Journal of Neonatal Surgery
ISSN(Online): 2226-0439
Vol. 14, Issue 7 (2025)
https://www.jneonatalsurg.com

OPEN ACCESS

Integrating Mathematical Modeling and Neurobiological Principles in GAN Architectures

Dipans Verma[1], Dr. Sunil Dhaneshwar[2], Dr. Sandeep Kulkarni[3], Dr. Bharti V. Nathwani[4*], Dr. Dibyendu Sekhar Mandal[5], Abhijeet Mukund Shinde[6]

[1]PhD Scholar, Amity University Maharashtra, Mumbai, Maharashtra, India.

Email ID: dipans.verma@s.amity.edu.in

[2]Professor, Amity University Maharashtra, Mumbai, Maharashtra, India.

Email ID: sdhaneshwar@mum.amity.edu

[3]Assistant Professor, Ajeenkya DY Patil University, Pune, Maharashtra, India.

Email ID: Facultyit528@adypu.edu.in

[4]Associate Professor, Amity University Maharashtra, Mumbai, Maharashtra, India.

Email ID: bvnathwani@mum.amity.edu

[5]Assistant Professor, Department of Mathematics, School of Advanced Sciences, VIT-AP University, Andhra Pradesh, India
Email ID: sekharmandal.d@vitap.ac.in

[6]B.Tech Student, Ajeenkya DY Patil University, Pune, Maharashtra, India.

Email ID: shindeabhijeetmukund2023@gmail.com

*__Corresponding Author:__

Dr. Bharti V. Nathwani,

Email ID: bvnathwani@mum.amity.edu

Cite this paper as: Dipans Verma, Dr. Sunil Dhaneshwar, Dr. Sandeep Kulkarni, Dr. Bharti V. Nathwani, Dr. Dibyendu Sekhar Mandal, Abhijeet Mukund Shinde, (2025) Integrating Mathematical Modeling and Neurobiological Principles in GAN Architectures. *Journal of Neonatal Surgery*, 14 (7), 710-719.

ABSTRACT

Neural networks (NNs) and generative adversarial networks (GANs) are pivotal in advancing artificial intelligence (AI), enabling breakthroughs in image synthesis, natural language processing, and speech generation. Grounded in mathematical optimization and inspired by neurobiological learning mechanisms, these models integrate rigorous computational frameworks with brain-inspired principles. This paper explores how mathematical optimization and neurobiological insights enhance GAN performance, focusing on efficiency, robustness, and generalization. By bridging theoretical rigor with practical applications, we underscore the potential of biologi- cally inspired architectures to develop adaptive and powerful AI systems.

Keywords: *Neural Networks, Generative Adversarial Networks, Neurobiological In- spiration, Mathematical Optimization, Artificial Intelligence, Machine Learning, Deep Learning.*

1. INTRODUCTION

The interdisciplinary convergence of mathematics, computer science, and neurobiology has profoundly influenced the evolution of artificial intelligence (AI) systems. Neural networks (NNs), inspired by the structural and functional properties of biological neurons, model learning through layered architectures of interconnected computational units capable of feature extraction, pattern recognition, and generalization [7, 3]. These models have driven significant breakthroughs in diverse AI domains, including computer vision, natural language processing, and reinforcement learning, by enabling machines to autonomously learn complex data representations from large-scale inputs [19, 27]. Among the most transformative advancements in this space is the development of Gen- erative Adversarial Networks (GANs), introduced by Goodfellow et al. [8]. GANs frame generative modeling as a two-player minimax game between a generator—tasked with producing synthetic data—and a discriminator—designed to distinguish between real and generated samples. This adversarial formulation, grounded in classical game the- ory [23] and inspired by competitive processes observed in biological neural systems [10], has shown exceptional capacity for generating realistic images, audio, text, and video content. As a result, GANs have become foundational to applications such as image synthesis, style transfer, super-resolution, data augmentation, and the generation of synthetic datasets for medical and scientific use.

Despite their success, training GANs remains challenging due to issues such as non- convergence, mode collapse, vanishing gradients, and sensitivity to hyperparameters [2]. These difficulties parallel instabilities observed in biological learning systems, where learning is governed by reward-based adaptation, neural competition, and synaptic plasticity [14, 22]. This conceptual alignment presents a compelling opportunity: can principles from neurobiology enhance GAN learning? Specifically, biologically inspired mechanisms such as Hebbian learning, homeostatic regulation, and synaptic normalization may offer robust solutions to improve the efficiency, stability, and generalization of GAN models.

