

Generative Artificial Intelligence in Biology and Medicine

Philip de Melo 

Department of Nursing and Allied Health, Norfolk State University, Norfolk, USA

Email: ferndemelo@gmail.com

How to cite this paper: de Melo, P. (2026) Generative Artificial Intelligence in Biology and Medicine. *Advances in Bioscience and Biotechnology*, 17, 185-198.
<https://doi.org/10.4236/abb.2026.175013>

Received: April 20, 2026

Accepted: May 26, 2026

Published: May 29, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Generative Artificial Intelligence in biology and medicine is emerging as a powerful approach for understanding complex biological systems and improving healthcare outcomes. This paper examines the role of generative models such as generative adversarial networks, variational autoencoders, and large language models in applications including drug discovery, protein structure prediction, medical imaging, and personalized medicine. In addition, the concept of PM GenAI, referring to predictive and personalized medicine powered by generative artificial intelligence, is explored as a key advancement that integrates patient specific data with generative modeling to support individualized diagnosis and treatment strategies. Generative AI systems are capable of learning from large scale biological and clinical datasets to produce realistic molecular structures, simulate biological processes, and enhance clinical decision making. In drug discovery, these models accelerate the identification of candidate compounds with desired therapeutic properties, significantly reducing development time and cost. In medical imaging, generative techniques improve image quality, enable robust data augmentation, and support early and accurate disease detection. Furthermore, applications in genomics and synthetic biology allow for the design of DNA sequences and prediction of functional biological outcomes. The paper also addresses important challenges including data bias, limited interpretability, ethical considerations, and regulatory constraints. The integration of GenAI introduces additional concerns related to data privacy and fairness in personalized healthcare systems. Addressing these challenges requires rigorous validation, transparent model design, and interdisciplinary collaboration between computational scientists, biologists, and clinicians. Overall, generative artificial intelligence and present significant opportunities to transform biological research and medical practice while necessitating careful oversight to ensure safe and equitable use.

Keywords

Generative Artificial Intelligence, PM GenAI, Predictive Medicine, Personalized Medicine, Drug Discovery, Protein Structure Prediction, Medical Imaging, Genomics, Synthetic Biology, Deep Learning, Generative Models, Clinical Decision Support, Healthcare Innovation, Bioinformatics, Ethical Considerations

1. Introduction

Recent advances in computational science and data availability have accelerated the integration of artificial intelligence into biology and medicine. Among these developments, generative artificial intelligence has emerged as a particularly transformative paradigm. Unlike traditional analytical models that focus on prediction or classification, generative models are designed to learn underlying data distributions and create new, realistic biological and medical data. This capability opens new possibilities for addressing complex challenges such as understanding molecular interactions, designing therapeutic compounds, and improving diagnostic systems.

In biological research, the increasing volume of high dimensional data generated through genomics, proteomics, and imaging technologies has created both opportunities and challenges. Generative artificial intelligence provides tools to extract meaningful patterns from these datasets while also enabling the simulation of biological processes that are difficult to observe experimentally. Techniques such as generative adversarial networks, variational autoencoders, and large language models have demonstrated significant potential in modeling biological sequences, predicting protein structures, and generating novel molecular designs with desired properties.

In the medical domain, generative artificial intelligence is reshaping clinical practice by enhancing diagnostic accuracy and supporting decision making. Applications in medical imaging allow for improved image reconstruction, noise reduction, and data augmentation, which are essential for training robust diagnostic models. Furthermore, generative systems are increasingly used to synthesize patient specific data, enabling more precise and individualized treatment strategies.

A key emerging concept in this context is a developed by de Melo and St. Rose the PM GenAI algorithm, which refers to the integration of generative artificial intelligence into predictive and personalized medicine. PM GenAI leverages patient specific clinical, genetic, and lifestyle data to generate tailored insights that support individualized healthcare. This approach has the potential to improve disease prevention, optimize treatment selection, and enhance patient outcomes by moving beyond one size fits all medical models.

