FUTURE TECHNO SCIENCE

Journal of Future Artificial Intelligence and Technologies

E-ISSN: 3048-3719

Research Article

Generative AI in Healthcare: An Analytical Review of Models, Clinical Applications, and Decision-Support Implications

Naglaa Fadul [1], Mona Fahad Alaskar [2], Kamal Bakari Jillahi [3,*], and Dalia Bassem El-Khaled [1]

[1] Faculty of Computer Studies, Arab Open University, Riyadh 11681, Saudi Arabia; e-mail : pt349.fcs@arabou.edu.sa; d.elkhaled@arabou.sa
[2] Faculty of Computer Science, Arab East College, Riyadh 11583, Saudi Arabia; e-mail: mfalaskar@arabeast.edu.sa
[3] Department of Computer Science, American University of Nigeria, Yola 800096, Nigeria; e-mail : kamal.bakari@aun.edu.ng
* Corresponding Author : Kamal Bakari Jillahi

Abstract: This review examines the rapidly expanding landscape of Generative Artificial Intelligence (GenAI) in healthcare, focusing on how models such as GANs, VAEs, diffusion models, and large language models are being explored across medical imaging, clinical documentation, synthetic data generation, drug discovery, and decision-support workflows. Despite GenAI's growing influence, persistent challenges, including limited annotated datasets, concerns over model generalizability, privacy risks, and the opacity of generative architectures, underscore the need for careful evaluation and governance. Accordingly, this study aims to map current applications, assess methodological and ethical constraints, and identify future research opportunities. Using a structured search across ScienceDirect, Scopus, and other sources, the study follows a structured narrative review complemented by quantitative descriptive analysis of the literature. The review also adopts PRISMA-guided screening and standardized data extraction, the review synthesizes evidence from 110 studies published up to October 2025. The findings indicate that the literature frequently reports improvements in imaging quality, data augmentation, molecular modeling workflows, and clinical documentation through generative approaches, particularly in technically constrained settings; however, evidence of clinically validated impact remains uneven across domains. While issues of bias, hallucination, and limited interpretability persist as significant obstacles, imaging-focused applications appear comparatively more mature than decision-support and patient-level modeling tasks. Across domains, diffusion models are commonly associated with higher visual fidelity in biomedical image generation, whereas LLMs demonstrate promise in narrative-oriented tasks but require stronger factual grounding and external verification mechanisms. Overall, the evidence suggests that GenAI's potential in healthcare is highly context-dependent and contingent upon robust validation frameworks, transparent governance, and human-in-the-loop oversight. The review concludes that responsible integration of GenAI—guided by ethical, legal, and clinical safeguards, will be essential for ensuring safe, equitable, and sustainable adoption in healthcare research, delivery, and policy.

Keywords: Clinical Decision Support; Ethical Implications; Generative AI; Healthcare Artificial Intelligence; Large Language Models; Medical Imaging; Synthetic Data; Trustworthy AI.

Received: November, 26th 2025
Revised: December, 10th 2025
Accepted: December, 16th 2025
Published: December, 31st 2025
Curr. Ver.: December, 31st 2025

1. Introduction

Generative AI (GenAI) continues to expand rapidly in healthcare, enhancing diagnostic precision, streamlining clinical workflows, and enabling personalized therapeutic strategies [1]–[3]. Amid growing clinical complexity and increasing patient expectations, AI offers transformative potential to improve patient outcomes, optimize resource utilization, and stimulate innovation in medical methodologies [4]–[6]. Within this landscape, GenAI has emerged as a disruptive paradigm capable of synthesizing realistic biomedical data, modeling complex biological systems, and automating narrative and decision-support functions [4], [7], [8].

Models such as GANs, VAEs, diffusion models, and LLMs address key challenges, including limited annotated datasets, privacy constraints, and the need for reliable simulation frameworks [1], [7], [9]–[12]. By generating clinically valid synthetic data, GenAI accelerates biomedical research, enhances diagnostic performance, and broadens access to emerging digital health technologies [1], [3], [4], [13].

This review critically examines GenAI's evolving applications in medical imaging, drug discovery, and electronic health record analytics, alongside the methodological, ethical, and regulatory dimensions shaping its integration into healthcare systems [1], [4], [6], [14]–[16]. It further outlines strategic directions for embedding GenAI responsibly within health systems to ensure outcomes that are equitable, transparent, and empirically grounded [1], [4], [17].

While several reviews have examined generative AI in healthcare, they tend to focus narrowly on specific model families, such as GAN-driven medical image synthesis, or on limited clinical domains, including radiology and drug discovery. Existing surveys often cover restricted time windows, omit multimodal generative architectures, or lack structured appraisals of methodological validity, explainability, and governance constraints. Importantly, prior reviews rarely integrate macro-level research trends or publication growth statistics, nor do they evaluate the maturity of GenAI applications across domains such as imaging, EHR synthesis, decision-support, molecular modeling, and digital twins. Moreover, the rapid emergence of international ethical and policy frameworks such as the WHO (2021) Ethics and Governance of AI in Health, the EU AI Act (2024), and evolving FDA/EMA guidance for adaptive AI/ML systems which has not been systematically linked to GenAI research trajectories in the existing literature. These gaps limit the ability of clinicians, policymakers, and researchers to understand how technical innovations intersect with governance, safety, and translational readiness.

This study conducts a structured narrative review, complemented by quantitative descriptive mapping of the literature, covering generative AI research in healthcare until October 2025. While PRISMA guidelines were adapted to inform the screening and selection process, the study does not aim to constitute a formal systematic literature review or meta-analysis, but rather to provide an integrative and cross-domain synthesis of existing work. In contrast to earlier reviews, it integrates (i) PRISMA-guided article selection, (ii) structured evidence mapping of GenAI applications, (iii) a taxonomy of model families, (iv) domain-wise maturity and methodological quality assessment, and (v) an evaluation of ethical, legal, and policy implications influencing clinical adoption. In addition, this review incorporates bibliometric-style trend visualizations to highlight the evolving research landscape across medical imaging, clinical documentation, synthetic data generation, molecular modeling, and decision-support systems.

The remainder of this paper is organized as follows: Section 2 describes the research protocol and PRISMA-based methodology. Section 3 introduces the foundational concepts and model families of generative AI. Section 4 contextualizes GenAI within healthcare systems, while Section 5 synthesizes application domains using a data-centric and patient-centric perspective. Sections 6 and 7 examine evaluation practices and ethical, legal, and social implications, respectively. Section 8 discusses key challenges and limitations, followed by future research directions in Section 9. Finally, Section 10 presents an integrative discussion and summary of key findings, and Section 11 concludes the paper.

2. Research Protocol

To systematically review the applications, benefits, risks, and future potential of generative AI technologies across healthcare domains, this review seeks to address the following research questions:

1. What are the current applications of generative AI in healthcare?
2. Which generative models are most commonly used in healthcare applications?
3. How are these applications evaluated and validated?
4. What ethical, technical, and clinical challenges arise from the application of these models?
5. What future research directions and opportunities can be identified?

2.1. Search Strategy for Reviewing Generative AI in Healthcare

The literature search was conducted across two major academic databases to ensure comprehensive coverage of both medical and technical perspectives:

- ScienceDirect: for clinically oriented research and medical applications
- Scopus: for interdisciplinary scholarly literature spanning computer science, engineering, and health sciences.
- Other Literature from sources such as preprints are also considered based on relevance and rigor.

2.2. Search String

The search strategy employed the following query string to retrieve relevant literature: "Generative AI" OR "Generative Artificial Intelligence" OR "GenAI" AND ("healthcare" OR "health care" OR "health system" OR "health").

2.3. Search Refinement

To ensure methodological rigor and relevance, the literature search was conducted across publications indexed up to October 2025, with primary analytical emphasis placed on recent advancements and emerging trends in generative AI research. Earlier studies were retained selectively to trace foundational concepts and methodological origins of specific generative AI paradigms. Only English-language sources were included. Citation chaining was employed to identify seminal and highly influential works, particularly those contributing to the historical development of generative models and their early adoption in healthcare contexts. Database-specific subject headings were used to improve search precision and coverage. These refinements ensured a focus on contemporary, high-impact research while maintaining continuity with foundational methodological literature.

2.4. Inclusion and Exclusion Criteria

The literature selection process prioritized peer-reviewed journal articles and systematic reviews to ensure methodological rigor and scholarly credibility. Non-technical commentaries, opinion pieces, and studies outside the healthcare context were excluded. The detailed eligibility framework, including inclusion and exclusion criteria, is summarized in Table 1.

Table 1. Eligibility criteria for inclusion and exclusion of studies.

Inclusion Criteria	Exclusion Criteria
Peer-reviewed articles, Conference papers and preprints	Non-scientific blog posts and news articles
Published between 1994 and 2025	Published before 1994
English-language publications	Non-English articles
Focused on healthcare applications of GenAI	Generic AI studies not specific to healthcare

To illustrate the outcome of applying these criteria, the overall search and screening process is summarized in Figure 1. A total of 394 records were initially identified and reduced to 250 after duplicate removal. Following title and abstract screening, 240 full-text articles were assessed for eligibility, of which 153 met the inclusion criteria. After excluding additional articles with methodological limitations or insufficient relevance, 110 studies were finalized for inclusion. Earlier publications were primarily retained to support historical tracing of generative AI paradigms, while the analytical synthesis emphasizes studies published during periods of higher research density.

2.5. Systematic Screening and Selection Methodology

The study selection process followed a three-stage screening protocol designed to ensure comprehensive and unbiased inclusion of relevant studies. Stage 1 involved two independent reviewers who evaluated each record against the predefined eligibility criteria and excluded clearly irrelevant studies. Stage 2 consisted of full-text screening to assess methodological soundness, relevance to generative AI applications in healthcare, and overall technical validity. Stage 3 applied PRISMA-compliant evaluation instruments to assess scientific merit, methodological transparency, and potential sources of bias. This multi-level screening process minimized selection bias and ensured the methodological integrity of the review.

2.6. Systematic Data Extraction and Synthesis Methodology

The included studies were systematically analyzed to capture both technical and clinical characteristics. A standardized data extraction template was developed to record model architecture, target healthcare domain, quantitative evaluation metrics, dataset attributes, clinical and technical outcomes, methodological limitations, and ethical considerations, as summarized in Table 5. The extracted data were organized across five thematic domains, enabling analysis of model performance, clinical validity, and implementation challenges. This integrated quantitative–qualitative synthesis facilitated the identification of cross-cutting trends, evidence gaps, and methodological best practices in healthcare-focused generative AI research, while maintaining transparency and reproducibility.

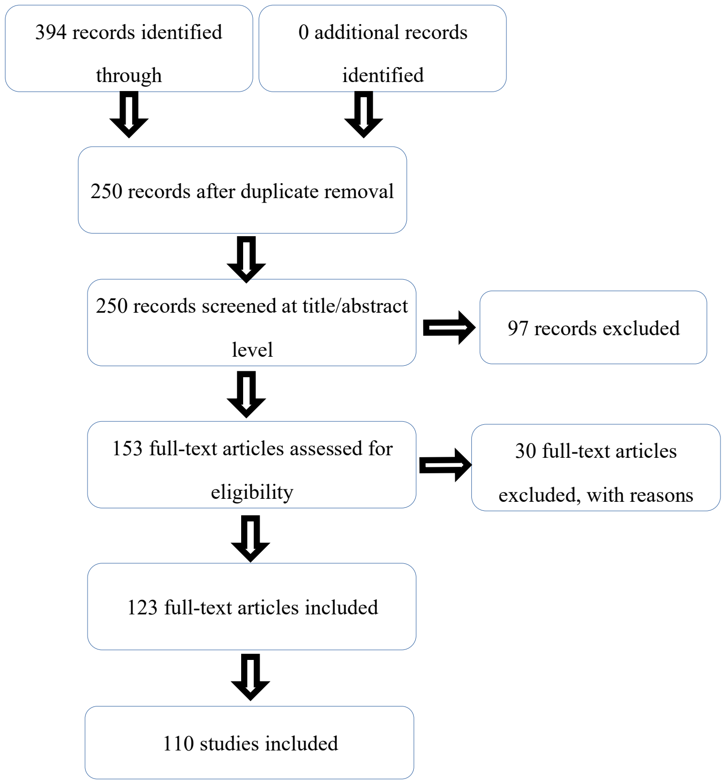

Figure 1. Literature search and assessment workflow

3. Foundations of Generative AI

Generative Artificial Intelligence (GenAI) refers to a class of algorithms designed to create new content such as images, text, audio, or structured data by learning statistical and semantic patterns from existing datasets [1], [4]. Its adoption in healthcare has been driven by advances in probabilistic modeling and unsupervised learning, extending earlier approaches such as Bayesian networks and principal component analysis [9], [18]. The introduction of Generative Adversarial Networks (GANs) by [9] demonstrated the capability of generator–discriminator architectures to synthesize data that are nearly indistinguishable from real clinical records, thereby enriching limited medical datasets [3], [19].

Figure 2 illustrates the evolving landscape of generative AI model adoption in healthcare, based on publications spanning the period 1994–2025. To ensure consistency with the overall bibliometric analysis and to account for sparse publication volume in early years, studies published prior to 2020 are aggregated into a single category (≤2020). As a result, the figure emphasizes temporal trends from 2020 onward, where publication density becomes sufficient to support year-level comparison. The stacked areas represent overlapping thematic classifications rather than mutually exclusive counts, reflecting the fact that individual studies may employ multiple generative model families. Early healthcare-focused generative AI research was dominated by image synthesis approaches based on GANs and VAEs. From 2022 onward, diffusion models and transformer-based large language models (LLMs) exhibit rapid growth, signaling a shift toward multimodal data generation and contextual reasoning. After 2023, hybrid architectures integrating LLMs with diffusion and VAE–GAN frameworks show accelerated adoption, indicating a transition toward unified generative systems capable of linking imaging, textual, and molecular data within a single learning paradigm.

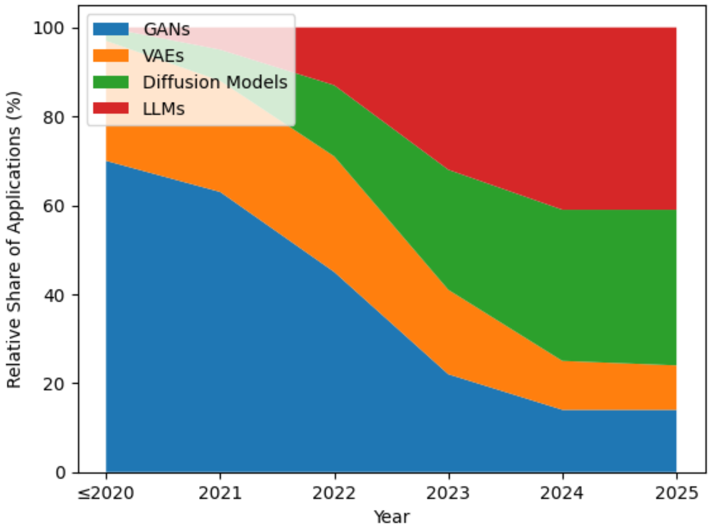

Figure 2. Shifting landscape of GenAI applications in healthcare.