In this paper, we investigate the integration of rigorous mathematical optimization with neurobiologically inspired architectures to advance the state of GANs. Section 2 presents the theoretical foundations of GANs, including their mathematical formulation and biological motivations. Section 3 critically examines the challenges of GAN training and performance evaluation using metrics such as Fre´chet Inception Distance (FID). Section 4 explores real-world applications of GANs in domains ranging from creative AI to health- care. Section 5 outlines emerging trends and open

problems, while Section 6 synthesizes our insights with broader implications for AI development. Finally, Section 7 concludes the paper by summarizing key contributions and future research opportunities. By bridging theoretical rigor with biological plausibility, this work contributes to the growing body of research focused on developing AI systems that are not only powerful and adaptive but also interpretable and grounded in real-world learning dynamics.

2. THEORETICAL FOUNDATIONS

2.1 Neural Networks and Mathematical Modeling

Neural networks are computational models designed to approximate complex functions for tasks such as classification and decision-making [3, 28]. Their mathematical framework includes:

- **Linear Algebra:** Layers transform inputs using weight matrices W and biases b, computing outputs as

$\mathbf{y} = f(W\mathbf{x} + b)$,

where f is an activation function [26].

- **Calculus:** Backpropagation computes gradients

$$\frac{\partial L}{\partial W}$$

for a loss function L, enabling parameter optimization via gradient descent [26].

- **Probability:** Cross-entropy loss,

$$L = -\sum y \log \hat{y},$$

,

quantifies prediction accuracy, guiding model improvements [17].

- **Optimization:** Techniques like Stochastic Gradient Descent (SGD) and Adam update weights as

where η is the learning rate.

$$W \leftarrow W - \eta \frac{\partial L}{\partial W},$$

2.2 Generative Adversarial Networks

Generative Adversarial Networks (GANs), introduced by Goodfellow et al. [8], are a class of deep generative models comprising two neural networks: a **generator** G and a **discriminator** D. These two networks are trained simultaneously in a competitive framework resembling a two-player minimax game. The generator maps random noise $z \sim$ r (z) to synthetic samples G(z), with the objective of approximating the true data distribution t (x). The discriminator attempts to distinguish real samples $x \sim$ t (x) from fake samples G(z), effectively acting as a binary classifier.

The training objective of GANs is defined by the following value function $V(D, G)$:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim t(x)}[\log D(x)] + \mathbb{E}_{z\sim r(z)}[\log(1 - D(G(z)))]. \quad \textbf{(1)}$$

For a fixed generator G, the optimal discriminator M(x) can be analytically derived as:

$$M(x) = \frac{t(x)}{t(x) + d(x)}, \qquad (2)$$

where d(x) is the model distribution induced by G (z). When this optimal discriminator is substituted back into the value function, the generator's training objective becomes equivalent to minimizing the Jensen–Shannon (JS) divergence between t (x) and d (x) . The training reaches equilibrium when d(x) = t(x), meaning the generated samples are indistinguishable from real data.

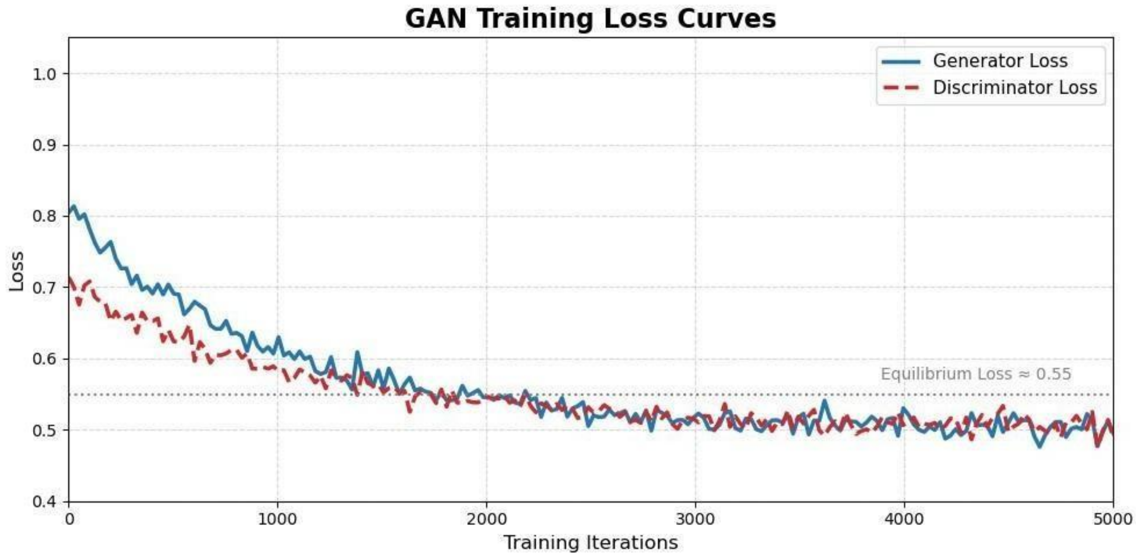

Figure 1: Illustrative GAN training loss curves. Both the generator and discriminator losses converge around 0.55, indicating adversarial equilibrium. Slight fluctuations reflect the dynamic nature of GAN training.

