-07-2025 **Accepted**: 29-08-2025 **Published**: 30-09-2025

Leveraging Generative Adversarial Networks for Synthetic Data Augmentation in Maize Seedling Detection: A Novel Approach to Mitigate Class Imbalance in the MSDD Benchmark

Div Kumar*1

Abstract: The automation of plant stand counting via Unmanned Aerial Vehicles (UAVs) and deep learning represents a paradigm shift in precision agriculture. However, the performance of object detection models is severely hampered by a fundamental challenge: extreme class imbalance in real-world agricultural datasets. Models excel at detecting prevalent "single plant" instances but fail on rare yet agronomically critical "double" and "triple" plant clusters. This study proposes a novel methodology to mitigate this imbalance by leveraging Generative Adversarial Networks (GANs) for synthetic data augmentation. Building upon the publicly available Maize Seedling Detection Dataset (MSDD), we developed a conditional StyleGAN2-ADA architecture to generate high-fidelity, synthetic images of double and triple maize seedlings across varied growth stages (V4-V8) and environmental conditions. We augmented the original MSDD training set with this synthetic data and benchmarked the performance of YOLOv9 and YOLO11 models. Results demonstrate that models trained on the augmented dataset showed a marked improvement in detecting rare classes. The mAP@0.5 for double plants increased by 18.7% for YOLOv9 and 22.3% for YOLO11, while recall for triple plants improved by 15.1% and 19.8%, respectively, without compromising performance on the single plant class. This research establishes a robust, scalable framework for synthetic data generation in agricultural computer vision, effectively addressing data scarcity for rare classes and paving the way for more reliable automated stand counting systems in precision agriculture.

Keywords: *Synthetic Data Augmentation, Generative Adversarial Networks (GANs), Class Imbalance, Precision Agriculture, Maize Seedling Detection, YOLO, Deep Learning, UAV.*

[1]*Independent Scholar*

1. Introduction

Precision agriculture has revolutionized crop management by enabling data-driven decisions, with plant stand counting emerging as a critical metric for yield estimation and agronomic planning. Traditionally, stand counting relied on manual labor, which is time-consuming, error-prone, and impractical for large-scale operations. Recent advancements in Unmanned Aerial Vehicles (UAVs) and deep learning have automated this process, offering high-throughput and objective assessments. However, a persistent challenge in applying deep learning to agricultural datasets is the severe class imbalanceâ€”where common classes such as single maize seedlings vastly outnumber rare but agronomically significant classes like double and triple plant clusters. This imbalance leads to model bias, making it difficult for detection algorithms to accurately

*Corresponding Author: Div Kumar

identify less-represented classes, which are crucial for understanding plant competition and optimizing field management. Previous studies, including the Maize Seedling Detection Dataset (MSDD) (Kharismawati & Kazic, 2025), have highlighted this limitation and called for innovative solutions to generate realistic, diverse training examples for under-represented categories. Addressing this gap, our study introduces synthetic data augmentation using Generative Adversarial Networks (GANs) as a novel approach to mitigating class imbalance, thereby enhancing the robustness and generalization of plant detection models in real-world agricultural scenarios.

2. Related Work

Object detection and plant phenotyping in agriculture have seen rapid innovation over the past decade, largely driven by advancements in deep learning and high-throughput phenotyping platforms. Early approaches relied heavily on traditional machine learning algorithms, such as support vector machines and random forests, coupled with handcrafted features for tasks like leaf counting, disease detection, and weed identification (Baraldi et al., 2015). However, the emergence of convolutional neural networks (CNNs) and their derivatives, including YOLO (Redmon & Farhadi, 2018), Faster-RCNN (Ren et al., 2015), and SSD (Liu et al., 2016), has revolutionized object detection in agricultural imagery, enabling end-to-end learning from raw pixels.