Advances in Variational Autoencoders (VAEs) and diffusion models have provided powerful tools for medical image generation, patient outcome simulation, and disease progression modeling. Consequently, GenAI has gained prominence across diverse healthcare applications, including medical imaging, electronic health record analytics, natural language processing, drug discovery, and personalized medicine [19], [20]. These developments illustrate how GenAI both reinforces established biomedical paradigms and enables new models of digital innovation in healthcare [9], [21], [22].

Figure 3 presents publication trends across major healthcare application domains of generative AI, using the same temporal aggregation strategy as Figure 2. Publications prior to 2020 are grouped into a single ≤2020 category due to limited volume, while annual trends from 2020 to 2025 are shown explicitly. The reported values represent overlapping thematic classifications rather than mutually exclusive categories, as individual studies may contribute to multiple application domains.

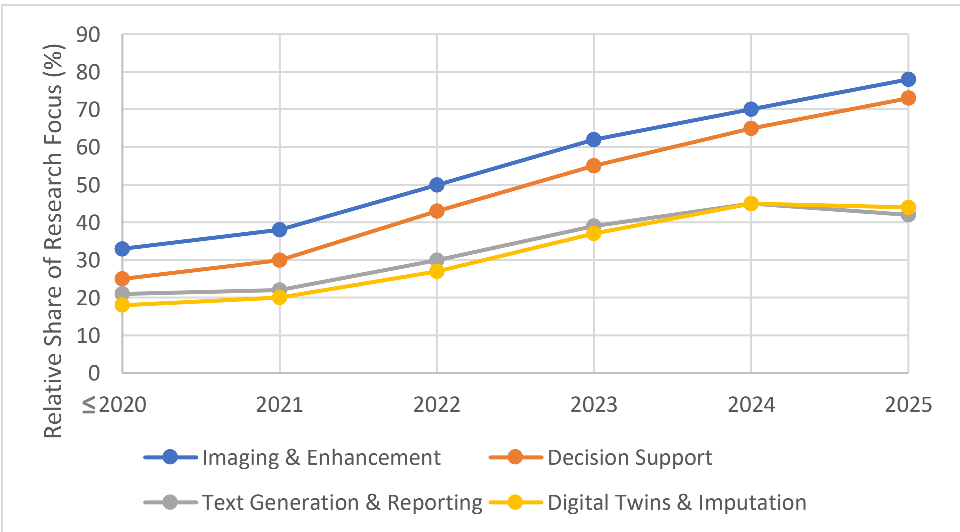

Figure 3. Publication trends by generative AI application domain.

As shown in Figure 3, medical imaging and image enhancement applications dominate the literature and exhibit rapid growth from 2020 onward, reflecting the maturation of GAN-, VAE-, and diffusion-based techniques for clinical imaging pipelines. Decision-support systems and clinical text generation demonstrate strong growth beginning in 2021, driven by increasing adoption of transformer-based LLMs and multimodal architectures. Digital twin and data-imputation applications remain comparatively less prevalent but show steady emergence after 2022. The apparent decline in early 2025 reflects partial-year indexing rather than a substantive reduction in research activity.

Taken together, Figures 2–4 provide a coherent bibliometric overview of the final corpus spanning 1994–2025. While early studies are aggregated due to sparse coverage, the period from 2020 onward reveals clear temporal and domain-specific research trajectories. Generative AI research in healthcare has transitioned from predominantly image-centric generation toward multimodal, language-driven, and decision-oriented applications, underscoring the increasing clinical integration and methodological diversification of generative models.

Figure 4 presents the distribution of publications included in the review by year. The reduced number of studies in 2025 reflects partial-year coverage, as the literature search was conducted in February. Overall, the trend highlights substantial growth in generative AI research compared with earlier years. To contextualize the representativeness of the final corpus (n = 110), a temporal distribution analysis was conducted by year of publication. Due to sparse coverage prior to 2020, studies published in 2018–2020 were aggregated into a single category (≤2020). 2023, 45 in 2024, and 21 in early 2025. Journal articles constituted approximately 82% of the corpus, while the remaining 18% comprised preprints (primarily arXiv and medRxiv).

Clinical text generation and automated reporting follow a similar upward trajectory, particularly with the widespread adoption of transformer-based LLMs after 2022. In contrast, digital twin and data-imputation applications remain emergent but demonstrate gradual growth, indicating increasing interest in simulation- and augmentation-driven clinical analytics. Overall, the upward trend through 2024 underscores expanding clinical integration of generative AI, while early 2025 figures likely underrepresent full-year output due to partial indexing.

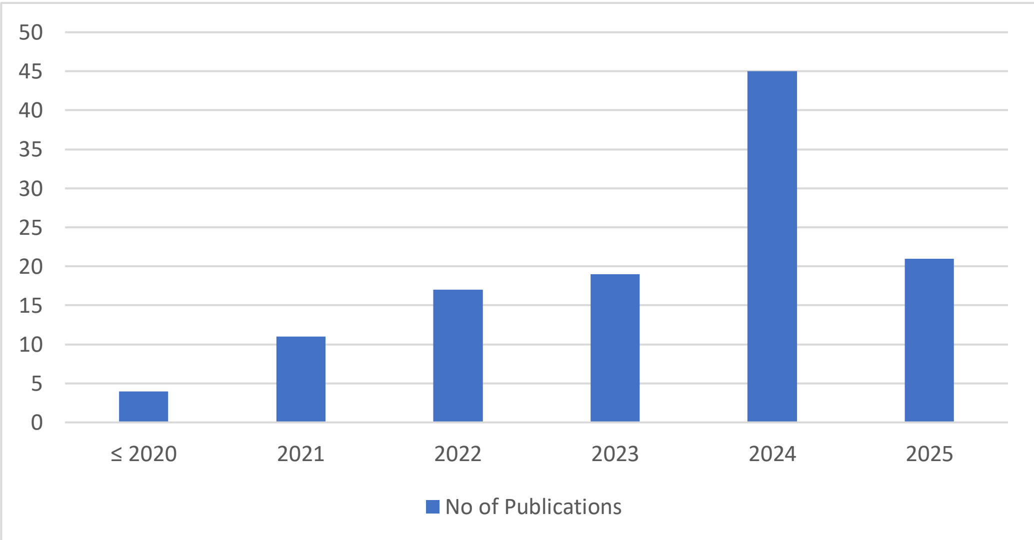

Figure 4. Number of publications included in the review by year.

Across model families, GANs remain the most widely adopted architecture, followed by VAEs and LLMs, while diffusion models and hybrid approaches are rapidly expanding the methodological diversity of GenAI applications in medical research [1], [4], [5]. The major GenAI model families identified in this review are discussed in the following subsections.

3.1 Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) are neural network architectures designed to synthesize artificial data instances [1], [4]. They have demonstrated significant efficacy in healthcare applications such as medical image synthesis, data augmentation, cross-domain translation, and the simulation of rare disease cases for diagnostic training [13], [23]. However, the deployment of GANs in clinical contexts presents challenges, including mode collapse, training instability, and limitations in clinical realism. Ongoing research addresses these issues through Wasserstein GANs (WGANs) and other stability-enhancing architectures [24]. Additionally, refinements in evaluation metrics and the incorporation of domain-specific constraints are advancing the development of clinically valid and ethically trustworthy synthetic datasets [13], [23]. Collectively, these innovations position GANs as foundational tools in the evolution of realistic, privacy-preserving, and diagnostically valuable healthcare data generation [12], [25].

3.2. Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs) are probabilistic generative models that encode input data into a structured latent space and reconstruct it to generate new, statistically similar instances [4]. They have shown promise in health informatics for generating synthetic electronic health records, reducing noise in medical imaging, and imputing missing clinical data [9], [14]. VAEs' well-organized latent representations facilitate anomaly detection and the modeling of patient-specific clinical trajectories [22]However, VAEs exhibit limitations in high-fidelity medical imaging, such as image blurriness and reduced structural sharpness, as summarized in Table 4 [24]. Hybrid approaches, including VAE–GAN architectures, integrate adversarial learning to enhance image realism while preserving the probabilistic strengths of VAEs [9]. Furthermore, disentangled VAE variants aim to improve the separability of clinical variables within the latent space, yielding more realistic, interpretable, and clinically useful synthetic data [9].

3.3. Diffusion Models

Diffusion models represent a newer class of generative frameworks capable of producing high-fidelity and realistic data by simulating a progressive denoising process [13], [23]. Inspired by principles from thermodynamics, these models generate detailed and lifelike imagery, offering substantial advances in diagnostic imaging, data augmentation, and training dataset expansion [9]. Nevertheless, diffusion models face computational and operational challenges, including high computational cost and slow sampling rates, which currently limit their routine deployment in clinical settings [19], [21]. Recent research efforts seek to mitigate these limitations through accelerated sampling techniques, such as Denoising Diffusion Implicit Models (DDIMs), as well as conditional diffusion frameworks [11], [25]. In addition, physics-informed diffusion models align synthetic outputs more closely with biological realism, thereby enhancing clinical validity and biomedical plausibility [3]. These developments position diffusion models as a transformative frontier for achieving high-resolution, interpretable, and trustworthy generative modeling in medical imaging and related healthcare applications [9].

3.4. Large Language Models (LLMs)

Large Language Models (LLMs) trained on medical and clinical corpora can generate human-like textual content for a range of healthcare applications, including automated discharge summaries, radiological report interpretation, clinical question answering, and patient interaction [9], [26]. Despite their strong performance, concerns related to accuracy and interpretability persist [3]. LLMs perform particularly well in decision-support and clinical documentation tasks; however, they are susceptible to hallucination, whereby plausible-sounding but incorrect statements are generated, posing risks in healthcare environments [9], [26]. The black-box nature of many LLM-based decision-support systems further complicates the verification of clinical recommendations and may increase the risk of inappropriate or misleading outputs [24]. Moreover, biases embedded in training datasets can be amplified by these models, potentially reinforcing inequities, instigating ethical conflicts, and adversely affecting patient outcomes [27], [28]. Table 2 summarizes model types, their key usage and insight drawn from studies on these models.

Table 2. Model-centric developments of GenAI models

Model Type	Key Contributions	Representative Studies	Key Insights
GANs	Data synthesis, denoising, medical image reconstruction	[19], [29]	Most mature and widely used for imaging; faces stability and bias challenges.
VAEs	Feature learning, molecular structure prediction	[30]–[32]	Effective for probabilistic reasoning; latent variables capture hidden patterns.
Diffusion Models	High-resolution generation, denoising diffusion probabilistic modeling	[33], [34]	Promising for drug and molecular synthesis; still limited in clinical deployment.
LLMs	Clinical documentation, question answering, summarization	[9], [35]	Rapidly expanding into radiology and EHR summarization; risks hallucination.

4. Premise of GenAI in Healthcare

Generative AI in healthcare has evolved from early probabilistic and unsupervised paradigms into a diversified toolkit encompassing GANs, VAEs, diffusion models, and large language models capable of generating clinically relevant images, text, and structured data [1], [4]. Historically, a key inflection point was the introduction of GANs, which demonstrated that generator–discriminator dynamics could yield synthetic samples approaching the distributional fidelity of real clinical data. Subsequent advances in VAEs and, more recently, diffusion models further expanded the design space for simulation, denoising, and patient-trajectory modeling [13], [25].

This progression is not merely technical; it directly responds to persistent healthcare constraints such as data scarcity, stringent privacy requirements, and the need for realistic in-silico environments for training and validation, as shown in Table 2. Collectively, these drivers explain the rapid diffusion of GenAI across imaging, EHR synthesis, and clinical documentation tasks [13].

Across clinical domains, the appeal of GenAI is largely pragmatic [11], [36]. In medical imaging, GANs and VAEs are widely used to synthesize MRI, CT, and PET data to support algorithm development where annotated labels are costly or scarce, while diffusion models and hybrid architectures are increasingly explored to overcome fidelity limitations in fine-grained anatomical detail [37]. In parallel, EHR-oriented research leverages generative approaches to construct de-identified cohorts and support downstream analytics without compromising patient confidentiality [16], [37]. These use cases position GenAI as an enabler of diagnostic development, privacy-preserving data sharing, and rare-disease research, even as open questions remain regarding artifact risks and the enforcement of domain-appropriate constraints for clinical reliability.

Figure 5 presents an evidence map of GenAI applications in healthcare, illustrating how different model classes align with foundational drivers, application domains, and evaluation and ethical considerations. As shown in Figure 5, the field still lacks harmonized validation norms. Imaging-focused studies frequently report metrics such as FID or PSNR, whereas text generation tasks rely on BLEU, ROUGE, or AUROC; however, consensus on what constitutes clinically meaningful validation remains limited. This fragmentation, exacerbated by siloed datasets and the scarcity of multimodal benchmarks, complicates reproducibility and cross-study comparability. It also underscores the need for cross-institutional evaluation protocols and standardized, ethically shareable testbeds (e.g., MIMIC-III, CheXpert) to move from technical promise toward scientific and clinical credibility.

Figure 6 depicts a taxonomic tree of GenAI models used in healthcare applications, organized across three levels: L1 (model categories), L2 (sub-categories), and L3 (representative examples or applications). Beyond technical classification, the ethics-and-governance substrate remains decisive for real-world adoption. Bias and fairness concerns arise when generative models are trained on demographically skewed corpora, while privacy-preserving strategies—such as differential privacy, k-anonymity, and federated learning—mitigate risks but also introduce trade-offs in fidelity, accountability, and interpretability. Regulatory guidance tailored specifically to generative systems remains emergent across jurisdictions, and, together with explainability gaps and workflow-integration challenges, continues to shape clinician trust and adoption.