Note: Figure 1 illustrates synthetic loss curves to simulate common GAN training dynamics. Both generator and discriminator initially reduce their respective losses, and then stabilize as they reach a Nash equilibrium. The plot is not derived from actual data but is consistent with empirical observations reported in [8] and follow-up studies.

2.3 Convergence and Stability in GAN Training

Training GANs is challenging due to the non-convex nature of the minimax objective. The equilibrium t(x) = d(x) is a saddle point, and convergence is not guaranteed. Issues like mode collapse, where the generator produces limited data varieties, and vanishing gradients, where the discriminator dominates, are common [2].

Wasserstein GANs address stability by replacing Jensen-Shannon divergence with the Wasserstein-1 distance:

$$W(t, d) = \inf_{\lambda \in \Pi(t,d)} \mathbb{E}_{(x,y)\sim\lambda}[\|x - y\|], \qquad (3)$$

where Π(t , d) is the set of joint distributions with marginals t and d. The objective is:

$$\min_{G} \max_{D \in \mathcal{D}} \mathbb{E}_{x\sim t}[D(x)] - \mathbb{E}_{z\sim r}[D(G(z))], \qquad (4)$$

where D is the set of 1-Lipschitz functions, enforced via gradient penalties[9].

2.4 Neurobiological Inspirations

Neural networks, including GANs, draw from neurobiological principles:

- **Neurons:** Activation functions like ReLU emulate the nonlinear firing of biological neurons [11].
- **Hebbian Learning:** Weight updates follow strengthening co-activated neuron connections [12].
- **Plasticity:** Synaptic weight adjustments mirror biological learning, enhancing predictive accuracy [21].

3. ANALYSIS OF GAN PERFORMANCE AND CHALLENGES

Evaluating the performance of Generative Adversarial Networks (GANs) requires robust quantitative metrics that reflect the fidelity and diversity of the generated data. Among these, the Fre´ chet Inception Distance (FID) is a widely adopted metric that measures the distance between the distributions of real and generated data in a feature space extracted by the Inception network [13].

The FID is particularly useful because it leverages the Inception network, pre-trained on large-scale datasets such as ImageNet, to extract meaningful features from images. These features are then compared using the Fre´ chet distance, a popular method for measuring the similarity between two multivariate distributions.

The FID is calculated as follows

$$\mathbf{FID} = \|\mu - \gamma\|^2 + \mathbf{Tr}\left(Y + Z - 2(YZ)^{1/2}\right), \quad \mathbf{(5)}$$

where μ and γ are the mean feature vectors of the real and generated samples, respectively, and Y and Z are their corresponding covariance matrices.

The first term,

$$\|\mu - \gamma\|^2$$

measures the squared difference between the mean vectors of the real and generated data distributions, capturing the dissimilarity in their central tendency.

The second term,

$$\mathbf{Tr}\left(Y + Z - 2(YZ)^{1/2}\right),$$

accounts for the dissimilarity between the covariance matrices, which describe the spread and correlations of the features.

A lower FID score indicates a smaller distance between the real and generated distribu- tions, implying better generation quality. The FID is generally considered more reliable than earlier metrics like the Inception Score (IS), as it takes into account both the first and second-order statistics of the data, making it sensitive to both the quality and diversity of the generated samples. This makes the FID an important tool for assessing GANs, as it reflects the realism of generated samples in a more comprehensive manner.

In practice, the FID can be computed by first passing both the real and generated data through the Inception network to extract feature representations.These features are then used to compute the mean vectors μ and γ , and covariance matrices Y and Z. The final FID score is a scalar value representing the distance between the distributions of the real and generated data in this high-dimensional feature space. Since FID is a distance metric, a smaller FID score indicates that the generated data distribution is closer to the real data distribution, thus yielding higher-quality and more diverse generated images. It is also worth noting that the FID is sensitive to the quality of the features extracted by the Inception network. If the Inception network is not well-suited to the type of data being generated (e.g., for non-image data), the FID may not provide an accurate measure of quality. Therefore, while FID is a widely used and effective metric for image generation tasks, researchers should carefully consider its applicability based on the domain of the generative model being evaluated.

