

Research Article

Generative Adversarial Networks for Synthetic Data Generation in Deep Learning Applications

Mohamed Islam Keskes*

Transilvania University of Brasov, Romania

Abstract

Generative Adversarial Networks (GANs) have emerged as a transformative approach for synthetic data generation in deep learning, addressing critical challenges such as data scarcity, privacy concerns, and algorithmic bias. This synthesis review provides a comprehensive analysis of GANs' role in creating high-fidelity synthetic data across diverse domains, including healthcare, finance, computer vision, and natural language processing. By leveraging an adversarial training process involving a generator and discriminator, GANs effectively capture intricate data distributions, generating realistic synthetic samples that enhance model robustness and generalization. The review explores foundational GAN principles, advanced architectures like DCGANs, cGANs, CycleGANs, and TimeGANs, and their applications in generating medical images, financial time-series, and tabular data. It also discusses the advantages of GANs, such as privacy preservation and cost-efficiency, alongside limitations, including training instability, mode collapse, and the lack of standardized evaluation metrics. Comparative analysis with other methods like Variational Autoencoders and traditional statistical approaches highlights GANs' superior realism for complex data types. Future research directions include improving training stability, developing robust evaluation benchmarks, and integrating privacy-enhancing techniques. This review underscores GANs' potential to revolutionize deep learning applications while emphasizing the need for ethical guidelines to mitigate misuse risks.

Introduction

The rapid growth of deep learning across various sectors has underscored the critical need for extensive, high-quality datasets to achieve top-tier model performance (LeCun, Bengio, & Hinton, 2015). These data-intensive models necessitate vast amounts of information to identify complex patterns and generalize effectively to new data [1]. However, obtaining real-world data is fraught with challenges, including data scarcity, high costs, and time-intensive processes for collection and annotation (Rolnick, et al. 2019) [2]. Furthermore, stringent privacy regulations governing sensitive data—such as personal, health, and financial records—along with legal and ethical restrictions, pose significant barriers (Veale & Binns, 2017). Real-world datasets may also contain algorithmic biases, further complicating their use (Barocas, Hardt, & Narayanan, 2019).

Synthetic data has become a powerful solution to address these challenges. By generating artificial datasets that replicate the statistical characteristics of real-world data, synthetic

data addresses data shortages, mitigates privacy concerns, and reduces biases [1,3]. It can be scaled and tailored for balanced class representation, making it especially useful for handling imbalanced datasets and simulating rare or complex scenarios [4]. These capabilities enhance model robustness and generalization in deep learning applications [5].

Generative Adversarial Networks (GANs) have become a cornerstone in synthetic data generation, gaining widespread recognition for their effectiveness [6]. GANs utilize an adversarial training framework involving two neural networks: a generator and a discriminator [7]. The generator learns the real data's underlying distribution to create synthetic samples that closely mimic it, while the discriminator works to distinguish real data from the generated samples [8]. This adversarial interplay, resembling a zero-sum game, pushes the generator to continuously improve, producing increasingly realistic synthetic data. The iterative feedback loop between the two networks ensures refined outputs, capturing complex data distributions with high fidelity [9,10].

More Information

***Address for correspondence:** Mohamed Islam Keskes, Transilvania University of Brasov, Romania, Email: mohamed.keskes@unitbv.ro

Submitted: July 05, 2025

Approved: July 25, 2025

Published: July 28, 2025

How to cite this article: Keskes MI. Generative Adversarial Networks for Synthetic Data Generation in Deep Learning Applications. J Artif Intell Res Innov. 2025; 1(1): 028-033. Available from: <https://dx.doi.org/10.29328/journal.jairi.1001004>

Copyright license: © 2025 Keskes MI. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Keywords: Generative adversarial networks; Synthetic data; Deep learning; Privacy preservation; Data scarcity