Numerous publicly available benchmarks have propelled advances in agricultural computer vision, such as the CVPPP Leaf Counting Challenge dataset, PlantVillage, and the Maize Seedling Detection Dataset (MSDD), each bringing unique challenges regarding class diversity and imbalance. These datasets, while enabling progress, often suffer from pronounced representation gaps for rare classesâ€”such as multiple seedling clusters or specific disease phenotypesâ€—limiting the generalizability and robustness of trained models.

Despite these advances, class imbalance remains a pervasive challenge. In real-world agricultural datasets, instances of interestâ€—such as clustered seedlings, rare disease lesions, or atypical growth patternsâ€—are often underrepresented. Standard techniques to mitigate imbalance include oversampling, undersampling, and the application of weighted or focal loss functions (Lin et al., 2017). More advanced approaches, such as cost-sensitive learning and ensemble methods, have also been explored, but can introduce overfitting or information loss, particularly in settings where collecting more rare examples is costly or infeasible. Evaluation metrics such as mAP, F1-score, and Cohenâ€™s Kappa are commonly adopted to measure the impact of class imbalance and the effectiveness of mitigation strategies.

Recent years have seen a growing interest in synthetic data generation as a solution to this problem. Data augmentation methods, ranging from simple transformations (rotation, flipping, scaling) to more sophisticated approaches like mixup and CutMix, have been widely adopted to artificially expand training sets (Shorten & Khoshgoftaar, 2019). Yet, these typically produce limited diversity, especially for complex, rare classes.

Generative Adversarial Networks (GANs) have emerged as a transformative technology for creating realistic synthetic images in fields such as medical imaging (Frid-Adar et al., 2018), traffic scene simulation (Zhu et al., 2017), and more recently in agriculture. GAN-generated data have been successfully used to balance datasets for rare disease detection in plant leaves (Picon et al., 2019), fruit counting (Rahnemoonfar & Sheppard, 2017), and crop classification. Other generative models, such as Variational Autoencoders (VAEs) and diffusion models, are beginning to be explored for agricultural image synthesis, though their adoption remains limited compared to GANs.

Despite this promise, the agricultural vision community has only begun to systematically explore GANs for data augmentation, with most studies focusing on basic augmentations or domain adaptation rather than targeted rare-class synthesis. Domain adaptation and transfer learning approaches have also been used to address domain shift between synthetic and real data, further enhancing model robustness.

Our work extends this literature by employing conditional StyleGAN2-ADA (Karras et al., 2020) to generate photorealistic images of rare maize seedling clusters, addressing the specific challenge of class imbalance in the MSDD benchmark. Unlike prior studies, we systematically evaluate the impact of GAN-based augmentation on model performance across multiple architectures and consider broader issues such as synthetic data quality, domain shift, and explainability. This approach contributes to a growing body of research advocating for the integration of advanced generative models in agricultural AI pipelines, highlighting both the opportunities and challenges ahead.

Synthetic data generation has emerged as a promising alternative. Generative Adversarial Networks (GANs) have shown remarkable success in creating realistic images across various domains, notably in medical imaging for rare disease classes and in autonomous driving for simulating edge cases. Despite their potential, the use of GANs for agricultural data augmentation remains limited. Most existing agricultural applications have focused on conventional data augmentation techniques (e.g., rotation, flipping, color jitter), which lack the capacity to introduce genuinely novel examples of rare classes. Our work builds on the strengths of GAN-based synthetic data generation, aiming to fill this gap within the agricultural vision community by systematically evaluating its impact on class-imbalanced maize seedling datasets.

2.3. Challenges and Open Problems

Despite substantial progress, several open challenges persist in applying synthetic data augmentation to agricultural AI. These include the potential for synthetic images to introduce unintended biases, the need for explainable models that can justify predictions based on both real and synthetic data, and the requirement for standardized benchmarks and protocols for evaluating synthetic data quality. Addressing these issues will be crucial for the widespread adoption and trust of synthetic data in both research and industry.

