

Generative Adversarial Networks (GANs): A Comprehensive Review of Architectures, Training Challenges, and Advancements

Balqis Nabilah Binti Ramanizan

Lecturer, School of AI Computing & Multimedia, Lincoln University College, Malaysia.

E-mail: balqisnabilah@lincoln.edu.my

Abstract

Amongst generative models, GANs have been one of the most powerful families that have exerted an influence in domains such as artificial intelligence and deep learning. Since their discovery by Ian Goodfellow back in 2014, GANs have generated realistic data, from images and videos up to audio, in a wide range of applications. The adversarial training paradigm allows GANs to learn sophisticated data distributions without explicit supervision, including a generator and a discriminator. However, despite those benefits, training GANs is inherently problematic since instability, mode collapse, and issues of convergence naturally pop up. This contribution reviews how GAN architecture has evolved, critically discusses the main challenges regarding GAN training, revisits the most promising developments toward enhancing stability and performance, and addresses rising trends such as diffusion models and hybrid frameworks. Furthermore, this paper points to the directions in which further research should be oriented with a view to the improvement of theoretical grounds, stability, and universality of GANs for practical use.

Keywords: GANs- Generative Adversarial Networks, Generative Models, GAN Architectures, Artificial Intelligence, Deep Learning, Adversarial Training.

1. Introduction

The intuition of generation in a model comes basically because one needs to have an idea about the underlying probability distribution of any dataset so that it generates new data samples similar to the actual ones. GANs, or Generative Adversarial Networks, gave quite a new adversarial solution to the given problem [1]. Figure 1 shows the Conceptual overview of GAN framework depicting the adversarial interaction between Generator (G) and Discriminator (D) as given below.

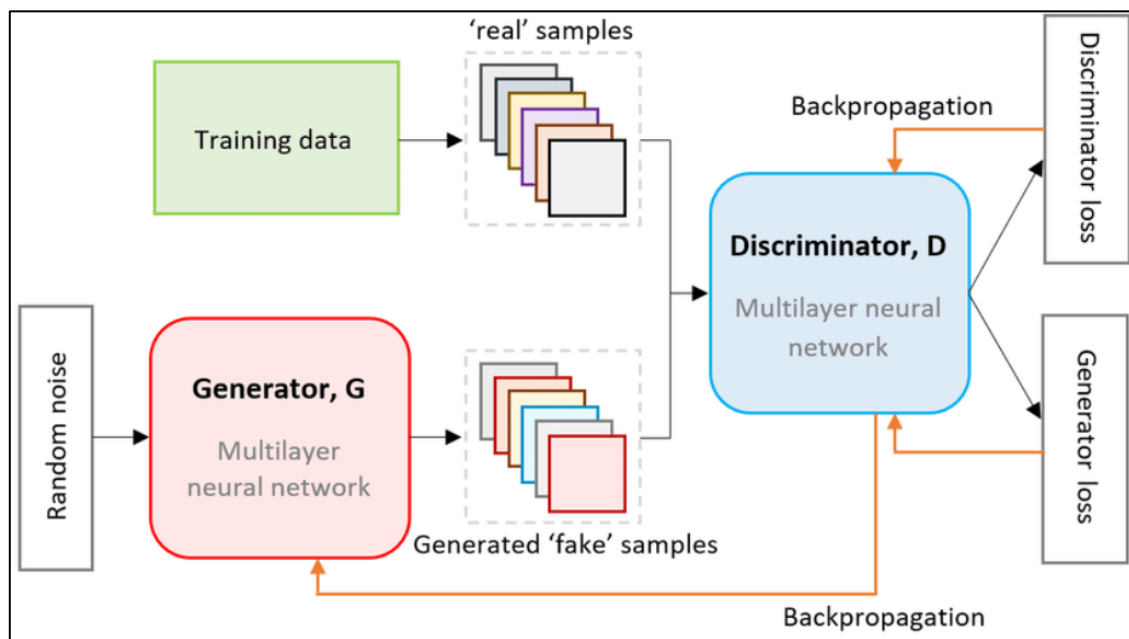


Figure 1. Conceptual Overview of GAN Framework Depicting the Adversarial Interaction Between Generator (G) and Discriminator (D) [16]

