

RESEARCH ARTICLE



Evolution of Generative Adversarial Networks (GANs) in Medicine: A Systematic Review of Architectures, Applications, and Implementation Challenges

Wan Aezwani Wan Abu Bakar^{1,*}, Nur Laila Najwa Josdi^{1,*}, Mustafa Man^{2,*}, Evizal Abdul Kadir³ and Bishwajeet Kumar Pandey⁴

¹ Faculty of Informatics and Computing, Universiti Sultan Zainal Abidin, Malaysia

² Faculty of Computer Science and Mathematics, Universiti Malaysia Terengganu, Malaysia

³ Department of Informatics Engineering, Universitas Islam Riau, Indonesia

⁴ Department of Computer Application, G.L. Bajaj Institute of Technology and Management, India

Abstract: Generative Adversarial Networks (GANs) have gained increasing attention in healthcare as a promising approach to addressing data scarcity, offering synthetic alternatives that support research while mitigating privacy risks. This review examines the landscape of GAN-based synthetic data generation in healthcare, with applications spanning medical imaging, electronic health records, genomics, and multimodal datasets. A systematic search guided by the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework identified 81 peer-reviewed studies published between 2014 and 2025, ensuring comprehensive coverage of methodological and translational developments. The review maps the diversity of GAN architectures employed, synthesizes evidence on evaluation strategies, and outlines ethical, privacy, and regulatory considerations that influence adoption. Results indicate that GANs often achieve strong fidelity and downstream utility, with emerging fairness-aware models addressing demographic bias. However, inconsistent validation practices, limited clinical integration, and unresolved ethical and governance challenges continue to hinder translation into real-world settings. Overall, the review consolidates methodological trends, barriers, and future directions, highlighting the potential of GANs to serve as viable tools to overcome data scarcity in healthcare research and practice.

Keywords: data fidelity, electronic health records (EHRs), generative adversarial networks (GANs), genomics, medical imaging, synthetic data

1. Introduction

The recent revolution in Artificial Intelligence (AI) has transformed medical diagnostics, enabling faster and more accurate disease prediction from multimodal healthcare datasets, such as medical imaging, biosignals, Electronic Health Records (EHRs), and genomics [1]. However, access to high-quality, diverse data remains limited—especially for rare conditions and specialized clinical use—and is further restricted by data protection regulations such as the Health Insurance Portability and Accountability Act (HIPAA) and General Data Protection Regulation (GDPR), alongside rising cybersecurity risks [2, 3].

Earlier efforts to address data scarcity, including oversampling methods such as the Synthetic Minority Oversampling Technique (SMOTE) and generative models, such as Variational AutoEncoders (VAEs), provided partial solutions but often failed to capture complex,

high-dimensional feature relationships, produced unrealistic samples, or lacked formal privacy safeguards [4, 5].

Generative Adversarial Networks (GANs) have emerged as a promising alternative, utilizing an adversarial process between a generator and a discriminator to create high-fidelity privacy-preserving datasets that closely mirror real clinical data [6]. These models support dataset augmentation, class balancing, rare disease research, and multimodal data integration while minimizing patient-data exposure. Evaluation of GAN-generated data typically benchmarks performance against real datasets or competing models using metrics of fidelity, utility, privacy, and clinical applicability [7].

Against this backdrop, this study examines the current state of GAN-based synthetic data generation in healthcare and its effectiveness in mitigating data scarcity. It reviews GAN architectures, evaluation approaches, ethical and regulatory considerations, and translational barriers to real-world adoption.

2. Literature Review

While synthetic data has been available for some time, its benefits have only recently garnered substantial recognition [8]. This growing interest stems from identified limitations in traditional privacy-preserving methods. To contextualize the role of GANs, we first review

***Corresponding authors:** Wan Aezwani Wan Abu Bakar, Faculty of Informatics and Computing, Universiti Sultan Zainal Abidin, Malaysia. Email: wanaezwani@unisza.edu.my; Nur Laila Najwa Josdi, Faculty of Informatics and Computing, Universiti Sultan Zainal Abidin, Malaysia. Email: s14189@putra.unisza.edu.my; Mustafa Man, Faculty of Computer Science and Mathematics, Universiti Malaysia Terengganu, Malaysia. Email: mustafaman@umt.edu.my

prior approaches to medical data handling, grouped into distinct methodological categories.

1) Privacy-Preserving Approaches

Data anonymization presents significant challenges, including high re-identification risks through external data linkages, considerable utility loss for Machine Learning (ML) applications, and difficulties preserving rare conditions or complex variable relationships [9]. Federated learning, while maintaining data locality, contends with privacy vulnerabilities in gradient exchanges, infrastructure complexities, and limitations in producing realistic datasets for development and testing purposes [10].

2) Real-World Data and Early Resources

Real-World Data (RWD) [11] has long been used in clinical research and health outcome analysis. However, challenges in data quality, standardization, and generalizability arise due to heterogeneity and missing values. Simulation-Based Training (SBT) has contributed to clinical education, while open datasets (e.g., Medical Information Mart for Intensive Care [MIMIC], PhysioNet) have provided benchmarks for research. Yet these approaches face high costs, limited realism, and issues of scope and representativeness [12, 13]. Similarly, data augmentation techniques enhanced ML pipelines, particularly for imaging tasks, but remain limited for tabular and structured data [14].

3) Emergence of Synthetic Data

These constraints have positioned synthetic data as a viable alternative that balances privacy protection with the need for accessible, diverse datasets to support research and innovation [15].

Table 1 illustrates the timeline of the evolution of medical data methods, highlighting the progression toward the adoption of synthetic data.

4) Generative Adversarial Networks

Introduced by Ian Goodfellow et al. [16] in 2014, GANs represent a pivotal advance in generative modeling. This breakthrough addressed two critical healthcare challenges: data scarcity and patient privacy. Unlike earlier approaches such as VAEs [17], GANs demonstrated superior performance in generating high-quality synthetic data while

preserving confidentiality capabilities. These strengths are particularly valuable in medicine, where complex data and strict privacy requirements create unique research barriers.

2.1. Applications of GANs in medical data synthesis

GAN-based applications in medicine can be grouped into seven categories reflecting different data modalities and clinical goals: (i) medical imaging, (ii) tabular clinical data, (iii) genomics, (iv) multimodal fusion, (v) pharmaceutical research and drug discovery, (vi) clinical decision support, and (vii) privacy-preserving infrastructures.

2.1.1. Medical imaging

GANs have become a powerful tool in medical imaging, addressing challenges in data augmentation, modality translation, and image enhancement. They enable the generation and refinement of imaging data to compensate for limited annotations and improve diagnostic accuracy. A notable example is [18], which used CycleGAN to translate chromatography (CT) to magnetic resonance imaging (MRI) scans while preserving anatomical structures—valuable given the superior soft-tissue contrast of MRI.

Their diagnostic potential is well established. Study [19] reviewed their use in enhancement, modality conversion, and dataset expansion, showing consistent gains in image quality, noise reduction, and task performance. Recent advances include a semi-supervised sequential GAN by [20], which achieved superior cross-modality translation for both paired and unpaired data, preserving anatomical and modality-specific features.

Further innovations extend the utility of GANs. Study [21] introduced a model that synthesizes high-fidelity images from limited datasets using specialized convolutional modules for multimodality augmentation. Similarly, [22] developed a low-noise discriminator (LND) for image denoising that improved peak signal-to-noise ratio by 9.75% and reduced processing time, demonstrating real-time clinical potential. Collectively, these works demonstrate how GANs enhance imaging quality, support cross-modality translation, and improve patient safety by reducing radiation exposure.

Table 1
Evolution of medical data methods

Category	Method	Applications	Limitations
Traditional privacy-preserving approaches	Data anonymization	Protects patient identity in datasets	<ul style="list-style-type: none"> • High re-identification risk • Reduced ML utility • Difficulties preserving rare conditions or complex relationships
Early data resources and augmentation methods	Federated Learning (FL)	Collaborative model training across distributed devices/institutions without sharing raw data	<ul style="list-style-type: none"> • Vulnerable to privacy leakage via gradients • Requires additional defenses to ensure data privacy
Emergence of synthetic data	Real-World Data (RWD)	Used in clinical research and health outcome analysis	<ul style="list-style-type: none"> • Challenges in data quality, standardization, and generalizability due to heterogeneity and missing data
	Simulation-Based Training (SBT)	Supports clinical education and skill development	<ul style="list-style-type: none"> • High costs, limited realism, and the need for specialized faculty training
	Open datasets (e.g., MIMIC, PhysioNet)	Provide benchmarks for research and model validation	<ul style="list-style-type: none"> • Limited scope, potential outdatedness, and the lack of representativeness
	Data augmentation techniques	Expand datasets for ML training, particularly imaging	<ul style="list-style-type: none"> • Primarily effective for images; limited applicability to tabular or structured data
	Synthetic data	Enables AI development, software testing, and research without compromising patient privacy	<ul style="list-style-type: none"> • Emerging best practices • Ongoing need for validation to ensure realism and utility

2.1.2. Tabular data

GANs are also widely applied to structured tabular data, including EHRs, clinical trials, and demographic records. This domain emphasizes specialized architectures designed to handle mixed data types, non-Gaussian distributions, and privacy constraints.

Conditional Tabular Generative Adversarial Network (CTGAN) [23] introduced conditional generation with non-uniform sampling, which effectively modeled continuous and categorical variables, producing realistic privacy-preserving synthetic patient records. Building on this, [24] presented an ML-based approach for generating realistic healthcare datasets suitable for validating clinical applications while ensuring privacy.

Enhanced models continue to emerge. CTAB-GAN+ [25] effectively manages mixed data and imbalanced class distributions, improving both quality and usability for robust AI development. An optimized Conditional GAN (cGAN) proposed by [26] generates synthetic cardiovascular data for mobile-based care, maintaining statistical accuracy and confidentiality.

MedGAN, initially designed for EHR generation, has also evolved. Study [27] demonstrated its capacity to reproduce feature distributions while enabling secure data sharing, reinforcing its relevance for privacy-sensitive medical research and inter-institutional collaboration.

2.1.3. Genomics

In genomics, GANs enable realistic synthesis of DNA, gene expression, and single-cell RNA sequencing (scRNA-seq) data, advancing precision medicine and integrative omics. Study [28] simulated artificial human genomes to improve genotype imputation for rare alleles, enhancing population-level analyses, while [29] achieved high statistical similarity in synthetic cat genome sequences.

Expanding into regulatory design, ExpressionGAN [30] generated DNA sequences inducing targeted mRNA expression, often surpassing natural controls, demonstrating potential in gene regulation research. For gene expression modeling, [31] proposed the Single-Cell Generative Adversarial Network (scGAN) and Conditional Single-Cell Generative Adversarial Network (cscGAN) for realistic scRNA-seq data generation, thereby improving marker gene detection and cell-type classification. MG-GAN [32] further enhanced cancer classification through data augmentation, and omicsGAN [33] integrated multi-omics datasets to improve cancer outcome prediction.

To address scRNA-seq sparsity and noise, [34] introduced scIGANs for accurate imputation while preserving gene–gene and cell–cell relationships. GroundGAN [35] incorporated gene regulatory networks for biologically grounded synthesis, and scGFT [36] introduced a train-free, Fourier-based architecture preserving intrinsic transcriptomic features. These works collectively demonstrate the growing role of GANs in generating realistic privacy-preserving genomic data for clinical and functional genomics research.