Despite its promise, the adoption of generative artificial intelligence in biology and medicine is accompanied by significant challenges. Issues related to data qual-

ity, model interpretability, ethical considerations, and regulatory compliance must be carefully addressed. In particular, the use of sensitive patient data in PM GenAI systems raises concerns about privacy, fairness, and accountability. Ensuring the safe and responsible deployment of these technologies requires collaboration across disciplines, as well as the development of transparent and robust validation frameworks.

Deep learning has rapidly transformed multiple scientific domains, particularly medicine and drug discovery. Advances in neural network architectures, generative models, and large-scale data processing have enabled breakthroughs ranging from disease diagnosis to molecular design. Foundational works such as Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) established the groundwork for modern generative AI, while more recent innovations have expanded these capabilities into real-world biomedical applications [1] [2].

The introduction of GANs marked a major milestone in generative modeling, enabling systems to produce realistic synthetic data through adversarial training [1]. Similarly, VAEs provided a probabilistic framework for learning latent representations of data [2].

Large-scale language models later demonstrated the effectiveness of transformer architectures in learning generalizable representations from massive datasets, enabling few-shot learning capabilities [3].

Deep learning has shown remarkable success in medical imaging tasks. Convolutional neural networks (CNNs) can achieve dermatologist-level accuracy in skin cancer classification [4]. Additionally, GAN-based approaches have been used to augment medical datasets, improving model performance in liver lesion classification tasks [5].

Deep learning has transformed drug discovery by accelerating multiple stages of development. Advanced models have been used for target identification, compound screening, and optimization, significantly improving efficiency [6].

Generative models have enabled the design of novel molecules with desired chemical properties, reducing reliance on traditional experimental approaches [7]. Neural networks combined with symbolic AI have also been used to plan chemical synthesis pathways, automating complex decision-making processes [8].

An AI-driven drug discovery approach identified potent DDR1 kinase inhibitors in a significantly reduced timeframe, showcasing the real-world impact of these techniques [9].

These approaches help address limited labeled data, although challenges related to interpretability and clinical integration remain.

One of the most significant achievements in computational biology is highly accurate protein structure prediction using deep learning [10]. Deep learning is also applied in genomics, proteomics, and systems biology, though challenges such as data limitations and reproducibility remain [11] [12].

Generative AI is increasingly applied beyond imaging and drug discovery. Text-to-image models can generate realistic images from textual descriptions [13],

while recent studies explore generative AI for disease classification, including diabetes prediction [14].

AI is expected to play a central role in augmenting clinical decision-making and improving patient outcomes [15] [16].

Despite significant progress, several challenges remain:

- Data limitations: Scarcity of high-quality labeled datasets.
- Interpretability: Models often function as “black boxes”.
- Ethical concerns: Bias, privacy, and accountability.
- Regulatory barriers: Need for rigorous clinical validation.

These challenges are widely recognized as barriers to adoption in healthcare [12].

The literature demonstrates that deep learning and generative AI have fundamentally reshaped medicine and drug discovery. From foundational generative models to advanced biological prediction systems, AI technologies are enabling unprecedented capabilities.

However, translating these advancements into clinical practice requires addressing interpretability, data quality, and ethical concerns. Future research should focus on integrating AI systems into healthcare while ensuring transparency and reliability.

2. Generative Artificial Intelligence

2.1. Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) are a class of deep learning models that have significantly advanced the field of generative artificial intelligence. Introduced by Ian Goodfellow in 2014, GANs are designed to generate new data samples that closely resemble real data. Their ability to model complex distributions makes them highly relevant for biological and medical applications, where data is often high dimensional, limited, and expensive to acquire.

In the context of biology and medicine, GANs have been applied to medical imaging, drug discovery, genomics, and personalized healthcare. They also play an important role in PM GenAI introduced by Philip de Melo by enabling the generation of patient specific data and supporting predictive and personalized treatment strategies.