The neural networks that would be involved in the proposed framework are one acting like a generator whose work it is to generate data as realistic as possible and another acting like the discriminator whereby it classifies data either real or fake. These normally train in a competitive environment where the objective is for the generator to successfully fool the discriminator while at the same time enabling the discriminator also to improve its discriminative features in finding fakes. Such a contradictory process is like a zero-sum game which results in an interactive learning process, and thus this hopes to yield high-quality samples from the generator [2]. Such adversarial interactions constitute a strong positive feedback loop. The Generator continuously improves in synthesizing more convincing fakes by learning from the successes of the Discriminator at spotting one of its current efforts.

On the other side of the spectrum, the Discriminator builds up his discriminative capability to recognize increasingly difficult fakes concocted by the Generator. While both neural networks keep improving together in synergy, the competition goes on. Training stops when the Generator is good enough and the Discriminator can't tell whether a given data point comes from the real or generated distribution beyond chance. By then, the Generator will have learned an effective model for the underlying probability distribution of the real data, capable of synthesizing new samples that statistically resemble the real dataset. With its elegant adversarial framework and pioneering generative powers, GANs stand right at the very forefront of enabling machines not only to understand but also create and act as agents in the world [3].

2. Literature Review

The study [4] examines the development of GANs, following their progression from Ian Goodfellow's research to their present-day state-of-the-art position. It explores the structure and training dynamics of GANs, emphasizing their use in a range of applications such as text-to-image conversion, style transfer, and image synthesis. Along with discussing issues like mode collapse and training instability, the paper offers a collection of recent GAN research that highlights creative methods in a range of fields.

Review [5] offers a comprehensive examination of the latest developments at this cutting-edge intersection of gene expression data and GANs, particularly from 2019 to 2023. Given the rapid advancement of deep learning technologies, thorough and inclusive evaluations of current procedures are essential for directing future research projects, exchanging knowledge, and spurring the field's ongoing development. By highlighting recent research and important works, this review helps professionals and scholars alike navigate the fascinating intersection of gene expression data systems and GANs.

The review [6] identifies limitations and potential research areas while highlighting the advantages of NAS in enhancing GAN performance, efficiency and stability. Key findings include the necessity of diverse datasets for evaluating GAN performance, the superiority of gradient-based approaches and evolutionary algorithms in specific contexts, and the significance of robust evaluation metrics beyond conventional scores like Fréchet Inception Distance (FID) and Inception Score (IS). This paper attempts to assist researchers in creating

more efficient NAS techniques and furthering the field of GANs by providing an organized comparison of current NAS-GAN approaches.

Survey [7] examine a number of training approaches put forth by various researchers in an effort to stabilize GAN training. It covers the original GAN model and its modified versions, a thorough examination of numerous GAN applications across various domains, and a thorough investigation of the different training challenges and training solutions for GANs. Result of the paper [8] shows that GANs are being used at various scales in the built environment, replacing more traditional approaches in some situations, and breaking new ground in issues that were previously disregarded. They are applied to a wide range of issues and data kinds, such as building design creation, spatiotemporal data privacy protection, vector data generation, and remote sensing data augmentation. The absence of excellent datasets selected especially for built environment issues is a prevalent problem, though. GANs might perform better with more data.

In the work [9] offer a framework for evaluating the relative benefits of GANs over conventional statistical techniques for creating synthetic data, considering the data's residual risk and usefulness. Here, we examine how GANs might be used to create artificial census microdata. We compare the data generated by tabular GANs with those generated by conventional data synthesis techniques using a disclosure risk metric and a set of utility metrics.

While numerous studies have proposed architectural and methodological innovations, few have unified the evaluation under a consistent theoretical framework. A critical analysis reveals that many comparative studies rely on subjective visual assessment rather than quantifiable metrics such as FID, IS, or KID. Moreover, benchmark inconsistencies and dataset bias often obscure the true advantages of one variant over another. This review synthesizes prior findings through both theoretical and empirical perspectives, establishing a cohesive understanding of GAN performance, convergence behavior, and generalization ability across diverse domains.