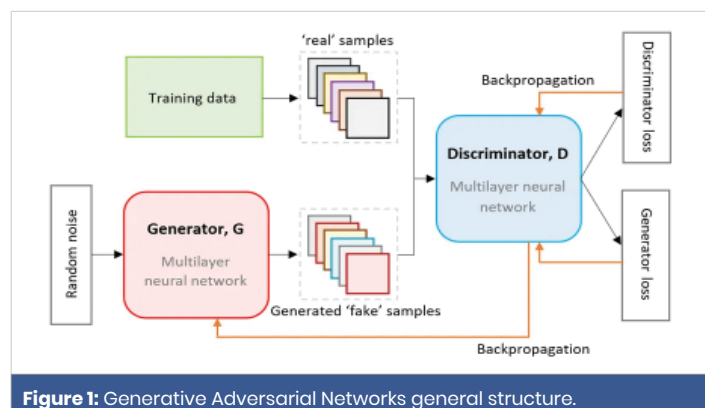
GANs outperform traditional generative models in generating synthetic data that closely mimics real-world data [11,12]. A key advantage is that the generator does not directly access the original data during training, which reduces the risk of data disclosure and potential privacy breaches [13]. This makes GANs particularly valuable for applications requiring privacy preservation while maintaining the quality and utility of synthetic datasets [6].

This review consolidates current research on utilizing GANs for synthetic data generation in deep learning. It explores GANs' core principles, applications across various fields, and advanced architectural innovations. The review compares GANs' strengths and limitations with other synthetic data generation methods and discusses quality assessment techniques. It provides a comprehensive overview for researchers and practitioners, highlighting challenges and future research directions in this rapidly evolving field.

Generative adversarial networks for synthetic data generation: A foundational overview

A GAN comprises two competing neural networks: the generator and the discriminator, both typically deep neural networks trained using backpropagation as shown in Figure 1 [11]. The generator learns the probability distribution of real-world data to produce synthetic samples that mimic the original data [14]. It takes a random noise vector, typically drawn from a Gaussian or uniform distribution, and transforms it into synthetic outputs such as images, tabular data, or time-series sequences [15]. The generator effectively learns a complex, non-linear mapping from a low-dimensional latent space to a higher-dimensional data space, generating new instances that are statistically similar to the training data [6].

In contrast, the discriminator acts as a binary classifier, distinguishing real data samples from synthetic ones produced by the generator (Wang, et al. 2025). It outputs a probability score indicating whether an input sample is real or fake, assigning high probabilities to real samples and low ones to synthetic ones [1]. This classification provides critical feedback to the generator, enabling it to refine its outputs and improve realism through the adversarial learning process [9].



GAN training is a competitive, zero-sum game. The generator aims to produce samples that fool the discriminator, while the discriminator strives to accurately classify real and fake samples [6]. Both networks are trained simultaneously in an iterative process: the discriminator improves its classification, and the generator refines its outputs to evade detection [16]. This adversarial loop continues until equilibrium, where the discriminator cannot reliably distinguish synthetic samples from real data, indicating the generator has approximated the true data distribution effectively [9,17].

The adversarial training process in GANs can be formally represented as a minimax optimization problem [7]. The objective function that governs this process typically involves the discriminator aiming to maximize the expected log-likelihood of correctly identifying real data and correctly identifying fake data (as fake). Simultaneously, the generator strives to minimize the expected log-likelihood of the discriminator correctly identifying its generated data as fake [11]. Mathematically, this can be expressed as:

$$\min_G \max_D V(D, G) = E_{X \sim P_{data}(X)} [D(X)] + E_{Z \sim P_z(Z)} [D(G(Z))]$$

Where:

- $D(x)$ represents the discriminator's output (the probability that x is real) for a real data sample x drawn from the real data distribution. $p_{data}(x)$.
- $G(z)$ represents the generator's output (a synthetic data sample) for a random noise vector z drawn from a noise distribution. $p_z(z)$.
- $D(G(z))$ represents the discriminator's output for the synthetic sample $G(z)$ (the probability that the synthetic sample is real).
- E denotes the expected value.

The discriminator seeks to maximize this value function by correctly classifying both real and fake data. The generator, on the other hand, aims to minimize this value function by producing synthetic data $G(z)$ that the discriminator is likely to classify as real (i.e., maximizing $D(G(z))$ or, equivalently, minimizing $1 - D(G(z))$). The equilibrium of this minimax game signifies that the generator has learned to produce synthetic data that is statistically indistinguishable from the real data [7,17].

Applications across diverse domains

The impact of GANs is expanding rapidly, transforming industries such as healthcare, finance, computer vision, and natural language processing. In healthcare, GANs are used to generate synthetic medical images, such as MRI and CT scans, assisting in diagnosis, treatment planning, and data augmentation for deep learning models [18,19]. They also enable data anonymization, addressing privacy concerns, with applications in brain imaging, cardiology, and oncology. For



example, models like medGAN generate synthetic electronic health records to mitigate data scarcity and preserve patient confidentiality [20].