GANs are based on a game theoretic framework involving two neural networks trained simultaneously:

- A generator G , which learns to produce synthetic data from random noise
- A discriminator D , which learns to distinguish between real and generated data

The objective of GAN training is commonly expressed as a minimax optimization problem:

$$\min_G \max_D E_{x \sim p_{data}(x)} [\log D(x)] + E_{z \sim p_z(z)} [\log(1 - D(G(z)))] \quad (1)$$

In this formulation, the generator attempts to minimize the objective while the

discriminator aims to maximize it. The training process continues until an equilibrium is reached, where the generated data becomes indistinguishable from real data.

Over time, several variants of GANs have been developed to improve performance and stability:

- Deep Convolutional GANs (DCGANs) use convolutional layers to generate high quality images, particularly useful in medical imaging.
- Conditional GANs (cGANs) incorporate additional information such as class labels or patient data to guide generation.
- Wasserstein GANs (WGANs) improve training stability by using a different loss function.
- CycleGANs enable image to image translation, such as converting MRI scans into CT images.

These variants expand the applicability of GANs across diverse biomedical tasks.

2.2. Applications in Biology and Medicine

GANs are widely used in medical imaging for image synthesis, enhancement, and reconstruction. They can generate realistic medical images, improve resolution, and reduce noise, thereby supporting more accurate diagnosis.

Figure 1 shows schematic representation of a Generative Adversarial Network (GAN) architecture illustrating the interaction between the generator and discriminator. The generator transforms random noise into synthetic images that resemble real medical images, while real images from the dataset are simultaneously provided to the discriminator. The discriminator evaluates both real and generated images and learns to distinguish between them, classifying each input as real or fake. Through this adversarial training process, the generator progressively improves its ability to produce highly realistic images, making GANs particularly effective for applications such as medical image synthesis, data augmentation, and diagnostic support in biomedical research.




Figure 1. Architecture of a Generative Adversarial Network (GAN) showing the interaction between the generator and discriminator, where real images from the dataset and synthetic images generated from random noise are evaluated to distinguish between real and fake data.

GANs have shown great potential in drug discovery by generating novel chemical compounds with desired properties. These models explore large chemical spaces efficiently, enabling the identification of candidate molecules for therapeutic development. This approach reduces both the time and cost associated with traditional drug discovery processes.

In genomics, GANs are used to generate DNA sequences and predict gene expression patterns. They assist in understanding regulatory mechanisms and enable the design of synthetic biological systems. These capabilities are important for advancing research in genetics and biotechnology.

GANs contribute to PM GenAI by enabling the generation of personalized biological and clinical data. They can simulate patient specific disease progression and predict treatment outcomes. This supports the development of individualized therapies and enhances decision making in clinical practice.

GANs offer several advantages in biological and medical applications:

- Ability to generate realistic and high-quality data.
- Effective in data scarce environments.
- Support for data augmentation and simulation.
- Capability to model complex biological systems.
- Contribution to personalized and predictive medicine.

Despite their strengths, GANs present several challenges:

- Training instability and convergence difficulties.
- Mode collapse, where diversity of generated data is limited.
- High computational requirements.
- Limited interpretability, which is critical in healthcare settings.
- Ethical concerns related to synthetic data and patient privacy.

These issues must be addressed to ensure safe and reliable use in medical applications.

Future research is focused on improving GAN stability, interpretability, and integration with other AI models. Combining GANs with large language models and multimodal systems may further enhance their capabilities. In PM GenAI, GANs are expected to play a central role in generating personalized insights and advancing precision medicine.

Figure 2 illustrates visualization of a Generative Adversarial Network (GAN) output showing the distribution of real data (blue) and generated data (red). The plot illustrates how the generator learns to approximate the underlying data distribution by producing synthetic samples that closely overlap with real data points. The similarity in spatial distribution and density between the two sets indicates effective adversarial training, where the discriminator's feedback enables the generator to progressively improve its output quality. This demonstrates the capability of GANs to replicate complex data patterns, making them valuable for applications such as data augmentation, simulation of biological datasets, and enhancement of machine learning models in biomedical research.