3. GAN Architecture

The core of a GAN lies in its adversarial, two-player game setup, involving two distinct neural networks

- Generator (G)

- Discriminator (D)

It learns the underlying distribution in real training data and generates synthetic samples that resemble the said distribution. It generally takes some input, often constituted of a random noise vector sampled from a Gaussian or uniform distribution, and transforms that input into a sample of data, say an image. At the beginning of training, the generated samples are usually poor, but a generator continuously refines its output. It is a binary classifier since it takes an input sample and tells whether it's a real sample drawn from the training dataset or a fake one created by the generator. Subsequently, it provides the probability within a range from 0 to 1 that the given sample is real. This is trained in such a way to give high scores for real data and low scores for generated ones. GAN training is essentially iterative and competitive in nature. While the generator tries to deceive the discriminator by generating data more and more like reality, the discriminator works at getting better at recognizing fakes. It is because of this kind of competitive movement that both of them move toward improvement. The objective for the generator will be to generate data that is convincing enough for the discriminator not to be able to tell the difference between it and actual data, that is, the discriminator outputs a probability close to 0.5 for generated samples.




Figure 2. Basic Architecture of GANs [10]

Figure 2 Basic architecture of the GAN model; the noisy input is projected via the generative network to produce synthetic samples, while the response from the discriminator provides feedback to the generator so as to learn better.

The training is guided by a minimax objective function, which can be expressed as:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]$$

Here:

- $D(x)$ is the discriminator's output for a real data sample x drawn from the real data distribution $p_{data}(x)$. The discriminator aims to maximize $\log D(x)$, meaning it wants to correctly identify real samples as real (output close to 1).
- $G(z)$ is the generated sample by the generator G from a random noise vector z sampled from a prior noise distribution $p_z(z)$.
- $D(G(z))$ is the output of discriminator for a fake sample $G(z)$. The aim of discriminator to maximize $\log(1-D(G(z)))$, meaning it wants to correctly identify fake samples as fake (output close to 0, so $1-D(G(z))$ is close to 1).
- The generator G aims to minimize this entire objective function, primarily by minimizing $\log(1-D(G(z)))$. This means the generator wants $D(G(z))$ to be close to 1, effectively fooling the discriminator.

In practice, the generator often maximizes $\log D(G(z))$ rather than minimizing $\log(1-D(G(z)))$ to avoid vanishing gradients in early training. In other words, what this means is that the generator learns from mistakes given by the discriminator, which changes with every different improvement the former makes. It is one big, continuous feedback loop that allows GANs to create both diverse and realistic synthetic data. However, this is an ideal balance that is hard to attain, with problems of non-convexity and instability in gradient flow. Advanced variants include Wasserstein GAN, where Jensen-Shannon divergence is replaced by Earth Mover's Distance for smoother gradients and provable convergence; and Least Squares GAN, which adds a quadratic cost for stabilization. The mathematical refinements reinforce the theoretical underpinning of adversarial training with more predictable convergence.

4. Training Challenges in GANs

GANs have become one of the most powerful generative modeling frameworks that enable synthesis of highly realistic data such as video, audio, and images. Besides their remarkable capabilities, GANs are notoriously hard to train owing to a number of inherent obstacles rooted in their adversarial learning setup; these may cause instability in the convergence and quality of generated outputs [11].




Figure 3. Training Challenges in GANs

Figure 3 summarizes some of the major difficulties in training GANs: non-convergence, mode collapse, vanishing gradients, and evaluation complexity. These are the major problems accompanying adversarial dynamics, and further work is needed in order to refine the loss functions and optimization strategies to mitigate these issues. Further details with respect to major training challenges for GANs are discussed below:

1. Non-Convergence and Instability: The biggest challenge with the training of GANs is to make them converge in a stable manner. Unlike other conventional deep learning models, which minimize a single loss function, GANs involve a two-player minimax game between a discriminator and a generator. The discriminator tries to distinguish between the generated and the real data, while the generator tries to generate data that can fool the discriminator. This adversarial nature often causes oscillations or divergence, rather than converging to a stable solution. Model architecture, initialization, and hyperparameter tuning also make a big difference in training.