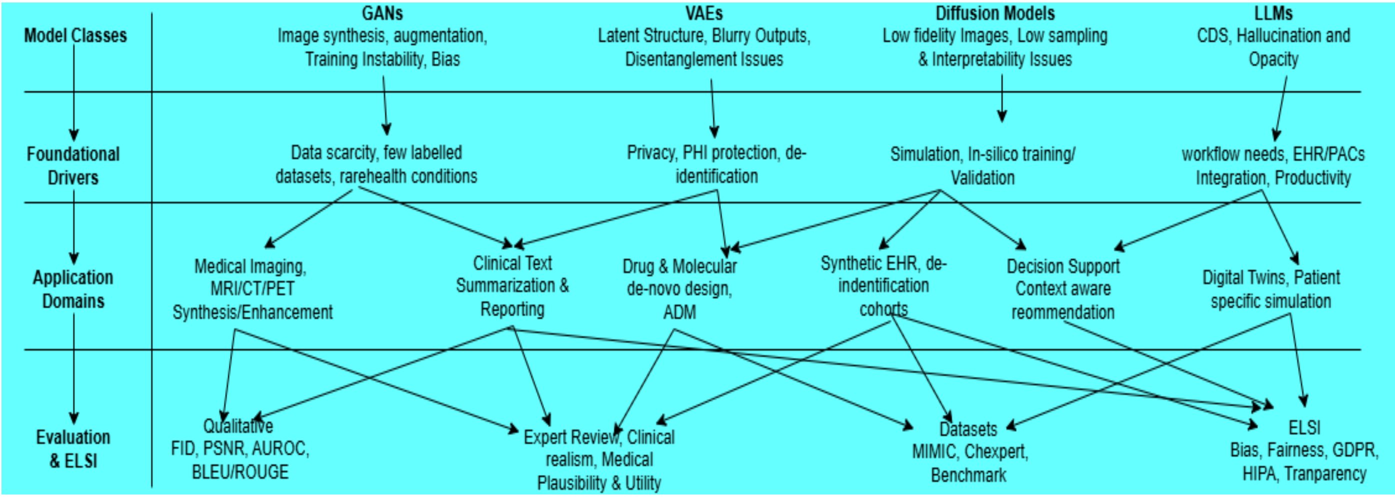

Figure 5. Evidence map of GenAI application in healthcare

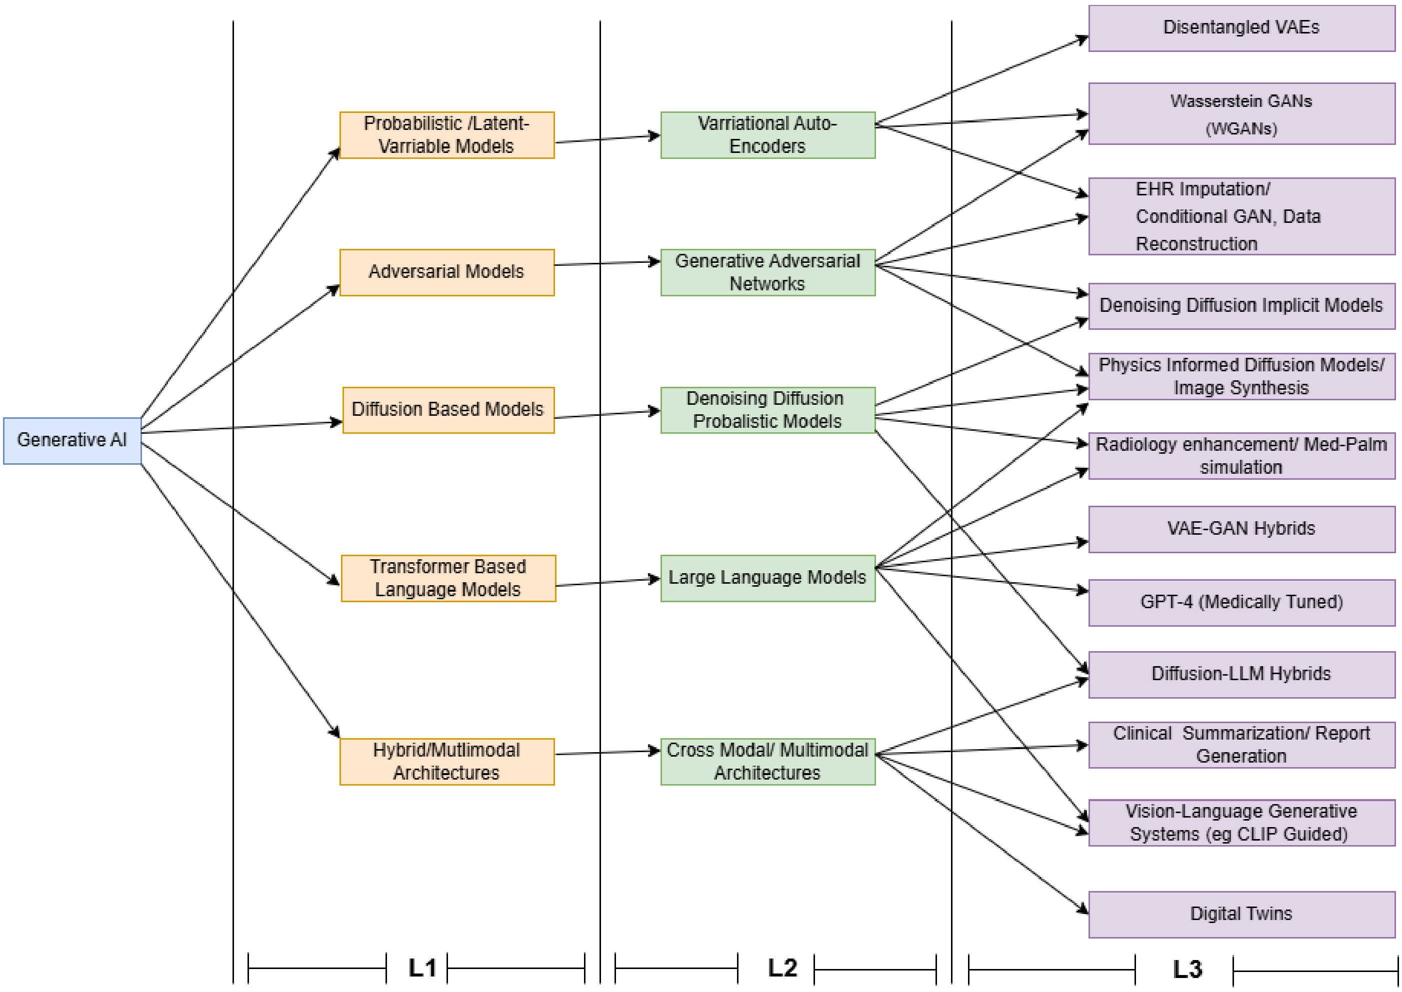

Figure 6. Taxonomic Tree of GenAI Models used in healthcare application. L1: Categories, L2: Sub-categories, L3: Examples/Applications

A balanced reading of the evidence suggests that GenAI's most defensible near-term value lies in augmentation rather than full automation. Reported gains in labeling efficiency, predictive accuracy, and human–AI synergy are promising, but must be interpreted through the lenses of equity and patient safety, as efficiency gains divorced from fairness risk entrenching existing disparities. As illustrated in Figure 6, the research frontier is converging toward multimodal generative architectures, privacy-preserving diffusion frameworks, and neuro-symbolic hybrids that integrate statistical generation with causal reasoning. These advances must operate within governance frameworks capable of auditing performance, tracing accountability, and sustaining clinician confidence. Related convergence trends are further discussed in Figure 12.

5. Domains of Healthcare Applications

5.1. Introduction to Application Categories

Generative AI applications in healthcare can be broadly organized into two overarching categories: (1) data-centric and (2) patient-centric applications. This distinction reflects the dual workflow of modern digital health systems, where upstream data generation, transformation, and enhancement processes support downstream clinical and patient-facing functionalities.

Data-centric applications primarily focus on improving, augmenting, or synthesizing structured and unstructured medical data. These include synthetic data generation, multimodal data fusion, missing data imputation, anomaly detection, and automated document processing. In contrast, patient-centric applications emphasize direct clinical utility, supporting diagnosis, prognosis, therapeutic planning, clinical decision-making, rehabilitation guidance, and patient engagement.

This conceptual division aligns with contemporary literature, which increasingly distinguishes between data-level innovation and clinically actionable systems. Structuring the review along these two dimensions enhances conceptual clarity, illustrates how foundational data processes underpin clinical intelligence, and provides a coherent basis for assessing the maturity of each application domain.

5.2. Data-Centric Applications

5.2.1. Medical Imaging

Generative AI models, particularly GANs and VAEs, are widely employed to synthesize realistic medical images in settings where annotated datasets are scarce, costly, or ethically constrained [38], [39]. These models support diagnostic algorithm development, rare disease research, and privacy-preserving data sharing [13], [25]. However, they often struggle to reproduce fine anatomical structures, potentially introducing artifacts that compromise diagnostic reliability [3], [27], [40].

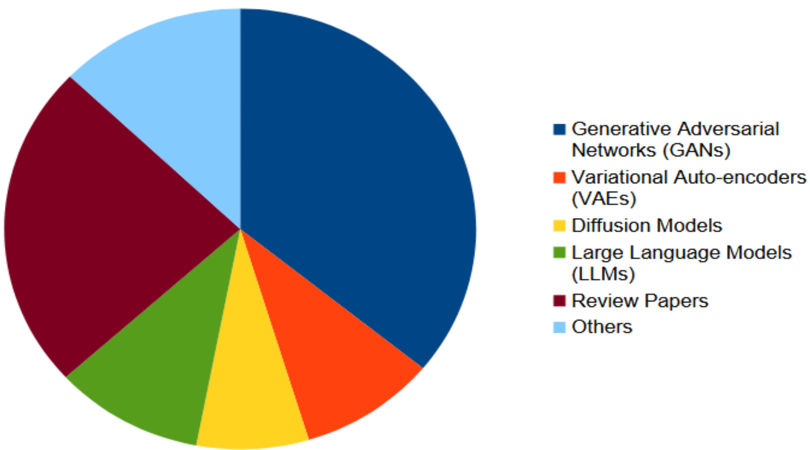

Figure 7. Categories of GenAI models used based on the papers under review. It isclear that GAN models are the most widely used due to their robust applicability followed by VAEs, then LLMs as the two later are comparably newer techniques

To address these limitations, recent studies increasingly integrate diffusion models and hybrid domain-constrained architectures to improve image fidelity and anatomical consistency [41]. Ethical considerations persist regarding the downstream use of synthetic images for clinical validation and decision-making [28]. Ongoing efforts aim to establish standardized benchmarks for assessing image quality, authenticity, and clinical reliability [39]. Medical imaging and enhancement remain the most extensively explored GenAI application domains, as reflected in the distribution of model usage shown in Figure 7.

5.2.2. Image Enhancement

Generative AI models, including GANs and VAEs, are also applied to enhance image resolution, contrast, and noise reduction, particularly in low-quality or low-dose imaging scenarios [39], [42]. These approaches improve diagnostic performance while preserving patient privacy [13]. Nonetheless, challenges remain in maintaining structural fidelity, which may affect diagnostic interpretation [2]. To improve robustness, recent studies explore diffusion-based and hybrid domain-aware architectures. Ethical concerns related to potential misuse of enhanced images in clinical decision-making and regulatory evaluation remain unresolved [3], [27].

5.2.3. Data Augmentation

Synthetic image generation is a critical technique for data augmentation in medical imaging, reducing overfitting and improving model generalizability across diverse patient populations [42]–[44]. The integration of real and synthetic data strengthens training pipelines, particularly under class imbalance or limited data conditions [44], [45]. However, synthetic augmentation may inadvertently reinforce latent biases present in the original datasets, resulting in overconfident but clinically flawed models [33]. Current research investigates fairness-aware generation, latent space interpolation, and style transfer techniques to enhance synthetic diversity and clinical plausibility [33], [42].

In parallel, Large Language Models (LLMs) are increasingly applied to automated structured reporting in radiology and pathology, translating visual findings into coherent narrative summaries [15], [33]. Despite these advances, challenges related to factual accuracy, contextual awareness, and regulatory compliance persist [2], [16], [46]. Emerging approaches, such as retrieval-augmented generation and multimodal reasoning frameworks, emphasize factual grounding and clinician oversight to improve reliability [15].

5.2.4. Data Imputation

Generative models are increasingly used to impute missing values in clinical databases by modeling complex inter-variable relationships [33], [42], [45]. Architectures such as VAEs and diffusion models generate plausible values for incomplete electronic health records, laboratory results, and longitudinal patient data [11], [47]. This approach improves data quality for downstream predictive modeling while avoiding biases associated with traditional imputation techniques (e.g., mean substitution or regression-based methods) [43].

Nevertheless, generative imputation can propagate hidden biases if missingness patterns are inaccurately modeled [27]. Recent studies explore attention-based mechanisms and adversarial validation to detect unreliable imputations [33], [45]. Hybrid models combining VAEs with symbolic reasoning show promise for improving interpretability [32].

5.2.5. Anomaly Detection

Generative models are applied to detect anomalous patterns in healthcare data, identifying potential fraud, reporting errors, or atypical clinical presentations [44], [48]. By learning normative data distributions, these models can identify outliers with high sensitivity [49]. Applications include billing anomaly detection, adverse drug reaction monitoring, and identification of subtle diagnostic inconsistencies, thereby enhancing clinical safety and organizational efficiency [37].

However, false positives remain a significant challenge, particularly in noisy clinical environments [3], [41]. Current research explores few-shot learning and multimodal fusion (e.g., combining imaging with laboratory trends) to improve specificity [3]. These trends are consistent with the application growth patterns illustrated in Table 3.

5.2.6. Predictive Models

Generative AI is increasingly employed in predictive modeling to forecast disease progression, treatment response, hospital readmission risk, and other clinical outcomes [50]–[52].

These models capture temporal and nonlinear patterns from longitudinal datasets, including EHRs, imaging studies, and genomic profiles [51], [53]. By anticipating changes in patient health trajectories, predictive generative models support proactive care planning, optimized resource allocation, and personalized interventions [35].

Nonetheless, temporal generative models often experience performance degradation over long prediction horizons due to error accumulation [37], [47], [54]. Recent work integrates uncertainty quantification and knowledge graphs to improve robustness [11]. Prospective clinical trials remain necessary to validate these models in real-world healthcare settings [42], [55].

5.2.7. Decision Support

Generative AI enhances clinical decision-support systems by producing context-aware recommendations tailored to patient profiles [56]. These systems integrate multimodal inputs including laboratory data, imaging, medications, and clinical history to generate diagnostic hypotheses, treatment suggestions, and risk assessments [24], [53]. Through natural language generation, LLMs also provide human-readable explanations that may reduce clinician cognitive burden and support informed decision-making [32], [57].

However, opaque reasoning processes may encourage over-reliance on automated recommendations [58]. Current research emphasizes interactive interfaces that allow clinicians to interrogate model outputs and hybrid architectures that link LLMs with curated medical knowledge bases [59]. Regulatory frameworks for adaptive AI-based decision-support systems remain under active development [15], [16], [25].

5.3. Patient-Centric Applications

Patient-centric applications of GenAI focus on direct clinical interaction, decision-making, and individualized patient support. Unlike data-centric applications, which primarily enhance underlying data infrastructure, patient-centric systems operate closer to the point of care and therefore face higher requirements for reliability, interpretability, and governance. As illustrated in Figure 8, patient-oriented use cases such as clinical summarization, decision support, and automated reporting represent a substantial proportion of current GenAI deployments, while more complex applications such as digital twins remain comparatively nascent.

5.3.1. Summarization of Clinical Notes

Generative AI techniques are increasingly used to summarize electronic health records (EHRs), producing concise and clinically relevant summaries for healthcare professionals [11]. Extractive and hybrid models identify critical patient information, thereby reducing cognitive load and improving clinical efficiency [19], [37]. These systems also support care coordination by transforming lengthy or fragmented clinical documentation into structured, easily interpretable summaries [37].

However, clinical summarization remains vulnerable to contextual and semantic errors, particularly in complex or inconsistently documented cases [43]. Recent approaches integrate structured EHR data with free-text narratives using graph-based or knowledge-augmented architectures to improve factual accuracy and coherence [46]. The growing demand for explainable summarization has further driven research into attention visualization, editable summaries, and clinician-in-the-loop validation interfaces [60], [61].