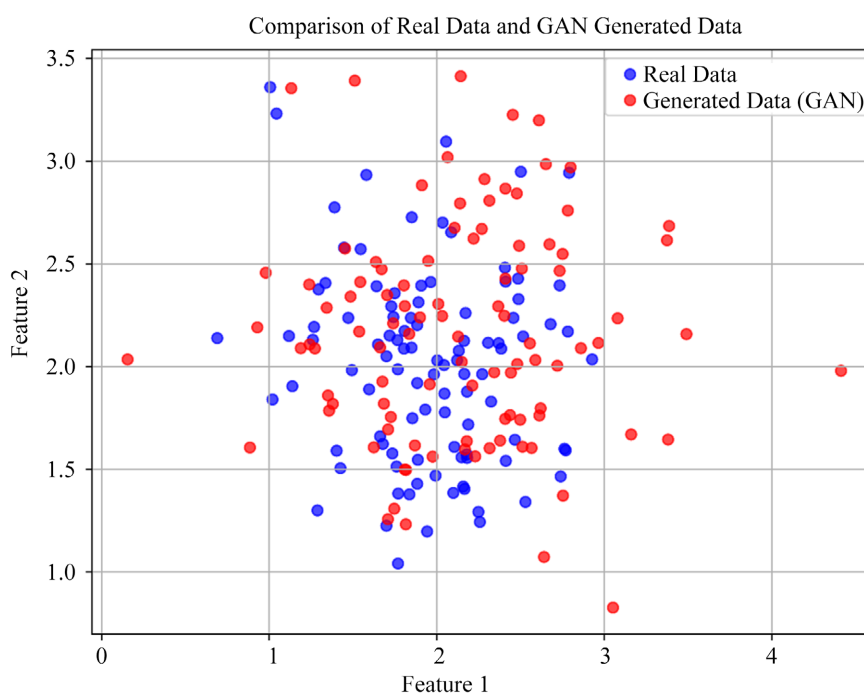


Figure 2. Scatter plot illustrating the comparison between real data (blue) and GAN generated data (red). The figure demonstrates how the generative adversarial network learns the underlying data distribution, producing synthetic samples that closely resemble the real data points.

2.3. Principal Model Generative Artificial Intelligence

Predictive and Personalized Medicine powered by Generative Artificial Intelligence, referred to as PM GenAI, represents a significant advancement in modern healthcare. The concept, introduced by Philip de Melo, integrates generative artificial intelligence with predictive and personalized medicine to enable more accurate diagnosis, individualized treatment, and improved patient outcomes. Unlike traditional healthcare models that rely on population based approaches, PM GenAI focuses on tailoring medical decisions to the unique biological, genetic, and environmental characteristics of each patient.

At its core, PM GenAI leverages advanced generative models such as Generative Adversarial Networks, variational autoencoders, and large language models to learn complex patterns from large scale biomedical datasets. These models are capable of generating realistic simulations of biological processes, predicting disease progression, and creating synthetic patient data that reflects individual variability. By combining genomic data, electronic health records, medical imaging, and lifestyle information, PM GenAI provides a comprehensive and personalized understanding of patient health.

One of the primary applications of PM GenAI is in early disease prediction and diagnosis. By analyzing patient specific data, generative models can identify subtle patterns and risk factors that may not be detectable through conventional methods. This enables early detection of diseases such as cancer, cardiovascular disorder-

ders, and neurological conditions, leading to more timely and effective interventions. Furthermore, these models can simulate different disease trajectories, offering valuable insights into how conditions may evolve in individual patients.

PM GenAI also plays a critical role in personalized treatment planning. Generative models can predict individual responses to different therapies, allowing clinicians to select the most effective treatment while minimizing adverse effects. In drug discovery, PM GenAI facilitates the design of targeted therapeutic compounds tailored to specific patient profiles, thereby advancing precision medicine. This reduces reliance on generalized treatment protocols and improves overall healthcare efficiency.