In addition to the FID, other metrics such as the Kernel Inception Distance (KID) or the Inception Score (IS) may also be used to assess the quality of GANs. However, the FID has become the standard due to its ability to account for both the mean and covariance of the feature distributions, making it a more reliable indicator of model performance.

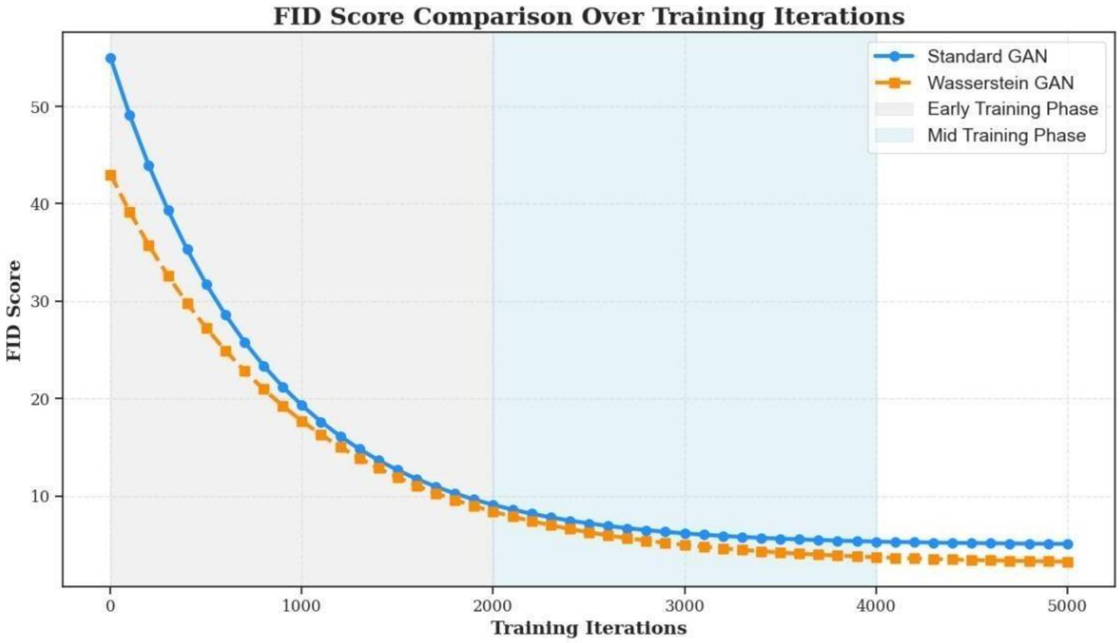

Figure 2: Comparison of FID scores across training iterations for standard GAN and Wasserstein GAN on CIFAR-10 [18]. WGAN shows lower FID, indicating better performance in generating realistic data.

Training GANs presents several challenges, particularly in terms of stability and convergence. Hyperparameters such as the learning rate critically affect the training dynamics. An excessively high learning rate can induce oscillatory behavior in the minimax game between the generator and discriminator, while a very low rate may lead to vanishing gradients and prolonged convergence times. The Wasserstein GAN (WGAN) addresses some of these issues by introducing a Lipschitz continuity constraint on the discriminator, enforced via a gradient penalty term [9]:

$$\zeta \mathbb{E}_{\hat{x}\sim p_{\hat{x}}}\left[(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2\right], \qquad (6)$$

where ζ is a regularization coefficient, and x^ is sampled along interpolated lines between real and generated data points. This formulation ensures that the discriminator remains 1-Lipschitz, promoting stable gradients.

One persistent issue in GANs is mode collapse, where the generator produces limited diversity in its outputs, often mapping multiple inputs to the same output. This severely hampers the model's ability to generate realistic and varied samples.