2.1.4. Multimodal fusion

Beyond single modalities, GANs enable the synthesis and integration of imaging, genomics, and clinical text, enhancing predictive modeling and survival analysis.

MedFusionGAN [37] fuses three-dimensional (3D) T1-weighted MRI and CT scans to improve brain tumor analysis and protocol generalization. cGANs [38] integrate imaging with genomic data for breast cancer mutation prediction, achieving superior Fréchet Inception Distance (FID) scores and higher accuracy when combined with real data.

The Mutual-Guided Cross-Modality Transformer (MGCT) [39] combines histopathology and genomic data to model tumor microenvironments, improving multi-cancer survival prediction. Similarly, HEALNet [40] integrates whole slide images with multi-

omics data while managing missing modalities, achieving state-of-the-art survival prediction.

Further multimodal innovations include a GAN framework [41] that generates multi-label discrete EHRs, modeling longitudinal trajectories with strong privacy preservation. Study [42] synthesized radiology reports by combining medical images, clinical knowledge, and text, improving diagnostic report generation.

2.1.5. Pharmaceutical research and drug discovery

The pharmaceutical sector is increasingly using GAN-generated synthetic data to accelerate research while maintaining privacy compliance. Syntegraph, in partnership with the Institute for Health Metrics, developed a synthetic data platform replicating over seven million de-identified EHRs [43]. This dataset supports real-world evidence studies and health economics research without exposing sensitive data. Likewise, Electronic Medical Record Bots (EMRBots) [44] demonstrated the practical use of synthetic EMRs in drug discovery and development.

2.1.6. Clinical decision support systems

GANs also strengthen clinical decision support by generating synthetic records and medical images for model training and validation. MDClone's synthetic data engine [45] enabled research on spine surgery outcomes by generating statistically equivalent patient records, thereby supporting risk analysis without using protected health information. In neuroimaging, 3D-StyleGAN [46] synthesizes high-resolution brain MRI scans, thereby expanding diagnostic datasets and improving model robustness.

2.1.7. Privacy-preserving data infrastructure

GANs contribute to privacy-preserving data infrastructures that enable secure, compliant data exchange and multi-institutional collaboration. HealthVerity [47] supported the National Cancer Institute's COVID-19 Real-World Data Infrastructure by integrating synthetic medical claims and laboratory and vaccination records using Privacy-Preserving Record Linkage, facilitating large-scale AI research while maintaining compliance. Similarly, ADS-GAN [48] demonstrates accurate EHR synthesis under complete anonymity, supporting privacy-conscious data sharing and large-scale analytics.

Table 2 summarizes these GAN applications across medical data domains, highlighting their key contributions, study years, and research impact.

2.2. Types of GANs used for medical data synthesis

Since the seminal work of Goodfellow et al. (2014) introduced the GAN framework, there has been substantial evolution in GAN architectures specifically designed for medical data synthesis. The study categorizes these models into three groups: (i) foundational GANs that established training stability and basic architectures, (ii) recent variants optimized for high-fidelity imaging and structured tabular synthesis, and (iii) hybrid models that combine multiple mechanisms to address the heterogeneity of medical data.

Although the original GAN remains the sole peer-reviewed model from the foundational 2014–2015 period, early influential preprints, such as the Conditional GAN (cGAN) and Deep Convolutional GAN (DCGAN), rapidly disseminated architectural innovations that were widely adopted in experimental settings.

However, it was not until 2016 that the field witnessed a surge of peer-reviewed advancements addressing key challenges in GAN training, namely instability, mode collapse, and limited interpretability. Although foundational models such as cGAN and DCGAN appeared earlier in the form of preprints or conference proceedings, to the best of the authors' knowledge, peer-reviewed applications relevant to medical

Table 2
GANs applications in medical data synthesis

Category	Application/Study	Year(s)	Impact/Notes
Medical imaging	CycleGAN (CT→MRI), semi-supervised sequential GAN, high-fidelity convolutional GAN	2018–2022	<ul style="list-style-type: none"> • Cross-modality synthesis and improved paired/unpaired translation • Preserves anatomical structures and features • Augments limited datasets across modalities
Tabular data (EHR, clinical)	LND denoising GAN, CTGAN / MedGAN / CTAB-GAN+, cGAN for mobile CV care	2018–2022	<ul style="list-style-type: none"> • Handles mixed/conditional data • Preserves statistical distributions and privacy • Enhances downstream utility (e.g., cardiovascular patient data) • +9.75% PSNR, reduces processing time
Genomics	GAN-based artificial genomes / cat genome synthesis, ExpressionGAN, scGAN / cscGAN / MG-GAN / omicsGAN, scIGANs / GroundGAN / scGFT	2019–2025	<ul style="list-style-type: none"> • Rare allele imputation, population-level studies • Synthetic regulatory DNA for targeted mRNA expression • Single-cell and multi-omics data synthesis • Preserves gene-gene/cell-cell relationships, improves downstream analysis
Multimodal data integration	MedFusionGAN, cGAN imaging+genomics, MGCT, HEALNet, multi-label EHR GAN, context-enhanced radiology GAN	2021–2023	<ul style="list-style-type: none"> • Cross-modality and multi-omics integration • Enhances diagnosis, mutation prediction, survival prediction, and trajectory modeling
Pharmaceutical research & drug discovery	Syntegra, EMRBots	2022	<ul style="list-style-type: none"> • Large-scale synthetic human/experimental records • Accelerates real-world evidence and drug research
Clinical decision support systems	MDClone, 3D-StyleGAN	2022–2023	<ul style="list-style-type: none"> • Synthetic patient records and MRI • Enhances outcome modeling and diagnostic tools
Privacy-preserving data infrastructure	HealthVerity, ADS-GAN	2024	<ul style="list-style-type: none"> • Secure synthetic EHR sharing • Regulatory compliant, supports multi-institution collaboration

data synthesis were limited or undocumented during 2014–2015.

Table 3 summarizes foundational GAN architectures, highlighting their main contributions and representative medical use cases. These models addressed early challenges such as training instability and mode collapse, enabling initial medical applications that include synthetic echocardiographs (ECGs) and MRI augmentation. For instance, InfoGAN [49] introduced disentangled latent representations, facilitating interpretable phenotype synthesis by controlling semantic features in the latent space. The DCGAN incorporated convolutional layers that significantly enhanced the spatial coherence and visual fidelity of synthetic medical images, such as brain MRI scans [50]. Meanwhile, WGAN [51] and its variant WGAN-GP [52] improved convergence behavior by replacing the Jensen-Shannon divergence with the Earth-Mover (Wasserstein) distance and introducing gradient penalties, which proved critical for stable training in biomedical imaging contexts. Spectral Normalization GAN (SNGAN) [53] further advanced stability by constraining the spectral norm of discriminator

weights, leading to improved convergence and higher-quality synthetic images. Additionally, cGANs [54] enabled label- or modality-conditioned generation, expanding applications to multi-contrast MRI synthesis and EHR generation conditioned on clinical labels.

Foundational GANs established the baseline for adversarial training and image synthesis, but their limited control over feature conditioning and challenges with structured data constrained broader clinical adoption. Models such as DCGAN and cGAN were pivotal for early medical imaging and label-conditioned EHR generation, respectively, highlighting how architectural modifications align with domain-specific requirements.

Building on foundational GANs, recent variants have advanced fidelity, diversity, and privacy preservation, with applications spanning imaging, EHRs, and genomics. StyleGAN [55] and its improved variants, including StyleGAN2-ADA [56], introduced style-based latent control and data-efficient augmentation techniques, enabling highly realistic facial phenotyping and modeling of rare diseases with

Table 3
Foundational GAN variants in medical data synthesis

Model	Year	Main Contribution	Medical Use Case
GAN	2014	Introduced adversarial training framework	Baseline synthetic ECGs, MRIs
InfoGAN	2016	Disentangled latent representations for interpretability	Phenotypic feature synthesis
DCGAN	2017	Convolutional layers for improved image coherence	Context-aware brain MRI synthesis
WGAN	2017	Earth-Mover distance for stable training	Histopathological image synthesis
WGAN-GP	2017	Gradient penalty for improved convergence	Brain MRI augmentation
SNGAN	2018	Spectral normalization to stabilize training	Medical image synthesis
cGAN	2019	Label- or modality-conditioned generation	Label-conditioned EHR generation

improved data efficiency. For structured data, tabular GANs, such as CTGAN [57], have become essential for synthesizing complex, heterogeneous datasets, such as EHRs and genomics, addressing the challenges posed by mixed data types and class imbalances. Privacy-preserving GAN [58] variants that incorporate differential privacy mechanisms have also emerged, facilitating the secure generation of synthetic medical data while maintaining patient confidentiality. Recently, Transformer-based GANs such as Trans-cGAN [59] have incorporated attention mechanisms and U-Net architectures to enhance cross-modality medical image synthesis, effectively capturing intricate anatomical details and improving clinical relevance.

Table 4 summarizes these modern GAN architectures and highlights their primary innovations and specific applications in medical data synthesis.

Recent variants demonstrate targeted optimizations: StyleGAN and StyleGAN2-ADA excel in high-fidelity image synthesis due to style-based latent control, while CTGAN dominates tabular data synthesis because it effectively handles mixed data types and class imbalances. Privacy-preserving GANs and Transformer-based models further extend applicability to sensitive EHRs and cross-modality imaging, reflecting the growing emphasis on clinical relevance and regulatory compliance.

While Section 2.1 mapped the diverse applications of GANs across medical domains, their success fundamentally depends on architectural innovations that address stability, interpretability, and clinical utility. This section surveys the evolution of GAN variants that underpin these applications. Some models, such as cGAN or CTGAN, appear across both application- and architecture-focused discussions; here we emphasize their methodological innovations rather than specific clinical deployments.

Hybrid GAN architectures have emerged to better capture the complex heterogeneity of medical data by combining complementary modeling strategies. MOSA [60] utilizes multi-omic synthetic data augmentation to study drug resistance in cancer. Hybrid GANs [61] integrate spatial and frequency domain features to synthesize histopathology images with both global structure and fine texture. EnhGAN [62] leverages conditional GANs with contrast enhancement to improve tumor subregion visibility and segmentation in brain MRI. HAGAN [63] employs hybrid attention mechanisms and hierarchical discriminators to generate anatomically consistent, realistic medical images. Recent models continue to integrate domain-specific strategies, such as spatial-frequency fusion and attention mechanisms, to enhance realism and clinical utility in challenging datasets.

Table 5 outlines these hybrid models, illustrating their specialized contributions and application domains.

Hybrid architectures combine complementary mechanisms to address heterogeneous and complex medical datasets. For instance, MOSA focuses on multi-omics data integration, while HAGAN and EnhGAN enhance anatomical consistency and segmentation accuracy in imaging. These design choices illustrate how domain-specific challenges drive the evolution toward specialized, high-utility GAN variants.

In summary, these developments chart a clear evolution in generative AI for medicine: a shift from general-purpose frameworks to specialized architectures fine-tuned for the complexities of medical imaging, genomics, and clinical records. This progression highlights the field's focused response to paramount challenges, including model stability, interpretability, data heterogeneity, and, most critically, clinical relevance. The outcome of this specialization is the ability to generate highly realistic and trustworthy synthetic datasets. As a result, GANs have matured into essential tools that not only augment data for training robust AI models but also protect patient privacy and facilitate groundbreaking research into rare and complex diseases.