In addition, PM GenAI contributes to advancements in medical imaging and genomics. Generative techniques can enhance image quality, generate missing data, and simulate patient specific imaging scenarios, improving diagnostic accuracy. In genomics and synthetic biology, these models enable the prediction of gene expression patterns and the design of personalized genetic interventions.

Despite its transformative potential, PM GenAI presents several challenges. The use of sensitive patient data raises concerns regarding privacy, security, and ethical responsibility. Ensuring fairness and minimizing bias in AI models is essential, particularly when applied to diverse populations. Moreover, the complexity of generative models can limit interpretability, which is a critical requirement in clinical decision making.

In conclusion, PM GenAI represents a paradigm shift toward more precise, data driven, and patient centered healthcare. By combining predictive analytics with generative modeling, it enables personalized medical solutions that have the potential to significantly improve patient outcomes. However, its successful implementation requires robust validation, ethical governance, and interdisciplinary collaboration across medicine, biology, and artificial intelligence.

Figure 3 depicts the comparison of real and synthetically generated datasets in a two-dimensional feature space. The scatter plot displays three distributions: real data (blue), data generated using Generative Adversarial Networks (GANs) (red), and data generated using a Probabilistic Model-based Generative AI approach (PM GenAI) (green). Each point represents an individual sample characterized by *Feature 1* (x -axis) and *Feature 2* (y -axis).

The real data form a moderately concentrated cluster centered approximately around (2.0, 2.0), reflecting the underlying true distribution. The GAN-generated data exhibit a wider spread with more extreme values, indicating higher variance and a tendency to produce outliers, which suggests that while GANs capture the general structure, they may overestimate distributional boundaries. In contrast, the PM GenAI-generated data show a distribution more tightly aligned with the real data cluster, with reduced dispersion and fewer outliers, indicating improved stability and closer approximation of the original data distribution.

Overall, the visualization highlights differences in generative performance: GANs provide diversity but with increased variability, whereas PM GenAI demonstrates better consistency and fidelity to the real data distribution.

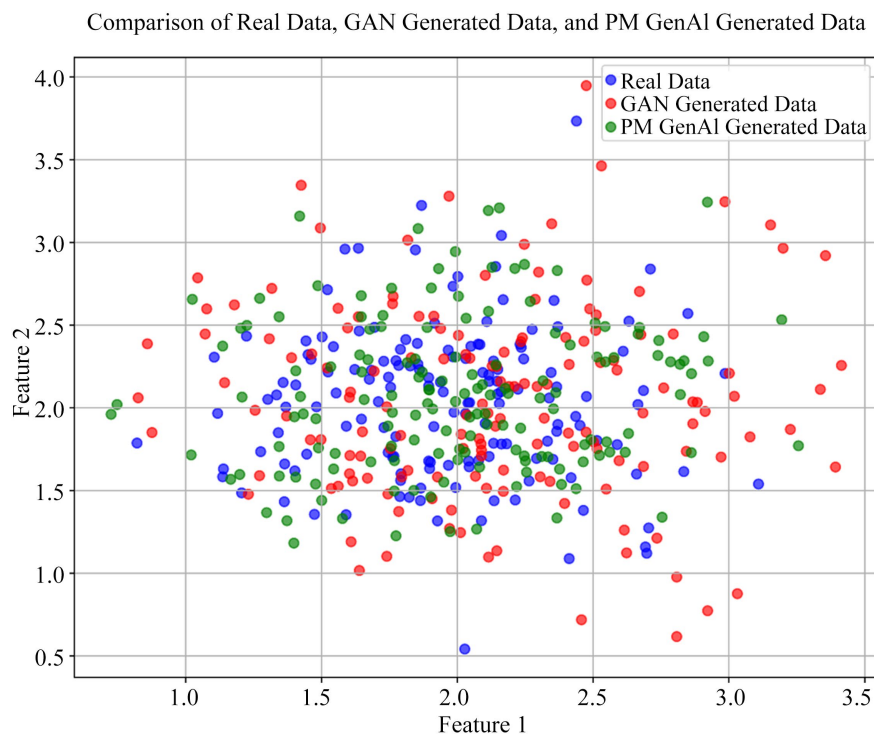


Figure 3. Scatter plot comparing real data, GAN generated data, and PM GenAI generated data. The PM GenAI samples are shown as a more personalized distribution that more closely follows the reference real data, while GAN generated samples illustrate a broader synthetic approximation of the underlying data distribution.