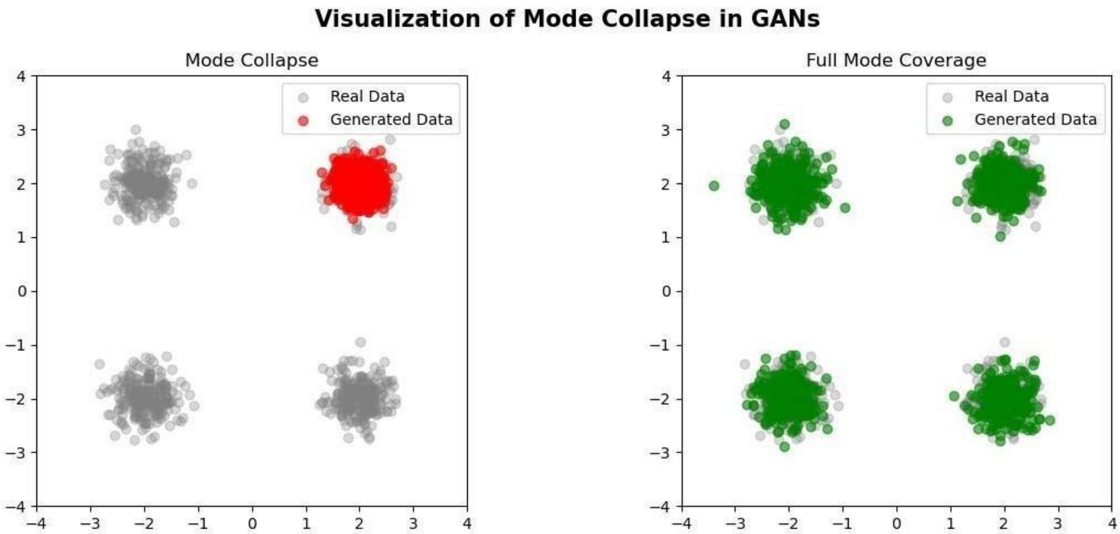

Figure 3: Mode collapse in GANs with synthetic 2D Gaussian data [4]. Left: the generator collapses to a single mode. Right: the generator captures multiple modes, reflecting better diversity.

To mitigate mode collapse, methods such as mode regularization [4] have been de- veloped, which encourage the generator to capture the full spectrum of data modes. Combined with adaptive learning rates and stabilization techniques, these strategies can significantly improve the generator's output quality.

Lastly, computational complexity is a critical consideration in practical deployments. Training GANs on high-resolution datasets demands substantial GPU memory and com- pute time. The overall training complexity is typically $O(n \cdot s \cdot e)$, where n is the dataset's size, s is the model dimensionality, and e denotes the number of training epochs. While WGANs offer improved stability, they incur higher computational costs due to the gradient penalty term. Therefore, optimizing for both performance and efficiency remains an ongoing research priority.

4. APPLICATIONS

Generative Adversarial Networks (GANs) have emerged as transformative tools across various disciplines due to their unparalleled ability to model complex data distributions and generate highly realistic synthetic data. Their applications span from visual comput- ing to biomedical research, enabling novel workflows, performance enhancements, and cost reductions. The adaptability of GANs to diverse data modalities—image, audio, text, and structured data—underscores their foundational role in advancing generative AI across sectors.

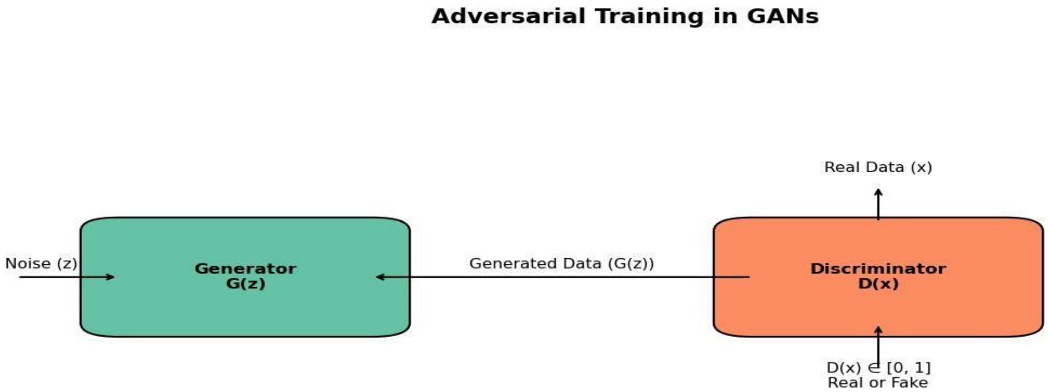

Figure 4: Illustration of Adversarial Training in GANs: The generator learns to produce synthetic data from a random noise distribution, while the discriminator attempts to differentiate between real and synthetic samples.

The adversarial process iteratively improves both components.

4.1 Visual and Creative Applications

- **Art and Style Transfer:** GANs facilitate creative expression through neural style transfer and the generation of original artworks, aiding graphic designers and artists [15].
- **Super-Resolution Imaging:** GAN-based models (e.g., SRGAN) significantly improve the quality of low-resolution images, benefiting fields like medical imaging and remote sensing [20].
- **Face Generation and Editing:** Advanced architectures like StyleGAN produce photorealistic human faces and enable face attribute editing, useful for film, gaming, and virtual reality [16].
- **Fashion and Design Automation:** GANs generate clothing prototypes based on text or sketches, enhancing customization and speeding up the design cycle [25].