In finance, GANs produce synthetic time-series data such as stock prices for fraud detection, forecasting, and trading strategy development [8]. TimeGAN is particularly effective at capturing temporal dependencies in financial data, enhancing risk modeling and algorithmic trading performance [21]. Similarly, FinGAN demonstrates GANs' capability in modeling complex financial distributions, supporting regulatory and market behavior analysis [22].

Computer vision has been a primary driver of GAN advancement, with applications in synthetic image and video generation, 3D model creation, and data augmentation for classification and detection tasks [11,23]. GANs also facilitate image-to-image translation (e.g., changing weather or season), super-resolution, and artistic generation such as anime character synthesis or photo restoration [12].

In NLP, while GANs are less prevalent than in vision, they generate synthetic text for tasks like text summarization and translation and support creative applications like poetry or story generation [24]. Though large language models dominate current NLP, GANs contribute to data augmentation and domain-specific text synthesis [24].

Beyond these, GANs are gaining traction in cybersecurity, fraud detection, and supply chain modeling. They generate synthetic fraudulent transactions or network traffic to improve model robustness and support imbalanced data training [25]. Models like table-GAN [26] and CTAB-GAN [27] highlight GANs' flexibility in structured data applications, underscoring their transformative potential across diverse fields.

Advanced GAN architectures and methodologies

Since the introduction of the original GAN framework, advanced architectures have been developed to overcome its limitations and enhance synthetic data generation [10]. These innovations address diverse data types and applications, improving the quality and utility of generated data [18].

Deep Convolutional GANs (DCGANs) marked a significant advancement by incorporating convolutional neural networks (CNNs) into the generator and discriminator [28]. This enabled DCGANs to synthesize realistic images by leveraging CNNs' ability to learn hierarchical spatial features, as seen in applications like generating fashion images from the Fashion MNIST dataset [23]. DCGANs set the stage for more sophisticated image synthesis models [29].

Conditional GANs (cGANs) introduced controlled generation by conditioning the process on supplementary information, such as class labels or textual descriptions [30]. This allows targeted data synthesis, particularly in healthcare, where cGANs generate medical images for specific pathologies

[31]. CTAB-GAN, a conditional tabular GAN, applies this principle to structured data, enhancing realism in synthetic tabular datasets [32].

CycleGANs facilitate unpaired image-to-image translation, learning mappings between domains without direct correspondences. Using cycle consistency loss, CycleGANs ensure reversible translations [33], proving valuable in healthcare for tasks like transforming MRI contrasts without paired data, showcasing GANs' ability to handle complex data mappings [34].

TimeGANs address time-series data generation by capturing temporal dependencies through a combination of supervised and unsupervised objectives. Applied to financial data like stock prices, TimeGAN outperforms alternatives by modeling realistic temporal dynamics, highlighting the need for specialized GAN designs for sequential data [21].

Tabular GANs address the challenges of structured data, which mixes numerical and categorical variables. Models like medGAN, CTAB-GAN, and table-GAN enhance synthetic tabular data generation. MedGAN uses an autoencoder-GAN hybrid for electronic health records [20], while CTAB-GAN and table-GAN incorporate classifiers to maintain semantic integrity [35]. Table-GAN employs hinge loss and classification loss to balance privacy and compatibility, demonstrating effectiveness in generating realistic tabular data (Xu, et al. 2019).

Training GANs requires careful optimization to overcome instability, vanishing gradients, and mode collapse. Strategies like hinge loss, gradient penalties, and hyperparameter tuning—such as adjusting epochs, batch sizes, and learning rates in FinGAN—stabilize training and improve outcomes (Xiaopeng, et al. 2020). For instance, FinGAN's adjustments enhanced its ability to capture complex financial patterns [36].

These advancements—DCGANs, cGANs, CycleGANs, TimeGANs, and tabular GANs—demonstrate the versatility of GANs in generating high-quality synthetic data across images, time-series, and structured datasets, with careful training ensuring their effectiveness in diverse applications [9] (Table 1).