2.4. Variational Autoencoders

Variational Autoencoders (VAEs) are a class of generative models that play a central role in modern generative artificial intelligence. Introduced by Diederik P. Kingma and Max Welling, VAEs provide a probabilistic framework for learning latent representations of complex data. Unlike Generative Adversarial Networks, which rely on adversarial training, VAEs are based on principles from Bayesian inference and variational optimization. This makes them particularly suitable for applications in biology and medicine, where uncertainty, variability, and interpretability are critical.

VAEs are widely used for modeling biological data such as gene expression, protein structures, and medical images. Their ability to learn meaningful latent spaces enables both data generation and representation learning, making them highly relevant for PM GenAI and personalized medicine.

A VAE consists of two main components:

- Encoder (Inference Network)

Maps input data x to a latent representation z , typically parameterized as a probability distribution.

- Decoder (Generative Network)

Reconstructs data from the latent variable z , generating outputs similar to the original data.

Unlike traditional autoencoders, VAEs do not map inputs to fixed points but to probability distributions, allowing them to capture uncertainty and variability in biological systems. The goal of a VAE is to model the data distribution $p(x)$ using latent variables z . This is achieved by maximizing the Evidence Lower Bound (ELBO):

$$\mathcal{L}(\theta, \phi; x) = E_{q_{\phi}(z|x)}[\log p_{\theta}(x|z)] - D_{KL}(q_{\phi}(z|x) \| p(z)) \quad (2)$$

where:

- $q_{\phi}(z|x)$ is the encoder (approximate posterior)
- $p(x|z)$ is the decoder (likelihood)
- $p(z)$ is the prior distribution (usually Gaussian)
- D_{KL} is the Kullback-Leibler divergence

This objective has two components:

1) Reconstruction Loss

Ensures generated data is similar to input data.

2) Regularization Term (KL Divergence)

Ensures latent space follows a structured distribution.

One of the most important features of VAEs is their continuous latent space. Each data point is mapped to a region in latent space, allowing smooth interpolation between samples. In biological contexts, this enables:

- Modeling transitions between cell states.
- Understanding disease progression.
- Exploring molecular variations.

VAEs are widely used to analyze high dimensional genomic data. They can:

- Reduce dimensionality of gene expression datasets.
- Identify patterns in genetic variation.
- Model cellular heterogeneity.

VAEs are applied to:

- Image reconstruction and denoising.
- Generation of synthetic medical images.
- Feature extraction for diagnostic models.

Figure 4 shows a scatter plot comparing real data (blue) with synthetic data generated by a Variational Autoencoder (green) and PM based GenAI approach (red). Both generative models capture the central tendency and overall spread of the real data distribution. Quantitatively, the similarity between distributions can be assessed using divergence and overlap metrics such as the Kullback-Leibler (KL) divergence and distributional overlap coefficients. Lower KL divergence values indicate that the generated distributions closely approximate the real data distribution, while higher overlap coefficients suggest strong agreement in feature space. In this context, the clustering and dispersion patterns imply that both VAE and PM GenAI achieve reasonable alignment with the real data, although subtle differences in density and variance suggest variations in how each model captures tail behavior and local structure.




Figure 4. Comparison of real data and synthetic data generated using Variational Autoencoder (VAE) and Principal Model GenAI (PM GenAI). The scatter plot illustrates the distribution of samples across two features, showing that both generative approaches capture the general structure of the real data.

3. Discussion

The comparative analysis of real data with synthetic datasets generated using Probabilistic Model GenAI (PM GenAI), Variational Autoencoders (VAE), and, by extension, Generative Adversarial Networks (GANs), highlights important differences in their ability to capture underlying data distributions. The scatter plot demonstrates that both PM GenAI and VAE models approximate the central structure of the real data, preserving the overall mean and variance across the two feature dimensions. However, subtle differences in dispersion, clustering density, and representation of extreme values reveal the inherent trade-offs between these generative approaches.