2.3. Types of GANs validation used for medical data synthesis

Evaluating the quality and utility of GAN-generated medical data is essential to determine whether it can effectively substitute or supplement real-world datasets. Validation ensures that synthetic data not only appears realistic but also replicates the statistical, structural, and clinically relevant properties of original data. While visual realism is valuable, rigorous technical and clinical evaluations are required to verify performance in downstream tasks. Broadly, validation approaches fall into two categories: (i) technical validation, emphasizing quantitative and algorithmic measures (e.g., distributional similarity, feature fidelity), and (ii) clinical validation, focusing on usability, expert review, and diagnostic performance.

2.3.1. Technical validation

Researchers assess synthetic data using quantitative metrics (e.g., FID, Inception Score [IS]), model-based testing (e.g., Train on Synthetic, Test on Real [TSTR]), and visualization tools (e.g., t-SNE, histograms). These methods evaluate alignment between real and synthetic data in terms of distribution, diversity, and feature representation.

FID [64] and IS [65] are common in medical imaging, quantifying perceptual realism and diversity through pretrained feature embeddings. They capture subtle distributional differences but depend on ImageNet-trained encoders, which may not fully represent clinical features. For structured data, the Maximum Mean Discrepancy (MMD) [66] directly compares statistical distributions and is model agnostic, although sensitive to kernel selection and dimensionality. TSTR [67] assesses predictive utility by measuring the performance of models trained on synthetic data on real data, although outcomes may conflate model bias with data quality. Dimensionality reduction methods such as t-SNE and PCA [68] offer qualitative visualization of distributional overlap but are parameter dependent and unsuitable for standalone validation.

Each technique highlights a different notion of quality: FID and IS capture perceptual fidelity, MMD measures structural similarity, TSTR reflects predictive value, and visualization aids interpretability. Their complementarity suggests that robust evaluation requires a combination of multiple statistical, task-based, and visualization approaches. Common techniques are summarized in Table 6.

Table 4
Recent GAN variants for high-fidelity and structured medical data synthesis

Model	Year	Main Contribution	Medical Use Case
StyleGAN	2019	Style-based latent control	Facial phenotyping in rare diseases
StyleGAN2-ADA	2020	Data-efficient augmentation	Rare disease modeling
CTGAN	2019	Handling mixed data types and imbalances	Synthetic EHR tables
cGAN with DP	2023	Privacy-preserving synthetic data generation	EHRs
Trans-cGAN	2023	Transformer–U-Net hybrid for image synthesis	Cross-modality MRI generation

Table 5
Hybrid GAN architectures for medical data synthesis

Model	Year	Main Contribution	Medical Use Case
MOSA	2024	Multi-omic synthetic data augmentation	Cancer research, drug resistance
HAGAN	2024	Hybrid attention and hierarchical discriminator	Medical image synthesis
Enhancement GAN (EnhGAN)	2025	Conditional GAN with contrast enhancement for segmentation	Brain tumor MRI synthesis and segmentation
Hybrid GAN (Spatial-Frequency)	2025	Fusion of spatial and frequency domain features	Histopathological image synthesis

Table 6
Technical validation on synthetic dataset

Metric / Method	Purpose	Example Use
Fréchet Inception Distance (FID)	Measures distributional similarity between real and synthetic images using deep feature embeddings	Evaluated synthetic medical images across various modalities
Inception Score (IS)	Assesses the quality and diversity of generated images based on classification confidence	Applied to synthetic chest X-rays and dermatoscopy images
Maximum Mean Discrepancy (MMD)	Measures the distance between distributions of real and synthetic data, applicable to both images and tabular data	Used in evaluating EHR tabular GANs such as CTGAN
Train on Synthetic, Test on Real (TSTR)	Evaluates how well models trained on synthetic data generalize to real data	Applied in synthetic EHR and genomics research
Dimensionality Reduction Visualization (e.g., t-SNE, PCA)	Visually assesses the overlap between real and synthetic data distributions	Used to evaluate structural similarity in synthetic longitudinal EHR datasets

2.3.2. Clinical validation

Clinical validation ensures that synthetic data preserves medically relevant features and remains suitable for real-world use. It involves comparing the characteristics and outcomes of synthetic datasets against real clinical data to confirm that key patterns, distributions, and relationships are retained. This step is essential for maintaining data integrity for applications such as disease modeling, diagnostic support, and epidemiological forecasting.

Several complementary methods address different aspects of clinical validity. Statistical validation [69] compares disease trajectories, comorbidity patterns, and visit distributions between real and synthetic EHRs, confirming temporal and cohort-level fidelity. Prospective clinical trial simulations [70] model treatment pathways and outcomes but rely on assumptions about synthetic patient responses. Downstream task performance [71] tests diagnostic or predictive models trained on synthetic data, although results are task dependent. Blinded diagnostic studies [72] provide strong perceptual validation but are resource intensive and limited in scale. Clinical usability studies [73] evaluate practical integration into workflows, emphasizing operational feasibility rather than statistical accuracy.

These approaches form a multi-layered validation framework: statistical and trial-based analyses ensure baseline clinical realism, task-based testing demonstrates utility, and expert-in-the-loop evaluations confirm real-world relevance. As summarized in Table 7, this combined evidence supports the reliability and translational potential of synthetic medical data while upholding privacy protections.

2.4. Challenges, limitations, and ethical considerations

2.4.1. Privacy risks and re-identification

Although GAN-generated medical datasets are designed to be privacy preserving, they remain susceptible to adversarial attacks. When models are overfitted or training data are scarce, synthetic records may inadvertently replicate real individuals, leading to the leakage of

sensitive information. Membership inference attacks have shown that adversaries can determine whether a particular record contributed to model training, especially when contrastive learning enhances attack precision [74].

Beyond membership inference, attribute inference and linkage attacks exploit correlations to reconstruct hidden traits or re-identify individuals by matching synthetic entries to real records [75, 76]. Similar vulnerabilities have been demonstrated in image-based GANs, where discriminators or black-box access can reveal training set membership, showing that risks extend beyond EHRs to other medical data modalities [77].

Ultimately, synthetic data generation involves a trade-off between utility and privacy. Techniques such as differential privacy can mitigate re-identification but often reduce data fidelity in high-dimensional healthcare settings. Table 8 summarizes key adversarial attack types and their implications for GAN-based medical data.

2.4.2. Bias amplification and fairness

Generative models may inadvertently amplify biases present in training data, posing an ongoing challenge for synthetic medical data. To address this, the Bias-Transforming GAN (BT-GAN) applies fairness constraints during generation to rebalance demographic and outcome disparities, such as unequal disease prevalence across gender or racial groups, while maintaining clinical validity [78].

In parallel, FairGAN promotes fairness by aligning distributions of protected attributes (e.g., race, gender), ensuring downstream classifiers trained on synthetic data perform equitably [79]. While FairGAN provides a general approach, BT-GAN adapts this concept for healthcare, preserving subgroup densities and clinical fidelity essential to medical analysis.

These models exemplify the tension between enforcing fairness and preserving data utility. Fairness-aware frameworks such as BT-GAN and FairGAN demonstrate that equity and realism can coexist, but achieving this balance requires careful optimization of fairness

Table 7
Clinical validation on synthetic dataset

Method	Purpose	Example Use
Hierarchical Autoregressive Language mOdel (HALO) for Longitudinal EHR Synthesis	To generate and validate high-dimensional, longitudinal synthetic EHRs that preserve clinical and temporal dependencies while maintaining privacy.	HALO-generated synthetic EHRs were compared with real patient data for disease prevalence and treatment trajectories, showing close alignment and predictive parity (AUC ≈ 0.94).
Prospective Clinical Trial Simulation	The use of synthetic patient populations to simulate clinical trials for drug safety and efficacy predictions.	Synthetic oncology patient cohorts are used to model trial outcomes and optimize study design before real-world trials.
Downstream Task Performance	Tests whether synthetic data supports clinical model training (e.g., segmentation, diagnosis).	GAN-generated CT scans evaluated for lung cancer detection accuracy.
Blinded Diagnostic Studies	Physicians diagnose cases from real and synthetic data to assess indistinguishability.	Used in validating CycleGAN-based cross-modality synthesis (e.g., CT-MRI).
Clinical Usability Studies	Assessment of synthetic data for medical education, software testing, or decision support.	Synthetic patient records are used in EMR system simulations or student training.

Table 8
Re-identification and privacy threats in synthetic healthcare data

Attack Type	Description	Example / Context
• Membership inference • Limitations of differential privacy	• Determines if a specific real record was used in training • Exploits statistical similarity between real and synthetic samples • DP reduces privacy risks but significantly degrades fidelity in high-dimensional healthcare data	Explicitly demonstrated against synthetic EHRs, enhanced by contrastive learning; Trade-off between data utility and privacy protection
Attribute inference	• Predicts unknown sensitive attributes from partially known traits	Attackers infer missing patient features (e.g., disease status) from correlations.
Linkage attacks	• Matches synthetic entries to real individuals using auxiliary datasets	Even without direct mapping, partial re-identification is possible.
Cross-domain inference	• GAN discriminators/black-box access used for training set membership detection	Image-based GANs (no explicit identity labels) are shown to leak membership info.

constraints within domain-specific limits. Table 9 summarizes the respective strategies, strengths, and applications of the models.

2.4.3. Data quality and fidelity

Ensuring that synthetic medical data adequately replicates the complexity, unpredictability, and subtle pathological aspects of real-world datasets remains a considerable task. Study [80] highlighted the trade-off between fidelity, accurately reflecting real data, and diversity, maintaining variability to prevent mode collapse and overfitting. Their GAN framework for retinal image synthesis addressed these issues by enhancing both visual quality and representational diversity, thereby improving trust in downstream diagnostic models.

Building on this, [81] proposed the Vessel and Style Guided GAN (VSG-GAN), which separates retinal image generation into vascular structure and background style components. Using style transformation and GAN inversion, VSG-GAN produces retinal images with diverse morphological patterns and superior realism across evaluation metrics.

Similarly, [82] reviewed synthetic data generation across healthcare domains, emphasizing the persistent challenge of maintaining fidelity and diversity in complex datasets. Robust GAN architectures must preserve fine pathological details while capturing real-world variability to ensure reliability and ethical deployment in clinical research and decision support. Table 10 summarizes representative models, their objectives, and comparative strengths and limitations.

Table 9
Fairness-aware GAN models for synthetic medical data

Model	Key Strategy	Remarks (Strengths and Limitations)	Application Context
BT-GAN	Bias-transforming constraints on demographic/outcome variables	• Reduces amplification of existing health data biases and preserves clinical validity • Requires careful tuning; fairness • Utility trade-off remains	Synthetic health datasets (e.g., disease prevalence across subgroups)
FairGAN	Adversarial fairness constraints enforcing parity in protected attributes	• Balances sensitive attributes (race, gender) during generation while maintaining data utility • Evaluated mainly on general tabular datasets, limited direct healthcare tests	• General tabular data • Adaptable to medical contexts

Table 10
Fidelity and diversity in GAN-generated medical data

Method / Model	Objective / Outcome	Remarks (Strengths & Limitations)
GAN framework for retinal images	<ul style="list-style-type: none"> • Improve fidelity and diversity in synthetic medical images • Retinal image synthesis balancing realism and variability 	<ul style="list-style-type: none"> • Preserves visual quality and diversity • Reduces mode collapse • May require careful tuning • Only demonstrated in the retinal domain
VSG-GAN	<ul style="list-style-type: none"> • Decouple vascular and background features to enhance diversity and fidelity • Retinal images with varied morphology 	<ul style="list-style-type: none"> • Generates more realistic and morphologically diverse images • Style-based control • Specialized retinal images • May not generalize to other modalities
Review of synthetic data in healthcare	<ul style="list-style-type: none"> • Assess fidelity/diversity challenges across datasets • Broad healthcare applications 	<ul style="list-style-type: none"> • Highlights common pitfalls • Informs robust GAN design • Conceptual only, does not provide implementation-level solutions

2.4.4. Ethical misuse and accountability

Study [83] highlights the growing risk of ethical misuse as generative AI becomes more accessible in clinical and research contexts. Without standardized accountability mechanisms, synthetic data may be misused, for example, in misleading publications, unauthorized data augmentation, or bypassing regulatory oversight. The authors call for clear ethical boundaries and institutional oversight to ensure transparency in how synthetic data are generated, validated, and applied, emphasizing traceability, consent, and shared responsibility among developers, clinicians, and policymakers.