5.3.2. Chatbots and Patient Communication

Chatbots powered by generative language models are increasingly employed to enhance patient engagement, particularly in remote or resource-limited healthcare settings [62]. Typical applications include symptom triage, appointment scheduling, post-discharge follow-up, and delivery of point-of-care health information [63], [64].

Despite their potential, chatbots require careful oversight to ensure medical accuracy, ethical triage behavior, and cultural sensitivity [65]. Limited contextual reasoning can lead to inaccurate or unsafe recommendations [64]. To mitigate these risks, recent work combines reinforcement learning with human feedback (RLHF), rule-based verification layers, and hybrid architectures that constrain generative responses using validated medical knowledge bases [61]. Ongoing research emphasizes governance mechanisms to ensure that chatbot-mediated interactions remain clinically safe and ethically compliant [25].

5.3.3. Drug Design

Generative models, including Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and reinforcement learning-based frameworks, are widely used in de novo drug design to generate novel drug-like molecules [11], [66]. These models explore large chemical spaces and optimize compounds for key pharmacological and physicochemical properties [67], [68]. By accelerating early-stage screening, GenAI reduces both the time and cost of traditional drug discovery pipelines and enables exploration of therapeutic candidates for rare or complex diseases [9]. However, many generated compounds exhibit limited synthetic feasibility or lack experimentally validated safety profiles [68]. To address these limitations, contemporary approaches integrate generative models with molecular dynamics simulations and in silico toxicity assessment to improve pharmacological realism and downstream translational viability [22].

5.3.4. Simulations (Molecular Binding and ADMET)

Generative models are increasingly applied to simulate molecular interactions, including ligand–receptor binding and absorption, distribution, metabolism, excretion, and toxicity (ADMET) properties [69], [70]. These simulations enable virtual prioritization of compounds prior to synthesis and laboratory testing, thereby accelerating preclinical development stages [71], [72]. By jointly modeling pharmacokinetics and pharmacodynamics, such approaches improve lead compound selection and reduce attrition rates in early drug development [73]. Recent research incorporates quantum-mechanical priors, federated learning strategies for proprietary data aggregation, and uncertainty-aware validation loops to enhance reliability [9], [11].

5.3.5. Protein Folding and Structure Prediction

Deep learning models, most notably AlphaFold, have achieved major breakthroughs in predicting protein three-dimensional structures from amino acid sequences [11]. These advances have significant implications for understanding protein–ligand interactions, guiding drug targeting, and identifying therapeutically relevant binding sites [74]

While current models excel at predicting static structures, they remain limited in modeling dynamic conformations, membrane proteins, and post-translational modifications. Ongoing research explores equivariant neural networks, protein–protein interaction modeling, and integration of cryo-EM data to address these challenges [67].

5.3.6. Synthetic Patient Data

Generative models such as VAEs and GANs are widely used to generate synthetic EHR datasets that replicate patient demographics, clinical histories, laboratory values, and treatment trajectories [14]. These datasets preserve key statistical properties of real patient populations while mitigating privacy risks, thereby enabling model development, benchmarking, and data sharing without exposing protected health information [70], [75] .

Synthetic EHRs are particularly valuable in settings where large annotated datasets are unavailable [3], [73]. However, generative approaches may inadvertently propagate or amplify biases present in the original data. As summarized in Table 3, current mitigation strategies include differential privacy mechanisms, causal generative modeling, and fairness-aware training protocols [15].

5.3.7. Digital Twins

Digital twins in healthcare are virtual representations of individual patients, constructed using generative models that integrate clinical, physiological, and behavioral data [35], [76]. These models simulate patient-specific disease trajectories and treatment responses, supporting precision medicine and personalized therapy optimization [31], [35], [65], [77].

Despite their conceptual appeal, patient-specific digital twins face significant challenges related to computational cost, sparse longitudinal data, and limited clinical validation. As reflected in the maturity assessment in Table 3, digital twins remain at an early stage of development, with ongoing research focusing on hybrid mechanistic–AI models, federated personalization, and real-world outcome validation [11], [68].

Across Sections 5.3.1–5.3.7, the maturity, benefits, and risks of patient-centric GenAI applications vary considerably. As summarized in Table 3, domains such as medical imaging, image enhancement, and clinical text processing demonstrate higher maturity, supported by structured data availability, established benchmarks, and partial regulatory validation. In contrast, complex patient-level simulations, particularly digital twins, remain at an early developmental stage due to high computational requirements, sparse longitudinal data, and limited real-world validation. This synthesis underscores that while patient-facing GenAI systems are increasingly explored, their readiness for routine clinical deployment depends strongly on domain-specific validation rigor, governance constraints, and explainability mechanisms.

Table 3. Domains of GenAI applications in healthcare.

Domain	Typical Use-Cases	Primary Benefits	Key Risks & Challenges	Maturity (as of 2025)
Medical Imaging	MRI/CT/PET synthesis; anomaly detection; multi-modal image fusion	Augments limited datasets; enhances privacy; enables rare-condition model training	Synthetic artifacts may mislead diagnostic models; validation gap between synthetic and clinical realism	High – GANs and diffusion models widely validated in imaging workflows
Image Enhancement	Super-resolution, denoising, and MRI→CT translation	Improves image clarity, reduces need for rescans, supports multimodal diagnostics	Risk of false anatomical structures; alignment and interpretability issues	High–Medium – Rapid progress, active regulatory and clinical validation
Clinical Text / Reporting	Automatic summarization and structured reporting for radiology and pathology	Reduces workload; standardizes documentation; supports multilingual reporting	Factual inconsistencies and hallucinations in generated reports; compliance with regulatory documentation standards	Medium – Advancing with hybrid retrieval-augmented and knowledge-grounded LLMs
Synthetic Data (EHR)	De-identified patient cohort generation; imputation of missing values	Enables data sharing and training without PHI exposure; mitigates class imbalance	Fidelity–privacy trade-offs; risk of bias reinforcement from original data	Medium – Growing adoption; dependent on privacy assurance mechanisms
Decision Support	Context-aware clinical recommendations and narrative explanations	Reduces clinician cognitive load; supports personalized and data-driven decisions	Over-reliance on opaque reasoning; need for human oversight and transparency	Medium – Hybrid human–AI systems emerging with interpretability focus
Predictive Models	Disease progression, treatment response, and readmission forecasting	Supports proactive and preventive care; improves resource planning	Propagation of uncertainty across long temporal horizons; lack of prospective validation	Medium – Active research phase; limited deployment in regulated settings
Digital Twins	Patient-specific physiological or disease simulation for therapy optimization	Enables individualized treatment planning and risk assessment	High computational cost; sparse longitudinal data; insufficient clinical validation	Low–Medium – Nascent domain with few large-scale demonstrations

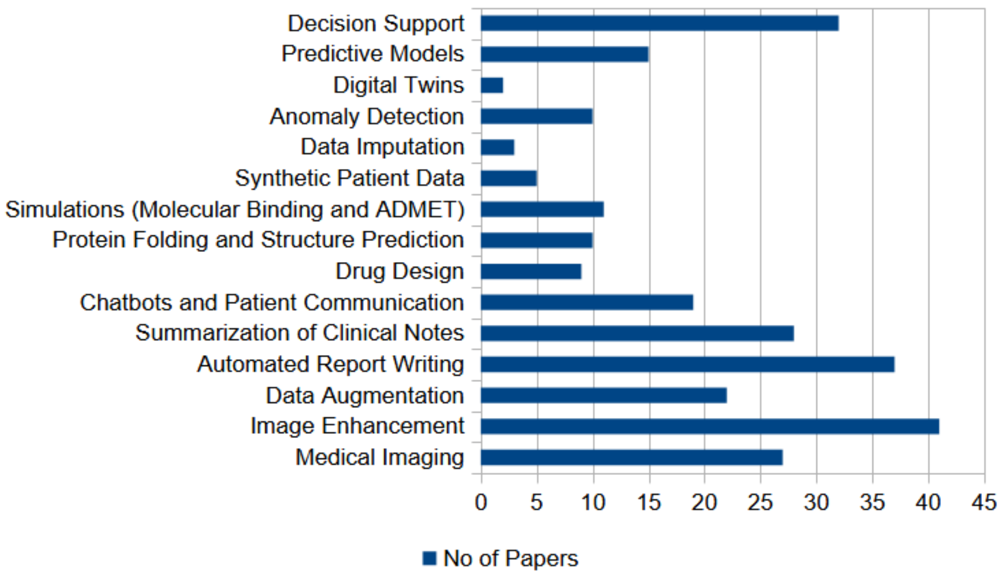

Figure 8. Areas of GenAI applications in healthcare in the papers under review. it can be seen that image enhancement and decision support are the most widely applied areas of healthcare for GenAI, while digital twins and data imputation are the least applied due to their newness.

Overall, patient-centric applications such as clinical summarization, decision support, and automated reporting represent the most actively deployed GenAI use cases, consistent with the publication distribution shown in Figure 8. In contrast, digital twins and advanced simulation-based systems remain less prevalent due to higher data, computational, and validation requirements. The relative strengths and limitations of the underlying generative models supporting these applications are further summarized in Table 4, highlighting trade-offs between realism, interpretability, computational cost, and clinical reliability.

Table 4. Comparative Capabilities and Limitations of GenAI models in healthcare.

Model Type	Core Mechanism	Primary Healthcare Applications	Strengths	Limitations	Representative Metrics	Illustrative Studies (1994–2025)
GANs	Adversarial training between generator and discriminator	Medical image synthesis, data augmentation, rare disease simulation	High realism; effective for small dataset augmentation	Mode collapse; instability; sensitive to hyperparameters	FID, SSIM, DSC	[13], [23]
VAEs	Latent-space probabilistic encoding and reconstruction	EHR data synthesis, noise reduction, missing data imputation	Stable training; interpretable latent features	Blurred outputs; limited fine detail	RMSE, AUC, KL divergence	[9], [14]
Diffusion Models	Gradual denoising of latent representations	High-fidelity medical imaging, synthetic pathology generation	Superior realism; noise-resistant; stable optimization	Computationally intensive; long training time	PSNR, FID	[45], [78]
LLMs	Transformer-based autoregressive generation	Clinical documentation, decision support, report summarization	Contextual fluency; multimodal potential	Hallucination; factual inconsistency; high data cost	BLEU, ROUGE-L, factual accuracy rate	[15], [46]
RLHF	Iterative alignment with expert feedback	Chatbots, triage assistants, adaptive diagnosis	Human-aligned responses; continual learning	Requires large-scale feedback loops	Reward accuracy, response safety	[3], [66]

6. Evaluation and Validation Methods of GenAI Models in Healthcare Applications

6.1. Quantitative Evaluation Methods

Quantitative evaluation of generative AI (GenAI) models in healthcare involves the systematic application of statistical and computational metrics to assess fidelity, diversity, and clinical relevance of generated outputs. Commonly used measures include the Fréchet Inception Distance (FID) and Peak Signal-to-Noise Ratio (PSNR), which quantify statistical similarity between real and synthetic medical images, enabling objective performance comparison across domains such as radiology and pathology [79].

For text-based generative models, including large language models (LLMs) and clinical summarization systems, evaluation metrics such as BLEU, ROUGE, and AUROC are employed to assess linguistic coherence, factual consistency, and diagnostic alignment against expert-validated reference corpora. Increasingly, multimodal evaluation metrics are being explored to capture integrated visual–textual outputs generated by hybrid GenAI architectures.

Despite their methodological rigor, quantitative evaluations remain constrained by the absence of standardized benchmarks, inconsistent metric interpretation across modalities, and weak correlation between numerical scores and downstream clinical utility. Consequently, the literature increasingly emphasizes the need for unified validation frameworks that integrate computational metrics with domain-specific clinical endpoints, ensuring that GenAI performance measures translate meaningfully into real-world healthcare outcomes.

As illustrated in Figure 9, GANs and diffusion models dominate imaging-based quantitative evaluation using metrics such as FID and PSNR, whereas LLMs primarily rely on BLEU and ROUGE scores for linguistic validation. This figure highlights the modality-dependent nature of quantitative assessment and the fragmentation of evaluation practices across GenAI model families.

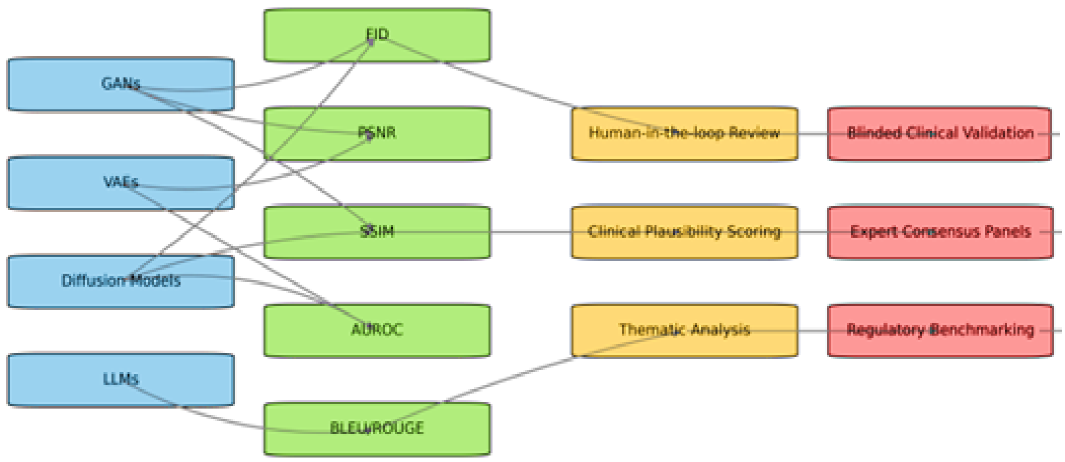

Figure 9. Evaluation and validation pathways for GenAI models in healthcare.

6.2. Qualitative Evaluation Methods

Qualitative evaluation of GenAI models emphasizes interpretability, usability, and contextual validity—dimensions that are often underrepresented in purely numerical assessments. These evaluations examine how effectively generated outputs reflect clinical reasoning, narrative coherence, and domain fidelity. Methods such as thematic content analysis, expert annotation comparison, and human-in-the-loop validation are commonly employed to assess transparency and reasoning quality in sensitive healthcare applications [13], [18], [26].

In medical imaging, qualitative evaluation typically involves clinician review of synthetic scans to assess anatomical plausibility, diagnostic relevance, and error sensitivity. For text-based applications, evaluators examine readability, factual grounding, ethical tone, and contextual appropriateness within generated clinical narratives. Thematic synthesis further supports the identification of latent biases, contextual misinterpretations, and cross-modal inconsistencies in multimodal outputs.

As emphasized across the reviewed literature, qualitative evaluation serves as an epistemic bridge between statistical performance and human trust. By contextualizing numerical metrics within clinical judgment, qualitative assessment provides a more holistic understanding of the reliability, safety, and ethical implications of generative AI systems in healthcare.