The VAE framework is grounded in probabilistic inference and learns a latent representation by optimizing a variational lower bound. This results in smooth and continuous latent spaces, which are advantageous for capturing global data structure and ensuring stable training. In the presented results, the VAE-generated data exhibits a relatively broad spread, suggesting that it effectively captures the overall variability of the dataset. However, VAEs are known to produce slightly over-smoothed distributions, which can lead to underestimation of sharp features or local density variations. This is consistent with the observed dispersion, where the VAE samples appear somewhat more diffuse compared to the real data.

In contrast, PM GenAI leverages explicit probabilistic modeling of system dynamics and data-generating mechanisms. This approach allows for greater interpretability and direct incorporation of domain knowledge, which is particularly valuable in biomedical contexts such as prostate cancer modeling. The PM GenAI-generated data demonstrates a closer alignment with the central clustering of the real dataset, suggesting that it may better capture the dominant modes of the distribution. Additionally, because PM GenAI is grounded in mechanistic or stochastic modeling, it provides a more transparent relationship between input parameters and generated outputs. However, this approach may be limited in capturing highly complex, high-dimensional patterns unless the underlying model is sufficiently expressive.

GANs represent another important class of generative models for comparison. GANs operate through an adversarial training process between a generator and a discriminator, enabling them to produce highly realistic samples that closely match the data distribution. In many applications, GANs outperform VAEs in terms of visual fidelity and sharpness of generated samples. However, this comes at the cost of training instability, mode collapse, and lack of explicit likelihood estimation. Unlike VAEs and PM GenAI, GANs do not provide a direct probabilistic interpretation of the data, making uncertainty quantification more challenging. In the context of biomedical modeling, this limitation can be significant, as understanding uncertainty and variability is critical for clinical decision-making.

Quantitatively, differences between these models can be assessed using metrics such as Kullback-Leibler divergence, Wasserstein distance, or Bhattacharyya coefficient. VAEs typically achieve moderate divergence values due to their smoothing effect, while GANs often achieve lower divergence in high-density regions but may fail to capture the full support of the distribution. PM GenAI models, depending on their formulation, can achieve strong alignment in regions governed by the underlying model assumptions but may deviate in unexplored or highly nonlinear regions. The visual overlap observed in the scatter plot suggests that both VAE and PM GenAI achieve reasonable fidelity, with PM GenAI potentially offering better alignment in central regions while GANs and VAE provide broader coverage.

Overall, the choice between PM GenAI, VAE, and GAN approaches depends on the specific objectives of the application. VAEs offer stability, interpretability of latent space, and ease of training, making them suitable for exploratory analysis and probabilistic modeling. GANs provide high-quality sample generation but require careful tuning and lack interpretability. PM GenAI stands out for its ability to incorporate domain knowledge and provide interpretable, physically meaningful outputs, making it particularly well-suited for applications such as tumor modeling and PSA dynamics. Future work should explore hybrid approaches that combine the strengths of these methods, such as integrating mechanistic models with deep generative frameworks to improve both realism and interpretability.

4. Conclusion

Generative Adversarial Networks (GANs) have become a powerful tool in generative artificial intelligence, with significant applications in biology and medicine. Their ability to generate realistic synthetic data has transformed medical imaging, where GAN-based augmentation has been used to improve liver lesion classification performance, as demonstrated by Matan Frid-Adar *et al.* (2018) [5]. One current limitation is that synthetic images may contain subtle artifacts and therefore require careful clinical validation before being used for diagnostic purposes. GANs have also significantly impacted drug discovery, where generative models are employed to design novel molecular structures with desired biological and chemical properties. However, a major limitation is that many generated compounds may not be biologically viable or practically synthesizable. Although GANs and related generative models have been applied to simulate biological data and support predictive modeling, important challenges remain, including the limited availability of high-quality datasets and the difficulty of interpreting model outputs in sensitive biomedical contexts. As part of the PM GenAI framework, GANs are contributing to the advancement of personalized healthcare. Continued research and responsible implementation will be essential to fully realize their potential.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