Similarly, [84] identifies ethical vulnerabilities in AI healthcare, such as data breaches, misuse of sensitive information, and unregulated commercialization, underscoring the need for robust governance frameworks. Recent studies further emphasize transparency, consent, and regulatory compliance as essential pillars of responsible AI. Study [85] advocates structured ethical frameworks for AI deployment, [86] explores the ethical–legal implications of consent and privacy, and [87] stresses institutional oversight and traceability to prevent misuse.

Table 11 summarizes major ethical challenges and best practices in synthetic medical data governance.

2.4.5. Toward unified ethical governance

Growing concerns over privacy, bias, and data fidelity have underscored the need for comprehensive ethical governance frameworks in synthetic medical data development. Study [88] introduces an ethical checklist for generative AI in healthcare, providing a practical guide for responsible model design and deployment, including GANs. The framework addresses interconnected risks such as privacy violations,

algorithmic bias, data fidelity loss, and lack of transparency, issues which are particularly critical when synthetic data influence diagnostic or therapeutic decisions. It emphasizes governance grounded in justice, accountability, and explainability, alongside technical robustness to maintain patient trust and uphold medical ethics.

In alignment, [89] advocates integrating ethical values with technical rigor, while [90] proposes a co-designed governance model tailored to healthcare institutions, enhancing stakeholder collaboration and real-world oversight. Study [91] further highlights the importance of secure infrastructure, strong data governance, and ethical guidelines for responsible use of AI. Collectively, these studies reinforce the need to establish a unified oversight mechanism to protect patient trust and align synthetic data practices with core medical principles. Table 12 summarizes notable contributions and key takeaways in this area.

3. Methodology

This systematic review was conducted in accordance with the PRISMA 2020 guidelines. The review question was formulated using the PECO (P Population, E Exposure, C Comparator, O Outcome) framework to ensure transparency and reproducibility. Table 13 summarizes the key components of the PECO framework applied in this study.

3.1. Eligibility criteria

Eligibility criteria were defined to operationalize the PECO framework into practical rules for study selection. Studies were included or excluded according to the criteria summarized in Table 14.

Table 11
Ethical issues and recommended mitigations in synthetic medical data

Ethical Issue	Description	Example / Implication	Recommended Mitigation
Misuse of synthetic data	Synthetic medical images or EHRs used in misleading publications or unauthorized augmentation	Using GAN-generated CT scans in studies without disclosing synthetic origin	Implement clear usage guidelines and disclosure policies
Lack of accountability	No standardized oversight for synthetic data generation and application	Circumventing regulatory requirements or bypassing ethical review	Establish institutional oversight and shared responsibility among developers, clinicians, and policymakers
Consent & privacy gaps	Patients may not consent to synthetic use of their data, and privacy risks remain	Using synthetic data derived from sensitive EHRs without informed consent	Integrate consent procedures and privacy safeguards
Traceability & transparency	Difficulty tracking synthetic outputs and their origin	Difficulty auditing or validating synthetic datasets in clinical pipelines	Maintain provenance records, logging generation process, and versioning

Table 12
Ethical governance reviews on GANs

Year	Main Contribution	Key Takeaways
2025	Structured ethical checklist for generative AI in healthcare	Practical guidance for responsible GAN deployment, addresses privacy, bias, data fidelity, and transparency
2024	Ethical principles for AI in medical research	Emphasizes justice, accountability, and explainability alongside technical rigor
2025	Co-designed AI governance framework for healthcare organizations	Promotes stakeholder collaboration and real-world oversight of AI systems
2024	Ethical implications of generative AI in clinical practice	Highlights the need for robust data governance, secure infrastructure, and ethical guidelines

Table 13
PECO framework for the systematic review

Component	Description
Population (P)	Healthcare datasets across imaging, EHR/tabular, genomics, and other biomedical domains
Exposure (E)	Use of GANs to generate synthetic data
Comparator (C)	Real-world datasets or alternative generative models, where available
Outcome (O)	Measures used to evaluate GAN-generated synthetic data, including fidelity (statistical similarity to real data), utility (performance in downstream ML tasks), privacy protection (resilience against re-identification), clinical or translational applicability, and any other metrics that benchmark or improve GAN performance

3.2. Search strategy

A comprehensive literature search was conducted to identify studies on GAN-based synthetic data generation in healthcare. To capture both biomedical and computer science perspectives, we searched PubMed, Scopus, Web of Science, and IEEE Xplore, covering publications from January 2014 to June 2025. While the primary focus of this review is on recent developments from 2020 onwards, studies published from 2014 to 2019 were also included to illustrate the evolution of GAN applications in healthcare. Searches were limited to English-language peer-reviewed articles and conference proceedings.

The selected databases reflect the dual publication patterns in this field. PubMed ensured coverage of medical and life science journals, while Scopus and Web of Science provided multidisciplinary indexing of high-impact outlets, including *Nature Communications*, *NPJ Digital*

Medicine, *IEEE Transactions on Medical Imaging*, and *Bioinformatics*. IEEE Xplore captured conference proceedings from leading venues, such as the Conference on Neural Information Processing Systems (NeurIPS), Medical Image Computing and Computer Assisted Intervention (MICCAI), Conference on Computer Vision and Pattern Recognition (CVPR), International Conference on Machine Learning (ICML), and IEEE International Conference on Bioinformatics and Biomedicine (IEEE BIBM), which are known for their innovations in technical GAN in healthcare.

The search was designed to capture studies meeting the inclusion criteria defined in Table 14. Keywords and controlled vocabulary were derived from the PECO framework:

Population (P): “medical,” “clinical,” “biomedical,” “EHR,” “genomic”

Exposure (E): “generative adversarial network,” “GAN,” “CTGAN,” “WGAN,” “StyleGAN”

Outcome (O): “synthetic data,” “data generation,” “data augmentation,” and evaluation metrics such as “FID,” “TSTR,” or “MMD”

An example Boolean search string used in Scopus was
("generative adversarial network" OR "GAN" OR "CTGAN" OR "WGAN" OR "StyleGAN")

AND ("synthetic data" OR "data generation" OR "data augmentation")

AND ("healthcare" OR "medical" OR "clinical" OR "biomedical" OR "genomic" OR "electronic health record" OR "EHR")

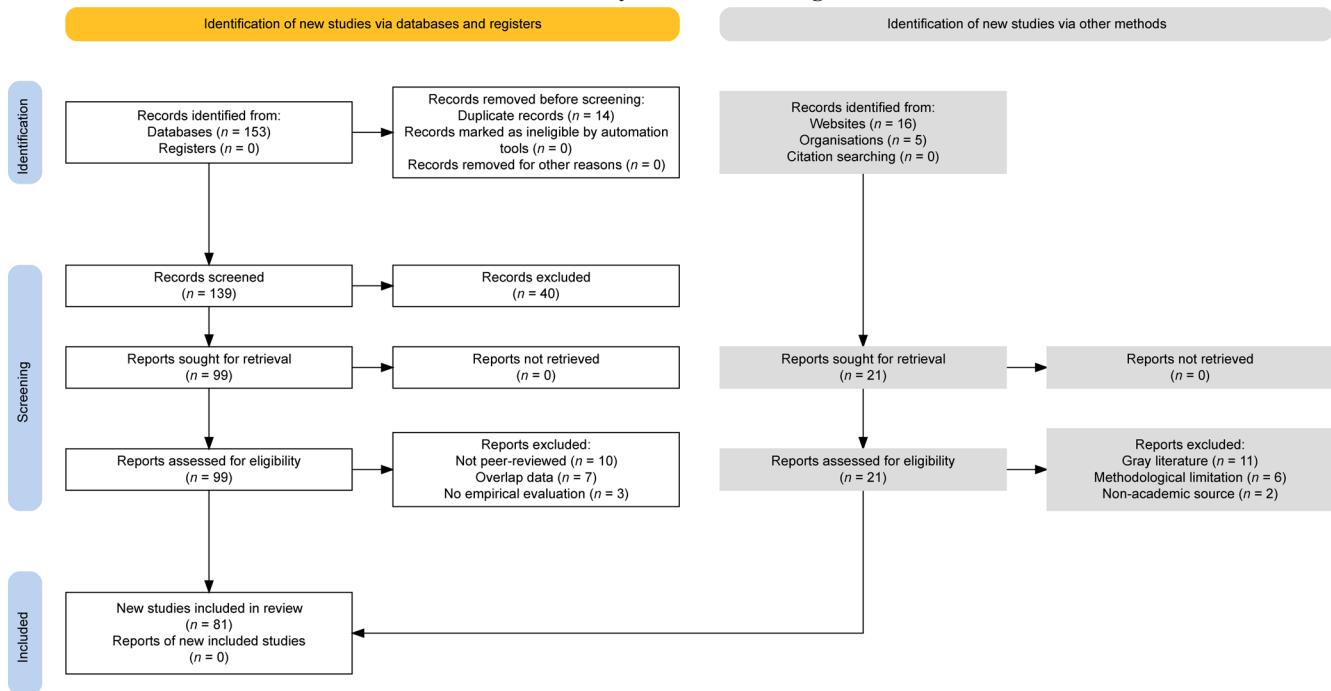
3.3. Study selection and bias considerations

All records retrieved from the database searches were manually screened and duplicates were removed through careful comparison. Titles and abstracts were evaluated against the eligibility criteria,

Table 14
Eligibility criteria for included studies

PECO Component / Criterion	Inclusion	Exclusion
Population (P)	• Healthcare datasets across imaging, EHR/tabular, genomics, and other biomedical domains	• Generic computer-vision benchmarks only
Exposure (E)	• Application of GANs for synthetic data generation (e.g., GAN, cGAN, WGAN, CycleGAN, StyleGAN, CTGAN)	• Methods papers with no empirical GAN application
Comparator (C)	• Real-world datasets or alternative generative models, where available	• Duplicate reports of the same dataset/model without novel analysis
Outcome (O)	• Studies reporting at least one validation of synthetic data (e.g., FID, IS, MMD, TSTR, clinical expert review, downstream task performance)	• Studies with no evaluation or empirical results of synthetic data
Time frame / language	• Published 2014–2025 • Written in English	• Non-English papers • Preprints or non-peer-reviewed sources

Figure 1
PRISMA study selection flow diagram



followed by full-text review of potentially relevant studies. Screening and selection were conducted systematically by a single reviewer to ensure consistency and transparency. This limitation is discussed further in Section 4.4. The overall selection process is summarized in the PRISMA 2020 flow diagram (Figure 1), which details the number of records identified, screened, excluded (with reasons), and ultimately included in the review.