2. Mode Collapse: Mode dropping or mode collapse is a phenomenon wherein the generator does not pay any attention to parts of the actual data distribution and generates output in a small range. Instead of generating a diverse set of samples, it focuses on a narrow set of outputs that consistently deceive the discriminator. This leads to poor generalization and limits the usefulness of the GAN, especially for tasks where diversity becomes an important aspect. In order to alleviate the issue, several techniques have

been suggested such as unrolled GANs, mini-batch discrimination, and Wasserstein GANs, though no universally accepted solution has been reached so far.

3. Vanishing Gradients: If, in GAN, the discriminator becomes too strong compared to the generator, it can classify the samples generated as fake with high confidence. Thus, the gradients of loss for the generator vanish since the generator receives very little feedback on how to do better. Therefore, it ultimately fails to learn from meaningful patterns of the data distribution. This is a particularly prevalent problem at the initial stages of training and hence hinders further progress. This issue is mitigated by using other variants of loss functions, such as those in WGANs or Least Squares GANs, which maintain the gradients meaningful throughout the training.

4. Balancing Generator and Discriminator: The most important insight is to keep the balance between the generator and discriminator. One should be careful that neither of the models becomes stronger than the other; otherwise, it will dominate the learning dynamics of the other model. A weak discriminator gives poor feedback, while one which is too powerful leaves no space for the generator to improve. This results in unstable training or suboptimal performance. Common strategies to deal with this balancing act include regular monitoring, dynamic adjustment of learning rates, and switching update frequencies.

5. Difficulty in Evaluation: Evaluation of GANs is not as straightforward as it is in the case of supervised learning models. There does not exist a single broadly accepted metric quantifying the quality and diversity of the generated samples. Most metrics are based on IS-Inception Score, FID-Frechet Inception Distance, and Precision-Recall for Distributions, all having their own limitations. Subjective "visual quality" is hard to quantify, and mode diversity is even more challenging.

6. High Sensitivity to Hyperparameters: The training of GANs is sensitive to specific hyper-parameter settings, optimization algorithms, model architectures, batch sizes, and learning rates. Smaller changes in these parameters may result in drastically different outcomes concerning the convergence and quality of generated outputs. This makes the GANs less robust and requires much experimentation and tuning, which may be computationally expensive and time-consuming.

7. Computational Cost: GAN training is computationally expensive, especially in applications that require synthesizing high-resolution images. This adversarial training mechanism will involve a large number of forward and backward passes through both generator and discriminator, which can be done over several iterations. This makes the time taken to train quite extensive and sometimes out of reach for many researchers.

8. Lack of Theoretical Understanding: Despite their empirical success, GANs still suffer from a lack of thorough theoretical investigation. Though adversarial training in GANs is based on game theory, in practice, the behavior of GANs often does not follow theoretical predictions. The gap makes it difficult to predict or guarantee the outcomes of training and limits the ability to design principled improvements to either the architecture or training procedures.

5. Advancements and Solutions

Though the original GAN architecture developed the foundation, its initial problems, such as unstable training and mode collapse-a scenario where the generator only produces a few varieties of output-led to many variants. Most variants include changes in the architecture, loss functions, or training procedures to circumvent these problems and further extend performance for a particular task at hand [12].

Variants of GANs




Figure 4. GAN Variants

Figure 4 summarizes major GAN variants, including DCGAN, CGAN, WGAN, LSGAN, PGGAN, and StyleGAN. Each of these introduced either an architectural or a mathematical innovation for better training stability, diversity, and quality of output.

DCGAN: The DCGAN is among the earliest and most important variants, setting architectural guidelines for the use of CNNs in both the discriminator and the generator. Key concepts that were involved included the use of batch normalization, the absence

of fully connected hidden layers, and the use of strided (discriminator) and fractional-strided (generator) convolutions instead of pooling; thus, much better training stability and higher-quality generated images were achieved.

