

Ganesan, Swathi and Somasiri, Nalinda (2026) A Narrative Review of GAN-Based Synthetic Data Generation in Disease Prediction. Journal of Data Science and Intelligent Systems.

Downloaded from: <https://ray.yorks.ac.uk/id/eprint/15256/>

The version presented here may differ from the published version or version of record. If you intend to cite from the work you are advised to consult the publisher's version:  
<https://doi.org/10.47852/bonviewJDSIS62028542>

Research at York St John (RaY) is an institutional repository. It supports the principles of open access by making the research outputs of the University available in digital form. Copyright of the items stored in RaY reside with the authors and/or other copyright owners. Users may access full text items free of charge, and may download a copy for private study or non-commercial research. For further reuse terms, see licence terms governing individual outputs. [Institutional Repositories Policy Statement](#)

# RaY

Research at the University of York St John

For more information please contact RaY at  
[ray@yorks.ac.uk](mailto:ray@yorks.ac.uk)

## REVIEW

# A Narrative Review of GAN-Based Synthetic Data Generation in Disease Prediction

Swathi Ganesan<sup>1,\*</sup> and Nalinda Somasiri<sup>1</sup><sup>1</sup>Department of Computer and Data Science, York St John University, UK

**Abstract:** Synthetic data generation has emerged as an important approach in the healthcare field to address data scarcity, disease class imbalance, and privacy restrictions that limit access to patient data. Among generative approaches, generative adversarial networks (GANs) have gained increasing attention, especially because of their ability to generate realistic data across complex data distributions such as medical imaging, electronic health records (EHRs), laboratory data, and phenotype codes. This narrative review focuses on the evolution of major GAN architectures and their application in disease prediction. The original GAN introduced the adversarial paradigm, while Deep Convolutional GAN advanced image generation and became widely used in MRI, CT, and histopathology tasks. Wasserstein GAN variants (WGAN and WGAN-GP) improve training stability and prove to be more suitable for tabular and structured healthcare data such as EHRs. More specialized architectures, including Conditional Tabular GAN and Medical GAN, further extended synthetic data generation to mixed-type datasets and sparse diagnostic records. The review also examines evaluation practices based on data fidelity, downstream utility, and privacy preservation, including differential privacy and resistance to membership inference attacks. Overall, the literature shows that GAN-generated synthetic data can support disease prediction research, but important challenges remain in benchmarking, reproducibility, interpretability, and ethical deployment. Emerging directions include hybrid GAN-diffusion models, federated training strategies, and standardized evaluation frameworks to support clinically reliable and privacy-preserving adoption.

**Keywords:** generative adversarial networks (GANs), synthetic data generation, electronic health records (EHR), AI in healthcare, disease prediction

## 1. Introduction

Artificial intelligence (AI) systems for disease prediction are often constrained by data scarcity, class imbalance, and privacy restrictions that limit access to patient-level data. These constraints can degrade model performance and hinder reproducibility across sites and populations. A growing body of work positions synthetic data as a practical complement to real data, expanding sample size, balancing minority classes, and enabling data sharing without exposing raw patient records. Recent reviews in healthcare underscore both the promise and the caveats of synthetic data across imaging, tabular, and time-series modalities [1].

Within generative approaches, generative adversarial networks (GANs) have become a central technique for synthesizing realistic samples by framing generation as a two-player game between a generator and a discriminator. The original GAN formulation established the adversarial learning paradigm and demonstrated compelling sample quality, catalyzing rapid adoption in medical AI [2]. Building on this, Wasserstein GAN (WGAN) introduced an Earth Mover's (Wasserstein-1) distance objective that improves training stability and correlates better with sample quality, an important property when synthesizing heterogeneous clinical data [3].

Healthcare spans diverse data types, including radiological images, pathology slides, electronic health records (EHR), laboratory measurements, and physiological time series, so GAN architectures have evolved along modality lines. Convolutional variants such as Deep Convolutional GAN (DCGAN) are widely used for image synthesis and augmentation in radiology and digital pathology, while EHR/tabular synthesis is commonly addressed with architectures tailored to mixed discrete-continuous variables (e.g., Medical GAN [MedGAN], Conditional Tabular GAN [CTGAN]) and with WGAN/WGAN-GP objectives to stabilize training on structured data. Collectively, this literature shows DCGAN-style models are effective for images, whereas WGAN-based and tabular-specific designs (CTGAN/MedGAN) are frequently preferred for structured health data [4].

Beyond feasibility, evaluation has matured from visual inspection to three complementary axes:

- 1) Fidelity – how closely synthetic distributions match real data (e.g., Wasserstein distance, Kolmogorov–Smirnov (KS) tests).
- 2) Utility – performance of downstream predictive models trained on or augmented with synthetic data.
- 3) Privacy – resilience to attacks such as membership inference and re-identification.

Recent studies and reviews emphasize the privacy–utility trade-off, noting that partially synthetic datasets may be more

\*Corresponding author: Swathi Ganesan, Department of Computer and Data Science, York St John University, UK. Email: [s.ganesan@yorksj.ac.uk](mailto:s.ganesan@yorksj.ac.uk)

vulnerable than fully synthetic ones and calling for standardized, multi-metric reporting in medical contexts [5].

Given this landscape, this literature review synthesizes GAN-based synthetic data generation for disease prediction across modalities, highlighting architectural choices (with attention to WGAN's stability advantages for tabular/EHR data), evaluation practices, and ethical/regulatory considerations. Our goal is to provide a theory-led, application-oriented overview that helps researchers select suitable GAN variants for specific medical data types and reporting standards that align with contemporary guidance on AI and data protection [3].

## 2. Scope of Literature Review

This review adopts a narrative (literature) review approach rather than a systematic review, allowing for a broader and more flexible synthesis of the most influential and representative studies on GANs for synthetic data generation in disease prediction. Unlike systematic reviews that exhaustively identify and appraise every eligible study, narrative reviews focus on key contributions, trends, and conceptual advances, offering a richer theoretical and methodological context.

### 2.1. Databases and search

To identify relevant literature, we consulted PubMed, IEEE Xplore, Scopus, Web of Science, and arXiv between January 2014 and January 2025, a period coinciding with the rapid development of GAN architectures and their first applications to medical AI. Search strings combined the terms “*synthetic data*,” “*GAN*,” “*disease prediction*,” “*healthcare*,” “*tabular data*,” and “*medical images*.” We also examined reference lists of recent reviews and key studies to ensure inclusion of foundational papers.

### 2.2. Selection criteria

Studies were included if they:

- 1) Described or applied a GAN architecture (including variants such as DCGAN, WGAN, MedGAN, CTGAN).
- 2) Generated synthetic data for any medical or health-related disease prediction task (tabular, image, or time-series).
- 3) Reported at least one evaluation metric (fidelity, utility, or privacy).

Excluded were papers focused solely on non-healthcare domains or on generative models other than GANs (e.g., Variational Autoencoders (VAEs), diffusion) unless used as a comparison to GANs.

Following this initial screening, the studies discussed in detail were curated purposively to reflect milestone contributions in the development and application of GANs for healthcare synthetic data generation. Papers were prioritized if they met one or more of the following criteria:

- 1) Introduced a foundational GAN architecture or a widely adopted variant relevant to healthcare data synthesis;
- 2) Demonstrated early or influential application of GANs in disease prediction;
- 3) Represented major healthcare data modalities, including tabular EHRs, medical imaging, and time-series data; and/or
- 4) Reported evaluation outcomes relevant to data fidelity, downstream utility, or privacy preservation.

As a narrative review, the intention was not to produce an exhaustive catalog of all published studies but to provide a balanced and conceptually representative synthesis of key methodological and applied developments in the field.

### 2.3. Scope and emphasis

This review focuses on the theoretical foundations of major GAN architectures, their subsequent extensions, and their applicability across key healthcare data modalities, including medical imaging, tabular EHRs, time-series biosignal data, and multi-modal disease prediction tasks. Particular attention is given to how different GAN variants address the technical challenges associated with these data types and how they support synthetic data generation for downstream predictive modeling.

While WGAN is highlighted as a leading approach for tabular/structured data, the review also examines DCGAN and other convolutional variants for medical imaging to provide a balanced view across modalities. This cross-modal perspective enables the review to capture the methodological evolution of GAN-based synthetic data generation from foundational architectures to more specialized healthcare applications. Rather than undertaking formal study-level bias assessment or exhaustive enumeration, the review emphasizes conceptual developments, comparative strengths, and practical relevance to disease prediction research.

## 3. Overview of Synthetic Data Generation in Healthcare

In this section, we survey how synthetic data generation (especially via GANs) is being used in healthcare, with attention to data types, application areas, and recent architectural advances.

### 3.1. Data modalities and application areas

Healthcare synthetic data span multiple modalities, each with distinct methodological challenges. These include tabular structured data such as EHRs and laboratory records, time-series and longitudinal data such as electrocardiogram (ECG) and continuous monitoring signals, and medical image data such as MRI, CT, ultrasound, and pathology images.

**Tabular structured data:** EHRs, clinical laboratory data, and risk factor datasets represent a major area of healthcare synthetic data generation. GANs are applied in this context to address minority-class underrepresentation, handle missing values, and enable privacy-protected data sharing. Ghosheh et al. reviewed GAN applications for EHRs and documented how fidelity and privacy metrics are used across such studies [6]. Beyond general data synthesis, more recent work has focused on improving the quality and completeness of structured healthcare records. Missing data remains one of the most critical challenges in real-world healthcare datasets, particularly in EHRs, where clinical variables are often irregularly recorded and highly sparse. To address this issue, a clinical conditional GAN (cGAN) framework has been proposed for imputing missing values by modeling nonlinear and multivariate relationships across patient records [7]. Unlike traditional imputation approaches that rely on linear assumptions or univariate statistics, this model conditions the generation process on available clinical features, enabling it to infer missing values while preserving complex dependencies between variables. This is particularly important in healthcare settings, where laboratory tests and clinical observations are not collected uniformly

over time. The study reported that the proposed framework significantly outperformed existing imputation methods, improving imputation accuracy by approximately 19.79% and enhancing downstream predictive performance on real-world multicenter diabetic datasets [7]. It also demonstrated robustness across varying levels of missingness, highlighting its practical relevance in clinical environments. These findings emphasize that GAN-based methods are valuable not only for synthetic data generation but also for improving data quality and completeness, both of which are essential for reliable machine learning-based disease prediction systems.