Advantages and benefits of using GANs for synthetic data

GANs provide substantial advantages in synthetic data generation across multiple domains. A key benefit is addressing data scarcity by creating vast amounts of realistic synthetic data, augmenting limited real-world datasets [37]. This is critical in fields like rare disease research or specialized industries where data collection is costly or challenging [20]. By generating diverse samples, GANs facilitate the training of robust deep learning models, mitigating overfitting and enhancing generalization [7].



Table 1: Comparative Table of GAN Architectures for Synthetic Data Generation.

Architecture	Key Features	Applications	Limitations
DCGAN	Incorporation of CNNs in the generator and discriminator.	Image synthesis, feature learning.	Can still suffer from training instability and mode collapse.
cGAN	Generation conditioned on additional input (e.g., labels, text).	Controlled data generation, image editing, and text-to-image synthesis.	Requires labeled or conditional data.
CycleGAN	Unpaired image-to-image translation using cycle consistency loss.	Style transfer, domain adaptation, image enhancement.	It can sometimes produce geometrically inconsistent results.
TimeGAN	Explicit modeling of temporal correlations for time-series data.	Synthetic financial data, healthcare time-series data.	Complexity in implementation and training.
medGAN	Combines an autoencoder with a GAN for mixed-type data (binary, continuous).	Synthetic electronic health records (EHR).	Originally designed for binary and continuous data, extensions are needed for multi-categorical data.
CTAB-GAN	Conditional GAN with a classifier to learn data semantics for tabular data.	Synthetic tabular data generation, handling mixed data types.	Evaluation metrics for tabular data can be inconsistent [12]
table-GAN	Adds a classifier network to enhance the semantic integrity of synthetic tables.	Synthetic tabular data generation, privacy preservation.	Performance can vary across different datasets and may not always capture all statistical nuances [13]

GANs also mitigate privacy risks by producing synthetic data that preserves the statistical properties of original datasets without exposing sensitive information. This facilitates data sharing and collaboration while complying with regulations like GDPR and HIPAA (Jordon, et al. 2019). Models like table-GAN are designed to synthesize tabular data, minimizing disclosure risks, which is vital in sensitive sectors such as healthcare and finance [26].

Moreover, GANs assist in addressing algorithmic bias in datasets. By generating synthetic data to balance imbalanced classes or represent underrepresented groups, GANs support the development of fairer machine learning models (Xu, et al. 2020). This is crucial for preventing AI systems from perpetuating societal biases, ensuring more equitable outcomes.

Synthetic data from GANs enhances model robustness and generalization by exposing models to a wider range of scenarios, including rare or complex cases that are not present in real data. This makes models more resilient to variations and noise, improving performance on unseen data. For example, GANs can simulate challenging conditions, enabling models to handle diverse real-world inputs effectively (Shmelkov, et al. 2018).

Finally, GANs offer cost and time efficiency compared to collecting and processing real-world data. Once trained, GANs can quickly generate large volumes of synthetic data, accelerating the development and deployment of deep learning models (Torfi, et al. 2020). This efficiency, combined with their ability to overcome data limitations, privacy concerns, and biases, makes GANs a transformative tool for advancing AI applications.

Challenges, limitations and considerations

Despite the many advantages of using GANs for synthetic data generation, various challenges, limitations, and considerations must be carefully addressed. A significant issue is the instability of the training process, often necessitating extensive hyperparameter tuning and sophisticated architectural design to achieve convergence [38]. Mode

collapse, where the generator produces limited sample variety, further complicates capturing the full diversity of real data and diminishes the utility of synthetic data (Srivastava, et al. 2017). Advanced training techniques and diligent monitoring are crucial to mitigate these issues.

Another limitation is the lack of standardized evaluation metrics to assess synthetic data quality and utility. Existing metrics often focus on specific aspects, like visual fidelity, but fail to capture the data’s usefulness for downstream tasks or preservation of complex relationships, hindering model comparisons (Borji, 2019). The high computational cost of training GANs, requiring powerful GPUs and significant time, poses a barrier for those with limited resources (Lucic, et al. 2018).

Privacy concerns arise when generators memorize training data patterns, risking information leakage. Balancing privacy and utility requires ongoing research into privacy-preserving techniques (Chen, et al. 2020). Additionally, the ability of GANs to create realistic synthetic data raises ethical concerns, including the potential for deepfakes and misinformation (Westerlund, 2019). Addressing these risks demands ethical guidelines, responsible data practices, and robust detection mechanisms to maintain trust in information sources.