CGAN: A CGAN allows targeted data generation, unlike the vanilla GAN that generates data randomly. This is achieved by providing conditioning information, such as class labels, text descriptions, or even another image, to both the generator and discriminator. That way, the GAN can generate specific types of outputs—for example, a cat image if conditioned on the "cat" label, or an image of a shoe based on its sketch.

Wasserstein GAN: WGAN replaces the JSD used in the loss function in the original GAN with the so-called Wasserstein distance or Earth Mover's Distance to solve the problem of training instability and mode collapse. That will provide a more stable and meaningful gradient signal even when the generated distribution is far from the target data distribution; thus, the training will be more robust with less mode collapse. WGANs often use "weight clipping" or "gradient penalty" (WGAN-GP) to enforce a Lipschitz constraint on the discriminator.

A key factor in why variants like Wasserstein GANs perform better than traditional GANs is because of their theoretical formulation. The original GAN minimizes Jensen–Shannon divergence. However, this saturates when distributions do not overlap. WGAN optimizes the Earth Mover's Distance:

$$W(p_r, p_g) = \inf_{\gamma \in \Pi(p_r, p_g)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|]$$

This metric provides continuous gradients even when the generated and real distributions are disjoint, thus providing more stable and meaningful updates. Beyond stability, WGAN's gradient penalty regularization imposes Lipschitz continuity, improving convergence predictability and sample diversity. Hence, its superior performance comes not merely from empirical stability but rather from stronger mathematical grounding.

LSGAN: The LSGAN suggests replacing the sigmoid cross-entropy loss, which is commonly used in traditional GANs, by a least-squares loss function for the discriminator. Concretely, this update tries to provide smoother and non-saturating

gradients, especially for samples far from the decision boundary, which should yield more stable training and higher-quality samples.

Progressive Growing GAN (PGGAN): PGGAN introduced a new training methodology where both the generator and discriminator are grown progressively during training. It started with small resolutions such as a 4×4 -pixel resolution and went up to 1024×1024 pixels by adding layers. This allows networks to learn coarser features first and then refine the finer details, thereby giving much better stability, faster training, and generation of incredibly high-quality images. StyleGAN is an extension of PGGAN. It proposed a mapping network along with AdaIN at every resolution of the generator after projecting the latent code into an intermediate latent space. This disentangles different artistic styles and features, such as pose, identity, hair color, and lighting, while generating realistic and controllable results for human faces. Its variants, StyleGAN2 and StyleGAN3, fix common artifacts and improve quality.

The following Table 1 provides a summary of the advantages associated with various GANs [13]

Table 1. Advantages of GANs

GAN Variant	Advantages
DCGAN	Improved training stability and better image quality.
CGAN	Enables targeted generation based on input conditions.
WGAN	More stable training; reduces mode collapse; works well when real and generated distributions are far apart.
LSGAN	Provides smoother gradients; improves training stability and output quality.
PGGAN	Enables learning from coarse to fine details; generates high-resolution, high-quality images.
StyleGAN	Allows fine-grained control over style attributes; generates highly realistic and customizable images.

6. Applications of GANs

From being a theoretical novelty, GANs came to be one of the most promising and broadly applied innovations within the scope of modern AI. Applications vary from the study of complex data distribution to synthesizing new samples that are strikingly realistic. Besides, its versatility allows it not only to create completely new content but also to transform existing data, enhance a dataset with better model training, and even contribute to solving some big challenges which come with changing society for preserving privacy or being prepared for climate change by generating synthetic data 14. Applications of GANs in various domains are depicted in Table 2 below.