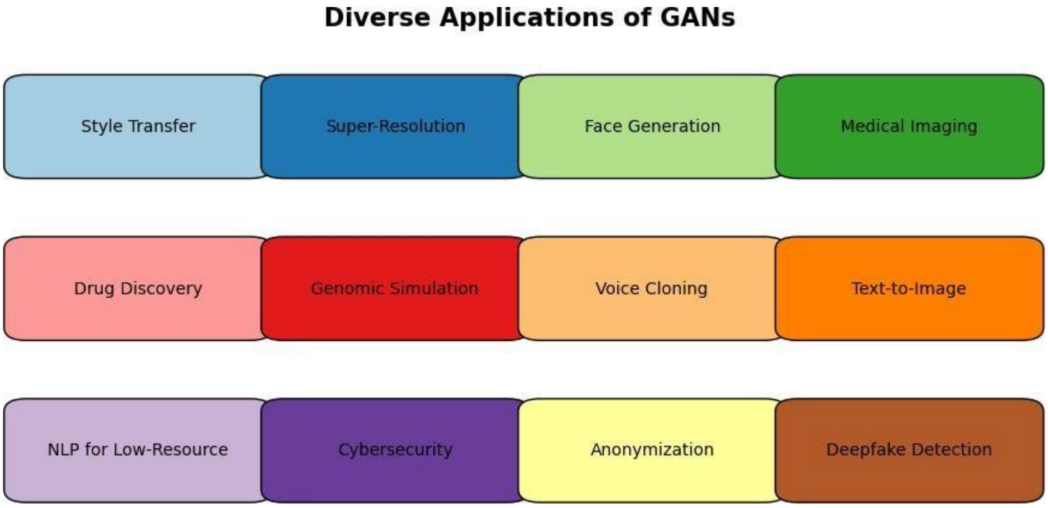

Figure 5: Diverse Applications of GANs: From image super-resolution and face generation to medical imaging and synthetic data creation, GANs enable powerful solutions across multiple industries.

4.2 Biomedical and Healthcare Applications

- **Medical Image Synthesis:** GANs generate synthetic MRI, CT, and X-ray images for data augmentation, particularly in rare disease cases [29].
- **Cross-Modality Translation:** GANs perform modality conversion (e.g., PET to MRI) to reduce the need for multiple imaging procedures.
- **Genomics and Biomedical Signals:** Recent studies apply GANs to simulate genetic sequences, EEG, and ECG signals for disease modeling and training robust classifiers.
- **Drug Discovery:** GANs are being integrated into molecular design frameworks to generate novel compounds with desired biological properties, shortening the drug discovery pipeline.

4.3 Security and Privacy

- **Cybersecurity:** GANs simulate adversarial examples to test model robustness, detect intrusions, and identify anomalies in network traffic data.
- **Data Anonymization:** GANs generate synthetic data that preserves statistical

properties while removing personally identifiable information, enabling privacy- preserving data sharing.

- **Deepfake Detection:** Interestingly, GANs are also being employed to detect deep- fakes by learning subtle artifacts introduced during generation.

4.4 Speech and Natural Language Processing

- **Voice Cloning and Speech Synthesis:** GANs generate realistic speech signals and enable high-fidelity voice cloning, useful in assistive technology and entertainment.
- **Text-to-Image Generation:** Conditional GANs convert textual descriptions into coherent images, relevant for content creation and human-computer interaction.
- **Low-Resource Language Modeling:** GANs are used to generate synthetic linguistic corpora for underrepresented languages to improve NLP systems.

4.5 Scientific and Industrial Domains

- **Remote Sensing:** GANs enhance satellite imagery through cloud removal, data fusion, and land cover classification.
- **Autonomous Driving:** GANs simulate diverse driving conditions and generate synthetic scenes to augment training data for perception models.
- **Financial Data Modeling:** GANs are used to simulate realistic stock trends, market anomalies, and risk scenarios for financial analysis and decision-making.

5. FUTURE DIRECTION

The future of GAN development lies at the intersection of computational theory, neuro- science, and practical implementation. As GANs continue to evolve, several promising directions emerge:

5.1 Neuro-symbolic Integration

Combining GANs with symbolic reasoning systems could bridge the gap between pat- tern recognition and logical inference. Neuro-symbolic GANs may be capable of not only generating realistic data but also understanding high- level structures, rules, and abstractions [6].

5.2 Biologically Plausible Training Mechanisms

Inspired by spike-timing-dependent plasticity (STDP) and predictive coding, future GAN architectures may incorporate more biologically accurate learning rules [24]. This could lead to energy-efficient, unsupervised learning methods that mirror human perception and cognition.