GANs in comparison to other synthetic data generation techniques

Generative Adversarial Networks are not the only method for synthetic data generation; other techniques like Variational Autoencoders (VAEs), Large Language Models (LLMs), and traditional statistical methods also play significant roles. VAEs, which encode data into a probabilistic latent space and decode it to generate new samples, offer greater training stability than GANs but often produce less realistic outputs, especially for complex data like images [39]. Both GAN-based (e.g., CTGAN) and VAE-based (e.g., TVAE) models are popular for tabular data [35].

LLMs excel in generating coherent, contextually relevant synthetic text, leveraging their training on vast text corpora, but are less versatile for structured data like images or tables

compared to GANs [40]. Traditional statistical methods, which model data properties like means and correlations, are computationally lighter and suitable for simpler tasks but struggle to capture complex, high-dimensional patterns [41]. These methods often require more manual domain expertise, unlike the automated learning of GANs. While GANs strike a strong balance of realism and fidelity for complex data such as images and time-series, the choice of technique depends on application requirements, data type, and trade-offs in training stability, computational cost, and data realism [40,42].

Conclusion

As a breakthrough in synthetic data generation, GANs generate highly realistic datasets, driving advancements in healthcare, finance, computer vision, and natural language processing. Their adversarial training process enables them to tackle data scarcity, mitigate privacy risks, and reduce algorithmic bias, advancing AI applications. However, challenges such as training instability, the absence of robust evaluation metrics, and ethical concerns remain. Ongoing research focuses on developing advanced GAN architectures, refining training techniques, and establishing standardized evaluation methods to address these issues. As these efforts advance, GANs promise to improve data sharing, model development, and the creation of fair, privacy-preserving, and robust deep learning solutions, with future improvements targeting stability, efficiency, controllability, and ethical deployment.