Table 2. Applications of GANs

Category	Specific Application	Description	Examples/Impact
Image Generation & Manipulation	Photorealistic Image Synthesis	Generating highly realistic images from scratch or text descriptions.	Creating faces of non-existent people, realistic landscapes, product mockups.
	Image-to-Image Translation	Transforming an image from one domain or style to another.	Sketch to photo, day to night, summer to winter, style transfer (e.g., CycleGAN).
	Super-Resolution	Enhancing the resolution and detail of low-quality images and videos.	Sharpening old photos, improving clarity of surveillance footage, medical imaging.
	Photo Inpainting/ Completion	Seamlessly filling in missing or damaged parts of an image.	Restoring old photographs, removing unwanted objects.
	Face Manipulation	Altering facial attributes (age, expression, identity), generating frontal views.	Age progression, virtual try-on for makeup/hair, deepfakes (controversial).
	Data Augmentation	Generating synthetic data to expand training datasets for other ML models.	Creating synthetic medical images, sensor data, or rare event data for training.

3D & Video	3D Object Generation	Creating realistic 3D models of objects and scenes.	Game assets, architectural visualization, VR/AR content.
	Video Generation & Prediction	Generating realistic video sequences or predicting future frames.	Film animation, virtual reality, autonomous vehicle simulation.
Healthcare & Medical	Medical Image Enhancement/ Synthesis	Improving quality of medical scans; generating synthetic patient data.	Better diagnostics from MRI/X-ray, training AI models with anonymized data.
	Anomaly Detection	Identifying unusual patterns or abnormalities in medical images.	Early detection of tumors, disease markers.
	Drug Discovery	Creating new molecular structures with the desired characteristics.	Accelerating new drug development, materials science.
Creative Arts & Ent.	Digital Art Generation	Creating unique and diverse artworks in various styles.	AI-generated paintings, abstract art, design concepts.
	Game Development	Automatically generating game assets (textures, characters, environments).	Faster game development, more diverse game worlds.
Cybersecurity	Adversarial Training	Generating "adversarial examples" to make AI models more robust against attacks.	Improving security of facial recognition, spam filters.
	Fraud Detection Data	Creating synthetic fraudulent transaction data for training detection systems.	Enhancing financial fraud detection.
Other Emerging	Text-to-Image Synthesis	Creating images directly from textual descriptions.	Creating visual content from written prompts (e.g., DALL-E, Midjourney principles).
	Audio Synthesis	Generating realistic music, speech, or sound effects.	AI-composed music, synthetic voiceovers, new sound designs.

These new GAN architectures are inexpensive, aiming at the goal of reducing computational and energy requirements without losses in generated quality. For this, several directions have been pursued: lightweight convolutional blocks, network pruning, knowledge distillation, and quantized inference can demonstrate that actual training costs can be reduced indeed by 50-70%. MobileGAN and TinyGAN generators and discriminators can fit into embedded and edge devices, respectively, and execute several synthesis tasks such as image translation and facial animation in real time. A number of recently proposed innovations in this direction democratize GAN research by lowering resource barriers while maintaining competitive visual fidelity. This will be particularly important for decentralized AI and on-device learning.

7. Future Directions

The future of GANs involves overcoming their limitations; hence, extending the domains they can be applied to. The importance here will be on the theoretical grounding for the dynamics of GAN training. A deeper understanding of the grounding in game-theoretic settings and the conditions under which convergence may occur could yield more robust methodologies for training. Other emerging trends are unsupervised and few-shot learning, whereby the GANs need to learn from little or no labeled data or adapt, with only a few examples, to new domains. Similarly, integration with multimodal learning frameworks is also likely to continue to grow, since these provide richer generative capabilities across images, audio, video, and text. Furthermore, hybrid generative models that merge GANs with other paradigms are also emerging, including combinations of VAEs, Normalizing Flows, and Diffusion Models. Such combinations will be needed to leverage the strengths of each for better sample quality, diversity, and interpretability. Ethical use of GANs is a big concern, and future work is especially needed regarding bias mitigation, output explainability, and responsible-use frameworks for preventing misuse in sensitive applications. Therefore, scalability, interpretability, and data efficiency are going to be driving the new generation of GANs.