A total of 174 records were identified, including 153 from electronic databases and 21 from other sources. After removing 14 duplicates, 160 unique records remained for title and abstract screening. Of these, 66 were excluded for not meeting the inclusion criteria. The full texts of 120 records were then assessed for eligibility (99 from databases and 21 from other sources). A total of 32 records were excluded: 13 from databases (due to methodological limitations, irrelevant focus, or non-peer-reviewed status) and 19 from other sources (Gray Literature [11], Methodological Limitations [6], Non-Academic Sources [2]). Finally, 81 studies were included in the review, comprising 79 from databases and 2 from other sources.

Potential sources of bias were considered both in the review process and in the included studies. Table 15 summarizes the main bias types, their sources, potential impacts, and the strategies employed to mitigate them.

3.4. Data extraction and synthesis

Data was systematically extracted to obtain key information from each included study. Extracted details included author and year, dataset type and domain, GAN architecture and configuration, evaluation metrics (fidelity, utility, privacy, interpretability), and main findings. This structured approach ensured consistency and transparency, and enabled a meaningful comparison across studies.

Given the heterogeneity of study designs, datasets, and evaluation metrics, a narrative synthesis approach was adopted. Studies were grouped thematically by data modality (imaging, EHR/tabular, genomics) and evaluation focus (fidelity, utility, privacy, clinical translation) to analyze patterns, strengths, and limitations, deriving insights into the current state and practical applicability of GAN-based synthetic data in healthcare.

4. Discussion

The analysis is structured around four key dimensions: Distribution of Included Studies by Publication Type and Article Category, Thematic Map of GAN-Based Synthetic Medical Data Literature, Technical Comparison of GAN Methods, and Limitations

Table 15
Bias considerations in the review process and included studies

Bias Type	Source	Description / Impact	Mitigation Strategy
Reviewer bias	Review process	<ul style="list-style-type: none"> Screening and data extraction performed by a single reviewer May introduce errors or inconsistencies 	<ul style="list-style-type: none"> Systematic search, predefined eligibility criteria, structured extraction template
Publication bias	Included studies	<ul style="list-style-type: none"> Only peer-reviewed English-language studies included May exclude null/negative results 	<ul style="list-style-type: none"> Acknowledge in synthesis and discussion
Methodological bias	Included studies	<ul style="list-style-type: none"> Small or single-source datasets, lack of external validation, selective reporting Limits generalizability 	<ul style="list-style-type: none"> Consider during synthesis Highlight limitations in discussion

and Future Directions. These perspectives illuminate both the current capabilities of GANs and the critical gaps that must be addressed for safe and effective clinical integration.

4.1. Distribution of included studies by publication type and article category

Table 16 categorizes the 81 included studies based on publication type and thematic focus. The majority are peer-reviewed journal articles ($n = 62$), including 46 original research studies and 16 conceptual works such as reviews, frameworks, and ethical discussions. Additionally, one peer-reviewed book was included. These serve as the foundation of the analysis, offering validated findings on GAN architectures, evaluation methods, and clinical applications.

Sixteen conference proceedings complement this core set, highlighting emerging developments in generative modeling. While some studies lack clinical validation, they represent important technical progress and reflect the rapidly evolving nature of the field.

Two additional sources were included after a careful credibility check: one technical industry report and one expert blog post. While the latter two are not peer reviewed, they provide practical perspectives that help bridge the gap between theory and real-world implementation.

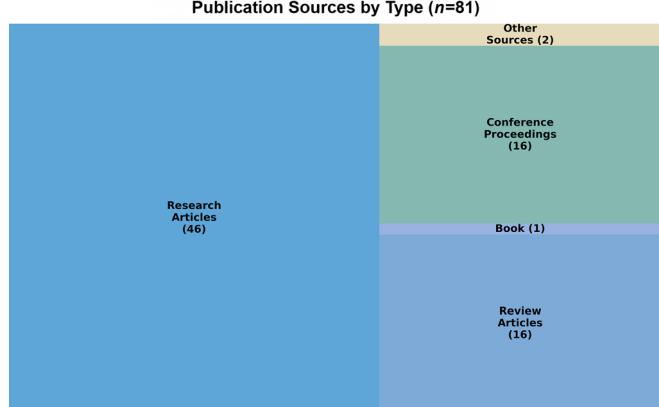
In total, the dataset consists of 62 original research studies, 17 review or perspective articles (including one book), and 2 vetted non-academic sources. This balanced mix allows for a robust and relevant synthesis, based on peer-reviewed research while incorporating diverse viewpoints from practice.

As shown in Figure 2, most included studies are peer reviewed, reflecting the strong academic foundation of GAN research in healthcare. Journal articles constitute the largest share, with 62 publications (76.5% of the total), including 46 research articles (56.8%) and 16 review articles (19.8%). Additionally, 1 book (1.2%) was included. These journals span high-impact outlets such as *Nature Communications*, *NPJ Digital Medicine*, *IEEE Transactions on Medical Imaging*, *Bioinformatics*, *Journal of Biomedical Informatics*, *Biophysical Journal*, *Computational and Structural Biotechnology Journal*, *Iranian Journal of Public Health*, *Preventing Chronic Disease*, *Frontiers in Artificial Intelligence*, *The Lancet Digital Health*, *Kosin Medical Journal*, *JMIR Research Protocols*, and *Journal of Medical Internet Research*, highlighting the interdisciplinary collaboration between AI and medical science.

The conference proceedings account for 16 studies (20% of the total), all of which are research articles from leading AI and healthcare venues, including NeurIPS, MICCAI, CVPR, and others such as the International Conference on Learning Representations (ICLR), ICML, IEEE BIBM, and *Proceedings on Privacy Enhancing Technologies* (PoPETs). Their presence demonstrates the active engagement of the ML community in addressing real-world medical challenges, including limited datasets, privacy, and annotation constraints, all of which utilize GAN-based solutions.

The remaining two sources (2%) are industry reports or blogs, providing additional perspective outside traditional academic publishing.

Figure 2
Distribution of publication type



Overall, this publication pattern reveals two main trajectories in GAN research: advancing core technical methodologies and addressing clinical demands for reliable, interpretable, and regulatory-compliant solutions, aligning closely with the objectives of this review to evaluate both technical innovation and practical applicability in healthcare.

4.2. Thematic map of GAN-based synthetic medical data literature

Figure 3 presents a thematic analysis of the included 81 studies, highlighting established research directions and emerging trends in GAN applications for medical data synthesis. Medical Imaging Synthesis is the largest category with 23 studies (28%), reflecting the suitability of GANs for pixel-based data generation in radiology. Tabular Data and EHR Synthesis follow with 18 studies (22%), demonstrating substantial interest in generating structured clinical data for research and decision support. Together, these two categories account for half of the reviewed literature, underscoring the technical focus of the field on clinically impactful applications.

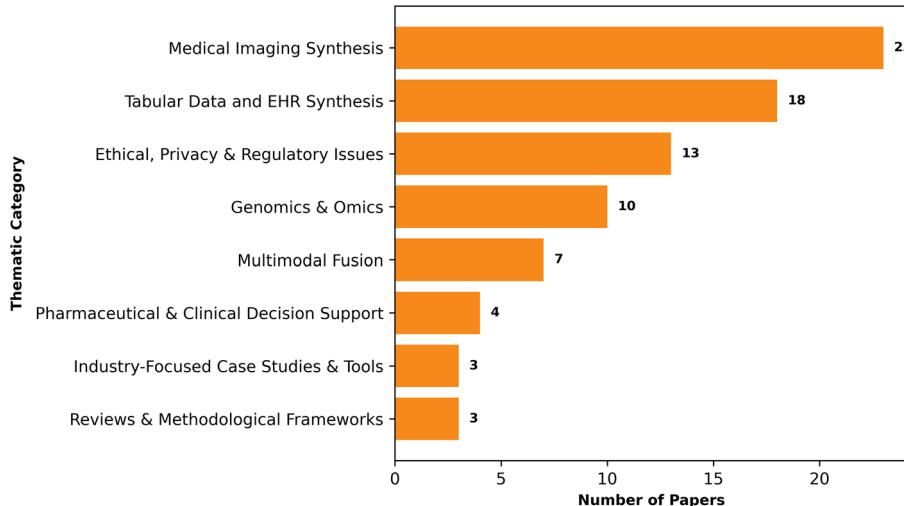
Emerging domains, though smaller in scale, are growing in representation: Genomics and Omics appear in 10 studies (12%), indicating their expanding application in biological data synthesis, while Multimodal Fusion is addressed in 7 studies (9%), representing innovative approaches to integrating heterogeneous clinical data. Ethical, Privacy, and Regulatory Issues are discussed in 13 studies (16%), and Pharmaceutical and Clinical Decision Support in 4 studies (5%), indicating a focus on implementation and translational challenges, despite these areas remaining underdeveloped relative to technical research. Reviews and Methodological Frameworks as well as Industry-Focused Case Studies appear in three studies each (4%), highlighting gaps in synthesizing best practices and capturing real-world deployment experiences.

Overall, this thematic distribution maps the current research landscape and reveals opportunities for future work, particularly in

Table 16
Classification of included studies by source type and article category

Source Type	Review (incl. Perspective)	Research Article	Other (Industry / Book)	Total by Source Type
Journal	16	46	1	63
Conference proceedings	0	16	0	16
Other (industry reports / press / blog)	0	0	2	2
Total by Publication Type	16	62	3	81

Figure 3
Proportional distribution of references by research theme in medical AI and data science



bridging technical innovation with clinical implementation, developing standardized evaluation frameworks, and addressing translational challenges in healthcare.

4.3. Technical comparison of GAN methods

To evaluate the methodological landscape of GAN-based synthetic medical data, Tables 17 and 18 summarize the key architectures identified in this review, organized by application domain. Table 17 focuses on GAN techniques applied to tabular EHR and clinical datasets, highlighting commonly used evaluation metrics, key findings, and reported strengths and limitations. Table 18 presents GAN techniques for imaging and multi-omics data, illustrating domain-specific performance and technical considerations.

These comparative tables provide a concise overview of the capabilities of each model, validation approaches, and practical constraints, facilitating identification of trends, domain-specific strengths, and remaining gaps in clinical translation. The following discussion interprets these trends in terms of fidelity, downstream utility,

fairness, and ethical considerations, highlighting critical observations relevant for real-world implementation.

The comparison of GAN techniques reveals several notable trends. Domain-specific strengths are evident: CTGAN, TVAE, and CopulaGAN excel at tabular EHR data, preserving statistical distributions and enabling downstream modeling, while imaging-focused GANs, such as CycleGAN, StyleGAN variants, and TranscGAN, achieve high visual fidelity for cross-modality or rare disease datasets. Hybrid architectures support multi-omics integration, providing flexible augmentation for complex datasets.