**Time-series and longitudinal data:** Physiological signals, such as ECG recordings and continuous monitoring data, as well as repeated longitudinal measurements, represent another important modality in healthcare synthetic data generation. The scoping review by Loni et al. identifies a growing body of work using adversarial models to generate time-series and longitudinal health records [8]. Beyond general sequence generation, GAN-based approaches have also been applied to clinically relevant signal restoration tasks, particularly in ECG denoising and reconstruction. A recent survey reports the use of architectures such as WGAN, Conditional Generative Adversarial Network (CGAN); CycleGAN, and Least Squares Generative Adversarial Network (LSGAN) for modeling ECG noise distributions and recovering clinically meaningful waveform patterns [9]. These methods have been shown to improve signal-to-noise ratio and reduce mean squared error relative to traditional signal processing approaches, while preserving diagnostically important waveform components such as the P wave, QRS complex, and T wave [9]. At the same time, the literature identifies ongoing challenges, including variability in noise sources, training instability, and computational cost, which continue to limit robust clinical deployment. Taken together, these findings reinforce the suitability of GAN-based

methods for time-series healthcare data, particularly in continuous disease monitoring and signal-based diagnostic applications.

**Image data:** Radiology, pathology, and medical imaging (MRI, CT, ultrasound). GANs (DCGAN, cGAN, etc.) have been heavily used for image augmentation, anomaly detection, and converting between image modalities. Ibrahim et al. (2024) provide a systematic survey covering image-based synthesis among other modalities [10]. More recent radiology-focused reviews further emphasize that synthetic medical imaging has become increasingly important for addressing data scarcity, class imbalance, and privacy constraints in image-based AI workflows, while also drawing attention to persistent challenges related to realism, validation, and clinical usability [11].

### 3.2. Key GAN architectures in use with evaluation metrics

Over the past several years, various GAN variants have been adapted for synthetic healthcare data. Several review studies have provided a broad overview of GAN architectures and their applications across domains, including healthcare. These works highlight the rapid growth of GAN-based models since 2016 and their use in tasks such as image processing, medical data analysis, and pattern generation, while also identifying key challenges such as mode collapse, training instability, and evaluation limitations [12]. Table 1 summarizes the key architectures and their roles.

### 3.3. Challenges in synthetic data generation

While GANs offer powerful tools, multiple challenges recur in the literature:

**Heterogeneity of tabular data:** Medical datasets often mix categorical, ordinal, continuous, and sometimes missing data. Many

**Table 1**  
Key GAN architectures in use with evaluation metrics

GAN variant	Main features/advantages	Typical use cases in healthcare	Evaluation metric
Original GAN/conditional GAN (cGAN)	Basic adversarial setup; conditional GANs incorporate class labels or attributes into both generator and discriminator through auxiliary-classifier mechanisms, enabling class-controlled sample generation (e.g., disease-specific synthesis)	Class-conditioned medical image synthesis and targeted minority-class augmentation for disease classification tasks [13]	Quality metrics: FID and IS are used to evaluate the visual quality and diversity of the generated medical images Computational considerations: Training stability and convergence time are important considerations given that medical images are of very high dimensions
DCGAN	Uses convolutional architectures; well-suited for images; stable in high-dimensional pixel spaces	Image augmentation (X-ray, CT, MRI), especially where real images are limited [10]	Quality metrics: FID is commonly used to assess realism, while statistical similarity measures (e.g., pixel-level distributions and correlation patterns) evaluate how closely synthetic images match real scans Utility metrics: Performance on downstream tasks (e.g., classification or detection)

(Continued)

**Table 1**  
(Continued)

GAN variant	Main features/advantages	Typical use cases in healthcare	Evaluation metric
Wasserstein GAN (WGAN/WGAN-GP, etc.)	Uses Wasserstein distance in the loss, which improves gradient behavior, reduces mode collapse, and yields more stable training. Especially relevant for heterogeneous, structured data	Tabular/EHR data generation; works that need stable output distributions; mixed categorical-continuous datasets. Ahmed et al. [14] compared multiple GAN variants, including WGAN, CTGAN, and CRAMER-GAN for tabular disease-related datasets and reported WGAN among the better performers	Privacy metrics: Re-identification risk, membership inference attack success rate, and formal differential privacy guarantees Utility metrics: Improvement in performing the tasks downstream, i.e., prediction tasks
Other variants/specialized GANs	e.g., MedGAN, CTGAN, DRAGAN, CRAMER-GAN, etc. These often incorporate modifications for mixed data types, privacy, stabilization, or better handling of discrete variables	Structured data, synthetic EHRs, privacy-aware tasks [15]	Privacy metrics: Re-identification risk, membership inference attack success rate, and formal differential privacy guarantees Utility metrics: Improvement in downstream prediction tasks

GANs struggle with preserving distributions across mixed variable types. Studies show that architectures or preprocessing that explicitly handle mixed data types (e.g., via specialized encoding) are more successful [16]. Recent work also shows that imbalance may exist within classes, not only between them. In medical image augmentation, GANs may improve class balance while still failing to capture sparse intra-class regions, limiting sample diversity and downstream performance [17].

**Training stability and mode collapse:** GANs can struggle with mode collapse (where many synthetic samples are similar), especially with limited data, which occurs commonly in rare disease contexts. To address this issue, the WGAN family, particularly WGAN-GP, is frequently used because of its improved training stability [10]. In addition, this review emphasizes evaluation across fidelity, utility, and privacy, underscoring the urgent need for standardized benchmarking protocols.

**Evaluation metrics:** There is no consensus on a standard set of metrics. Many works report fidelity (distribution similarity, statistical distance), utility (performance of downstream prediction), and sometimes privacy (risk of re-identification). However, metrics vary widely, making comparison difficult. Reviews call for more standardized benchmarks [10].

**Privacy–utility trade-offs:** Synthetic data can leak information if privacy is not handled properly; yet, stronger privacy constraints can degrade utility. Many studies highlight this trade-off [8].

**Generalizability and clinical validity:** Even when a synthetic generator works well on a given dataset, issues emerge when applying it in new settings or external validation. Ensuring that synthetic disease patterns are clinically plausible and not just statistically accurate is critical [18].

### 3.4. Recent empirical comparisons

Some recent works provide direct empirical comparisons between GAN variants for healthcare synthetic data:

- 1) Ahmed et al. [14] generated synthetic data using six GAN variations, including GAN, cGAN, CTGAN, Cramer-GAN, DRAGAN, and WGAN on healthcare tabular datasets (Breast Cancer, Lung Cancer, CTG). The authors compared classification performance (XGBoost, Support Vector

Machine (SVM), statistical/correlation analyses, and hybrid datasets (real + synthetic) versus purely synthetic datasets. Their results suggest that CTGAN and cGAN often perform well for classification utility on tabular data, while WGAN shows stable statistical fidelity.

- 2) The survey “A Comprehensive Survey of Synthetic Tabular Data Generation” by Shi et al. [16] compares a broad range of generative paradigms (GANs included) for tabular data, highlighting trade-offs in fidelity, privacy, and heterogeneity handling.

## 4. GAN Architectures and Applications

GANs have become a key tool for addressing several recurring challenges in healthcare data, particularly data scarcity, class imbalance, and privacy constraints. By learning to approximate complex data distributions, GAN-based generators can synthesize realistic medical samples that supplement or partially replace sensitive real-world data. This is especially valuable in disease prediction tasks where annotated datasets are limited, highly imbalanced, or expensive to collect.

Different GAN variants have been introduced to mitigate specific technical limitations of the original framework, such as training instability, mode collapse, lack of control over the generated outputs, or difficulty modeling heterogeneous tabular data. Architectural refinements, including conditioning mechanisms, convolutional layers, and alternative loss functions, have enabled tailored solutions for applications ranging from tumor detection in imaging to chronic disease risk prediction in structured EHRs.

Taken together, these architectures form a family of models that can be matched to the characteristics of the target data (image, tabular, time-series) and the downstream prediction task in disease modeling.

### 4.1. Variants of GANs

Although all GANs share the same generator–discriminator backbone, several variants have been developed to tackle specific drawbacks of the original formulation and to adapt the adversarial idea to diverse healthcare data modalities. In each case,

the generator aims to produce synthetic samples that are indistinguishable from the real data, while the discriminator evaluates their authenticity. What differs between variants is how the generator is guided, which loss function is optimized, and how the architecture is adapted to image, tabular, or sequential data.

#### 4.1.1. Original GAN

The original GAN introduced by Goodfellow et al. formalized the adversarial training paradigm as a minimax game between generator and discriminator [2]. The generator receives random noise as input and learns to map it into synthetic samples that mimic the real data distribution, whereas the discriminator attempts to distinguish real samples from generated ones. Over successive training iterations, the two networks co-evolve: the generator improves at fooling the discriminator, and the discriminator becomes more accurate at detecting fakes.

Despite its conceptual elegance, the original GAN is prone to training instability and mode collapse, particularly in high-dimensional settings typical of medical data [2]. Early healthcare studies often used this baseline architecture as a proof of concept, for instance, to augment small medical imaging datasets with synthetic scans, demonstrating that even simple GANs can improve classifier performance when real data are scarce.