6.3. Expert Review and Clinical Validation

Expert review remains the definitive standard for validating GenAI systems in healthcare contexts. This process involves structured evaluation by domain specialists, such as clinicians, radiologists, pathologists, and biomedical informaticians, who assess clinical soundness, safety, and practical applicability of generated outputs [47]. Expert evaluations frequently adopt blinded or comparative study designs, wherein synthetic and real outputs are assessed against criteria including diagnostic realism, interpretive value, and decision-support usability.

Beyond surface-level assessment, expert validation examines clinical concordance by determining whether GenAI-generated recommendations, reconstructions, or summaries align with established medical guidelines and treatment protocols [26]. Importantly, expert feedback also drives iterative model refinement, supporting adaptive improvement and alignment with regulatory expectations from authorities such as the FDA and EMA. The literature consistently demonstrates that embedding expert-in-the-loop validation enhances transparency, accountability, and trustworthiness, key prerequisites for clinical adoption. In this sense, expert review transforms GenAI evaluation from a static performance check into a continuous epistemic dialogue between algorithmic systems and human clinical expertise.

7. Ethical, Legal, and Social Implications (ELSI)

The literature consistently highlights the interdependent relationship between ethical, legal, and social (ELS) dimensions in the deployment of generative AI within healthcare systems. Ethical principles—such as fairness, bias mitigation, transparency, and accountability—interact closely with legal frameworks governing liability, responsibility, and fitness-for-purpose. At the same time, social factors, including public trust, perceived utility, and equitable access, strongly influence how generative AI systems are accepted, adopted, or contested by stakeholders and end users [27], [28], [80].

These three dimensions collectively inform governance mechanisms, which integrate ethical norms, legal requirements, and social expectations into enforceable oversight structures, including regulatory bodies, ethics committees, and institutional review processes [81], [82]. Governance, in turn, feeds back into the ethical, legal, and social domains, enabling continuous alignment between technological innovation, public values, and regulatory safeguards. Figure 10. illustrates this interdependence, showing how ethical principles, legal considerations, and social factors mutually reinforce one another and converge into a centralized governance layer that shapes responsible GenAI deployment in healthcare.

Figure 10 depicts the bidirectional relationships among ethical principles, legal considerations, and social factors, and their convergence into governance structures such as regulatory bodies, ethics boards, and institutional oversight mechanisms. These feedback loops ensure that generative AI systems remain aligned with clinical safety, public trust, and regulatory compliance.

While this section addresses overarching ELSI considerations, these issues become more concrete when linked to the technical distinctions discussed earlier. Generative imaging models, including GANs and diffusion models, raise risks related to synthetic image authenticity, diagnostic misrepresentation, and provenance loss, stemming from their pixel-level generative capacity. In contrast, LLM-based systems introduce semantic and reasoning risks, such as hallucinations, fabricated citations, biased narratives, and unsafe clinical recommendations [27], [28]. These distinctions underscore the need for model-specific ethical guidance: imaging models require provenance tracking, watermarking, and validation safeguards, whereas LLM-based systems demand truthfulness controls, retrieval augmentation, hallucination monitoring, and rigorous human oversight.

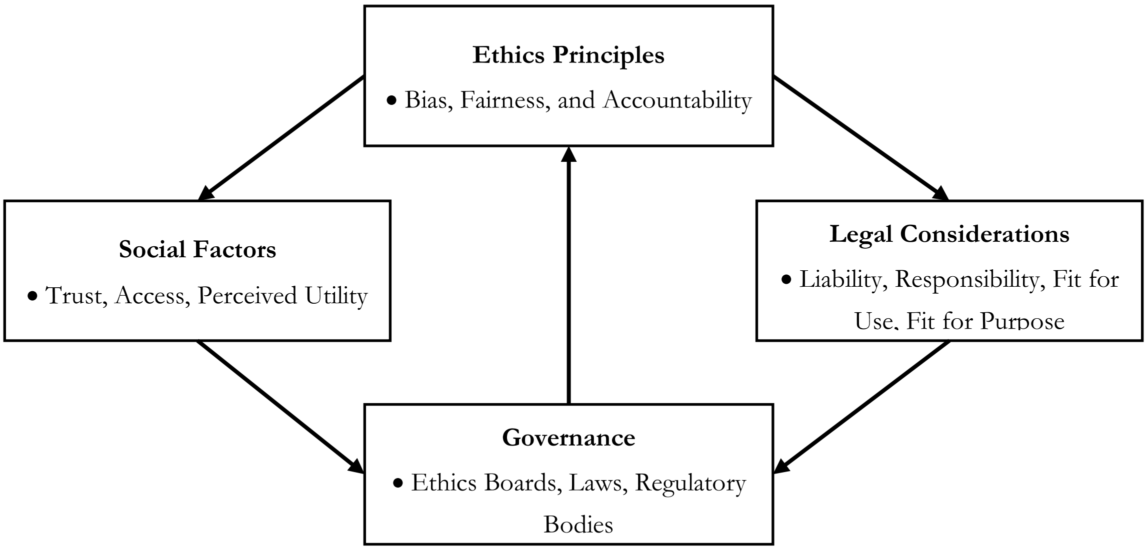

Figure 10. Interdependence of Ethical, Legal, and Social Dimensions in Generative AI Governance

7.1. Privacy

The generation of synthetic healthcare data must ensure that fabricated datasets cannot be reverse-engineered or linked to identifiable individuals, as such leakage would constitute a serious breach of patient confidentiality [70], [75]. State-of-the-art privacy-preserving approaches—including differential privacy, k-anonymity, and privacy-aware GAN variants—are increasingly employed to mitigate re-identification risks [70].

Even low-probability privacy breaches can undermine ethical obligations and regulatory compliance, reinforcing the need for robust anonymization pipelines and adversarial testing [15], [16]. Accordingly, validation procedures increasingly incorporate statistical similarity testing, membership inference attacks, and re-identification risk assessments to ensure that synthetic records cannot be exploited to disclose protected health information (PHI) [27], [43].

7.2. Bias and Fairness

Training generative models on imbalanced or non-representative clinical datasets risks amplifying systemic biases and producing inequitable outcomes for underrepresented populations [10], [27], [41], [54]. For example, datasets dominated by specific demographic groups may yield synthetic outputs that inadequately represent minority populations, thereby reinforcing disparities in downstream clinical decision-making [62].

Mitigation strategies include fairness-aware learning methods such as reweighting, adversarial debiasing, and synthetic minority oversampling, combined with continuous auditing using fairness metrics (e.g., demographic parity and equalized odds) [73], [83]. Ensuring fairness across populations is thus a central ethical requirement for trustworthy GenAI deployment in healthcare.

7.3. Transparency and Accountability

The use of generative AI in healthcare necessitates interpretable and auditable models that allow stakeholders to understand and scrutinize how outputs are generated [83]–[85]. Black-box architectures, particularly deep neural networks, often lack transparency, which can impede regulatory approval and erode clinician and patient trust [26], [62]. Interpretability techniques such as SHAP, LIME, and attention visualization have been increasingly applied to enhance transparency and support accountability [15], [16]. In parallel, comprehensive documentation of data sources, model architectures, training procedures, and validation protocols strengthens traceability, facilitates regulatory compliance, and enables corrective action when biases or errors are identified [16], [37].

7.4. Compliance

Synthetic healthcare data must comply with stringent legal and regulatory frameworks, including the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and applicable Food and Drug Administration (FDA) guidelines [15], [86]. The GDPR mandates data minimization and purpose limitation, requiring synthetic datasets to be non-reidentifiable while remaining functionally informative [43], [86]. HIPAA imposes strict access controls and audit requirements for data usage [43].

For FDA-regulated use cases—such as AI-assisted diagnosis—robust validation is required to demonstrate equivalence between synthetic data and real-world clinical evidence [15], [16]. Failure to meet these standards can result in legal consequences and undermine the credibility of AI-driven healthcare solutions [24]. Table 5 summarizes the core ethical, regulatory, and governance concerns associated with GenAI deployment in healthcare, highlighting key focus areas, representative studies, and their practical implications for clinical adoption.

Table 5. Ethical, regulatory, and governance considerations of GenAI in healthcare

Concern	Focus	Notable Studies	Implications
Bias and Fairness	Population imbalance	[27]	GenAI requires equitable dataset design
Privacy	Synthetic data for anonymity	[33]	Differential privacy and federated learning emerging
Explainability	Clinician trust	[81], [83]	Drives demand for interpretable AI
Regulation	FDA, GDPR alignment	[15], [21]	Lack of harmonized global framework

8. Perils of GenAI Applications in Healthcare (Challenges and Limitations)

8.1. Data Scarcity

The development of effective generative models in the healthcare domain is significantly constrained by the scarcity of high-quality, well-annotated medical datasets [11]. Unlike broader application domains, medical data acquisition is limited by strict privacy regulations, institutional ethics review requirements, and the need for specialized domain expertise for accurate labeling [45], [87]. Rare diseases and longitudinal patient records are particularly underrepresented, resulting in fragmented and incomplete training corpora [38]. Although approaches such as semi-supervised learning, federated learning, and transfer learning from related domains offer partial mitigation, they often introduce trade-offs in data fidelity and representativeness [88]. In the absence of sufficient labeled data, generative models may produce clinically implausible or statistically inconsistent outputs, thereby undermining downstream analytical and clinical applications [43], [88].

8.2. Generalization

A persistent challenge in medical generative AI is limited generalization beyond the training distribution [84], [89]. Overfitting, where synthetic outputs replicate training data rather than capturing true biological variability, compromises external validity [69]. For instance, models trained on patient cohorts from large urban hospitals may not generalize to smaller institutions or demographically diverse populations [38], [89]. Techniques such as domain adaptation, adversarial validation, and controlled noise injection into synthetic data generation have been proposed to improve robustness [62]. Nevertheless, rigorous cross-institutional validation remains essential to demonstrate that generative outputs preserve epidemiological and pathophysiological validity across heterogeneous clinical settings [3], [43].

8.3. Explainability

The black-box nature of deep generative architectures, including GANs and diffusion models, presents a significant barrier to clinical adoption, as synthetic outputs cannot be easily traced back to interpretable sources [81], [83], [85]. Unlike discriminative models, where feature importance can often be quantified, generative models lack intuitive interpretability paradigms [1]. This opacity raises concerns that synthetic data may propagate latent biases or artifacts into diagnostic or decision-support systems [13], [59]. Emerging solutions include hybrid architectures—such as VAEs augmented with attention mechanisms—as well as post-hoc interpretability techniques, including latent space traversal and attribution-based analysis [58], [90]. Regulatory and governance bodies increasingly emphasize interpretability requirements to prevent the misuse of synthetic data in biomedical research and clinical care [91] [92].

8.4. Workflow Integration

Integrating generative models into healthcare information systems introduces substantial technical and organizational challenges, including interoperability with electronic health records (EHRs), clinician trust, and limited access to high-performance computing infrastructure [79]. Synthetic data pipelines must align with clinical workflows, requiring temporally consistent EHR trajectories and Picture Archiving and Communication System (PACS)-compatible imaging formats [93]. Real-time deployment is further complicated by system latency, care-continuity disruptions, and model drift over time [24], [32]. Effective integration therefore demands close collaboration among clinicians, data engineers, and regulatory stakeholders to ensure operational feasibility, user acceptance, and compliance with healthcare IT standards [15], [16].

8.5. GenAI for Healthcare Analytics and Decision-Making

Generative AI is increasingly embedded in operational and policy-level healthcare analytics, extending its role beyond individual clinical prediction to system-wide decision support. At the operational level, GenAI enables the construction of synthetic EHR cohorts that retain statistical fidelity while protecting patient privacy, supporting fine-grained analyses for quality monitoring, risk stratification, and service optimization [41], [71]. These datasets also underpin demand-forecasting models that anticipate patient flow, operating-theatre utilization, readmission risk, and bed occupancy, informing staffing and resource allocation decisions in high-pressure healthcare environments [26], [94].

In clinical operations, LLM-based documentation tools and automated report generation reduce clinician cognitive burden and improve throughput in radiology, pathology, and emergency care settings [13], [15]. In imaging services, generative enhancement techniques—such as super-resolution and denoising—reduce repeat scans and accelerate interpretation pipelines, indirectly supporting downstream operational planning [23], [95].

At the policy and public health level, GenAI facilitates equity-aware simulations using synthetic or augmented population datasets to evaluate interventions across sociodemographic subgroups and underserved regions [87], [96]. When combined with explainable AI techniques, these models provide transparent and auditable evidence to support regulatory review, guideline formulation, and clinical governance. Moreover, GenAI supports multi-institutional epidemiological collaboration by enabling privacy-preserving data sharing through federated learning and synthetic-data intermediaries, thereby improving reproducibility in surveillance and system-level performance assessment [45]. Collectively, these developments shift GenAI from a purely predictive paradigm toward a broader framework of actionable, trustworthy, and governance-aligned intelligence.

Robust safeguards remain essential for safe deployment. Human-in-the-loop oversight, explainability mechanisms (e.g., SHAP, saliency maps, and attention-based interpretation), fairness testing across patient subgroups, and seamless integration with EHR and PACS systems are critical to ensuring that generative analytics translate into accountable and clinically sound decision-making [13].

9. Promise of GenAI Application in Healthcare (Future Directions)

9.1. Multimodal Models

The integration of heterogeneous data modalities—such as clinical notes, medical images, and physiological signals—offers substantial potential for improving diagnostic accuracy and holistic patient characterization [3], [23]. By leveraging cross-modal correspondences, multimodal generative models can produce context-aware synthetic data that more closely reflect the complexity of real-world clinical practice [84]. For example, jointly modeling radiological images and their associated diagnostic reports enables the extraction of clinically meaningful associations, thereby reducing ambiguity in downstream interpretation tasks [3], [23].

Despite these advantages, challenges remain, including modality alignment, heterogeneous data formats, and increased computational requirements [97]. Emerging paradigms such as cross-modal transformers and contrastive learning frameworks are being explored to unify diverse data sources while preserving semantic integrity, enabling synthetic datasets to support multimodal diagnostic and decision-support systems more effectively [43], [98].

9.2. Federated Learning

Federated learning (FL) provides a decentralized paradigm for training generative models, allowing healthcare institutions to collaboratively improve algorithms without sharing raw patient data [65], [70]. By distributing training across sites and aggregating only model updates, FL mitigates privacy risks and aligns with data protection regulations such as GDPR and HIPAA [28]. This approach is particularly valuable for rare diseases and underrepresented populations, where individual institutions often lack sufficient data volume [98].

However, FL introduces technical challenges, including communication overhead, statistical heterogeneity due to non-IID data distributions, and vulnerability to adversarial attacks. These issues necessitate robust solutions such as secure aggregation protocols (e.g., homomorphic encryption) and adaptive optimization strategies [99]. Overall, FL's capacity to democratize access to high-quality synthetic data while preserving confidentiality positions it as a critical enabler of scalable and collaborative GenAI development in healthcare [5], [100].