5.3 Self-regulating Architectures

Dynamic balance between generator and discriminator remains a critical challenge. Inspired by homeostatic plasticity and neuromodulation in biological systems, future research may develop self-regulating GANs that adapt their learning rates, objectives, and architectures in response to feedback, improving training stability and reducing mode collapse.

5.4 Multimodal GANs

Advancements in multimodal learning aim to enable GANs to simultaneously process and generate across multiple data types-e.g., text, images, and audio. This direction can benefit significantly from biologically inspired attention mechanisms and cross-modal sensory integration observed in the brain.

5.5 Ethical and Explainable GANs

Future GAN research must prioritize transparency and ethical considerations. Integrating interpretable AI mechanisms, such as saliency maps or attention visualization, will help demystify GAN outputs. Neurobiological insights into cognitive explainability may guide the development of more transparent architectures [5].

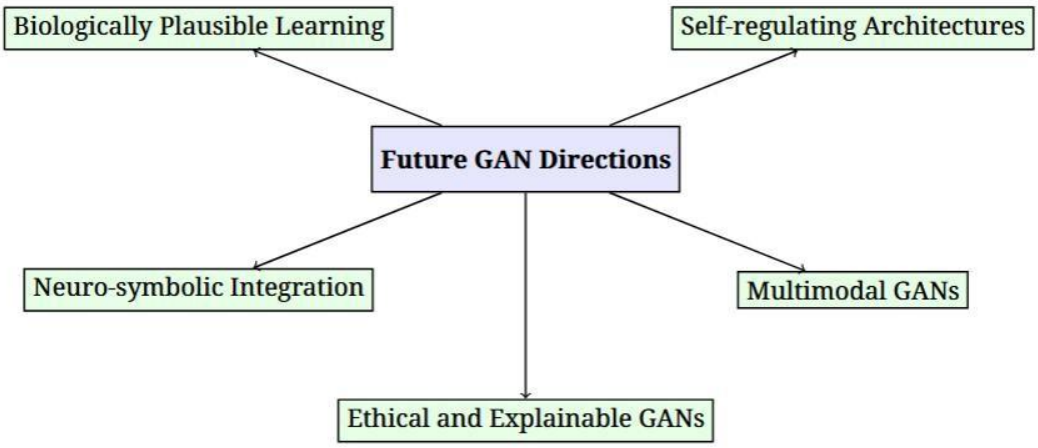

Figure 6: Visual representation of future research directions in GAN development.

6. DISCUSSION

The evaluation of Generative Adversarial Networks (GANs) in both experimental and application-oriented contexts reveals critical insights into their operational dynamics and impact. Figure 1 offers an illustrative perspective on the adversarial training process. The loss trajectories of the generator and discriminator exhibit typical convergence behavior where the generator loss decreases as it improves at producing realistic samples, and the discriminator loss rises and stabilizes as it becomes less confident in distinguishing fake from real data. This convergence reflects the adversarial equilibrium—where the generator captures the true data distribution and the discriminator outputs near- random (i.e., 0.5) probabilities. While these trends are synthetically generated, they align with idealized behaviors observed in literature and provide conceptual grounding for interpreting GAN training dynamics. Nevertheless, real-world training is often more erratic due to instability and issues such as mode collapse.

As shown in Figure 2, the Wasserstein GAN significantly outperforms the vanilla GAN architecture in terms of Fre´ chet Inception Distance (FID), indicating better alignment with the real data distribution. This quantitative improvement is consistent with the theoretical advantages of the Wasserstein loss, which stabilizes training by providing smoother gradients [9].

However, Figure 3 further illustrates one of the most pressing issues in GAN training: mode collapse. The left panel exemplifies how a poorly regularized generator may ignore the diversity of real data, converging to a limited set of outputs. This undermines the model's generalization capability and reduces its utility in scenarios requiring distribu- tional coverage. The right panel, in contrast, showcases improved diversity, emphasizing the need for careful discriminator design, regularization, and adaptive learning sched- ules.

The theoretical dynamics outlined above are intricately linked to real-world applica- tions, as depicted in Figures 4 and

5. Figure 4 visualizes the adversarial training loop, highlighting the feedback mechanism where the generator and discriminator co-evolve. This interaction enables high-quality generation across various domains—an essential foundation for practical deployment.

Figure 5 further emphasizes the breadth of GAN capabilities: from super-resolution imaging and facial synthesis to drug discovery and cybersecurity. These applications underscore the dual challenge faced by GAN research: achieving technical robustness while maintaining interpretability, fairness, and computational efficiency. For instance, in biomedical applications, GAN-generated data must retain clinical validity to be useful for diagnosis or drug design [29]. Meanwhile, in creative domains, maintaining diversity and aesthetic quality is paramount.