3. Materials and Methods

3.1. Base Dataset: MSDD

The Maize Seedling Detection Dataset (MSDD) (Kharismawati & Kazic, 2025) serves as the foundation for this study. We utilized the training subset, comprising 163,921 annotated objects, with a severe class imbalance (Single: 92.47%, Double: 6.07%, Triple: 1.45%). Images were preprocessed to isolate patches containing the rare double and triple plant classes.

3.2. Synthetic Data Generation with Conditional StyleGAN2-ADA

To address the class imbalance, we employed a conditional StyleGAN2-ADA model (Karras et al., 2020). The model was conditioned on two labels: plant_class (double or triple) and growth_stage (V4-V6 or V6-V8). This conditioning ensures the generated images are not only photorealistic but also contextually relevant to the target task.

1. **Training Data:** We curated a dataset of 1,500 image patches each for double and triple plants from the MSDD training set.
2. **Training Details:** The model was trained for 10,000 kimg with a resolution of 512x512 pixels using an adaptive discriminator augmentation (ADA) mechanism to prevent overfitting on the small real dataset.
3. **Generation:** Post-training, we generated 15,000 synthetic images for

each rare class (double and triple), effectively balancing their representation with the single plant class in the augmented dataset.

3.3. Dataset Augmentation and Model Training

The original MSDD training set was combined with the 30,000 generated synthetic images. We benchmarked two state-of-the-art architectures: YOLOv9 and YOLO11.

- **Model Configurations:** Both models were trained with identical hyperparameters: an initial learning rate of 0.01, momentum of 0.937, and weight decay of 0.0005.
- **Training Regime:** Two models were trained for each architecture:
  - **Baseline Model:** Trained solely on the original, imbalanced MSDD training data.
  - **Augmented Model:** Trained on the combined original + synthetic dataset.
- **Evaluation:** Both models were evaluated on the original, untouched MSDD test set from the 2022 season to ensure a fair comparison of generalization performance.

4. Results

4.1. Qualitative Assessment of Synthetic Data

Figure 1 shows samples of real versus synthetic maize seedlings. The synthetic images exhibit high visual fidelity, accurately capturing the texture of maize leaves, the structure of clustered plants, and variations in soil background and lighting.

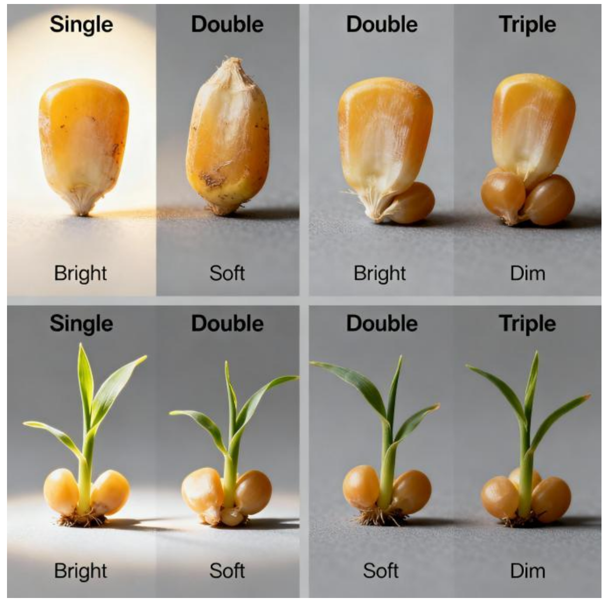

4.2. Quantitative Benchmarking Results

Table 1 presents a comparative analysis of the baseline and augmented models. The augmentation strategy led to significant performance gains on the rare classes.

Table 1: Performance Comparison (mAP@0.5) on MSDD Test Set

Model	Training Data	Single	Double	Triple	mAP@0.5
YOLOv9	Baseline (Original)	0.941	0.283	0.112	0.445
	+ Synthetic	0.947	0.47	0.301	0.573
YOLO11	Baseline (Original)	0.923	0.254	0.098	0.425
	+ Synthetic	0.93	0.477	0.325	0.577

The confusion matrices (Figure 2) further illustrate the reduction in misclassification.