Several objective metrics have been employed to assess algorithm performance across these domains. For tabular datasets, MMD and TSTR evaluate distribution similarity and downstream task utility, respectively. For imaging applications, FID and IS quantify visual fidelity and diversity, often complemented by dimensionality reduction techniques such as t-SNE or PCA to visualize the overlap between real and synthetic data. Some studies have also incorporated clinical or expert validation, including blinded diagnostic assessments or simulated clinical trials, to provide a real-world evaluation of synthetic

Table 17
GAN techniques for tabular EHR and clinical data

GAN Technique / Model	Evaluation Metrics	Key Findings / Performance	Strengths / Limitations
CTGAN	MMD, TSTR	<ul style="list-style-type: none"> Preserves categorical correlations High fidelity in mixed-type tabular data 	<ul style="list-style-type: none"> Effective for structured data Limited for imaging Widely adopted for clinical datasets
MedGAN / CTAB-GAN+	MMD, TSTR	<ul style="list-style-type: none"> Captures conditional distributions Good downstream task performance 	<ul style="list-style-type: none"> Conditional generation improves realism May require large datasets
TVAE	MMD, FID	<ul style="list-style-type: none"> Comparable to CTGAN on small datasets Preserves statistical patterns 	<ul style="list-style-type: none"> Effective for smaller datasets Slightly less robust on highly imbalanced features
CopulaGAN	MMD	<ul style="list-style-type: none"> Strong preservation of distributional characteristics 	<ul style="list-style-type: none"> Limited handling of categorical variables Mainly for numeric tabular data
Fairness-aware GANs (BT-GAN, FairGAN)	MMD, TSTR, fairness metrics	<ul style="list-style-type: none"> Reduces demographic bias Maintains clinical validity 	<ul style="list-style-type: none"> Fairness–utility trade-off Limited real-world validation

Table 18
GAN techniques for imaging and multi-omics data

GAN Technique / Model	Application Domain	Evaluation Metrics	Key Findings / Performance	Strengths / Limitations
CycleGAN	Imaging (CT ↔ MRI)	FID, IS, blinded diagnostic studies	<ul style="list-style-type: none"> High-quality cross-modality synthesis Preserves anatomical structures 	<ul style="list-style-type: none"> Effective for paired/unpaired translation Primarily for imaging Strong visual realism Not applicable to tabular data Preclinical validation only
STNG / Trans-cGAN	Imaging	FID, IS	<ul style="list-style-type: none"> High-fidelity image synthesis Captures complex anatomical features 	<ul style="list-style-type: none"> High visual fidelity Specialized for image-based tasks Limited tabular application Combines multiple architectures May require complex training pipelines
StyleGAN / StyleGAN2-ADA	Imaging / rare disease phenotyping	FID, IS	<ul style="list-style-type: none"> Generates visually realistic images Style-based latent control 	<ul style="list-style-type: none"> High visual fidelity Specialized for image-based tasks Limited tabular application Combines multiple architectures May require complex training pipelines
Hybrid GANs (MOSA, HAGAN, EnhGAN, Spatial-Frequency)	Multi-omic, imaging	MMD, TSTR, FID	<ul style="list-style-type: none"> Multi-domain augmentation Preserves structural and frequency features 	<ul style="list-style-type: none"> High visual fidelity Specialized for image-based tasks Limited tabular application Combines multiple architectures May require complex training pipelines

data. Together, these metrics provide a multifaceted assessment encompassing statistical, visual, and functional performance.

Despite these validations, clinical evaluation remains limited; only a few studies perform blinded diagnostic assessments, trial simulations, or downstream task evaluation, highlighting the ongoing “bench-to-bedside” gap. While GANs generally demonstrate high fidelity and structural realism, limitations persist in generalizability, categorical feature handling, and dataset-specific tuning, with preclinical evaluation predominating and no methods yet fully transitioning to large-scale clinical trials. Fairness-aware GANs partially mitigate bias in tabular datasets but require careful trade-offs between data utility and demographic parity.

Moreover, regulatory pathways and ethical governance for GAN-generated data are underdeveloped, and privacy risks, including membership inference and linkage attacks, as well as potential misuse, must be addressed alongside considerations of fidelity and fairness. Integrating ethical checklists, consent frameworks, and transparency measures will be critical for real-world translation. Overall, the technical performance of state-of-the-art GANs is promising, particularly for domain-specific applications; however, robust clinical validation, standardized evaluation, and ethical safeguards remain essential for their safe and effective deployment in healthcare research and practice.

4.4. Limitations and future directions

This systematic review provides a comprehensive review of GAN applications in medical data generation, but several significant constraints still warrant attention. First, the rapid development of GAN research means that, despite a rigorous selection process, the dataset may not fully capture the latest architectural advancements or novel variants that emerge after the deadline. This is an inherent challenge in any rapidly evolving field, suggesting that continuous updates or dynamic systematic reviews are crucial to maintain an up-to-date understanding of the state-of-the-art.

Second, limiting research to English publications and peer-reviewed sources may have excluded relevant studies, including null or negative results, thus introducing potential publication bias. Future research could mitigate this limitation by collaborating with multilingual teams or professional translation services, thereby expanding inclusion criteria and ensuring a more globally representative perspective.

Third, the screening and data extraction were performed by a single reviewer, which could introduce reviewer bias due to errors or

inconsistencies. While structured extraction templates and predefined eligibility criteria were used to mitigate this risk, future multi-reviewer validation would further strengthen the reliability of the study.

Fourth, significant methodological heterogeneity exists across the studies, including differences in dataset characteristics, model architectures, and evaluation metrics, which complicates direct comparison and synthesis of findings. Many studies relied on small or single-source datasets, lacked external validation, or selectively reported results, thus limiting generalizability. The lack of widely accepted benchmarks and standardized reporting practices remains a key obstacle. Moving forward, the development and adoption of community standards and the sharing of datasets are essential to facilitate more consistent and transparent assessment of GAN models in the medical field.

Fifth, despite technical advances, the majority of reviewed studies remain at the preclinical stage. Only a small number included preliminary clinical validation, and none have progressed to full-scale clinical trials, highlighting a significant “bench-to-bedside” gap. Regulatory pathways for integrating GAN-synthesized data into clinical practice have not yet been established. Agencies such as the FDA may require extensive validation to ensure both fidelity and patient safety, representing a key barrier to clinical translation. Addressing these gaps is crucial for translating synthetic data from methodological research into safe and effective clinical applications.

Lastly, the stringent inclusion criteria prioritized methodological quality and relevance but excluded studies without empirical GAN applications, purely methodological papers, duplicate datasets, or studies lacking validation metrics. While this approach ensures rigor, it may overlook preliminary or exploratory work that has the potential to introduce impactful innovations. Future work can incorporate more nuanced, tiered review approach to include such early-stage research to track emerging trends and assess their maturation over time.

5. Conclusion

This review has examined the rapidly evolving landscape of GAN-based synthetic data generation in healthcare, highlighting its transformative potential and the challenges that shape its current trajectory. Evidence from imaging, tabular health records, and emerging applications such as genomics suggests that GANs can meaningfully augment or even substitute real-world datasets by preserving statistical fidelity, enhancing privacy, and enabling downstream analysis where data scarcity would otherwise hinder progress. At the same time, novel adaptations such as CTGAN variants and fairness-aware models signal

an ongoing shift from proof-of-concept studies toward more mature, problem-oriented implementations.

Despite these advances, significant obstacles remain. The predominance of single-source or small-scale datasets, along with methodological variability in model architectures and evaluation metrics, limits the comparability of results and the generalizability of insights. Many reports of progress emphasize technical similarity rather than clinical relevance, raising concerns about its translation into real-world practical application. Ethical issues, including privacy guarantees, fairness, and the environmental cost of training, remain underexplored, particularly when considering deployment in sensitive clinical settings.

Taken together, these findings underscore both the promise and the fragility of current GAN-based approaches. To realize their potential, future research should adopt standardized, transparent benchmarks that extend beyond technical accuracy to include fairness, robustness, and clinical utility. Broader collaboration across institutions and disciplines will be critical to ensure diversity of datasets and to capture insights that single studies cannot provide. Furthermore, integrating ethical and regulatory considerations from the outset, rather than as afterthoughts, will be essential for building trust and fostering adoption.

Ultimately, GANs represent more than a technical innovation; they embody a paradigm shift in how healthcare data may be generated, shared, and applied. Their future impact will depend not only on advances in architectural design but also on the willingness of the research community to embrace inclusivity, standardization, and responsible innovation. By addressing these dimensions, GAN-based synthetic data generation can progress from a promising research tool to a cornerstone of equitable, scalable, and secure healthcare research and practice.

Recommendations

To advance GAN-based synthetic medical data research, future reviews should adopt dynamic and regularly updated literature collection strategies to capture emerging developments while preserving connections to foundational work. Broadening the scope to include non-English and regional publications can enhance the inclusivity and global relevance of findings. Additionally, promoting methodological standardization in datasets, architectures, and evaluation metrics would improve study comparability. Future syntheses should also track and evaluate promising yet understudied GAN variants to uncover novel directions, while fostering interdisciplinary collaboration to integrate insights from fields such as bioinformatics and AI ethics for a more comprehensive understanding of synthetic medical data applications.

Acknowledgement

We extend our sincere gratitude to the Center of Research Excellence and Innovation Management (CREIM) of Universiti Sultan Zainal Abidin (UniSZA) for their generous financial support toward the publication of this work. Special appreciation is also due to the dedicated project team at UniSZA, particularly Nur Laila Najwa Josdi, a PhD candidate, for her technical contributions, and Dr. Wan Aezwani Wan Abu Bakar, who served as her supervisor, for providing conceptual guidance and technical proofreading. Our thanks are also extended to Assoc. Prof. Ts. Dr. Mustafa Man from Universiti Malaysia Terengganu (UMT) for facilitating network linkages and offering valuable feedback as a reviewer. Additionally, we are grateful for the international collaboration and insights provided by Dr. Evizal Abdul Kadir, Research Fellow at Universitas Islam Riau, Indonesia, and Dr. Bishwajeet Pandey from G.L. Bajaj Institute of Technology and Management, India whose contributions enriched this study.

Funding Support

This work is co-sponsored by Center for Research Excellence and Innovation Management (CREIM) of Universiti Sultan Zainal Abidin (UniSZA) and Research Management Innovation Center (RMIC) of Universiti Malaysia Terengganu (UMT).

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

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Author Contribution Statement

Wan Aezwani Wan Abu Bakar: Conceptualization, Methodology, Formal analysis, Resources, Visualization, Supervision, Funding acquisition. **Nur Laila Najwa Josdi:** Software, Data curation, Writing – original draft, Writing – review & editing, Funding acquisition. **Mustafa Man:** Validation, Investigation, Supervision, Project administration. **Evizal Abdul Kadir:** Resources. **Bishwajeet Kumar Pandey:** Project administration.