9.3. Explainable AI

The adoption of synthetic data and generative models in clinical environments depends critically on transparency and interpretability, as opaque systems can undermine trust among clinicians and regulators [92], [94]. Explainable AI (XAI) techniques—including attention mechanisms, saliency mapping, and counterfactual explanations—play a central role in elucidating how synthetic outputs are generated and in validating their clinical plausibility [3]. For instance, in radiology, generative models should provide interpretable evidence that synthetic lesions exhibit biologically plausible characteristics rather than artifacts introduced during training [5].

Regulatory agencies increasingly emphasize the need for explainability in the certification of AI-driven medical tools, prompting calls for standardized reporting of data provenance, generation methods, and potential biases [57], [97]. Transparent generative models not only facilitate regulatory approval but also enable continuous refinement through human-in-the-loop feedback, strengthening alignment with clinical practice [59], [94].

9.4. Standardization

A major barrier to the widespread adoption of synthetic data in healthcare is the absence of standardized guidelines for its generation and validation [43]. Regulatory bodies and professional organizations must establish frameworks that assess synthetic data quality across multiple dimensions, including realism (fidelity), representativeness (diversity), and non-reidentifiability (privacy preservation) [69], [75]. Institutions such as the FDA, EMA, and IEEE are actively developing evaluation protocols that combine quantitative similarity metrics with clinician-in-the-loop validation to ensure compliance with clinical and ethical standards [3], [27], [43].

Furthermore, interoperability standards, such as HL7 FHIR for electronic health record integration, are essential to enable seamless adoption of synthetic data across healthcare systems [43]. Achieving consensus on these standards will be pivotal for fostering innovation while safeguarding patient interests and maintaining data integrity across institutions [9], [83].

10. Expanded Literature Synthesis

The transformative role of GenAI in healthcare has evolved from a predominantly conceptual paradigm into a multidimensional applied science spanning imaging, clinical text, molecular modeling, and healthcare systems analytics [23], [98], [101]. Over the past decade, the field has transitioned from classical statistical modeling toward generative architectures capable of learning complex, high-dimensional representations from heterogeneous and multimodal clinical data.

Early GenAI applications (1994–2020) were largely dominated by GANs, primarily addressing data scarcity in medical imaging through realistic synthetic data generation and data augmentation strategies that improved diagnostic model robustness. These early advances established a methodological foundation upon which subsequent diversification occurred. VAEs and transformer-based models extended generative modeling beyond imaging into structured clinical records, molecular representations, and narrative medical text [97], [102].

More recent developments, particularly diffusion models and LLMs, have introduced a new generation of contextual, stable, and explainable generative systems capable of synthesizing data while supporting reasoning, abstraction, and cross-domain generalization. This shift marks a conceptual transition from data generation alone toward interpretable and adaptive knowledge systems designed to support precision medicine and decision-making workflows [45], [67], [71].

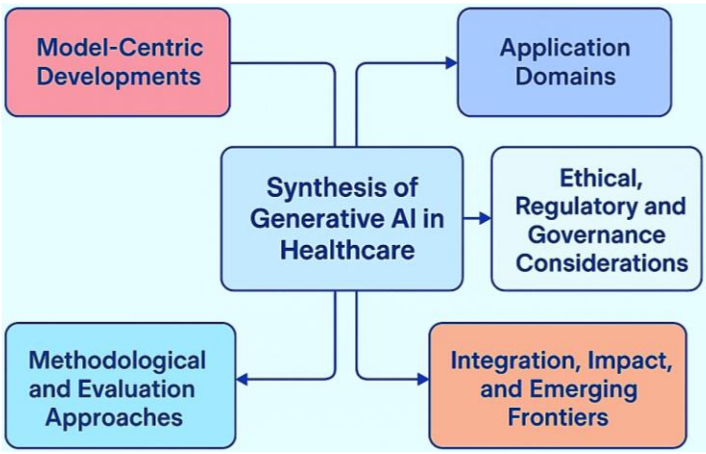

Figure 11. Interdependence of GenAI in healthcare

As illustrated in Figure 11, the literature reveals a strong interdependence between model maturity, application scope, evaluation rigor, and governance readiness. Model-centric advances directly influence feasible application domains, while ethical, regulatory, and governance constraints increasingly shape how generative systems are designed, validated, and deployed in clinical environments.

From a technical perspective, distinct trade-offs persist across major GenAI model families [92], [103]. GANs remain the most mature and widely validated approach for medical image synthesis, anomaly detection, and data augmentation, offering high perceptual realism but continuing to face challenges related to training instability and mode collapse [31], [64], [69]. VAEs excel in structured latent representations and probabilistic reasoning, making them suitable for molecular design and patient phenotype modeling, although they often produce outputs with reduced visual or structural fidelity [32]. Diffusion models demonstrate superior stability and high-resolution synthesis for biomedical imaging and molecular generation, albeit at substantial computational cost. Meanwhile, LLMs, such as BioGPT, Med-PaLM, and GPT-4, extend generative capability into clinical documentation, summarization, and knowledge reasoning, while introducing concerns related to hallucination, verification, and traceability [33], [57], [95].

Across these architectures, a recurring methodological gap is evident: the lack of unified benchmarks, standardized evaluation metrics, and reproducible validation protocols continues to limit cross-study comparability and external validity, particularly across institutions and modalities. This limitation, previously highlighted in Figure 5, remains a central barrier to large-scale clinical translation [104]

Methodologically, recent studies increasingly adopt hybrid learning paradigms that combine generative modeling with reinforcement learning, explainability tools, and domain constraints, as summarized in Table 6. Techniques such as SHAP, Grad-CAM, and LIME are frequently reported as post-hoc interpretability aids; however, relatively few studies integrate these methods into real-world clinical pipelines in a systematic or longitudinal manner. Evaluation practices remain fragmented: imaging studies predominantly report FID or PSNR, while text-based applications rely on BLEU, ROUGE, or AUROC, with limited consensus on what constitutes clinically meaningful validation. A consolidated overview of evaluation strategies and their alignment with model classes is presented in Figure 9.

Table 6. Emerging frontiers of GenAI applications in healthcare.

Frontier	Description	Potential Impact
Multimodal Generative Models	Unified image–text–signal learning	Foundational for precision medicine
Federated Diffusion Models	Secure decentralized model training	Privacy-compliant collaboration
Neurosymbolic Integration	Hybrid reasoning and generative modeling	Transparent decision-making
Governance Frameworks	AI auditability and regulation	Enhances accountability and trust

Beyond technical performance, an expanding body of literature emphasizes the ethical, regulatory, and socio-technical dimensions of GenAI adoption [15], [28]. Bias and fairness remain persistent concerns, as generative models trained on demographically skewed datasets risk amplifying existing inequities in diagnosis and treatment [10], [27]. Although synthetic data generation and federated learning offer promising privacy-preserving alternatives, they introduce new challenges related to fidelity, accountability, and auditability [69], [75]. Regulatory guidance from institutions such as the FDA and EMA, alongside frameworks such as GDPR and HIPAA, remains fragmented and insufficiently tailored to generative systems, leaving much of current deployment guided by voluntary or provisional standards [105]. These uncertainties compound challenges related to clinician trust, explainability, and workflow integration [103].

Despite these constraints, empirical evidence consistently reports tangible benefits of GenAI adoption, including accelerated data labeling, improved predictive accuracy, and enhanced human–AI collaboration [96], [101], [106], [107]. These gains are most pronounced in radiology, oncology, and pharmacogenomics, where GenAI systems function as augmentative tools rather than autonomous decision-makers [88], [101]. However, the literature repeatedly cautions that efficiency gains must be interpreted through ethical and equity lenses, as performance improvements without governance safeguards risk reinforcing systemic disparities.

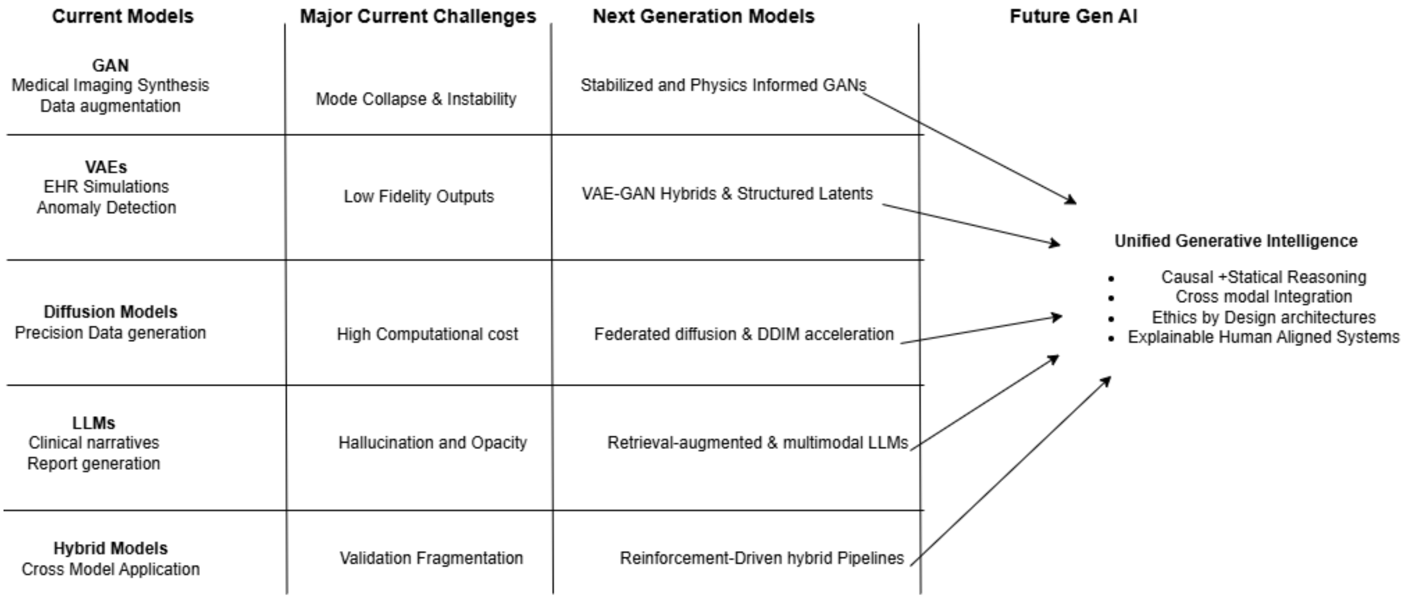

Figure 12. Models, challenges, and future directions in GenAI for healthcare

Looking forward, the literature anticipates convergence among multimodal generative architectures, federated diffusion models, and neurosymbolic AI frameworks that combine statistical generation with causal reasoning [31], [32]. These paradigms aim to address current trade-offs by enabling scalable, interpretable, and privacy-preserving collaboration across heterogeneous datasets [108]. As summarized in Tables 7 and 8, the success of GenAI in healthcare will ultimately depend not only on algorithmic sophistication but also on institutional readiness, regulatory clarity, and human-centered governance frameworks [21], [103]. In this sense, GenAI is redefining healthcare intelligence by shifting the discipline from reactive analytics toward proactive, explainable, and ethically aligned computational reasoning [51], [104].

Table 7. Impact and performance of GenAI in healthcare.

Parameter	Empirical Findings	Implication
Efficiency	20–40% faster data labeling and synthesis	Workflow optimization
Accuracy	~15% performance gains over baselines	Context-sensitive improvement
Human–AI Synergy	Improved productivity and satisfaction	Augmentation, not replacement
Accessibility	Broader access to healthcare AI tools	Risk of misuse and inequity

Table 8. Integration and implementation pathways.

Focus	Integration Pathway	Obstacles	Observed Trends
Workflow Adoption	EHR-integrated GenAI systems	Interoperability, clinician trust	Gradual uptake
Infrastructure	Cloud and federated models	Cost, latency	Shift toward hybrid computing
Evaluation in Practice	Prospective clinical trials	Limited validation	Pilot deployments

10.1. Policy Implications

The effective integration of generative AI into healthcare requires governance frameworks that explicitly prioritize interpretability, safety, and equity. First, explainability must be institutionalized: clinical- and policy-facing GenAI tools should incorporate feature-attribution maps, uncertainty profiles, and transparent data provenance to enable auditability, clinical oversight, and conformity with practice guidelines [15], [93]. Second, policymakers should promote privacy-by-design architectures, combining federated learning with clinically validated synthetic datasets to enable secure cross-institutional analytics while remaining compliant with data protection regulations [1], [4]. Such approaches reduce direct data exposure while preserving analytic utility for system-level planning and research.

Third, health systems must implement continuous equity and performance monitoring, including fairness audits prior to deployment and drift detection mechanisms after implementation. These processes should be supported by subgroup-level performance dashboards that guide corrective governance actions when disparities emerge [3], [16]. Finally, procurement and adoption pathways should align with transparency and interoperability standards. This includes the use of interoperable evaluation metrics, model cards, and full lifecycle documentation provided by vendors or developers, as summarized in Table 8 [4]. Together, these policy measures ensure that GenAI adoption advances clinical value without compromising accountability or public trust.

10.2. Research Implications

Future research should prioritize the development of unified multimodal evaluation benchmarks that integrate conventional technical metrics (e.g., FID, PSNR, BLEU/ROUGE) with clinically meaningful endpoints such as diagnostic concordance, time-to-decision, or downstream patient outcomes [86], [100]. Bridging this gap is essential to align algorithmic performance with real-world clinical impact. Advancing hybrid neurosymbolic and multimodal architectures represents another critical research direction, as these approaches can improve causal interpretability and robustness across heterogeneous institutional settings [101]. Such models may better support generalization while enabling transparent reasoning over complex clinical contexts.

A further priority is prospective, multi-site external validation, ideally incorporating equity-aware sampling strategies to ensure that reported performance generalizes beyond controlled research environments [109]. Without such validation, claims of robustness and safety remain incomplete. Finally, translation-oriented research should increasingly integrate implementation science, examining cost-effectiveness, workflow disruption, alert fatigue, and human–AI teaming dynamics. These factors are central to establishing the real-world value of generative models in both operational and policy contexts [101], [109].

11. Conclusions

Generative AI (GenAI) is increasingly transforming the healthcare landscape through applications in medical image synthesis, predictive diagnosis, personalized treatment planning, and molecular discovery [14]. Leveraging deep learning architectures such as diffusion models and GANs, these technologies can generate clinically valuable representations, expand limited datasets, and accelerate biomedical research workflows [3]. In radiology, generative models enhance image resolution and replicate rare pathologies for training and education, while in molecular science they enable in silico compound design and reduce reliance on costly wet-lab experimentation [11].

Despite these advances, clinical deployment must adhere to stringent validation requirements for safety, effectiveness, and ethical fitness for purpose [37], [50]. Algorithmic bias, data privacy risks, and limited interpretability necessitate robust governance mechanisms and sustained transdisciplinary collaboration, as highlighted in Table 5 [10], [27]. Close coordination among clinicians, data scientists, and regulators is essential to ensure that GenAI systems conform to standardized evaluation and approval pathways [41], [95].