Figure 1 illustrates the adversarial training process between the generator and discriminator. The generator receives random noise as input and attempts to produce synthetic samples that resemble the real data distribution. The discriminator processes both real data and generated samples, learning to classify them as real or fake. During training, adversarial feedback flows from the discriminator back to the generator, enabling the generator to iteratively improve its ability to create realistic outputs. This adversarial loop forms the foundation for all GAN variants used in subsequent sections.

In the medical domain, the original GAN framework established the basis for later healthcare-specific generative models by demonstrating how adversarial training could be used to generate realistic synthetic data from limited datasets [19]. Recent reviews in medical imaging note that this foundational generator-discriminator structure has supported a wide range of downstream applications, including synthetic image generation, image enhancement, segmentation, detection, classification, reconstruction, and augmentation [19]. Although later variants

were introduced to improve stability, conditioning, and cross-modal performance, the original GAN remains conceptually important because it introduced the adversarial learning mechanism that underpins subsequent developments in medical image synthesis and analysis.

#### 4.1.2. Conditional GAN

The cGAN extends the original GAN framework by introducing class-conditioning information into both the generator and discriminator [13]. Instead of relying solely on random noise, the generator also receives an auxiliary variable such as a class label, disease status, or other patient attributes that guide the type of sample it should produce. The discriminator is trained with the same conditioning variable, so it can assess whether a sample is both realistic and consistent with the specified condition. This conditioning mechanism makes cGANs highly effective for healthcare applications where the goal is to generate class-specific or attribute-controlled samples. For example, cGANs can synthesize tumor-specific MRI images, minority-class patient records, or disease-positive time-series sequences, thereby addressing data imbalance in prediction tasks [1]. Because the model explicitly incorporates clinically relevant labels, cGANs serve as a powerful augmentation strategy for disease classification pipelines where certain diagnostic categories are underrepresented.

Figure 2 illustrates the cGAN workflow, where both the generator and discriminator receive the same conditioning variable (e.g., class label). The generator uses this information to produce synthetic samples aligned with the specified class, while the discriminator checks whether the input sample is real or synthetic. The discriminator also checks whether the input sample matches the provided condition. This architecture enables targeted synthetic data generation, making cGANs particularly suited for imbalanced medical datasets [20].

Beyond healthcare, cGANs have also been used in other structured image synthesis tasks, illustrating the flexibility of conditioning mechanisms in controlling generated outputs [21]. However, in medical applications, their value lies more specifically in enabling class-guided synthesis, minority-class augmentation, and disease-targeted data generation.

#### 4.1.3. Deep Convolutional GAN

DCGAN represents one of the most influential architectural refinements of the original GAN framework. It replaces fully connected layers with convolutional and transposed-convolutional (deconvolutional) layers, enabling the model to learn hierarchical spatial representations far more effectively than a multilayer perceptron. Jenkins and Roy [22] describe how this architecture leverages convolutional feature extraction, batch normalization, and nonlinear activation functions to construct stable training pipelines capable of generating high-fidelity synthetic images.

In the DCGAN generator, a latent noise vector is progressively upsampled through stacked transpose-convolutional layers

Figure 1  
General GAN architecture





Figure 2  
Conditional GAN architecture



to form a synthetic image. Each layer learns increasingly complex spatial patterns, beginning with coarse structural outlines and refining into fine-grained textures. Conversely, the discriminator uses standard convolutional layers to downsample real and generated images into feature maps, enabling more robust discrimination than fully connected counterparts.

These architectural features make DCGAN particularly effective for medical imaging tasks, where spatial coherence and anatomical consistency are essential. The model has been widely applied for:

- 1) Image augmentation, generating synthetic CT, MRI, and X-ray scans to expand limited datasets for tumor detection or lung disease classification.
- 2) Anatomical structure synthesis, where DCGAN learns to replicate patterns such as lesions, organ boundaries, and tissue texture with high fidelity.
- 3) Support for modality translation pipelines, in which DCGAN often acts as the backbone generator for more complex GAN variants in MRI-CT mapping tasks.
- 4) Super-resolution, serving as a foundational component in frameworks that enhance undersampled or low-contrast medical images.

Beyond healthcare imaging, DCGAN has also demonstrated strong performance in modeling fine-grained spatial patterns in other domains, such as deepfake generation and biometric synthesis. For instance, Ganesan et al. [23] illustrate how DCGAN architectures can capture subtle structural variations in human facial features, an insight that directly translates to generating anatomically coherent synthetic medical images.

Figure 3 shows a typical DCGAN configuration, where the generator upsamples a noise vector into an image using transpose-convolutional layers and the discriminator uses convolutional filters for downsampling. This layered design enables DCGAN to learn spatial hierarchies essential for high-fidelity medical image synthesis. Despite the rise of more specialized image-focused GANs, DCGAN remains a widely used baseline because of its stability, interpretability, and strong spatial feature learning, making it a reliable option for disease prediction applications where high-quality synthetic images are needed.

Extensions of DCGAN have also explored supervised learning settings to improve performance on multi-class image datasets. A supervised DCGAN framework has been proposed to incorporate label information directly into the training process, enabling

more effective learning across multiple categories and improving classification performance [24]. Such approaches highlight the potential of GANs not only for data generation but also for enhancing supervised learning tasks, particularly when labeled data are available.

#### 4.1.4. Wasserstein GAN

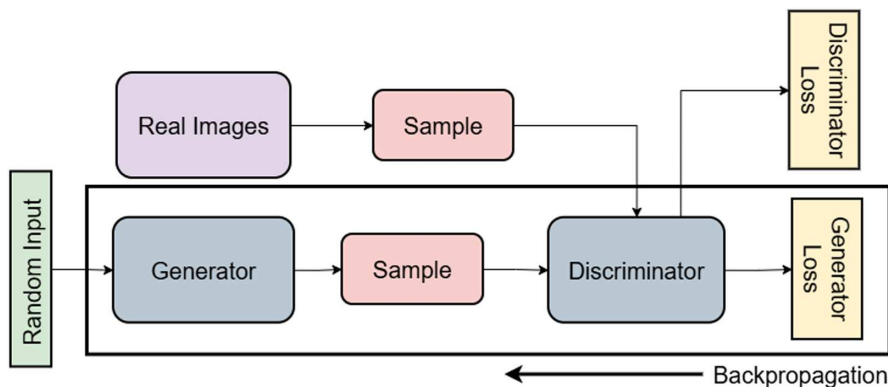
The WGAN was introduced to address key challenges associated with the original GAN framework, such as unstable gradients, sensitivity to hyperparameters, and mode collapse. Rather than using the Jensen–Shannon divergence, WGAN optimizes the Wasserstein (Earth Mover’s) distance, which provides smoother and more meaningful gradients even when the generated and real distributions have little overlap [3]. As a result, WGAN offers more stable and reliable convergence behavior compared to earlier GAN variants.

A fundamental change in WGAN is the replacement of the discriminator with a critic, which outputs a continuous score indicating how close a sample is to the real data distribution. The critic is trained to assign higher scores to real data and lower scores to generated samples, while the generator is optimized to increase the critic’s score for synthetic data. This formulation yields a more informative training signal and mitigates common issues such as mode collapse.

The rapid evolution of GAN architectures has led to extensive comparative studies evaluating their performance across different tasks and datasets. A recent comparative analysis demonstrates that while standard GANs are capable of generating realistic samples, they often suffer from instability, mode collapse, and inconsistent convergence [25]. In contrast, WGAN introduces the Wasserstein distance as a loss function, resulting in smoother gradient behavior and improved training stability. Further enhancement through WGAN-GP addresses limitations associated with weight clipping by incorporating a gradient penalty term, leading to more stable convergence and higher-quality synthetic outputs. Experimental findings indicate that WGAN-GP consistently produces clearer and more structurally coherent outputs with fewer artifacts compared to earlier variants. These insights are particularly relevant for healthcare applications, reinforcing the preference for WGAN-based architectures in modeling complex and high-dimensional data distributions.

The effectiveness of WGAN in medical image augmentation has been demonstrated in brain MRI studies. A WGAN-GP framework was used to generate synthetic brain MR images

Figure 3  
Deep Convolutional GAN architecture



from healthy datasets to address data scarcity and support age classification tasks [26]. When combined with real images, the synthetic data improved Convolutional Neural Network (CNN) classification accuracy from 95.14% to 98.37%, indicating that WGAN-based augmentation can enhance robustness and generalization in medical image analysis, particularly when labeled data are limited [26].

Beyond its theoretical contributions, WGAN has been effectively used in practice across various data modalities. For example, the work by Ebner and Eltelt [27] demonstrates how a standard adversarial framework can incorporate stability principles that conceptually align with WGAN-like training improvements. Their approach applies these stability concepts to audio inpainting tasks, where the model learns to reconstruct missing segments of audio signals. Although their application domain differs from healthcare, the emphasis on improved gradient behavior and training stability remains relevant to any domain requiring robust generation of structured, high-dimensional data. This highlights the versatility of WGAN principles and their adaptability beyond image synthesis.

Figure 4 [27] illustrates a typical WGAN architecture, where the critic evaluates real and synthetic samples using a continuous scoring function, and the generator seeks to maximize this score for its outputs. This configuration enhances training stability and allows WGAN to model complex data distributions more effectively than classical GANs [27].

In healthcare, WGAN is especially valuable for modeling tabular medical datasets, such as EHRs, laboratory values, and mixed-type clinical variables. These datasets often contain complex dependencies, non-Gaussian distributions, and significant class imbalance that benefit from the stable gradient properties of WGAN. A study by Ganesan and Somasiri [28] further demonstrates this capability by using WGAN to generate synthetic records for underrepresented patient groups while maintaining distributional fidelity. Their findings reinforce WGAN’s suitability for disease prediction scenarios where real clinical data are limited, imbalanced, or privacy-sensitive.