8. Conclusion

GANs were revolutionary in the field of generative modeling by providing an efficient adversarial training mechanism. While the original GAN architecture was great, most of the stability and scalability regarding quality in generated data were achieved with subsequent

innovations. However, attempts to train GANs remain challenging despite the progress during the last years due to mode collapse, non-convergence, and gradient instability problems. But GANs are pushing the edges of what is possible in AI-from new architectures to better loss functions and further advanced training strategies. It thus concerns anything from image synthesis and medical imaging to finance and digital arts. As new studies along these research courses continue to improve, GANs would be able to play even more significant roles in the future of AI if some challenges about training dynamics and ethical concern issues can be overcome.

References

- [1] Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial networks." *Communications of the ACM* 63, no. 11 (2020): 139-144.
- [2] Creswell, Antonia, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A. Bharath. "Generative adversarial networks: An overview." *IEEE signal processing magazine* 35, no. 1 (2018): 53-65.
- [3] Sharma, Preeti, Manoj Kumar, Hitesh Kumar Sharma, and Soly Mathew Biju. "Generative adversarial networks (GANs): introduction, taxonomy, variants, limitations, and applications." *Multimedia Tools and Applications* (2024): 1-48.
- [4] Pawar, D. R., and Pravin Yannawar. "Advancements and Applications of Generative Adversarial Networks: A Comprehensive Review." *Int. J. Res. Appl. Sci. Eng. Technol* 12 (2024).
- [5] Lee, Minhyeok. "Recent advances in generative adversarial networks for gene expression data: a comprehensive review." *Mathematics* 11, no. 14 (2023): 3055.
- [6] Alotaibi, Abrar, and Moataz Ahmed. "Neural Architecture Search for Generative Adversarial Networks: A Comprehensive Review and Critical Analysis." *Applied Sciences* 15, no. 7 (2025): 3623.
- [7] Jabbar, Abdul, Xi Li, and Bourahla Omar. "A survey on generative adversarial networks: Variants, applications, and training." *ACM Computing Surveys (CSUR)* 54, no. 8 (2021): 1-49.

- [8] Wu, Abraham Noah, Rudi Stouffs, and Filip Biljecki. "Generative Adversarial Networks in the built environment: A comprehensive review of the application of GANs across data types and scales." *Building and Environment* 223 (2022): 109477.
- [9] Little, Claire, Mark Elliot, Richard Allmendinger, and Sahel Shariati Samani. "Generative adversarial networks for synthetic data generation: a comparative study." *arXiv preprint arXiv:2112.01925* (2021).
- [10] Saxena, Divya, and Jiannong Cao. "Generative adversarial networks (GANs) challenges, solutions, and future directions." *ACM Computing Surveys (CSUR)* 54, no. 3 (2021): 1-42.
- [11] Nayak, Ankitha A., P. S. Venugopala, and B. Ashwini. "A Systematic Review on Generative Adversarial Network (GAN): Challenges and Future Directions." *Archives of Computational Methods in Engineering* (2024): 1-34.
- [12] Durgadevi, M. "Generative Adversarial Network (GAN): A general review on different variants of GAN and applications." In *2021 6th International Conference on Communication and Electronics Systems (ICCES)*, pp. 1-8. IEEE, 2021.
- [13] Hong, Yongjun, Uiwon Hwang, Jaeyoon Yoo, and Sungroh Yoon. "How generative adversarial networks and their variants work: An overview." *ACM Computing Surveys (CSUR)* 52, no. 1 (2019): 1-43.
- [14] Alqahtani, Hamed, Manolya Kavakli-Thorne, and Gulshan Kumar. "Applications of generative adversarial networks (gans): An updated review." *Archives of Computational Methods in Engineering* 28 (2021): 525-552.
- [15] Chakraborty, Tanujit, Ujjwal Reddy KS, Shraddha M. Naik, Madhurima Panja, and Bayapureddy Manvitha. "Ten years of generative adversarial nets (GANs): a survey of the state-of-the-art." *Machine Learning: Science and Technology* 5, no. 1 (2024): 011001.
- [16] https://commons.wikimedia.org/wiki/File:GANs_Technique.webp