Looking ahead, GenAI is poised for deeper integration into clinical workflows and research pipelines [110]. Promising directions include patient-specific synthetic datasets for rare disease modeling and improved generalization of AI-driven diagnostic systems across diverse populations [1]. In medical education, GenAI is expected to play an increasing role in the creation of realistic virtual patients and simulation environments for healthcare training [81], [107].

At the same time, strong momentum is building toward making generative models transparent, interpretable, and ethically governed, reflecting the sensitive nature of medical data and clinical decision-making [58], [110]. Regulatory paradigms will be decisive in shaping the safe and responsible use of synthetic data and AI-generated clinical recommendations. As the field matures, healthcare stands to benefit substantially from GenAI's capacity to accelerate discovery, augment decision-making, and ultimately improve patient outcomes at scale [25].

Author Contributions: Conceptualization: N.F. and M.F.A.; methodology: N.F. and K.B.J.; validation: N.F., M.F.A., and F.S.I.; investigation: N.F. and K.B.J.; resources: K.B.J.; writing—original draft preparation: N.F. and K.B.J.; writing—review and editing: M.F.A. and F.S.I.; visualization: K.B.J.; supervision: N.F. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: The authors acknowledge the use of the AI-assisted tool DeepSeek during the writing process to support content structuring and drafting. All content generated with the assistance of this tool was carefully reviewed, verified, and edited by the authors, who take full responsibility for the accuracy, integrity, and originality of the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

References

[1] V. Capraro *et al.*, "The impact of generative artificial intelligence on socioeconomic inequalities and policy making," *PNAS Nexus*, vol. 3, no. 6, May 2024, doi: 10.1093/pnasnexus/pgae191.

[2] S. J. Aitken *et al.*, "Codesign of health technology interventions to support best-practice perioperative care and surgical waitlist management," *BMJ Heal. Care Informatics*, vol. 31, no. 1, p. e100928, Mar. 2024, doi: 10.1136/bmjhci-2023-100928.

[3] A. Wilson, H. Saeed, C. Pringle, I. Eleftheriou, P. A. Bromiley, and A. Brass, "Artificial intelligence projects in healthcare: 10 practical tips for success in a clinical environment," *BMJ Heal. Care Informatics*, vol. 28, no. 1, p. e100323, Jul. 2021, doi: 10.1136/bmjhci-2021-100323.

[4] S. Singh, S. Singh, S. Kraus, A. Sharma, and S. Dhir, "Characterizing generative artificial intelligence applications: Text-mining-enabled technology roadmapping," *J. Innov. Knowl.*, vol. 9, no. 3, p. 100531, Jul. 2024, doi: 10.1016/j.jik.2024.100531.

[5] D. L. Roberts and M. Candi, "Artificial intelligence and innovation management: Charting the evolving landscape," *Technovation*, vol. 136, p. 103081, Aug. 2024, doi: 10.1016/j.technovation.2024.103081.

[6] S. L. Chan, E. Lum, M. E. H. Ong, and N. Graves, "Implementation science: A critical but undervalued part of the healthcare innovation ecosystem," *Heal. Care Sci.*, vol. 1, no. 3, pp. 160–165, Dec. 2022, doi: 10.1002/hcs2.22.

[7] Y. Zhai and X. Du, "Differential effectiveness of COVID-19 health behaviors: The role of mental health conditions in mask-wearing, social distancing, and hygiene practice," *Heal. Care Sci.*, vol. 2, no. 4, pp. 286–290, Aug. 2023, doi: 10.1002/hcs2.60.

[8] X. Wang and D. Zhang, "Choices of medical institutions and associated factors in older patients with multimorbidity in stabilization period in China: A study based on logistic regression and decision tree model," *Heal. Care Sci.*, vol. 2, no. 6, pp. 359–369, Dec. 2023, doi: 10.1002/hcs2.73.

[9] D. Shin *et al.*, "Artificial intelligence versus clinical judgement: how accurately do generative models reflect CNS guidelines for chiari malformation?," *Clin. Neurol. Neurosurg.*, vol. 248, p. 108662, Jan. 2025, doi: 10.1016/j.clineuro.2024.108662.

[10] I. Straw and H. Wu, "Investigating for bias in healthcare algorithms: a sex-stratified analysis of supervised machine learning models in liver disease prediction," *BMJ Heal. Care Informatics*, vol. 29, no. 1, p. e100457, Apr. 2022, doi: 10.1136/bmjhci-2021-100457.

[11] A. Beheshti and W. Mansoor, "Exploring the convergence of Internet of things and big data technologies in the age of generative artificial intelligence," in *Empowering IoT with Big Data Analytics*, Elsevier, 2025, pp. 333–354. doi: 10.1016/B978-0-443-21640-4.00004-1.

[12] M. Heikkilä, O. Heino, and P. Rautiainen, "System's Crisis Resilience as a Societal Crisis: Knowledge Structure and Gaze of the Finnish Health Care System," *Heal. Care Anal.*, Jan. 2024, doi: 10.1007/s10728-023-00479-3.

[13] P. Sarvari, Z. Al-fagih, A. Ghuwel, and O. Al-fagih, "A systematic evaluation of the performance of GPT-4 and PaLM2 to diagnose comorbidities in MIMIC-IV patients," *Heal. Care Sci.*, vol. 3, no. 1, pp. 3–18, Feb. 2024, doi: 10.1002/hcs2.79.

[14] H. Kamil, R. Rachmah, I. Irvanizam, and E. Wardani, "Exploring Health Professionals' Perceptions on Health-ID, an Electronic Integrated Patient Progress Documentation System: A Qualitative Study in Indonesia," *J. Multidiscip. Healthc.*, vol. Volume 13, pp. 1649–1656, Nov. 2020, doi: 10.2147/JMDH.S270740.

[15] L. Pantanowitz *et al.*, "Regulatory Aspects of Artificial Intelligence and Machine Learning," *Mod. Pathol.*, vol. 37, no. 12, p. 100609, Dec. 2024, doi: 10.1016/j.modpat.2024.100609.

[16] K. H. Zou *et al.*, "Harnessing Real-World Data for Regulatory Use and Applying Innovative Applications," *J. Multidiscip. Healthc.*, vol. Volume 13, pp. 671–679, Jul. 2020, doi: 10.2147/JMDH.S262776.

[17] D. Chen, Y. Ma, H. Xiao, and Z. Yan, "Development trends of etiological research contents and methods of noncommunicable diseases," *Heal. Care Sci.*, vol. 2, no. 5, pp. 352–357, Oct. 2023, doi: 10.1002/hcs2.69.

[18] T. Saheb, M. Sidaoui, and B. Schmarzo, "Convergence of artificial intelligence with social media: A bibliometric & qualitative analysis," *Telemat. Informatics Reports*, vol. 14, p. 100146, Jun. 2024, doi: 10.1016/j.teler.2024.100146.

[19] C. Modica, J. Lewis, and R. Bay, "Advancing Virtual at-Home Care for Community Health Center Patients Using Patient Self-Care Tools, Technology, and Education," *J. Multidiscip. Healthc.*, vol. Volume 17, pp. 521–531, Feb. 2024, doi: 10.2147/JMDH.S443973.

[20] K. D. Graham, A. Steel, and J. Wardle, "Embracing the Complexity of Primary Health Care: System-Based Tools and Strategies for Researching the Case Management Process," *J. Multidiscip. Healthc.*, vol. Volume 14, pp. 2817–2826, Oct. 2021, doi: 10.2147/JMDH.S327260.
[21] G. Fisher *et al.*, "Barriers and facilitators to learning health systems in primary care: a framework analysis," *BMJ Heal. Care Informatics*, vol. 31, no. 1, p. e100946, Jun. 2024, doi: 10.1136/bmjhci-2023-100946.
[22] M. Batool, O. Sanumi, and J. Jankovic, "Application of artificial intelligence in the materials science, with a special focus on fuel cells and electrolyzers," *Energy AI*, vol. 18, p. 100424, Dec. 2024, doi: 10.1016/j.egyai.2024.100424.
[23] Y. Wu *et al.*, "Evolution and major changes of the diagnosis and treatment protocol for COVID-19 patients in China 2020–2023," *Heal. Care Sci.*, vol. 2, no. 3, pp. 135–152, Jun. 2023, doi: 10.1002/hcs2.45.
[24] K. Seangpraw and P. Ong-Artborirak, "Knowledge and Behaviors Toward Health Care for Elderly Patients with Hypertension, and Quality of Life Among Informal Caregivers in Northern Thailand," *J. Multidiscip. Healthc.*, vol. Volume 13, pp. 1771–1780, Dec. 2020, doi: 10.2147/JMDH.S284886.
[25] I. D. Mienye, Y. Sun, and E. Ileberi, "Artificial intelligence and sustainable development in Africa: A comprehensive review," *Mach. Learn. with Appl.*, vol. 18, p. 100591, Dec. 2024, doi: 10.1016/j.mlwa.2024.100591.
[26] K. Sadeghi R., D. Ojha, P. Kaur, R. V. Mahto, and A. Dhir, "Explainable artificial intelligence and agile decision-making in supply chain cyber resilience," *Decis. Support Syst.*, vol. 180, p. 114194, May 2024, doi: 10.1016/j.dss.2024.114194.
[27] M. G. Hanna *et al.*, "Ethical and Bias Considerations in Artificial Intelligence/Machine Learning," *Mod. Pathol.*, vol. 38, no. 3, p. 100686, Mar. 2025, doi: 10.1016/j.modpat.2024.100686.
[28] T. P. Raju, "Artificial Intelligence: From Theory to Application and Ethical Considerations," *Int. J. Res. Appl. Sci. Eng. Technol.*, vol. 12, no. 5, pp. 639–641, May 2024, doi: 10.22214/ijraset.2024.59463.
[29] D. M. Hartley *et al.*, "An overview of Internet biosurveillance," *Clin. Microbiol. Infect.*, vol. 19, no. 11, pp. 1006–1013, Nov. 2013, doi: 10.1111/1469-0691.12273.
[30] E. Vázquez-Cano, M.-P. Quicios-García, J. Fombona, and J. Rodríguez-Arce, "Latent factors on the design and adoption of gamified apps in primary education," *Educ. Inf. Technol.*, vol. 28, no. 11, pp. 15093–15123, Nov. 2023, doi: 10.1007/s10639-023-11797-3.
[31] L. Zhang, W. Wang, T. Fang, and Y. Song, "ConKE: Conceptualization-Augmented Knowledge Editing in Large Language Models for Commonsense Reasoning," *arXiv*. May 28, 2025. [Online]. Available: http://arxiv.org/abs/2412.11418
[32] Q. He, J. Yu, and W. Wang, "Large Language Model-Enhanced Symbolic Reasoning for Knowledge Base Completion," *arXiv*. Jan. 02, 2025. [Online]. Available: http://arxiv.org/abs/2501.01246
[33] Z. Tan *et al.*, "Large Language Models for Data Annotation and Synthesis: A Survey," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2024, pp. 930–957. doi: 10.18653/v1/2024.emnlp-main.54.
[34] T. R. Noviandy, G. M. Idroes, and I. Hardi, "An Interpretable Machine Learning Strategy for Antimalarial Drug Discovery with LightGBM and SHAP," *J. Futur. Artif. Intell. Technol.*, vol. 1, no. 2, pp. 84–95, Aug. 2024, doi: 10.62411/faith.2024-16.
[35] H. Abdollahi *et al.*, "Theranostic digital twins: Concept, framework and roadmap towards personalized radiopharmaceutical therapies," *Theranostics*, vol. 14, no. 9, pp. 3404–3422, 2024, doi: 10.7150/thno.93973.
[36] A. Mohanty and S. Mishra, "A Comprehensive Study of Explainable Artificial Intelligence in Healthcare," in *Augmented Intelligence in Healthcare: A Pragmatic and Integrated Analysis*, 2022, pp. 475–502. doi: 10.1007/978-981-19-1076-0_25.
[37] D. D. AlThubaity and A. Y. Mahdy Shalby, "Perception of Health Teams on the Implementation of Strategies to Decrease Nursing Errors and Enhance Patient Safety," *J. Multidiscip. Healthc.*, vol. Volume 16, pp. 693–706, Mar. 2023, doi: 10.2147/JMDH.S401966.
[38] T. Dhar, N. Dey, S. Borra, and R. S. Sherratt, "Challenges of Deep Learning in Medical Image Analysis—Improving Explainability and Trust," *IEEE Trans. Technol. Soc.*, vol. 4, no. 1, pp. 68–75, Mar. 2023, doi: 10.1109/TTS.2023.3234203.
[39] L. Pinto-Coelho, "How Artificial Intelligence Is Shaping Medical Imaging Technology: A Survey of Innovations and Applications," *Bioengineering*, vol. 10, no. 12, p. 1435, Dec. 2023, doi: 10.3390/bioengineering10121435.
[40] L. Zhang *et al.*, "Generative AI enables medical image segmentation in ultra low-data regimes," *Nat. Commun.*, vol. 16, no. 1, p. 6486, Jul. 2025, doi: 10.1038/s41467-025-61754-6.
[41] F. Abdelfattah, H. Madi, M. Al-Washahi, A. AlAraimi, S. Nagi, and M. Abbas, "Harnessing Artificial Intelligence, Business Intelligence, and Digital Technologies for achieving supply chain excellence in Oman: Investigating the mediating role of predictive analytics," *J. Open Innov. Technol. Mark. Complex.*, vol. 11, no. 4, p. 100644, Dec. 2025, doi: 10.1016/j.joitmc.2025.100644.
[42] J. P. Bona, F. W. Prior, M. N. Zozus, and M. Brochhausen, "Enhancing Clinical Data and Clinical Research Data with Biomedical Ontologies - Insights from the Knowledge Representation Perspective," *Yearb. Med. Inform.*, vol. 28, no. 01, pp. 140–151, Aug. 2019, doi: 10.1055/s-0039-1677912.
[43] F. J. Santos Arteaga *et al.*, "On the capacity of artificial intelligence techniques and statistical methods to deal with low-quality data in medical supply chain environments," *Eng. Appl. Artif. Intell.*, vol. 133, p. 108610, Jul. 2024, doi: 10.1016/j.engappai.2024.108610.
[44] M. Atzmueller, T. Bohne, and P. Windler, "Knowledge-Augmented Explainable and Interpretable Learning for Anomaly Detection and Diagnosis," in *Handbook on Neurosymbolic AI and Knowledge Graphs*, 2025. doi: 10.3233/FAIA250234.
[45] M. Dion, P. AbdelMalik, and A. Mawudeku, "Big Data and the Global Public Health Intelligence Network (GPHIN)," *Canada Commun. Dis. Rep.*, vol. 41, no. 9, pp. 209–214, Sep. 2015, doi: 10.14745/ccdr.v41i09a02.
[46] R. Karuppal, "The impact of artificial intelligence on medical article writing: A boon or a bane?," *J. Orthop.*, vol. 63, pp. 98–100, May 2025, doi: 10.1016/j.jor.2024.10.045.
[47] M. K. Makwero *et al.*, "Characterising the performance measurement and management system in the primary health care systems of Malawi," *African J. Prim. Heal. Care Fam. Med.*, vol. 16, no. 1, Jan. 2024, doi: 10.4102/phcfm.v16i1.4007.
[48] F.-M. Liou, Y.-C. Tang, and J.-Y. Chen, "Detecting hospital fraud and claim abuse through diabetic outpatient services," *Health Care Manag. Sci.*, vol. 11, no. 4, pp. 353–358, Dec. 2008, doi: 10.1007/s10729-008-9054-y.
[49] E. Jin *et al.*, "LogicAD: Explainable Anomaly Detection via VLM-based Text Feature Extraction," *Proc. AAAI Conf. Artif. Intell.*, vol. 39, no. 4, pp. 4129–4137, Apr. 2025, doi: 10.1609/aaai.v39i4.32433.