- [1] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. (2014) Generative Adversarial Networks. *Advances in Neural Information Processing Systems*, **27**, 2672-2680. <https://arxiv.org/abs/1406.2661>
- [2] Kingma, D.P. and Welling, M. (2013) Auto-Encoding Variational Bayes. <https://arxiv.org/abs/1312.6114>
- [3] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., *et al.*, (2020) Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, **33**, 1877-1901,
- [4] Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., *et al.* (2017) Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks. *Nature*, **542**, 115-118. <https://doi.org/10.1038/nature21056>
- [5] Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J. and Greenspan, H. (2018) Gan-Based Synthetic Medical Image Augmentation for Increased CNN Performance in Liver Lesion Classification. *Neurocomputing*, **321**, 321-331. <https://doi.org/10.1016/j.neucom.2018.09.013>
- [6] Chen, H., Engkvist, O., Wang, Y., Olivecrona, M. and Blaschke, T. (2018) The Rise of Deep Learning in Drug Discovery. *Drug Discovery Today*, **23**, 1241-1250. <https://doi.org/10.1016/j.drudis.2018.01.039>
- [7] Kwon, Y., Kim, S., Lee, J. and Kang, S. (2021) Generative Models for Molecular Design in Drug Discovery. *Journal of Chemical Information and Modeling*, **61**, 3775-3789.

- [8] Segler, M.H.S., Preuss, M. and Waller, M.P. (2018) Planning Chemical Syntheses with Deep Neural Networks and Symbolic Ai. *Nature*, **555**, 604-610.
<https://doi.org/10.1038/nature25978>
- [9] Zhavoronkov, A., Ivanenkov, Y.A., Aliper, A., Veselov, M.S., Aladinskiy, V.A., Aladinskaya, A.V., *et al.* (2019) Deep Learning Enables Rapid Identification of Potent DDR1 Kinase Inhibitors. *Nature Biotechnology*, **37**, 1038-1040.
<https://doi.org/10.1038/s41587-019-0224-x>
- [10] Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., *et al.* (2021) Highly Accurate Protein Structure Prediction with AlphaFold. *Nature*, **596**, 583-589. <https://doi.org/10.1038/s41586-021-03819-2>
- [11] Angermueller, C., Pärnamaa, T., Parts, L. and Stegle, O. (2016) Deep Learning for Computational Biology. *Molecular Systems Biology*, **12**, 878-890.
<https://doi.org/10.15252/msb.20156651>
- [12] Ching, T., Himmelstein, D.S., Beaulieu-Jones, B.K., Kalinin, A.A., Do, B.T., Way, G.P., *et al.* (2018) Opportunities and Obstacles for Deep Learning in Biology and Medicine. *Journal of The Royal Society Interface*, **15**, Article 20170387.
<https://doi.org/10.1098/rsif.2017.0387>
- [13] Ramesh, A., Dhariwal, P., Nichol, A., Chu, C. and Chen, M. (2022) Hierarchical Text-Conditional Image Generation with CLIP Latents.
<https://arxiv.org/abs/2204.06125>
- [14] Melo, P.D. and Rose, M.S. (2025) Accurate Classification of Diabetes via PM Generative Ai. *Advances in Bioscience and Biotechnology*, **16**, 379-409.
<https://doi.org/10.4236/abb.2025.169025>
- [15] Topol, E.J. (2019) High-Performance Medicine: The Convergence of Human and Artificial Intelligence. *Nature Medicine*, **25**, 44-56.
<https://doi.org/10.1038/s41591-018-0300-7>
- [16] Rajkomar, A., Dean, J. and Kohane, I. (2019) Machine Learning in Medicine. *New England Journal of Medicine*, **380**, 1347-1358.
<https://doi.org/10.1056/nejmra1814259>