Recent work in medical imaging further supports the effectiveness of WGAN-GP for synthetic image generation. In chest X-ray synthesis, a modified WGAN-GP incorporating adaptive  $\lambda$  and the AdaBelief optimizer has been shown to improve training stability and image quality compared with DCGAN, WGAN, and standard WGAN-GP, as assessed using Fréchet inception distance (FID) and Inception Score (IS) [29]. These findings reinforce the suitability of WGAN-based approaches for privacy-conscious data augmentation and synthetic medical image generation.

#### 4.1.5. Conditional Tabular GAN

CTGAN was introduced to address the unique challenges associated with modeling tabular datasets, which are common in healthcare. Unlike images, which have spatial regularity, medical tabular data typically consist of mixed variable types (continuous, categorical, ordinal), highly imbalanced distributions, and complex interdependencies that classical GANs struggle to learn. CTGAN overcomes these limitations by incorporating a conditional generator and a training-by-sampling strategy that explicitly targets underrepresented categories and non-Gaussian variable distributions [30].

A central innovation in CTGAN is its ability to handle categorical variables through a conditional mechanism. During each training iteration, the model selects a discrete attribute (e.g., diagnosis category) and samples real data that belong to that attribute. The generator then receives both random noise and the selected condition, producing synthetic rows consistent with that attribute. This conditional process ensures that CTGAN learns the distribution of each subgroup more accurately, making it well-suited for medical datasets with diagnostic labels, comorbidities, or patient stratification markers.

In addition, CTGAN uses a mode-specific sampling strategy, which focuses on difficult or underrepresented values of continuous features. This improves its ability to model skewed clinical variables such as laboratory biomarkers, physiological measurements, or rare comorbidity patterns, variables that standard GANs or even WGAN variants often struggle to reproduce.

As shown in Figure 5 [31], the work by Asimopoulos et al. [31] depicts a typical CTGAN workflow, where real patient records are grouped by discrete attributes during training, and the generator synthesizes new samples conditioned on these attributes. The discriminator evaluates whether each record (real or synthetic) matches the corresponding condition, enabling CTGAN to model both feature distributions and inter-feature dependencies within tabular healthcare data. CTGAN has been widely adopted in healthcare studies requiring the synthesis of realistic patient-level records. It has been applied to generate:

- 1) Synthetic EHR rows, including demographics, laboratory values, diagnosis codes, and treatment indicators;
- 2) Minority-class patient cohorts, supporting predictive modeling tasks with imbalanced clinical outcomes;
- 3) Tabular clinical datasets that require preserving multi-feature interactions for risk prediction models;
- 4) Synthetic disease phenotypes, enabling testing and validation of machine learning models without exposing sensitive patient data [32].

Figure 4  
Wasserstein GAN architecture

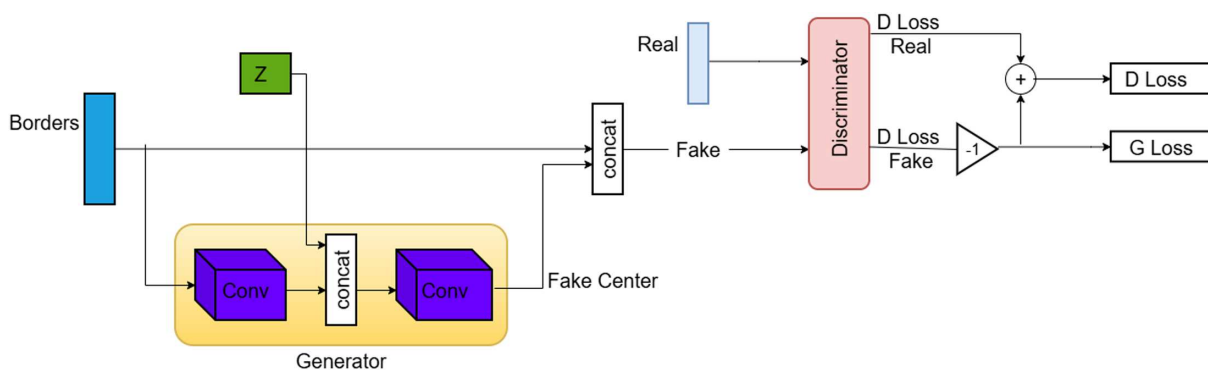



Figure 5  
CTGAN architecture



Because of its focus on conditional modeling and mixed-type variables, CTGAN is frequently used when the downstream task involves disease prediction from structured medical datasets, especially when real-world data are limited or contain rare disease categories. Its ability to replicate complex tabular distributions makes it a strong complement to image-focused GANs such as DCGAN and a specialized solution for structured data where WGAN may still struggle without additional transformations.

Recent research has extended GAN-based tabular synthesis in healthcare, particularly for cardiovascular datasets containing mixed continuous and categorical variables. Enhanced cGAN architectures have been proposed to better preserve heterogeneous feature distributions and inter-feature dependencies through specialized handling of continuous and categorical data [33]. Experimental results using metrics such as the Kolmogorov–Smirnov statistic and Jaccard similarity suggest improved preservation of marginal distributions and feature relationships compared with baseline methods such as CTGAN and Tabular Variational Autoencoder (TVAE) [33]. These findings reinforce the importance of GAN-based synthetic tabular data for privacy-preserving healthcare analytics and robust predictive modeling.

#### 4.1.6. Medical GAN


MedGAN is one of the earliest GAN architectures designed specifically for high-dimensional, sparse, and discrete medical data, particularly EHRs. Traditional GANs struggle with multi-hot encoded diagnosis codes, procedure vectors, or binary clinical indicators because these inputs lack the continuous structure found in image or audio data. MedGAN overcomes this challenge by integrating a denoising autoencoder with the GAN framework, allowing it to model sparse clinical vectors more effectively [34].

In MedGAN, the autoencoder first learns a compressed latent representation of the high-dimensional EHR data. The generator is then trained to produce synthetic latent vectors, which are passed through the decoder to reconstruct full-dimensional clinical records. This hybrid approach enables MedGAN to handle discrete and binary medical attributes, as the decoder learns to map latent variables back into structured multi-hot representations. Meanwhile, the discriminator works at the data-space level, distinguishing between real and decoded synthetic records, which encourages the generator to produce latent vectors that decode into realistic patient profiles.

A key advantage of MedGAN is its ability to preserve co-occurrence patterns across diagnosis and procedure codes, a crucial requirement for disease prediction tasks. Medical datasets often contain strong dependencies between conditions (e.g., diabetes–hypertension comorbidities), and MedGAN has been shown to reproduce these patterns more effectively than baseline GANs. This makes it highly suitable for generating synthetic cohorts for risk modeling, readmission prediction, and comorbidity analysis.

Figure 6 [35] presents the imaging-based MedGAN workflow proposed by Guo et al. [35]. This model differs from the original EHR-focused MedGAN introduced by Choi et al. [34], which was designed for sparse, multi-hot clinical records. In the imaging-based version, MedGAN is described as an adaptive extension of the GAN framework designed to address training instability, mode collapse, and poor convergence in medical image generation tasks. By incorporating dynamic training strategies and Wasserstein-based convergence measures, this approach improves adversarial stability and supports downstream tasks including disease classification and lesion localization in data-limited imaging scenarios [35]. Although both models share the name MedGAN, they are distinct and correspond to different healthcare data modalities.

Figure 6  
MedGAN architecture



MedGAN has been applied to generate:

- 1) Synthetic diagnosis and procedure vectors, supporting predictive modeling when real data are sparse or privacy-restricted;
- 2) High-dimensional EHR profiles, including thousands of binary indicators;
- 3) Rare disease patient records, enhancing classifier performance in imbalanced prediction tasks;
- 4) Comorbidity structures, useful for disease correlation analysis and simulation studies [34, 36].

Because MedGAN directly addresses the discrete nature of clinical coding systems, it occupies a distinct position among GAN variants. Whereas WGAN and CTGAN focus on tabular mixed-type data, MedGAN is optimized for binary, sparse, high-dimensional vectors, making it particularly valuable for applications involving ICD codes, procedure codes, medication vectors, and other structured EHR fields.

Although GAN variants have substantially expanded the applicability of synthetic data generation in healthcare, each architecture also presents important limitations and is best suited to particular clinical scenarios. The original GAN is conceptually foundational but is often affected by training instability and mode collapse, especially in high-dimensional medical settings. cGANs improve control over generated outputs and are useful when class labels or disease categories are available, although their performance depends on the quality and balance of the conditioning

information. DCGAN is highly effective for image-based applications because of its convolutional design, but it is less suitable for structured tabular records. WGAN improves training stability and is often preferred for heterogeneous or imbalanced healthcare datasets, although it introduces added training complexity and computational demands. More specialized models such as CTGAN and MedGAN are better aligned with mixed-type tabular data and sparse clinical records, respectively, but they may still struggle with rare event preservation or complex inter-feature dependencies. Therefore, the selection of a GAN variant should be guided by the nature of the target data, the downstream medical objective, and the balance required between fidelity, stability, and interpretability.

Table 2 summarizes the key limitations of common GAN variants and their most appropriate healthcare application scenarios, providing a practical guide to model selection across different medical data types.

#### 4.2. Applications across disease prediction

GANs have increasingly been adopted in healthcare research as prominent tools for synthesizing multimodal data, including images, tabular records, and time-series signals. Disease prediction has, in turn, benefited from the diversity of GAN architectures and the adaptability of adversarial training across different data types.