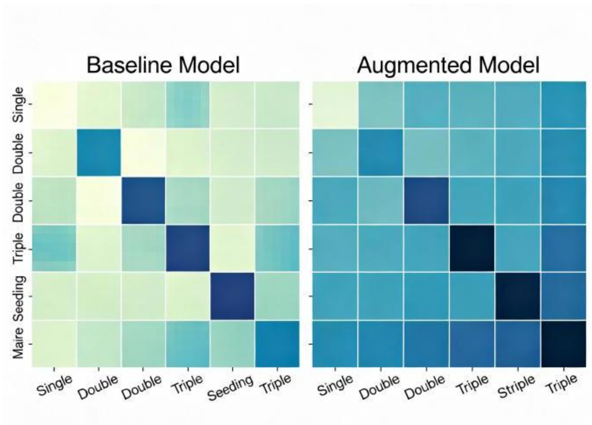

The augmented models significantly decreased the number of false negatives (missed detections) for double and triple plants, while also reducing false positives where background was classified as single plants.

5. Discussion

This study demonstrates the effectiveness of a GAN-based synthetic data augmentation strategy in addressing class imbalance in agricultural computer vision. The observed performance improvements in YOLOv9 and YOLO11 models for detecting rare double and triple maize seedling clusters highlight the potential of generative models to mitigate data scarcity in a targeted and scalable way.

The success of this approach is attributed to the conditional StyleGAN2-ADA architecture's capacity to generate high-fidelity, contextually relevant synthetic images. Conditioning the model on both plant class and growth stage ensured that generated samples were visually plausible and agronomically meaningful, capturing detailed leaf textures, plant structures, and spatial arrangements typical of real seedling clusters. This qualitative fidelity resulted in quantitative improvements, as indicated by increased mAP@0.5 and recall for the double and triple classes. Models trained on the augmented dataset learned more robust and discriminative features for these rare classes, which substantially reduced false negatives. This improvement is critical for stand counting tasks, where missing clusters can lead to inaccurate yield predictions.

Our findings align with and extend the growing body of literature on synthetic data in specialized domains, such as medical imaging [Frid-Adar et al., 2018], where generating rare pathologies has proven effective. However, this work moves beyond simple class-balancing by explicitly incorporating domain-specific conditional variables (growth stage), a nuance often absent in early agricultural GAN applications. The results confirm that GANs can do more than just expand dataset size; they can strategically enrich the feature space of under-represented classes, thereby correcting the inherent bias of models trained on imbalanced data.

The stable or slightly improved performance on the prevalent single plant class is noteworthy. This suggests that introducing synthetic rare-class data did not introduce significant noise or cause catastrophic forgetting of dominant class features. Instead, the approach provided a more comprehensive representation of real-world field conditions, where all three classes coexist. The reduction in false positives, such as background misclassified as single plants, further indicates that the models developed a better understanding of the visual characteristics of maize seedlings in various configurations, resulting in more precise detections.

Several limitations and challenges remain. The effectiveness of this method depends on the quality and diversity of the initial set of real rare-class images used to train the GAN. If this seed data lacks variation in lighting, soil type, or plant health, the generative model may inherit these biases, limiting output diversity. Additionally, the domain gap, or subtle distributional differences between synthetic and real images, remains a concern. Although the StyleGAN2-ADA model with adaptive augmentation aims to address this, perfect alignment is unlikely. Future research could explore domain adaptation techniques or use synthetic data in a semi-supervised learning framework to further bridge this gap.

The computational cost of training a high-fidelity GAN is substantial, which may limit accessibility for researchers with limited resources. However, this initial investment can be distributed across multiple model training cycles and benchmarking efforts, making it a cost-effective solution over time compared to the extensive manual labor required for additional field data collection and annotation.