References

- Hei Z, Sun W, Yang H, Zhong M, Li Y, Kumar A, et al. Novel domain-adaptive Wasserstein generative adversarial networks for early bearing fault diagnosis under various conditions. *Reliab Eng Syst Saf*. 2025;257:110847. Available from: <https://doi.org/10.1016/j.res.2025.110847>
- Keskes MI, Nita MD. Developing an AI tool for forest monitoring: Introducing SylvaMind AI. *Bull Transilv Univ Brasov Ser II For Wood Ind Agric Food Eng*. 2024;17(66):39–54. Available from: <https://doi.org/10.31926/but.fwi.2024.17.66.2.3>
- Goncalves A, Ray P, Soper B, Stevens J, Coyle L, Sales AP. Generation and evaluation of synthetic patient data. *BMC Med Res Methodol*. 2020;20(1):108. Available from: <https://doi.org/10.1186/s12874-020-00977-1>
- Frid-Adar M, Diamant I, Klang E, Amitai M, Goldberger J, Greenspan H. GAN-based synthetic medical image augmentation for improved CNN performance in liver lesion classification. *Neurocomputing*. 2018;321:321–31. Available from: <https://doi.org/10.1016/j.neucom.2018.09.013>
- Xu L, Veeramachaneni K. Synthesizing tabular data using generative adversarial networks. *arXiv [Preprint]*. 2018 [cited 2025 Jul 25]. Available from: <https://arxiv.org/abs/1811.11264>
- Li Y, Bai F, Lyu C, Qu X, Liu Y. A systematic review of generative adversarial networks for traffic state prediction: Overview, taxonomy, and prospects. *Inf Fusion*. 2025;117:102915. Available from: <https://doi.org/10.1016/j.inffus.2024.102915>
- Wang X, Jiang H, Mu M, Dong Y. A trackable multi-domain collaborative generative adversarial network for rotating machinery fault diagnosis. *Mech Syst Signal Process*. 2025;224:111950. Available from: <https://doi.org/10.1016/j.ymssp.2024.111950>
- Chen Y, Yang XH, Wei Z, Heidari AA, Zheng N, Li Z, et al. Generative adversarial networks in medical image augmentation: A review. *Comput Biol Med*. 2022;144:105382. Available from: <https://doi.org/10.1016/j.compbiomed.2022.105382>
- Creswell A, White T, Dumoulin V, Arulkumaran K, Sengupta B, Bharath AA. Generative adversarial networks: An overview. *IEEE Signal Process Mag*. 2018;35(1):53–65. Available from: <https://doi.org/10.1109/MSP.2017.2765202>
- Gui J, Sun Z, Wen Y, Tao D, Ye J. A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE Trans Knowl Data Eng*. 2021;35(4):3313–31. Available from: <https://doi.org/10.1109/TKDE.2021.3130191>
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial nets. In: *Advances in Neural Information Processing Systems*. 2014;27:2672–80. Available from: https://proceedings.neurips.cc/paper_files/paper/2014/hash/f033ed80deb0234979a61f95710dbe25-Abstract.html
- Karras T, Laine S, Aila T. A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE; 2019;4401–10. Available from: <https://doi.org/10.1109/CVPR.2019.00453>
- Esteban C, Hyland SL, Rätsch G. Real-valued medical time series generation with recurrent conditional GANs. *arXiv [Preprint]*. 2017. Available from: <https://arxiv.org/abs/1706.02633>
- Keskes MI. Review of the current state of deep learning applications in agriculture. *Preprints [Preprint]*. 2025. Available from: <https://doi.org/10.20944/preprints202504.1290.v2>
- Radford A, Metz L, Chintala S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv [Preprint]*. 2016. Available from: <https://arxiv.org/abs/1511.06434>
- Wang D, Qin X, Song F, Cheng L. Stabilizing training of generative adversarial nets via Langevin Stein variational gradient descent. *IEEE Trans Neural Netw Learn Syst*. 2022;33(7):2768–80. Available from: <https://doi.org/10.1109/TNNLS.2020.3045082>
- Goodfellow I. NIPS 2016 tutorial: Generative adversarial networks. *arXiv [Preprint]*. 2016. Available from: <https://arxiv.org/abs/1701.00160>
- Pan Z, Yu W, Yi X, Khan A, Yuan Y, Zheng Y. Recent progress on generative adversarial networks (GANs): A survey. *IEEE Access*. 2019;7:36322–33. Available from: <https://doi.org/10.1109/ACCESS.2019.2905015>
- Singh NK, Raza K. Medical image generation using generative adversarial networks: A review. In: *Studies in Computational Intelligence*. Vol. 937. Springer; 2021;77–96. Available from: https://doi.org/10.1007/978-981-15-9735-0_5
- Armanious K, Jiang C, Fischer M, Küstner T, Hepp T, Nikolaou K, et al. MedGAN: Medical image translation using GANs. *Comput Med Imaging Graph*. 2020;79:101684. Available from: <https://doi.org/10.1016/j.compmedimag.2019.101684>
- Yoon J, Jarrett D, van der Schaar M. Time-series generative adversarial networks. In: *Advances in Neural Information Processing Systems*. 2019;32:5508–18. Available from: https://proceedings.neurips.cc/paper_files/paper/2019/file/c9efe5f26cd17ba6216bbe2a7d26d490-Paper.pdf
- Wiese M, Knobloch R, Korn R, Kretschmer T. Quant GANs: Deep generation of financial time series. *Quant Finance*. 2020;20(9):1419–40. Available from: <https://doi.org/10.1080/14697688.2020.1730426>
- Isola P, Zhu JY, Zhou T, Efros AA. Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE; 2017;1125–34. Available from: <https://doi.org/10.1109/CVPR.2017.632>
- Yu L, Zhang W, Wang J, Yu Y. SeqGAN: Sequence generative adversarial nets with policy gradient. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 2017;31(1). Available from: <https://doi.org/10.1609/aaai.v31i1.10604>
- Lin Y, Wang H, Liu C. IDSGAN: Generative adversarial networks for attack generation against intrusion detection. *IEEE Access*. 2020;8:80086–96. Available from: <https://doi.org/10.1109/ACCESS.2020.2989151>