Table 2  
Limitations of common GAN variants and their suitable healthcare scenarios

GAN variant	Main limitation	Most suitable healthcare scenario
Original GAN	Prone to training instability and mode collapse; limited control over outputs	Baseline or proof-of-concept synthetic data generation
cGAN	Performance depends on high-quality labels and can be affected by class imbalance	Class-specific generation, minority-class augmentation, labeled disease categories
DCGAN	Less suitable for tabular or mixed-type clinical variables	Medical imaging, such as MRI, CT, X-ray, ultrasound, and pathology images
WGAN/WGAN-GP	More complex optimization and higher tuning requirements	Structured or heterogeneous healthcare data, including EHR and imbalanced tabular datasets
CTGAN	May struggle with rare categories and highly complex dependencies	Mixed-type tabular healthcare datasets and disease prediction records
MedGAN	Less suitable for image data and complex temporal structures	Sparse EHRs, diagnosis codes, and multi-label clinical records

#### 4.2.1. Tumor detection and classification

In oncology, GANs are widely used because medical imaging datasets are often limited and class-imbalanced. cGANs and DCGANs can generate synthetic histopathology slides or brain tumor MRI scans to augment training datasets and improve classifier performance [1, 4]. These augmented samples have been shown to reduce overfitting and improve sensitivity in tumor classification tasks.

#### 4.2.2. Chronic disease prediction

Structured tabular datasets are more common than large imaging datasets, mostly in cases of cardiovascular disease and diabetes. WGANs and CTGANs are particularly suitable for modeling such data because they can capture correlations that might exist between heterogeneous variables (e.g., lab results, demographics, and lifestyle factors) without collapsing the modes [28]. Risk assessment models had been developed on these synthetic datasets; hence, there is an indirect stratification of patients based on risk, which may still be applicable in low-resource settings.

#### 4.2.3. Rare disease studies

With this approach, patients suffering from rare diseases and having a few cases can be very much at a disadvantage. cGANs and CTGANs generate patient records for rare conditions such as pulmonary hypertension or genetic disorders realistically, oversampling minority classes [1, 32]. This increases the generalizability of the classifiers and supports early diagnostic tool developments.

#### 4.2.4. Image modality translation

Multimodal imaging plays an important role in medical practice. GANs, particularly DCGANs and pix2pix form cGANs, have been used for modality translation tasks for generating CT images from MRI and MRI from Positron Emission Tomography (PET), thereby reducing additional imaging procedures and supporting multimodal prediction pipelines [10]. These methods enhance cross-modality tumor detection and facilitate prognostic modeling.

#### 4.2.5. Time-series disease monitoring

Besides image and tabular data, GANs have also been extended toward time-series signals, ECG, and Electroencephalography (EEG) data. By creating realistic sequences, adversarial systems help in training deep models for arrhythmia detection, seizure prediction, and sleep-stage classification [36]. The value of such methods is higher for continuous disease monitoring, where it becomes costly to obtain large, annotated datasets.

#### 4.2.6. Comorbidity and disease progression analysis

MedGAN and its hybrid derivatives have been strong performers in synthetic EHRs, where mixtures of diagnostic codes and medications must be accounted for together [34, 36]. Synthetic EHRs were also used for predicting disease progression (such as complications arising from diabetes) and comorbidity patterns (e.g., hypertension with obesity), supporting clinical decision-making without disclosing sensitive patient information.

### 4.3. Summary

GAN architectures for healthcare retain a shared foundation—a generator synthesizing data and a discriminator or critic assessing its quality, but the key innovations lie in domain-specific challenges.

- 1) cGANs provide control by conditioning on labels so that one can target the minority disease classes for synthesis.
- 2) DCGANs use convolutional filters to detect spatial hierarchies, thus fitting well to medical images.
- 3) WGANs offer the advantage of providing a more stable criterion for learning, especially in areas such as structured/tabular data like EHRs.
- 4) CTGAN and MedGAN further extend this landscape by addressing mixed-type tabular records and sparse discrete clinical codes, respectively.

GANs have been successfully applied across medical applications for image augmentation in tumor detection, risk prediction for chronic diseases, oversampling for rare diseases, multimodality image translation, time-series signal generation, and EHR synthesis for comorbidity analysis. Architectural innovation in tandem with the data modality images (DCGAN/cGAN), structured EHRs (WGAN/CTGAN/MedGAN), and signals (time-series GANs) has allowed researchers to extend the inference power of machine learning models and provide remedies for data scarcity and privacy-constrained environments.

## 5. Evaluation of GAN-Generated Data

GANs represent a significant advancement in generative modeling because they can learn complex data distributions and produce realistic synthetic data. However, no single evaluation metric is sufficient to assess GAN performance comprehensively. Instead, the literature supports a multidimensional evaluation framework based on quality, utility, privacy, and computational considerations. Recent systematic reviews further highlight the growing role of GANs in medical image reconstruction and enhancement tasks. Across a wide range of imaging modalities, GAN-based approaches have been applied to improve image quality, reconstruct incomplete data, and support diagnostic workflows [37]. These studies also show that GAN performance is influenced by factors such as network architecture, dataset characteristics, and loss function design, while emphasizing ongoing challenges related to training stability and evaluation consistency in medical imaging applications [37].

### 5.1. Quality metrics

The metrics indicate the degree of similarity between generated data and real data distributions in terms of realism and diversity. Typically, these metrics are used to compare different GAN architectures and to monitor training progress.

**Statistical similarity:** This metric computes the difference in the statistical properties between the synthetic and the real datasets. For tabular data, such differences may be checked in terms of, for example, column-wise distributions, means, and correlation matrices. However, for images, a more abstract statistical comparison is generally employed in a feature space.

**Distribution overlap:** These measures assess how closely the distribution of generated data matches that of the real data. Commonly used metrics include:

- 1) FID [38]
- 2) IS [39]

While metrics such as FID explicitly compare real and generated data distributions, others such as IS evaluate output quality and diversity.

**Wasserstein distance:** The Wasserstein distance is a metric used in optimal transport theory. It measures the weighted cost of transforming one probability distribution into another according to some cost function. With the Wasserstein distance, one gets meaningful, continuous gradients even when the distributions have zero overlap, unlike the Jensen–Shannon divergence used in the original GAN. This feature made the learning process in the WGAN much more stable. A lower distance indicates better alignment of the two distributions [3].

Recent review work has further emphasized that GAN evaluation remains an open challenge, as commonly used automated metrics cannot fully capture both objective fidelity and subjective quality [40]. This is highly relevant in healthcare settings, where synthetic outputs may appear statistically plausible yet still fail to reflect clinically meaningful structure. As a result, more balanced evaluation frameworks that combine automated scores with human-informed assessment may be necessary when judging the quality and usability of medical synthetic data.

## 5.2. Utility metrics

These metrics assess the functional value of synthetic data, specifically whether it can effectively replace real data in practical applications.

**Improvement in downstream prediction tasks:** One of the typical approaches is to train a machine learning model on synthetic data and use a separate, real-world test set to evaluate its performance. If the model performs well, then the synthetic data are deemed highly useful. This approach often constitutes a more task-specific and practical evaluation than general quality metrics.

**AUC and F1-score:** These are just standard metrics used for evaluating the performance of a classifier created on synthetic data. The area under the receiver operating characteristic curve (AUC-ROC) checks the ability of a classifier to classify an instance into one or another class, whereas the F1-score gives a balanced estimate of precision and recall over imbalanced datasets [39].

## 5.3. Privacy metrics

Privacy evaluation is paramount for applications involving sensitive data such as health records. These metrics quantify the risk of information leakage of datasets in training.

**Re-identification risk:** This is a chance that a synthetic data record is linked to a real individual who was part of the original training data. Researchers attempt to “re-identify” records using known information or adversarial attacks [41].

**Differential privacy (DP):** A formal, mathematical framework for quantifying privacy guarantees. GANs can be trained with differentially private mechanisms that limit the impact of any single training record on the success of the final generated distribution, thus producing a provable privacy guarantee [19].

**Membership inference:** This attack attempts to determine whether a point of data was used by the machine learning model in its training. If a membership inference attack succeeds at a reasonably high frequency, then the model has memorized exact training data and may indicate privacy leakage [42].

## 5.4. Computational considerations

These metrics are centered on attributes that could be observed during each GAN training session, which are considered quite challenging.

**Training stability:** In GAN training methods, instability is commonly observed, underscored by the adversarial minimax game shared between the generator and discriminator. Metrics or observations attributing stable and non-collapsing loss curves to a good behavior of GANs have been offered [43]. Typical modes of failure, from mode collapse (loss of diversity on the generator’s side) to oscillations of the loss, are key signs of instability.

**Convergence time:** Convergence time is the amount of time it takes for the GAN to achieve a stable equilibrium, where it can produce high-quality samples. While difficult to define mathematically, faster convergence practically requires quicker model development.

## 6. Discussion

The findings of this narrative review highlight the growing relevance of GAN-based synthetic data generation across diverse areas of disease prediction and healthcare analytics. While each GAN variant addresses specific challenges such as modeling spatial complexity in medical images, handling heterogeneous tabular features, or stabilizing training on limited datasets, their collective evolution demonstrates a clear trend toward architectures tailored for clinical data scarcity, privacy preservation, and class imbalance. The reviewed studies consistently show that synthetic data, when generated using suitably chosen GAN models, can enhance model robustness, support downstream diagnostic tasks, and mitigate constraints associated with real-world clinical datasets. However, despite these advancements, challenges remain regarding evaluation standards, clinical validity, and the interpretability of synthetic data, all of which warrant closer examination in subsequent subsections.

### 6.1. Trends and insights

The application of GANs in disease prediction highlights serious, clear-cut trends whereby GAN architectures are reviewed concerning the data modality. DCGANs remain the most widely used model for generating medical imagery of high quality. The convolutional architecture enables them to effectively capture spatial hierarchies and fine-grained details present in radiology images, pathological images, and dermatology images. With a wealth of publications and open-source implementations around the concept of DCGANs, their superiority in this realm is well established, making them the preferred toolkit for image augmentation and synthesis.

Conversely, WGANs and their variants (e.g., WGAN-GP) have gained much interest in handling tabular data, such as EHRs and lab values. Regular GANs have great difficulties in managing the heterogeneity of tabular data, which contains combinations of continuous, categorical, and discrete variables. WGANs, by using the Wasserstein distance as a loss function, allow a more stable training process and avoid mode collapse that poses an obstacle to the generation of a wide and representative variety of patient records [3]. In the case of difficult mixed-type datasets, WGANs are preferably chosen because of this kind of stability, as concluded by Ahmed et al. [14].