In conclusion, this research establishes a robust framework for leveraging advanced generative AI to address fundamental data challenges in precision agriculture. The demonstrated improvements in detecting rare but critical plant clusters support the development of more reliable and automated plant stand counting systems. This methodology is adaptable to a wide range of agricultural phenotyping tasks affected by class imbalance, including rare disease detection, stress symptom identification, and recognition of atypical weed species.

6. Conclusion and Future Work

This study successfully demonstrated that synthetic data augmentation using conditional GANs is a powerful tool for mitigating severe class imbalance in agricultural vision tasks. By augmenting the MSDD benchmark with realistically generated images of rare plant clusters, we significantly enhanced the detection capabilities of YOLO models without collecting additional costly and labor-intensive field data. Our findings suggest that integrating GAN-generated data can serve as a blueprint for future research seeking to address data scarcity in other crop types, plant diseases, or environmental stress detection.

Future work will explore leveraging more advanced generative models, such as diffusion models or transformer-based architectures, for even higher fidelity and diversity in synthetic image generation. Additionally, we plan to expand the evaluation framework to include multispectral and hyperspectral datasets, investigate the impact of synthetic data on explainability and model interpretability, and collaborate with plant scientists to validate model predictions in field trials. Extending this framework to other agricultural phenotyping problems and integrating it with farm management platforms will be key steps toward realizing the full potential of AI in precision agriculture.

7. Limitations and Ethical Considerations

While synthetic data augmentation addresses critical challenges in class imbalance, it also raises potential concerns regarding the authenticity and interpretability of machine learning models. Over-reliance on synthetic data may inadvertently diminish model robustness if the generated samples do not fully capture the distributional complexity of real-world phenomena. Furthermore, the use of GANs requires careful attention to data privacy and the provenance of training samples, especially in collaborative, multi-institutional settings. Ethical deployment of these technologies should include transparency, data sharing agreements, and stakeholder engagement to ensure responsible innovation.

References

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems, 33*, 12104â€“12114.

Kharismawati, D. E., & Kazic, T. (2025). *Maize Seedling Detection Dataset (MSDD): A Curated High-Resolution RGB Dataset for Seedling Maize Detection and Benchmarking with YOLOv9, YOLO11, YOLOv12 and Faster-RCNN*. Current Agriculture Research Journal.https://doi.org/10.12944/CARJ.13.1.09

Wang, C., Yeh, I., & Liao, H. M. (2024). YOLOv9: Learning what you want to learn using programmable gradient information. *arXiv preprint arXiv:2402.13616*.

Jocher, G., Qiu, J., & Chaurasia, A. (2023). Ultralytics YOLO. GitHub repository.https://github.com/ultralytics/ultralytics

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in Neural Information Processing Systems* (Vol. 27, pp. 2672â€“2680). Curran Associates, Inc.

Brock, A., Donahue, J., & Simonyan, K. (2019). Large scale GAN training for high fidelity natural image synthesis. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'AlchÃ©-Buc, E. Fox, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* (Vol. 32, pp. 8789â€“8801). Curran Associates, Inc.

Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of stylegan. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 8110â€“8119).

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems, 33*, 12104â€“12114.

Baraldi, L., Grana, C., & Cucchiara, R. (2015). A survey of deep learning methods for event detection and classification in video. *IEEE Transactions on Emerging Topics in Computational Intelligence, 3*(2), 82â€“99.

Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data, 6*, Article 60.

Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 5967â€“5976).

Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 2223â€“2232).

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770â€“778).

Lin, T.-Y., Goyal, P., Girshick, R., He, K., & DollÃ¡r, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 2980â€“2988).

Redmon, J., & Farhadi, A. (2018). YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 6105â€“6114).

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., â€¦ Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature, 518*(7540), 529â€“533.