References

- Al-Antari, M. A. (2023). Artificial intelligence for medical diagnostics—Existing and future AI technology. *Diagnostics*, 13(4), 688. <https://doi.org/10.3390/diagnostics13040688>
- Woerner, S., & Baumgartner, C. F. (2024). Navigating data scarcity using foundation models: A benchmark of few-shot and zero-shot learning approaches in medical imaging. In *International Workshop on Foundation Models for General Medical AI*, 30–39. https://doi.org/10.1007/978-3-031-73471-7_4
- Abgrall, G., Monnet, X., & Arora, A. (2025). Synthetic data and health privacy. *JAMA*, 333(7), 567–568. <https://doi.org/10.1001/jama.2024.25821>
- Sufi, F. (2024). Addressing data scarcity in the medical domain: a GPT-based approach for synthetic data generation and feature extraction. *Information*, 15(5), 264. <https://doi.org/10.3390/info15050264>
- Arora, A., & Arora, A. (2022). Generative adversarial networks and synthetic patient data: Current challenges and future perspectives. *Future Healthcare Journal*, 9(2), 190–193. <https://doi.org/10.7861/fhj.2022-0013>
- Isasa, I., Hernandez, M., Epelde, G., Londoño, F., Beristain, A., Larrea, X., ... & Konstantinidis, E. (2024). Comparative assessment of synthetic time series generation approaches in healthcare: Leveraging patient metadata for accurate data synthesis. *BMC Medical Informatics and Decision Making*, 24(1), 27. <https://doi.org/10.1186/s12911-024-02427-0>
- Susser, D., Schiff, D. S., Gerke, S., Cabrera, L. Y., Cohen, I. G., Doerr, M., ... & Wagner, J. K. (2024). Synthetic health data: Real ethical promise and peril. *Hastings Center Report*, 54(5), 8–13. <https://doi.org/10.1002/hast.4911>

[8] van Kesteren, E. J. (2024). To democratize research with sensitive data, we should make synthetic data more accessible. *Patterns*, 5(9). <https://doi.org/10.1016/j.patter.2024.101049>

[9] Stadler, T., Oprisanu, B., & Troncoso, C. (2022). Synthetic data-anonymisation groundhog day. In *31st USENIX Security Symposium*, 1451–1468.

[10] Wei, W., Liu, L., Wu, Y., Su, G., & Iyengar, A. (2021). Gradient-leakage resilient federated learning. In *41st International Conference on Distributed Computing Systems*, 797–807. <https://doi.org/10.1109/ICDCS51616.2021.00081>

[11] De Cristofaro, E. (2024). Synthetic data: Methods, use cases, and risks. *IEEE Security & Privacy*, 22(3), 62–67. <https://doi.org/10.1109/MSEC.2024.3371505>

[12] Liu, F., & Panagiotakos, D. (2022). Real-world data: A brief review of the methods, applications, challenges and opportunities. *BMC Medical Research Methodology*, 22(1), 287. <https://doi.org/10.1186/s12874-022-01768-6>

[13] Elendu, C., Amaechi, D. C., Okatta, A. U., Amaechi, E. C., Elendu, T. C., Ezeh, C. P., & Elendu, I. D. (2024). The impact of simulation-based training in medical education: A review. *Medicine*, 103(27), e38813. <https://doi.org/10.1097/MD.00000000000038813>

[14] Chen, D. (2020). Open data: Implications on privacy in health-care research. *Blockchain in Healthcare Today*, 3, 10-30953. <https://doi.org/10.30953/bhtry.v3.144>

[15] Islam, T., Hafiz, M. S., Jim, J. R., Kabir, M. M., & Mridha, M. F. (2024). A systematic review of deep learning data augmentation in medical imaging: Recent advances and future research directions. *Healthcare Analytics*, 5, 100340. <https://doi.org/10.1016/j.health.2024.100340>

[16] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.

[17] Nikolentzos, G., Vazirgiannis, M., Xypolopoulos, C., Lingman, M., & Brandt, E. G. (2023). Synthetic electronic health records generated with variational graph autoencoders. *NPJ Digital Medicine*, 6(1), 83. <https://doi.org/10.1038/s41746-023-00822-x>

[18] Zhang, Z., Yang, L., & Zheng, Y. (2018). Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9242–9251. <https://doi.org/10.1109/CVPR.2018.00963>

[19] Wang, T., Lei, Y., Fu, Y., Wynne, J. F., Curran, W. J., Liu, T., & Yang, X. (2021). A review on medical imaging synthesis using deep learning and its clinical applications. *Journal of Applied Clinical Medical Physics*, 22(1), 11–36. <https://doi.org/10.1002/acm2.13121>

[20] Yang, X., Lin, Y., Wang, Z., Li, X., & Cheng, K. T. (2019). Bi-modality medical image synthesis using semi-supervised sequential generative adversarial networks. *IEEE Journal of Biomedical and Health Informatics*, 24(3), 855–865. <https://doi.org/10.1109/JBHI.2019.2922986>

[21] Feng, Y., Zhang, B., Xiao, L., Yang, Y., Gegen, T., & Chen, Z. (2024). Enhancing medical imaging with gans synthesizing realistic images from limited data. In *International Conference on Electronic Technology, Communication and Information*, 1192–1197. <https://doi.org/10.1109/ICETCI61221.2024.10594540>

[22] Alanazi, T. M., & Mercorelli, P. (2024). Precision denoising in medical imaging via generative adversarial network-aided low-noise discriminator technique. *Mathematics*, 12(23), 2227–7390. <https://doi.org/10.3390/math12233705>

[23] Xu, L., Skoulikidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional gan. *Advances in Neural Information Processing Systems*, 32.

[24] Arvanitis, T. N., White, S., Harrison, S., Chaplin, R., & Despotou, G. (2022). A method for machine learning generation of realistic synthetic datasets for validating healthcare applications. *Health Informatics Journal*, 28(2), 14604582221077000. <https://doi.org/10.1177/14604582221077000>

[25] Zhao, Z., Kunar, A., Birke, R., Van der Scheer, H., & Chen, L. Y. (2024). Ctab-gan+: Enhancing tabular data synthesis. *Frontiers in Big Data*, 6, 1296508. <https://doi.org/10.3389/fdata.2023.1296508>

[26] Alqlaity, 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>

[27] Baowaly, M. K., Lin, C. C., Liu, C. L., & Chen, K. T. (2019). Synthesizing electronic health records using improved generative adversarial networks. *Journal of the American Medical Informatics Association*, 26(3), 228–241. <https://doi.org/10.1093/jamia/ocy142>

[28] Yelmen, B., Decelle, A., Ongaro, L., Marnetto, D., Tallec, C., Montinaro, F., ... & Jay, F. (2021). Creating artificial human genomes using generative neural networks. *PLoS Genetics*, 17(2), e1009303. <https://doi.org/10.1371/journal.pgen.1009303>

[29] Hazra, D., Kim, M. R., & Byun, Y. C. (2022). Generative adversarial networks for creating synthetic nucleic acid sequences of cat genome. *International Journal of Molecular Sciences*, 23(7), 3701. <https://doi.org/10.3390/ijms23073701>

[30] Zrimec, J., Fu, X., Muhammad, A. S., Skrekas, C., Jauniskis, V., Speicher, N. K., ... & Zelezniak, A. (2022). Controlling gene expression with deep generative design of regulatory DNA. *Nature Communications*, 13(1), 5099. <https://doi.org/10.1038/s41467-022-32818-8>

[31] Marouf, M., Machart, P., Bansal, V., Kilian, C., Magruder, D. S., Krebs, C. F., & Bonn, S. (2020). Realistic in silico generation and augmentation of single-cell RNA-seq data using generative adversarial networks. *Nature Communications*, 11(1), 166. <https://doi.org/10.1038/s41467-019-14018-z>

[32] Chaudhari, P., Agrawal, H., & Kotecha, K. (2020). Data augmentation using MG-GAN for improved cancer classification on gene expression data. *Soft Computing*, 24(15), 11381–11391. <https://doi.org/10.1007/s00500-019-04602-2>

[33] Ahmed, K. T., Sun, J., Cheng, S., Yong, J., & Zhang, W. (2022). Multi-omics data integration by generative adversarial network. *Bioinformatics*, 38(1), 179–186. <https://doi.org/10.1093/bioinformatics/btab608>

[34] Xu, Y., Zhang, Z., You, L., Liu, J., Fan, Z., & Zhou, X. (2020). ScIGANs: Single-cell RNA-seq imputation using generative adversarial networks. *Nucleic Acids Research*, 48(15), e85–e85. <https://doi.org/10.1093/nar/gkaa506>

[35] Zinati, Y., Takiddeen, A., & Emad, A. (2024). GRouNdGAN: GRN-guided simulation of single-cell RNA-seq data using causal generative adversarial networks. *Nature Communications*, 15(1), 4055. <https://doi.org/10.1038/s41467-024-48516-6>

[36] Nouri, N. (2025). Single-cell RNA-seq data augmentation using generative Fourier transformer. *Communications Biology*, 8(1), 113. <https://doi.org/10.1038/s42003-025-07552-8>

[37] Safari, M., Fatemi, A., & Archambault, L. (2023). MedFusionGAN: Multimodal medical image fusion using an unsupervised deep generative adversarial network. *BMC Medical Imaging*, 23(1), 203. <https://doi.org/10.1186/s12880-023-01160-w>

[38] Huang, Z. H., Chen, L., Sun, Y., Liu, Q., & Hu, P. (2024). Conditional generative adversarial network driven radiomic prediction of mutation status based on magnetic resonance imaging of breast cancer. *Journal of Translational Medicine*, 22(1), 226. <https://doi.org/10.1186/s12967-024-05018-9>

[39] Liu, M., Liu, Y., Cui, H., Li, C., & Ma, J. (2023). Mgct: Mutual-guided cross-modality transformer for survival outcome prediction using integrative histopathology-genomic features. In *International Conference on Bioinformatics and Biomedicine*, 1306–1312. <https://doi.org/10.1109/BIBM58861.2023.10385897>

[40] Hemker, K., Simidjevski, N., & Jamnik, M. (2024). Healnet: Multimodal fusion for heterogeneous biomedical data. *Advances in Neural Information Processing Systems*, 37, 64479–64498.

[41] 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 *Machine Learning for Healthcare Conference*, 286–305.

[42] Li, H., Wang, H., Sun, X., He, H., & Feng, J. (2025). Context-enhanced framework for medical image report generation using multimodal contexts. *Knowledge-Based Systems*, 310, 112913. <https://doi.org/10.1016/j.knosys.2024.112913>

[43] Syntegra. (2022). Syntegra and the institute for health metrics expand strategic partnership to bring privacy-guaranteed healthcare data to life sciences. Retrieved from: <https://www.prnewswire.com/news-releases/syntegra-and-the-institute-for-health-metrics-expand-strategic-partnership-to-bring-privacy-guaranteed-healthcare-data-to-life-sciences-301652498.html>

[44] Kartoun, U. (2019). Advancing informatics with electronic medical records bots (EMRBots). *Software Impacts*, 2, 100006. <https://doi.org/10.1016/j.simpa.2019.100006>

[45] Greenberg, J. K., Landman, J. M., Kelly, M. P., Pennicooke, B. H., Molina, C. A., Foraker, R. E., & Ray, W. Z. (2023). Leveraging artificial intelligence and synthetic data derivatives for spine surgery research. *Global Spine Journal*, 13(8), 2409–2421. <https://doi.org/10.1177/21925682221085535>

[46] Hong, S., Marinescu, R., Dalca, A. V., Bonkhoff, A. K., Bretzner, M., Rost, N. S., & Golland, P. (2021). 3D-StyleGAN: A style-based generative adversarial network for generative modeling of three-dimensional medical images. In S. Engelhardt, I. Oksuz, D. Zhu, Y. Yuan, A. Mukhopadhyay, N. Heller, ..., & Y. Xue (Eds.), *Deep Generative Models, and Data Augmentation, Labelling, and Imperfections* (pp. 24–34). Springer. https://doi.org/10.1007/978-3-030-88210-5_3

[47] HealthVerity. (2024). Advancing COVID-19 research with the National Cancer Institute. Retrieved from: <https://blog.healthverity.com/case-study-partnering-with-nci>

[48] Ghebrehiwet, I., Zaki, N., Damseh, R., & Mohamad, M. S. (2024). Revolutionizing personalized medicine with generative AI: A systematic review. *Artificial Intelligence Review*, 57(5), 128. <https://doi.org/10.1007/s10462-024-10768-5>

[49] Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., & Abbeel, P. (2016). Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in Neural Information Processing Systems*, 29.