Domain-specific GANs, such as CTGAN and MedGAN, attempt to solve the peculiarities of tabular healthcare data: long-tail distributions and multi-label conditions. As a result, they are able to control their impulses and are skilled at solving problems that general-purpose GANs struggle with. Earlier reviews that focused primarily on GAN applications in either imaging or tabular data alone [1, 10], whereas this review integrates findings across multiple modalities including tabular, time-series, and imaging. This review also explicitly compares key GAN architectures side by side, highlighting WGAN's distinct advantages for structured health data. The synthesis presented here can guide researchers in selecting GAN architectures tailored to their data modality: DCGAN for high-dimensional imaging tasks, WGAN for heterogeneous tabular EHRs, and CTGAN/MedGAN for mixed-type records requiring fine-grained control over discrete attributes.

## 6.2. Strengths and weaknesses of GAN approaches

Despite their promise, the implementation of GANs in healthcare is loaded with many implementation issues, posing significant challenges. Knowing their strengths and weaknesses is an operative mechanism for their deployment with utmost responsibility.

### 6.2.1. Strengths

**Data augmentation:** GANs can generate large volumes of synthetic data. This is particularly useful for rare diseases or conditions where patient data are scarce, as GANs can almost generate realistic examples. This ability helps to address the "small data" problem, allowing the training of more robust and generalized machine learning models that would otherwise risk overfitting on a limited dataset [44].

**Privacy preservation:** One of the most compelling advantages of GANs in the healthcare industry is that they generate synthetic data with statistical maturity like that of real data but without compromising any original patient information. Hence, researchers would be able to freely share realistic datasets for collaborative studies and algorithm development while drastically reducing patient privacy concerns. This bridges the gap that arises from the legal and ethical issues tied to sharing very sensitive data like EHRs [36].

### 6.2.2. Weaknesses

**Training instability:** Training the GAN is an infamously difficult task. The instability in the adversarial minimax game between the generator and the discriminator often leads to the very common modes of failure. One can expect the loss functions to bounce around between local minima with no convergence, never telling when the model has become "good enough." Such instability then makes the training difficult, leading at times to extensive hyperparameter tuning or the use of specialized architectures to handle it [39].

**Model collapse:** Another major weakness of many GAN architectures is mode collapse. During mode collapse, the generator learns to produce only a few selected samples out of the real data distribution and disregards most of the data. For example, a GAN trained on a rich set of medical images may learn only to generate images of tumors of a single type and ignore others. The resulting data thus become very poor candidates for training robust models [39].

**Privacy concerns:** With the potential to preserve privacy, GANs create an additional privacy attack vector. Research,

therefore, has shown that certain kinds of GANs are in fact vulnerable to privacy attacks such as membership inference attacks. An attacker may gain confidence in whether someone was used as training data for the GAN. This calls for even more mitigation tools, such as DP, which provide stronger guarantees to the system [42].

By framing GAN-generated synthetic data within the context of privacy, fairness, and clinical acceptability, this review adds a layer of ethical and regulatory awareness often absent from purely technical surveys.

## 6.3. Implications for disease prediction

Disease predictive models developed by AI require synthetic data produced by GANs for model generalization and fairness. Providing synthetic data from augmented, small, and imbalanced datasets or those under private restrictions could nearly improve a model's ability to generalize to new, unseen patient populations. This is particularly important for rare diseases or conditions, where real data are virtually non-existent. A larger training set with greater variety helps avoid overfitting to limited examples [44]. On the other hand, synthetic data might also provide a solution to ameliorate unfairness from data bias. Real-life medical datasets often have demographic groups that are underrepresented (say, certain ethnicities or gender groups), which in turn makes the existing models perform poorly on those populations. GANs can artificially generate data from these underrepresented groups upon which real datasets shall be leveled, thus making sure that models trained for all patients are more unbiased and more accurate [45]. If there is a way to reduce algorithmic bias, then that is one aspect that goes a fair way toward equitable healthcare outcomes. The cross-modal perspective taken here points toward hybrid models (e.g., GAN-diffusion or GAN-federated approaches) as promising next steps for scalable, privacy-preserving synthetic data generation in healthcare.

## 6.4. Limitations of this narrative review

This review adopts a narrative rather than a systematic approach, which has implications for completeness and reproducibility. Although multiple databases were searched and recent surveys were consulted to identify influential work on GAN-based synthetic data generation in healthcare, the search was not exhaustive, and some relevant studies may have been missed. No formal study-level quality appraisal or risk-of-bias assessment was performed, as the emphasis was on architectural themes, evaluation practices, and application patterns rather than on comparing individual study outcomes. In addition, the field of generative AI is evolving rapidly, so papers published after the search window (up to early 2025) are not captured. These limitations mean that the findings should be interpreted as a conceptual and methodological synthesis rather than as a definitive or exhaustive account of all GAN applications in disease prediction.

## 6.5. Future research directions

Although the applications of GANs have been demonstrated for disease prediction, medical imaging, and synthetic EHR generation, certain open directions remain critical for progressing the field.

**Hybrid GANs and GAN+ diffusion models:** One of the prominent trends in the research is the development of hybrid generative

market systems, where GANs are combined with other architectures to leverage their complementary strengths. Perhaps the most notable example would be the combination of GANs with diffusion models [46]. While GANs can quickly synthesize high-fidelity images, they often suffer from unstable training procedures and mode collapse, thereby restricting output diversity. Diffusion models, by contrast, are famous for their stable training and the ability to generate a wide variety of high-quality and diverse samples. However, they come with a high computational cost and exhibit a slow image generation speed. Hybrid models such as SupResDiffGAN aim to inherit these advantages in producing quality images with high stability and speed in generation for tasks like super-resolution of medical imagery.

**Federated learning (FL):** Data silos and privacy regulations present major barriers; therefore, GANs and FL integration is a huge research field. The FL paradigm prescribes that several institutions (hospitals) collaborate in training the GAN model without sharing their raw patient data. Each institution trains a local discriminator using its own data. An update is transmitted to a central server where the global generator is trained. Such a general structure allows for a robust and generalized model to be trained on a larger and more diverse dataset, which complies with stricter patient privacy and security regulations. Research works on making the frameworks more efficient and secure, some of which prevent the leakage of the generator from the server to clients [41].

**Standardized evaluation frameworks:** The absence of a universally accepted evaluation framework remains a major challenge in GAN-based synthetic data research. Future work should focus on developing standardized evaluation frameworks that enable results to be reported consistently and compared reliably across studies and application domains. Such frameworks should extend beyond traditional image-quality metrics, such as FID and IS, to include broader dimensions such as downstream utility, privacy preservation, and clinical acceptability. They should also incorporate rigorous assessments of re-identification risk and clinician-blind evaluation to determine whether generated images are clinically indistinguishable from real data. Establishing such comprehensive evaluation standards would strengthen trust, transparency, and confidence in the clinical use of GAN-generated data [47].

## 7. Ethical Considerations

The ethical use of GANs in healthcare is complex, involving critical components of privacy, fairness, and transparency that are essential for achieving clinical acceptability. Even though GANs are more suitable for generating synthetic data for research purposes under privacy protection, they remain vulnerable to attacks. Some of the prime concerns raised by researchers relate to re-identification and membership inference attacks, where an adversary might identify whether data belonging to an individual was used to train a GAN. Accordingly, the dominant research trends have centered around integrating DP, which means that noise added during training ensures that no single data point can significantly influence the final model. As a result, it provides strong privacy protection [19].

GANs may also exacerbate biases found in genuine healthcare datasets, potentially leading to unfair outcomes for minorities [45]. Ensuring fairness requires devising bias-aware GAN architectures and then auditing the synthetic data so that it adequately and fairly represents the disparate patient demographics. Moreover, existing reviews emphasize that successful adoption of machine

learning in healthcare depends not only on technical performance but also on ethical safeguards, transparency, and clinician trust, all of which are critical when deploying synthetic data systems [48].

Beyond statistical fidelity and downstream predictive performance, recent studies increasingly emphasize the importance of clinical validity through domain-expert assessment. Clinical validity in this context refers not only to preserving overall data distributions but also to ensuring that generated records, images, or temporal patterns remain medically plausible, internally coherent, and reflective of meaningful disease relationships. Emerging evaluation approaches therefore advocate complementing algorithmic metrics with clinician-centered validation, such as blind expert review and assessment of diagnostic realism. For example, recent work has proposed hybrid evaluation frameworks that combine quantitative metrics with expert assessment to ensure that synthetic data are clinically relevant and usable in real-world settings [49]. More broadly, recent reviews of generative AI in healthcare emphasize that evaluation must extend beyond technical performance to include real-world clinical utility, safety, explainability, and responsible integration into healthcare workflows [50]. This is particularly important in healthcare, where synthetic samples may appear statistically convincing while still failing to capture subtle but clinically important dependencies.

These concerns are closely linked to one of the most prominent barriers to GAN adoption: the lack of transparency and interpretability. Most GANs are still regarded as black-box systems, which can reduce clinician confidence in models trained on synthetic data. It is therefore essential to develop validation and interpretability strategies that demonstrate not only the realism of GAN-generated data but also their clinical plausibility and reliability [47]. Addressing these ethical and practical limitations will be critical to building the confidence of both medical practitioners and patients in the use of GAN-based synthetic data.