26. Park N, Mohammadi M, Gorde K, Jajodia S, Park H, Kim Y. Data synthesis based on generative adversarial networks. *Proc VLDB Endow.* 2018;11(10):1071–83. Available from: <https://doi.org/10.14778/3231751.3231757>
27. Zhou J, Han X, Zhang L, Fan Y. CTAB-GAN: Effective table data synthesizing. *arXiv [Preprint]*. 2020 [cited 2025 Jul 25]. Available from: <https://arxiv.org/abs/2010.01906>
28. Zhang H, Xu T, Li H, Zhang S, Wang X, Huang X, et al. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. IEEE; 2017; 5907–15. Available from: <https://doi.org/10.1109/ICCV.2017.629>
29. Karras T, Aila T, Laine S, Lehtinen J. Progressive growing of GANs for improved quality, stability, and variation. *arXiv [Preprint]*. 2018. Available from: <https://arxiv.org/abs/1710.10196>
30. Reed S, Akhtar Z, Yan X, Du L, Chintala S. Generative adversarial text-to-image synthesis. In: *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. 2016;1060–9.
31. Frid-Adar M, Klang E, Amitai M, Goldberger J, Greenspan H. Synthetic data augmentation using GAN for improved liver lesion classification. *Neurocomputing*. 2018;321:321–31. Available from: <https://doi.org/10.1016/j.neucom.2018.09.013>
32. Zhao Z, Kunar A, Birke R, Chen LY. CTAB-GAN: Effective table data synthesizing. *arXiv [Preprint]*. 2021. Available from: <https://arxiv.org/abs/2102.07669>
33. Zhu JY, Park T, Isola P, Efros AA. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. IEEE; 2017;2223–32. Available from: <https://doi.org/10.1109/ICCV.2017.244>
34. Yang Q, Li N, Zhao Z, Fan X, Chang EI, Xu Y. MRI cross-modality image-to-image translation. *Sci Rep*. 2020;10(1):3753. Available from: <https://doi.org/10.1038/s41598-020-60520-6>
35. Majeed A, Hwang SO. Moving conditional GAN close to data: Synthetic tabular data generation and its experimental evaluation. *IEEE Trans Big Data*. 2024; Available from: <https://doi.org/10.1109/TBDATA.2024.3442534>
36. Takahashi S, Chen Y, Tanaka-Ishii K. Modeling financial time-series with generative adversarial networks. *Physica A*. 2019;527:121261. Available from: <https://doi.org/10.1016/j.physa.2019.121261>
37. Guo K, Chen J, Qiu T, Guo S, Luo T, Chen T, et al. MedGAN: An adaptive GAN approach for medical image generation. *Comput Biol Med*. 2023;163:107119. Available from: <https://doi.org/10.1016/j.combiomed.2023.107119>
38. Keskes MI. Review of the current state of deep learning applications in agriculture. *Preprints [Preprint]*. 2025. Available from: <https://doi.org/10.20944/preprints202504.1290.v2>
39. Kingma DP, Welling M. Auto-encoding variational Bayes. *arXiv [Preprint]*. 2014. Available from: <https://arxiv.org/abs/1312.6114>
40. Miletic M, Sariyar M. Assessing the potential of LLMs and GANs as state-of-the-art tabular synthetic data generation methods. In: *Lecture Notes in Computer Science*. Vol. 14975. Springer; 2024;374–89. Available from: https://doi.org/10.1007/978-3-031-69651-0_25
41. Du Y, Luo D, Yan R, Wang X, Liu H, Zhu H, et al. Enhancing job recommendation through LLM-based generative adversarial networks. *Proc AAAI Conf Artif Intell*. 2024;38(8):8363–71. Available from: <https://doi.org/10.1609/aaai.v38i8.28678>
42. Deng M, Chen L. CDGFD: Cross-domain generalization in ethnic fashion design using LLMs and GANs: A symbolic and geometric approach. *IEEE Access*. 2025;13:7192–207. Available from: <https://doi.org/10.1109/ACCESS.2024.3524444>
43. Chao X, Cao J, Yuqin L, Dai Q. Improved training of spectral normalization generative adversarial networks. In: *2020 2nd World Symposium on Artificial Intelligence (WSAI)*. IEEE; 2020;24–8. Available from: <https://doi.org/10.1109/WSAI49636.2020.9143310>
44. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial networks. *Commun ACM*. 2020;63(11):139–44. Available from: <https://doi.org/10.1145/3422622>
45. Hinton GE, Srivastava N, Krizhevsky A, Sutskever I, Salakhutdinov RR. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv [Preprint]*. 2012. Available from: <https://arxiv.org/abs/1207.0580>