[50] Nie, D., Trullo, R., Lian, J., Petitjean, C., Ruan, S., Wang, Q., & Shen, D. (2017). Medical image synthesis with context-aware generative adversarial networks. In M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, & S. Duchesne (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017* (pp. 417–425). Springer. https://doi.org/10.1007/978-3-319-66179-7_48

[51] Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. *International Conference on Machine Learning*, 214–223.

[52] Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. (2017). Improved training of wasserstein gans. *Advances in Neural Information Processing Systems*, 30.

[53] Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*. <https://doi.org/10.48550/arXiv.1802.05957>

[54] Dar, S. U., Yurt, M., Karacan, L., Erdem, A., Erdem, E., & Cukur, T. (2019). Image synthesis in multi-contrast MRI with conditional generative adversarial networks. *IEEE Transactions on Medical Imaging*, 38(10), 2375–2388. <https://doi.org/10.1109/TMI.2019.2901750>

[55] Karras, T., Laine, S., & Aila, T. (2019). A style-based generator architecture for generative adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4401–4410. <https://doi.org/10.1109/CVPR.2019.00453>

[56] Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33, 12104–12114.

[57] Xu, L., Skoulikidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional gan. *Advances in Neural Information Processing Systems*, 32.

[58] Sun, C., van Soest, J., & Dumontier, M. (2023). Generating synthetic personal health data using conditional generative adversarial networks combining with differential privacy. *Journal of Biomedical Informatics*, 143, 104404. <https://doi.org/10.1016/j.jbi.2023.104404>

[59] Li, Y., Han, N., Qin, Y., Zhang, J., & Su, J. (2023). Trans-cGAN: Transformer-Unet-based generative adversarial networks for cross-modality magnetic resonance image synthesis. *Statistics and Computing*, 33(5), 113. <https://doi.org/10.1007/s11222-023-10282-8>

[60] Cai, Z., Apolinário, S., Baião, A. R., Pacini, C., Sousa, M. D., Vinga, S., ... & Gonçalves, E. (2024). Synthetic augmentation of cancer cell line multi-omic datasets using unsupervised deep learning. *Nature Communications*, 15(1), 10390. <https://doi.org/10.1038/s41467-024-54771-4>

[61] Liu, Q., Zhou, T., Cheng, C., Ma, J., & Hoque Tania, M. (2025). Hybrid generative adversarial network based on frequency and spatial domain for histopathological image synthesis. *BMC Bioinformatics*, 26(1), 29. <https://doi.org/10.1186/s12859-025-06057-9>

[62] Hamghalam, M., & Simpson, A. L. (2024). Medical image synthesis via conditional GANs: Application to segmenting brain tumours. *Computers in Biology and Medicine*, 170, 107982. <https://doi.org/10.1016/j.combiomed.2024.107982>

[63] Ju, Z., Zhou, W., Kong, L., Chen, Y., Li, Y., Sun, Z., & Shan, C. (2025). Hagan: Hybrid augmented generative adversarial network for medical image synthesis. *Machine Intelligence Research*, 1–14. <https://doi.org/10.1007/s11633-024-1528-y>

[64] Abdusalomov, A. B., Nasimov, R., Nasimova, N., Muminov, B., & Whangbo, T. K. (2023). Evaluating synthetic medical images using artificial intelligence with the GAN algorithm. *Sensors*, 23(7), 3440. <https://doi.org/10.3390/s23073440>

[65] Wang, J., Tabassum, N., Toma, T. T., Wang, Y., Gahlmann, A., & Acton, S. T. (2022). 3D GAN image synthesis and dataset quality

assessment for bacterial biofilm. *Bioinformatics*, 38(19), 4598–4604. <https://doi.org/10.1093/bioinformatics/btac529>

[66] 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>

[67] Rashidian, S., Wang, F., Moffitt, R., Garcia, V., Dutt, A., Chang, W., ... & Saltz, J. (2020). SMOOTH-GAN: Towards sharp and smooth synthetic EHR data generation. In M. Michalowski, & R. Moskovitch (Eds.), *Artificial Intelligence in Medicine* (pp. 37–48). Springer. https://doi.org/10.1007/978-3-030-59137-3_4

[68] Achterberg, J. L., Haas, M. R., & Spruit, M. R. (2024). On the evaluation of synthetic longitudinal electronic health records. *BMC Medical Research Methodology*, 24(1), 181. <https://doi.org/10.1186/s12874-024-02304-4>

[69] Theodorou, B., Xiao, C., & Sun, J. (2023). Synthesize high-dimensional longitudinal electronic health records via hierarchical autoregressive language model. *Nature Communications*, 14(1), 5305. <https://doi.org/10.1038/s41467-023-41093-0>

[70] Eckardt, J. N., Hahn, W., Röllig, C., Stasik, S., Platzbecker, U., Müller-Tidow, C., ... & Middeke, J. M. (2024). Mimicking clinical trials with synthetic acute myeloid leukemia patients using generative artificial intelligence. *NPJ Digital Medicine*, 7(1), 76. <https://doi.org/10.1038/s41746-024-01076-x>

[71] Yang, S., Kim, K. D., Ariji, E., Takata, N., & Kise, Y. (2023). Evaluating the performance of generative adversarial network-synthesized periapical images in classifying C-shaped root canals. *Scientific Reports*, 13(1), 18038. <https://doi.org/10.1038/s41598-023-45290-1>

[72] Acharyya, M., Moharana, B., Jain, S., & Tandon, M. (2024). A double-blinded study for quantifiable assessment of the diagnostic accuracy of AI tool “ADVEN-i” in identifying diseased fundus images including diabetic retinopathy on a retrospective data. *Indian Journal of Ophthalmology*, 72, S46-S52. https://doi.org/10.4103/IJO.IJO_3342_22

[73] El Emam, K., Mosquera, L., Fang, X., & El-Hussuna, A. (2022). Utility metrics for evaluating synthetic health data generation methods: Validation study. *JMIR Medical Informatics*, 10(4), e35734. <https://doi.org/10.2196/35734>

[74] Alshanti, A., Rasheed, A., & Westad, F. (2025). Privacy re-identification attacks on tabular GANs. *Security and Privacy*, 8(1), e469. <https://doi.org/10.1002/spr.2469>

[75] 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>

[76] Yoon, J., Mizrahi, M., Ghalaty, N. F., Jarvinen, T., Ravi, A. S., Brune, P., ... & Pfister, T. (2023). EHR-Safe: Generating high-fidelity and privacy-preserving synthetic electronic health records. *NPJ Digital Medicine*, 6(1), 141. <https://doi.org/10.1038/s41746-023-00888-7>

[77] Hayes, J., Melis, L., Danezis, G., & De Cristofaro, E. (2017). Logan: Membership inference attacks against generative models. *arXiv Preprint: 1705.07663*. <https://doi.org/10.48550/arXiv.1705.07663>

[78] Ramachandranpillai, R., Sikder, M. F., Bergström, D., & Heintz, F. (2024). Bt-GAN: Generating fair synthetic healthdata via bias-transforming generative adversarial networks. *Journal of Artificial Intelligence Research*, 79, 1313–1341. <https://doi.org/10.1613/jair.1.15317>

[79] Wu, X., Xu, D., Yuan, S., & Zhang, L. (2022). Fair data generation and machine learning through generative adversarial networks.

In R. Razavi-Far, A. Ruiz-Carcia, V. Palade, & J. Schmidhuber (Eds.), *Generative adversarial learning: Architectures and applications* (pp. 31–55). Springer International Publishing. https://doi.org/10.1007/978-3-030-91390-8_3

[80] Costa, P., Galdran, A., Meyer, M. I., Niemeijer, M., Abràmoff, M., Mendonça, A. M., & Campilho, A. (2017). End-to-end adversarial retinal image synthesis. *IEEE Transactions on Medical Imaging*, 37(3), 781–791. <https://doi.org/10.1109/TMI.2017.2759102>

[81] Liu, J., Xu, S., He, P., Wu, S., Luo, X., Deng, Y., & Huang, H. (2024). VSG-GAN: A high-fidelity image synthesis method with semantic manipulation in retinal fundus image. *Biophysical Journal*, 123(17), 2815–2829. <https://doi.org/10.1016/j.bpj.2024.02.019>

[82] Pezoulas, V. C., Zardis, D. I., Mylona, E., Androutsov, 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>

[83] Hasan, S. S., Fury, M. S., Woo, J. J., Kunze, K. N., & Ramkumar, P. N. (2025). Ethical application of generative artificial intelligence in medicine. *Arthroscopy: The Journal of Arthroscopic & Related Surgery*, 41(4), 874–885. <https://doi.org/10.1016/j.artthro.2024.12.011>

[84] Farhud, D. D., & Zokaei, S. (2021). Ethical issues of artificial intelligence in medicine and healthcare. *Iranian Journal of Public Health*, 50(11), 1–5. <https://doi.org/10.18502/ijph.v50i11.7600>

[85] Dankwa-Mullan, I. (2024). Health equity and ethical considerations in using artificial intelligence in public health and medicine. *Preventing Chronic Disease*, 21, E64. <https://doi.org/10.5888/pcd21.240245>

[86] Nasir, M., Siddiqui, K., & Ahmed, S. (2025). Ethical-legal implications of AI-powered healthcare in critical perspective. *Frontiers in Artificial Intelligence*, 8, 1619463. <https://doi.org/10.3389/frai.2025.1619463>

[87] Boraschi, D., van der Schaar, M., Costa, A., & Milne, R. (2025). Governing synthetic data in medical research: The time is now. *The Lancet Digital Health*, 7(4), e233–e234. <https://doi.org/10.1016/j.landig.2025.01.012>

[88] Ning, Y., Teixayavong, S., Shang, Y., Savulescu, J., Nagaraj, V., Miao, D., ... & Liu, N. (2024). Generative artificial intelligence and ethical considerations in health care: A scoping review and ethics checklist. *The Lancet Digital Health*, 6(11), e848–e856. [https://doi.org/10.1016/S2589-7500\(24\)00143-2](https://doi.org/10.1016/S2589-7500(24)00143-2)

[89] Yu, S., Lee, S. S., & Hwang, H. (2024). The ethics of using artificial intelligence in medical research. *Kosin Medical Journal*, 39(4), 229–237. <https://doi.org/10.7180/kmj.24.140>

[90] Freeman, S., Wang, A., Saraf, S., Potts, E., McKimm, A., Coiera, E., & Magrabi, F. (2025). Developing an AI governance framework for safe and responsible AI in health care organizations: protocol for a multimethod study. *JMIR Research Protocols*, 14(1), e75702. <https://doi.org/10.2196/75702>

[91] Chen, Y., & Esmaeilzadeh, P. (2024). Generative AI in medical practice: in-depth exploration of privacy and security challenges. *Journal of Medical Internet Research*, 26, e53008. <https://doi.org/10.2196/53008>

How to Cite: Bakar, W. A. W. A., Josdi, N. L. N., Man, M., Kadir, E. A., & Pandey, B. K. (2025). Evolution of Generative Adversarial Networks (GANs) in Medicine: A Systematic Review of Architectures, Applications, and Implementation Challenges. *Artificial Intelligence and Applications*. <https://doi.org/10.47852/bonviewAIA52026216>