## 8. Conclusion

With the emergence of GANs, synthetic data generation has become a feasible and increasingly valuable approach for disease prediction research. GANs help address limited data availability, class imbalance, and privacy restrictions by enabling the generation of realistic medical images, EHRs, and laboratory data. In robust predictive modeling, especially in populations where collecting and distributing large-scale datasets is nearly impossible, the value of GAN-generated synthetic datasets is unmatched; they offer diverse and representative data. Hence, from these applications, GAN methods aid the advancement of machine learning in healthcare while somewhat sidestepping the ethical and logistical issues pertinent to real patient data. This review consolidates and evaluates the rapidly expanding literature on GAN-based synthetic data generation for disease prediction, providing researchers with a single reference point for architectural choices, evaluation practices, and ethical considerations.

Among the different types of GANs, the WGANs (WGAN and WGAN-GP) hold merit in handling tabular medical data that are often heterogeneous, sparse, and imbalanced. These models replace the Jensen–Shannon divergence with the Wasserstein distance, which carries with it smoother gradients with respect to critic parameters and more stable convergence behavior, hence reduced mode collapse between other models. Stability is important in generating structured EHRs and laboratory results, as traditional GANs often are not able to grasp the full depth of the interaction among the categories and continuous variables.

Hence, WGAN-based methods have thus far shown promising results, for example, in synthesizing patient records that retain statistical soundness and thus would be preferable for downstream prediction tasks. By highlighting the strengths and weaknesses of specific GAN variants, the paper offers practical insights for investigators seeking to enhance data availability, model robustness, and fairness in medical AI.

While the work and use of GANs in healthcare AI create good prospects, they simultaneously bring challenges. Future work may explore hybrid GANs and diffusion-based approaches to improve data fidelity, as well as federated and privacy-preserving GANs, allowing for secure collaboration across institutions. Future progress will depend on community-wide efforts to establish standardized benchmarks and cross-institutional collaborations, particularly those integrating GANs with FL, to ensure both scientific rigor and privacy compliance.

Meanwhile, the community needs to address reproducibility, evaluation, and ethical deployment, ensuring that synthetic data supports not only improved prediction performance but also patient privacy and clinical soundness. Striking a balance between innovation and responsibility would be instrumental in unleashing the full potential of synthetic data generated by GANs in support of next-generation disease prediction and healthcare analytics. As highlighted in prior work on machine learning integration in healthcare, ethical and operational considerations must complement technical advancements to ensure clinical trust and adoption [48]. If such steps are taken, GAN-generated synthetic data can become a cornerstone of next-generation disease prediction, enabling equitable, privacy-preserving, and clinically reliable AI tools across diverse healthcare settings.

## Acknowledgment

The authors are grateful to the academics and research assistants at York St John University for their support in sourcing and collecting the review papers from various databases, which substantially contributed to this study.

## Funding Support

This research is supported and funded by York St. John University, Research School under QR Project QR24-25-22, AY24-25.

## Ethical Statement

This study does not contain any studies with human or animal subjects performed by any of the authors.

## Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

## Data Availability Statement

No new primary datasets were generated or analyzed in this study. This article is a narrative literature review based on previously published studies. The datasets discussed in the manuscript are available from the original publications and associated repositories cited below.

The structured EHR datasets reviewed in this article are available at <https://doi.org/10.1145/3636424>, reference number [6],

and <https://doi.org/10.1016/j.compbmed.2023.107188>, reference number [7]. The multicenter diabetic EHR dataset is available at <https://doi.org/10.1016/j.compbmed.2023.107188>, reference number [7]. Time-series and longitudinal healthcare datasets, including ECG-related datasets, are available at <https://doi.org/10.1038/s41746-024-01409-w>, reference number [8], and <https://doi.org/10.1109/icoa62581.2024.10753742>, reference number [9]. Medical imaging datasets used in GAN-based synthetic image generation, augmentation, and reconstruction studies are available at <https://doi.org/10.1016/j.compbmed.2025.109834>, reference number [10]; <https://doi.org/10.1148/radiol.232471>, reference number [11]; <https://doi.org/10.1016/j.compbmed.2022.105382>, reference number [20]; <https://doi.org/10.1109/tiptekno63488.2024.10755233>, reference number [26]; [https://doi.org/10.1007/978-981-96-1185-0\\_38](https://doi.org/10.1007/978-981-96-1185-0_38), reference number [29]; and <https://doi.org/10.1016/j.compbmed.2025.110094>, reference number [37].

The Breast Cancer Wisconsin (Diagnostic), Lung Cancer Patient, and Fetal Cardiotocography (CTG) datasets are available at <https://doi.org/10.1007/s41060-025-00816-w>, reference number [14], and [https://github.com/Halal-Abdulrahman-Ahmed/MedSynth\\_GANVariants](https://github.com/Halal-Abdulrahman-Ahmed/MedSynth_GANVariants). The Breast Cancer Wisconsin (Diagnostic) dataset is also available from the UCI Machine Learning Repository at <https://archive.ics.uci.edu/dataset/17/breast+cancer+wisconsin+diagnostic>, and the Fetal Cardiotocography dataset is available from the UCI Machine Learning Repository at <https://archive.ics.uci.edu/dataset/193/cardiocography>. The cardiovascular dataset was obtained from King Faisal Specialist Hospital and Research Centre and is subject to access restrictions as stated in the original article at <https://doi.org/10.3390/s24237673>, reference number [33].

The brain MRI dataset is available at <https://brain-development.org/ixi-dataset/>. Additional synthetic EHR and longitudinal health record datasets discussed in relation to MedGAN and mixed-type EHR generation are available at <https://proceedings.mlr.press/v68/choi17a.html> and <https://doi.org/10.1038/s41746-023-00834-7>, reference number [36]. Where datasets are not separately hosted in an open repository, access conditions and availability are governed by the original cited publications.

## Author Contribution Statement

**Swathi Ganesan:** Conceptualization, Methodology, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition. **Nalinda Somasiri:** Visualization, Supervision, Project administration, Funding acquisition.

## References

- [1] Pezoulas, V. C., Zaridis, D. I., Eugenia, M., Androutsos, C., Apostolidis, K., Tachos, N. S., & Fotiadis, D. I. (2024). Synthetic data generation methods in healthcare: A review on open-source tools and methods. *Computational and Structural Biotechnology Journal*, 23, 2892–2910. <https://doi.org/10.1016/j.csbj.2024.07.005>
- [2] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., . . . , & Bengio, Y. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144. <https://doi.org/10.1145/3422622>
- [3] Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, 70, 214–223.