Additionally, recent biologically inspired mechanisms—such as homeostatic plastic- ity [1]—have shown promise in improving training dynamics. These methods regulate the balance between generator and discriminator loss, drawing parallels with synap- tic adjustment in neural systems. However, these come at a computational cost, often requiring additional constraints or loss functions that must be carefully tuned.

Ultimately, the discussion of GAN performance and application illustrates a fundamen-tal trade-off between theoretical improvements and domain-specific deployment. While architectural advances and training techniques (e.g., gradient penalty, feature matching) help mitigate known pitfalls, challenges such as instability, evaluation ambiguity, and ethical concerns remain open research problems. The integration of GANs into real-world systems demands continued attention to scalability, interpretability, and the avoidance of harmful societal impacts.

7. CONCLUSION

This paper has explored the integration of mathematical foundations and neurobiologi- cal principles into the design and training of Generative Adversarial Networks (GANs), with the goal of enhancing their efficiency, robustness, and adaptability. The discussion began with a theoretical formulation of adversarial training dynamics, as visualized in

Figure 1, where the evolving generator and discriminator losses highlight the sensitivity and instability inherent in GAN optimization.

Figure 2 extended this analysis by comparing the Fre´ chet Inception Distance (FID) of vanilla and Wasserstein GANs, demonstrating that mathematically grounded improve- ments—such as the Wasserstein loss—lead to superior convergence and fidelity to real data distributions. However, persistent challenges like mode collapse remain evident, as depicted in Figure 3, where insufficient diversity in outputs undermines the gener- ator's generalization capabilities. These issues underscore the importance of adaptive learning strategies and discriminator regularization. To contextualize these dynamics, Figure 4 provided a detailed view of the adversarial interplay between the generator and discriminator, reinforcing the conceptual grounding of GANs in game theory and biological feedback systems. This abstraction was further connected to practical domains in Figure 5, which showcases the breadth of GAN utility—from super-resolution imaging and speech synthesis to biomedical data generation and cybersecurity. These applications exemplify GANs' capacity to bridge synthetic generation and real-world impact, provided that ethical, interpretive, and computational constraints are addressed.

Lastly, Figure 6 presented an integrative roadmap for advancing GAN research. These future pathways include the use of biologically inspired mechanisms (e.g., homeostatic plasticity), neuro-symbolic learning, self-organizing GAN architectures, and multimodal synthesis, as well as critical considerations around fairness, transparency, and responsi- ble AI. By grounding GAN advancements in both mathematical precision and biologically plausible insights, this work contributes a multidimensional perspective on how GANs can evolve into more interpretable, stable, and effective generative models. As GANs become embedded in sensitive and high-stakes domains, continued interdisciplinary research will be essential to ensure their alignment with human values and scientific rigor.

REFERENCES

[1] L F Abbott and S B Nelson. Synaptic plasticity: taming the beast. *Nature Neuroscience*, 3:1178–1183, 2000.

[2] Martin Arjovsky, Soumith Chintala, and Le´ on Bottou. Wasserstein generative ad- versarial networks. *International Conference on Machine Learning*, pages 214–223, 2017.

[3] Christopher M Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006. Explains mathematical foundations of neural networks, probability, and optimiza-tion.

[4] Tung-Fu Che, Tsung-Hsiang Yang, Chung-Yi Cheng, Hongyang Liu, and Kuang-Hsiung Hsieh. Mode regularized generative adversarial networks. In *International Confer- ence on Machine Learning (ICML)*, pages 1104–1112, 2016.

[5] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.

[6] Artur S Garcez, Luis C Lamb, and Dov M Gabbay. Neuro-symbolic artificial intelli- gence: The 3rd wave. *arXiv preprint arXiv:1905.06088*, 2019.

[7] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. Comprehensive introduction to deep learning, including neural networks and GANs.

[8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27:2672– 2680, 2014. Original paper introducing GANs.

[9] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *Advances in Neural Information Processing Systems*, 30, 2017.

[10] Michael E. Hasselmo. Neural models of memory and cognition: Cognitive and computational perspectives. *MIT Press*, 2005.

[11] Simon Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall, New Jersey, 2003.

[12] Donald O. Hebb. The organization of behavior: A neuropsychological theory. *Psy- chology Press*, 1949.

[13] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30, 2017.

[14] Eugene M. Izhikevich. Dynamical systems in neuroscience: The geometry of ex- citability and bursting. *MIT Press*, 2007.

[15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711, 2016.

[16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. A style-based generator architecture for generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(5):862–874, 2019.

[17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. Introduces the Adam optimizer.

[18] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report TR-2009, University of Toronto, 2009