- [4] Akpınar, M. H., Sengur, A., Salvi, M., Seoni, S., Faust, O., Mir, H., . . . , & Acharya, U. R. (2025). Synthetic data generation via generative adversarial networks in healthcare: A systematic review of image- and signal-based studies. *IEEE Open Journal of Engineering in Medicine and Biology*, 6, 183–192. <https://doi.org/10.1109/OJEMB.2024.3508472>
- [5] Adams, T., Birkenbihl, C., Otte, K., Ng, H. G., Rieling, J. A., Näher, A.-F., . . . , & Fröhlich, H. (2025). On the fidelity versus privacy and utility trade-off of synthetic patient data. *iScience*, 28(5), 112382. <https://doi.org/10.1016/j.isci.2025.112382>
- [6] Ghosheh, G. O., Li, J., & Zhu, T. (2024). A survey of generative adversarial networks for synthesizing structured electronic health records. *ACM Computing Surveys*, 56(6), 147. <https://doi.org/10.1145/3636424>
- [7] Bernardini, M., Doinychko, A., Romeo, L., Frontoni, E., & Amini, M.-R. (2023). A novel missing data imputation approach based on clinical conditional Generative Adversarial Networks applied to EHR datasets. *Computers in Biology and Medicine*, 163, 107188. <https://doi.org/10.1016/j.combiomed.2023.107188>
- [8] Loni, M., Poursalim, F., Asadi, M., & Gharehbaghi, A. (2025). A review on generative AI models for synthetic medical text, time series, and longitudinal data. *npj Digital Medicine*, 8(1), 281. <https://doi.org/10.1038/s41746-024-01409-w>
- [9] Sraitih, M., El Hassani, A. H., Jabrane, Y., & Andres, E. (2024). Generative adversarial networks in ECG signal denoising: A survey. In *2024 10th International Conference on Optimization and Applications*, 1–7. <https://doi.org/10.1109/icoa62581.2024.10753742>
- [10] Ibrahim, M., Khalil, Y. A., Amirrajab, S., Sun, C., Breeuwer, M., Pluim, J., . . . , & Dumontier, M. (2025). Generative AI for synthetic data across multiple medical modalities: A systematic review of recent developments and challenges. *Computers in Biology and Medicine*, 189, 109834. <https://doi.org/10.1016/j.combiomed.2025.109834>
- [11] Koetzier, L. R., Wu, J., Mastrodicasa, D., Lutz, A., Chung, M., Koszek, W. A., . . . , & Willeminck, M. J. (2024). Generating synthetic data for medical imaging. *Radiology*, 312(3), e232471. <https://doi.org/10.1148/radiol.232471>
- [12] Aggarwal, A., Mittal, M., & Battineni, G. (2021). Generative adversarial network: An overview of theory and applications. *International Journal of Information Management Data Insights*, 1(1), 100004. <https://doi.org/10.1016/j.ijmei.2020.100004>
- [13] Odena, A., Olah, C., & Shlens, J. (2017). Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning*, 70, 2642–2651.
- [14] Ahmed, H. A., Nepomuceno, J. A., Vega-Márquez, B., & Nepomuceno-Chamorro, I. A. (2025). Synthetic data generation for healthcare: Exploring generative adversarial networks variants for medical tabular data. *International Journal of Data Science and Analytics*, 20(6), 5739–5754. <https://doi.org/10.1007/s41060-025-00816-w>
- [15] Petrakos, N. Z., Moodie, E. E. M., & Savy, N. (2025). A framework for generating realistic synthetic tabular data in a randomized controlled trial setting. *Statistics in Medicine*, 44(18–19), e70227. <https://doi.org/10.1002/sim.70227>
- [16] Shi, R., Wang, Y., Du, M., Shen, X., Chang, Y., & Wang, X. (2025). *A comprehensive survey of synthetic tabular data generation*. arXiv. <https://doi.org/10.48550/arXiv.2504.16506>
- [17] Ding, H., Huang, N., Wu, Y., & Cui, X. (2025). Improving imbalanced medical image classification through GAN-based data augmentation methods. *Pattern Recognition*, 166, 111680. <https://doi.org/10.1016/j.patcog.2025.111680>
- [18] Yoon, J., Drumright, L. N., & van der Schaar, M. (2020). Anonymization through data synthesis using generative adversarial networks (ADS-GAN). *IEEE Journal of Biomedical and Health Informatics*, 24(8), 2378–2388. <https://doi.org/10.1109/jbhi.2020.2980262>
- [19] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318. <https://doi.org/10.1145/2976749.2978318>
- [20] Chen, Y., Yang, X.-H., Wei, Z., Heidari, A. A., Zheng, N., Li, Z., . . . , & Guan, Q. (2022). Generative adversarial networks in medical image augmentation: A review. *Computers in Biology and Medicine*, 144, 105382. <https://doi.org/10.1016/j.combiomed.2022.105382>
- [21] Singh, T. P., Singh, T. M., & G, D. K. (2025). Using conditional GAN for facial recognition enhancement. In *2025 International Conference on Intelligent Computing and Control Systems*, 1511–1515. <https://doi.org/10.1109/ICICCS65191.2025.10985656>
- [22] Jenkins, J., & Roy, K. (2024). Exploring deep convolutional generative adversarial networks (DCGAN) in biometric systems: A survey study. *Discover Artificial Intelligence*, 4(1), 42. <https://doi.org/10.1007/s44163-024-00138-z>
- [23] Ganesan, S., Pokhrel, S., & Somasiri, N. (2025). Navigating the ethics and legality of deepfake technology: Advancements, implications and responsible deployment. In K. Ramamurthy, S. Kulanthaivelu, S. Anand, & T. Murugan (Eds.), *Generative AI unleashed: Advancements, transformative applications and future frontiers* (pp. 159–176). Institution of Engineering & Technology. [https://doi.org/10.1049/pbpc076e\\_ch9](https://doi.org/10.1049/pbpc076e_ch9)
- [24] Öcal, A., & Özbakır, L. (2021). Supervised deep convolutional generative adversarial networks. *Neurocomputing*, 449, 389–398. <https://doi.org/10.1016/j.neucom.2021.03.125>
- [25] Shi, Z., Teng, J., Zheng, S., & Guo, K. (2023). Exploring the effects of various generative adversarial networks techniques on image generation. In *2023 IEEE 11th Joint International Information Technology and Artificial Intelligence Conference*, 1796–1799. <https://doi.org/10.1109/itaic58329.2023.10409102>
- [26] Yaman, B., Yılmaz, O. Z., Darici, M. B., & Ozmen, A. (2024). Age classification by WGAN brain MR image augmentation. In *2024 Medical Technologies Congress*, 1–4. <https://doi.org/10.1109/tiptekno63488.2024.10755233>
- [27] Ebner, P. P., & Eltelt, A. (2020). *Audio inpainting with generative adversarial network*. arXiv. <https://doi.org/10.48550/arXiv.2003.07704>
- [28] Ganesan, S., & Somasiri, N. (2025). Mitigating data scarcity in healthcare through Wasserstein generative adversarial network—A case study in medical application. In *2025 10th International Conference on Machine Learning Technologies*, 278–286. <https://doi.org/10.1109/ICMLT65785.2025.11193164>
- [29] Senthil Murugan, K. R., & Swapna, T. R. (2025). Synthetic data generation of chest X-ray images using modified WGAN GP algorithm. In *Data Science and Applications: Proceedings of ICDSA 2024*, 5, 493–505. [https://doi.org/10.1007/978-981-96-1185-0\\_38](https://doi.org/10.1007/978-981-96-1185-0_38)
- [30] Majeed, A., & Hwang, S. O. (2025). Moving conditional GAN close to data: Synthetic tabular data generation

- and its experimental evaluation. *IEEE Transactions on Big Data*, 11(3), 1188–1205. <https://doi.org/10.1109/TBDDATA.2024.3442534>
- [31] Asimopoulos, D. C., Radoglou-Grammatikis, P., Makris, I., Mladenov, V., Psannis, K. E., Goudos, S., & Sarigiannidis, P. (2023). Breaching the defense: Investigating FGSM and CTGAN adversarial attacks on IEC 60870-5-104 AI-enabled intrusion detection systems. In *Proceedings of the 18th International Conference on Availability, Reliability and Security*, 52, 1–8. <https://doi.org/10.1145/3600160.3605163>
- [32] Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., & Kasneci, G. (2024). Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 35(6), 7499–7519. <https://doi.org/10.1109/TNNLS.2022.3229161>
- [33] Alqulaity, M., & Yang, P. (2024). Enhanced conditional GAN for high-quality synthetic tabular data generation in mobile-based cardiovascular healthcare. *Sensors*, 24(23), 7673. <https://doi.org/10.3390/s24237673>
- [34] Choi, E., Biswal, S., Malin, B., Duke, J., Stewart, W. F., & Sun, J. (2017). Generating multi-label discrete patient records using generative adversarial networks. In *Proceedings of the 2nd Machine Learning for Healthcare Conference*, 68, 286–305.
- [35] Guo, K., Chen, J., Qiu, T., Guo, S., Luo, T., Chen, T., & Ren, S. (2023). MedGAN: An adaptive GAN approach for medical image generation. *Computers in Biology and Medicine*, 163, 107119. <https://doi.org/10.1016/j.compbiomed.2023.107119>
- [36] Li, J., Cairns, B. J., Li, J., & Zhu, T. (2023). Generating synthetic mixed-type longitudinal electronic health records for artificial intelligent applications. *npj Digital Medicine*, 6(1), 98. <https://doi.org/10.1038/s41746-023-00834-7>
- [37] Hussain, J., B ath, M., & Ivarsson, J. (2025). Generative adversarial networks in medical image reconstruction: A systematic literature review. *Computers in Biology and Medicine*, 191, 110094. <https://doi.org/10.1016/j.compbiomed.2025.110094>
- [38] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S. (2017). GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *31st Conference on Neural Information Processing System*.
- [39] Borji, A. (2022). Pros and cons of GAN evaluation measures: New developments. *Computer Vision and Image Understanding*, 215, 103329. <https://doi.org/10.1016/j.cviu.2021.103329>
- [40] Shah, A., Rakesh, N., Gulhane, M., Agrawal, P., Kaur, M., & Dixit, S. (2025). Evaluating generative models: A review of challenges and opportunities in assessment of generative adversarial networks. In *2025 2nd International Conference on Computational Intelligence, Communication Technology and Networking*, 12–16. <https://doi.org/10.1109/CICTN64563.2025.10932502>
- [41] Jordon, J., Yoon, J., & van der Schaar, M. (2019). PATE-GAN: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*, 1–21.
- [42] Zhang, Z., Yan, C., & Malin, B. A. (2022). Membership inference attacks against synthetic health data. *Journal of Biomedical Informatics*, 125, 103977. <https://doi.org/10.1016/j.jbi.2021.103977>
- [43] Mtetwa, J. T., Ogudo, K. A., & Pudaruth, S. (2025). Adaptive gradient penalty for Wasserstein GANs: Theory and applications. *Mathematics*, 13(16), 2651. <https://doi.org/10.3390/math13162651>
- [44] Sandfort, V., Yan, K., Pickhardt, P. J., & Summers, R. M. (2019). Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Scientific Reports*, 9(1), 16884. <https://doi.org/10.1038/s41598-019-52737-x>
- [45] Chen, C., Qin, C., Ouyang, C., Li, Z., Wang, S., Qiu, H., . . . , & Rueckert, D. (2022). Enhancing MR image segmentation with realistic adversarial data augmentation. *Medical Image Analysis*, 82, 102597. <https://doi.org/10.1016/j.media.2022.102597>
- [46] Islam, S., Aziz, M. T., Nabil, H. R., Jim, J. R., Mridha, M. F., . . . , & Shin, J. (2024). Generative adversarial networks (GANs) in medical imaging: Advancements, applications, and challenges. *IEEE Access*, 12, 35728–35753. <https://doi.org/10.1109/ACCESS.2024.3370848>
- [47] Schuit, G., Parra, D., & Besa, C. (2026). Perceptual evaluation of GANs and diffusion models for generating X-rays. In *Human-AI Collaboration: First International Workshop*, 93–101. [https://doi.org/10.1007/978-3-032-08970-0\\_9](https://doi.org/10.1007/978-3-032-08970-0_9)
- [48] Ganesan, S., & Somasiri, N. (2025). Navigating the integration of machine learning in healthcare: Challenges, strategies, and ethical considerations. *Journal of Computational and Cognitive Engineering*, 4(1), 8–23. <https://doi.org/10.47852/bonviewJCCE42023600>
- [49] Luschi, A., Tognetti, L., Cartocci, A., Cevenini, G., Rubegni, P., & Iadanza, E. (2025). Advancing synthetic data for dermatology: GAN comparison with multi-metric and expert validation approach. *Health and Technology*, 15(3), 553–562. <https://doi.org/10.1007/s12553-025-00971-x>
- [50] Partap, A., Mubeen, U. S., Mohan, N. C., Chikkanna, Y., & Iyengar, A. (2025). The role of generative AI in transforming healthcare services. In *2025 25th International Conference on Software Quality, Reliability, and Security Companion*, 33–39. <https://doi.org/10.1109/QRS-C65679.2025.00015>

**How to Cite:** Ganesan, S., & Somasiri, N. (2026). A Narrative Review of GAN-Based Synthetic Data Generation in Disease Prediction. *Journal of Data Science and Intelligent Systems*. <https://doi.org/10.47852/bonviewJDSIS62028542>