[50] M. M. Alsaleh *et al.*, "Prediction of disease comorbidity using explainable artificial intelligence and machine learning techniques: A systematic review," *Int. J. Med. Inform.*, vol. 175, p. 105088, Jul. 2023, doi: 10.1016/j.ijmedinf.2023.105088.
[51] S. Ardabili *et al.*, "COVID-19 Outbreak Prediction with Machine Learning," *Algorithms*, vol. 13, no. 10, p. 249, Oct. 2020, doi: 10.3390/a13100249.
[52] A. P. Zhao *et al.*, "AI for science: Predicting infectious diseases," *J. Saf. Sci. Resil.*, vol. 5, no. 2, pp. 130–146, Jun. 2024, doi: 10.1016/j.jnlssr.2024.02.002.
[53] Y. Wang, J. Zhong, and R. Kumar, "A Systematic Review of Machine Learning Applications in Infectious Disease Prediction, Diagnosis, and Outbreak Forecasting," *Prerpints.org*. Apr. 15, 2025. doi: 10.20944/preprints202504.1250.v1.
[54] Z. L. Teo *et al.*, "Generative artificial intelligence in medicine," *Nat. Med.*, vol. 31, no. 10, pp. 3270–3282, Oct. 2025, doi: 10.1038/s41591-025-03983-2.
[55] E. C. Daniel, S. Tirunagari, K. Batth, D. Windridge, and Y. Balla, "Interpretable Machine Learning for Predicting Multiple Sclerosis Conversion from Clinically Isolated Syndrome," *medRxiv*. Jul. 19, 2024. doi: 10.1101/2024.07.18.24310578.
[56] A. Akkari and M. Radhouane, "Learning in emergencies contexts: From the building of the concept to multifaced actions in the field," *Prospects*, vol. 53, no. 1–2, pp. 15–28, May 2023, doi: 10.1007/s11125-023-09639-8.
[57] F. Nooralahzadeh, Y. Zhang, J. Furst, and K. Stockinger, "Multi-Modal Data Exploration via Language Agents," *arXiv*. Nov. 25, 2025. [Online]. Available: http://arxiv.org/abs/2412.18428
[58] A. Budhkar, Q. Song, J. Su, and X. Zhang, "Demystifying the black box: A survey on explainable artificial intelligence (XAI) in bioinformatics," *Comput. Struct. Biotechnol. J.*, vol. 27, pp. 346–359, 2025, doi: 10.1016/j.csbj.2024.12.027.
[59] J. O. Ferstad, E. B. Fox, D. Scheinker, and R. Johari, "Learning Explainable Treatment Policies with Clinician-Informed Representations: A Practical Approach," *arXiv*. Nov. 26, 2024. [Online]. Available: http://arxiv.org/abs/2411.17570
[60] W. Samek and K.-R. Müller, "Towards Explainable Artificial Intelligence," in *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, 2019, pp. 5–22. doi: 10.1007/978-3-030-28954-6_1.
[61] Y. Chen *et al.*, "An Explainable Recommendation Method for Artificial Intelligence of Things Based on Reinforcement Learning With Knowledge Graph Inference," *IEEE Trans. Consum. Electron.*, vol. 71, no. 1, pp. 1980–1991, Feb. 2025, doi: 10.1109/TCE.2024.3521435.
[62] M. Kinney, M. Anastasiadou, M. Naranjo-Zolotov, and V. Santos, "Expectation management in AI: A framework for understanding stakeholder trust and acceptance of artificial intelligence systems," *Heliyon*, vol. 10, no. 7, p. e28562, Apr. 2024, doi: 10.1016/j.heliyon.2024.e28562.
[63] C. Smith, "First, do no harm: institutional betrayal and trust in health care organizations," *J. Multidiscip. Healthc.*, vol. Volume 10, pp. 133–144, Apr. 2017, doi: 10.2147/JMDH.S125885.
[64] H. Dweik and M. Nawahdah, "Bridging reality and AI: The role of augmented reality chatbots in interactive learning," in *19th International Technology, Education and Development Conference*, Mar. 2025, pp. 4136–4141. doi: 10.21125/inted.2025.1048.
[65] S. T. Tshitenge and J. M. Nthitu, "COVID-19 frontline primary health care professionals' perspectives on health system preparedness and response to the pandemic in the Mahalapye Health District, Botswana," *African J. Prim. Heal. Care Fam. Med.*, vol. 14, no. 1, Apr. 2022, doi: 10.4102/phcfm.v14i1.3166.
[66] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artif. Intell.*, vol. 267, pp. 1–38, Feb. 2019, doi: 10.1016/j.artint.2018.07.007.
[67] S. Pengpid and K. Peltzer, "Health Care Responsiveness by Conventional, Traditional and Complementary Medicine Providers in a National Sample of Middle-Aged and Older Adults in India in 2017–2018," *J. Multidiscip. Healthc.*, vol. Volume 15, pp. 773–782, Apr. 2022, doi: 10.2147/JMDH.S357761.
[68] H. Li *et al.*, "A systematic review on the cost-effectiveness of the computer-assisted orthopedic system," *Heal. Care Sci.*, vol. 1, no. 3, pp. 173–185, Dec. 2022, doi: 10.1002/hcs2.23.
[69] L. Reddy *et al.*, "Classifying simulated gait impairments using privacy-preserving explainable artificial intelligence and mobile phone videos," *PLOS Digit. Heal.*, vol. 4, no. 9, p. e0001004, Sep. 2025, doi: 10.1371/journal.pdig.0001004.
[70] M. Kirienko *et al.*, "Distributed learning: a reliable privacy-preserving strategy to change multicenter collaborations using AI," *Eur. J. Nucl. Med. Mol. Imaging*, vol. 48, no. 12, pp. 3791–3804, Nov. 2021, doi: 10.1007/s00259-021-05339-7.
[71] V. Lakshminarayanan, A. Ravikumar, H. Sriraman, S. Alla, and V. K. Chattu, "Health Care Equity Through Intelligent Edge Computing and Augmented Reality/Virtual Reality: A Systematic Review," *J. Multidiscip. Healthc.*, vol. Volume 16, pp. 2839–2859, Sep. 2023, doi: 10.2147/JMDH.S419923.
[72] R. E. -, "Revolutionizing Knowledge Synthesis: Introducing AI-Enhanced Collaborative Theory (AIECT)," *Int. J. Multidiscip. Res.*, vol. 6, no. 5, Oct. 2024, doi: 10.36948/ijfmr.2024.v06i05.28203.
[73] J. Verma, L. Sandys, A. Matthews, and S. Goel, "Readiness of artificial intelligence technology for managing energy demands from renewable sources," *Eng. Appl. Artif. Intell.*, vol. 135, p. 108831, Sep. 2024, doi: 10.1016/j.engappai.2024.108831.
[74] R.-Y. Ong, F.-M. Lum, and L. F. Ng, "The Fine Line between Protection and Pathology in Neurotropic Flavivirus and Alphavirus Infections," *Future Virol.*, vol. 9, no. 3, pp. 313–330, Mar. 2014, doi: 10.2217/fvl.14.6.
[75] C.-N. Hang, Y.-Z. Tsai, P.-D. Yu, J. Chen, and C.-W. Tan, "Privacy-Enhancing Digital Contact Tracing with Machine Learning for Pandemic Response: A Comprehensive Review," *Big Data Cogn. Comput.*, vol. 7, no. 2, p. 108, Jun. 2023, doi: 10.3390/bdcc7020108.
[76] F. Grignaffini, P. Simeoni, A. Alisi, and F. Frezza, "Computer-Aided Diagnosis Systems for Automatic Malaria Parasite Detection and Classification: A Systematic Review," *Electronics*, vol. 13, no. 16, p. 3174, Aug. 2024, doi: 10.3390/electronics13163174.
[77] E. Lamptey, "Building sustainable capacity for better access to diabetes care in low-resource settings: A critical review of global efforts and integrated strategies," *Heal. Care Sci.*, vol. 3, no. 2, pp. 131–139, Apr. 2024, doi: 10.1002/hcs2.89.
[78] L. M. Thimoteo, M. M. Vellasco, J. Amaral, K. Figueiredo, C. L. Yokoyama, and E. Marques, "Explainable Artificial Intelligence for COVID-19 Diagnosis Through Blood Test Variables," *J. Control. Autom. Electr. Syst.*, vol. 33, no. 2, pp. 625–644, Apr. 2022, doi: 10.1007/s40313-021-00858-y.

[79] L. M. Mupara, J. J. O. Mogaka, W. R. Brieger, and J. M. Tsoka-Gwegweni, "Scorecard metrics for assessing the extent of integration of community health worker programmes into national health systems," *African J. Prim. Heal. Care Fam. Med.*, vol. 13, no. 1, Nov. 2021, doi: 10.4102/phcfm.v13i1.2691.
[80] R. E. -, "The Holistic Intelligent Healthcare Theory (HIHT): Integrating AI for Ethical, Transparent, and Human-Centered Healthcare Innovation," *Int. J. Multidiscip. Res.*, vol. 6, no. 5, Oct. 2024, doi: 10.36948/ijfmr.2024.v06i05.28846.
[81] A. Holzinger, C. Biemann, C. S. Pattichis, and D. B. Kell, "What do we need to build explainable AI systems for the medical domain?," *arXiv*. Dec. 28, 2017. [Online]. Available: http://arxiv.org/abs/1712.09923
[82] C. Pesquita, "Towards Semantic Integration for Explainable Artificial Intelligence in the Biomedical Domain," in *Proceedings of the 14th International Joint Conference on Biomedical Engineering Systems and Technologies*, 2021, pp. 747–753. doi: 10.5220/0010389707470753.
[83] U. Ehsan, Q. V. Liao, M. Muller, M. O. Riedl, and J. D. Weisz, "Expanding Explainability: Towards Social Transparency in AI systems," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, May 2021, pp. 1–19. doi: 10.1145/3411764.3445188.
[84] P. Gupta, B. Ding, C. Guan, and D. Ding, "Generative AI: A systematic review using topic modelling techniques," *Data Inf. Manag.*, vol. 8, no. 2, p. 100066, Jun. 2024, doi: 10.1016/j.dim.2024.100066.
[85] F. Giuste *et al.*, "Explainable Artificial Intelligence Methods in Combating Pandemics: A Systematic Review," *IEEE Rev. Biomed. Eng.*, vol. 16, pp. 5–21, 2023, doi: 10.1109/RBME.2022.3185953.
[86] L. Chen and X. Xu, "Effect Evaluation of the Long-Term Care Insurance (LTCI) System on the Health Care of the Elderly: A Review," *J. Multidiscip. Healthc.*, vol. Volume 13, pp. 863–875, Aug. 2020, doi: 10.2147/JMDH.S270454.
[87] Y. Bouba, E. K. Tsinda, M. D. M. Fonkou, G. S. Mmbando, N. L. Bragazzi, and J. D. Kong, "The Determinants of the Low COVID-19 Transmission and Mortality Rates in Africa: A Cross-Country Analysis," *Front. Public Heal.*, vol. 9, Oct. 2021, doi: 10.3389/fpubh.2021.751197.
[88] R. L. Berkelman, R. T. Bryan, M. T. Osterholm, J. W. LeDuc, and J. M. Hughes, "Infectious Disease Surveillance: a Crumbling Foundation," *Science (80-.).*, vol. 264, no. 5157, pp. 368–370, Apr. 1994, doi: 10.1126/science.8153621.
[89] A. Abd-Alrazaq *et al.*, "Artificial Intelligence in the Fight Against the COVID-19 Pandemic: Opportunities and Challenges," in *Multiple Perspectives on Artificial Intelligence in Healthcare*, 2021, pp. 185–196. doi: 10.1007/978-3-030-67303-1_15.
[90] A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018, doi: 10.1109/ACCESS.2018.2870052.
[91] G. Carloni, A. Berti, and S. Colantonio, "The Role of Causality in Explainable Artificial Intelligence," *WIREs Data Min. Knowl. Discov.*, vol. 15, no. 2, Jun. 2025, doi: 10.1002/widm.70015.
[92] M. Belghachi, "A Review on Explainable Artificial Intelligence Methods, Applications, and Challenges," *Indones. J. Electr. Eng. Informatics*, vol. 11, no. 4, Dec. 2023, doi: 10.52549/ijeei.v11i4.5151.
[93] S. Chen, J. Yu, S. Chamouni, Y. Wang, and Y. Li, "Integrating machine learning and artificial intelligence in life-course epidemiology: pathways to innovative public health solutions," *BMC Med.*, vol. 22, no. 1, p. 354, Sep. 2024, doi: 10.1186/s12916-024-03566-x.
[94] N. A. Aziz, A. Manzoor, M. D. Mazhar Qureshi, M. A. Qureshi, and W. Rashwan, "Unveiling Explainable AI in Healthcare: Current Trends, Challenges, and Future Directions," *medRxiv*. Aug. 10, 2024. doi: 10.1101/2024.08.10.24311735.
[95] D. Xu *et al.*, "Harnessing Large Language Models for Knowledge Graph Question Answering via Adaptive Multi-Aspect Retrieval-Augmentation," *Proc. AAAI Conf. Artif. Intell.*, vol. 39, no. 24, pp. 25570–25578, Apr. 2025, doi: 10.1609/aaai.v39i24.34747.
[96] A. Broadbent and T. Grote, "Can Robots Do Epidemiology? Machine Learning, Causal Inference, and Predicting the Outcomes of Public Health Interventions," *Philos. Technol.*, vol. 35, no. 1, p. 14, Mar. 2022, doi: 10.1007/s13347-022-00509-3.
[97] Y. Dang *et al.*, "Explainable and Interpretable Multimodal Large Language Models: A Comprehensive Survey," *arXiv*. Dec. 03, 2024. [Online]. Available: http://arxiv.org/abs/2412.02104
[98] M. Madanchian, H. Taherdoost, M. Vincenti, and N. Mohamed, "Transforming Leadership Practices through Artificial Intelligence," *Procedia Comput. Sci.*, vol. 235, pp. 2101–2111, 2024, doi: 10.1016/j.procs.2024.04.199.
[99] F. Armknecht *et al.*, "A Guide to Fully Homomorphic Encryption," *Cryptology ePrint Archive*. 2015. [Online]. Available: https://eprint.iacr.org/2015/1192
[100] X. Hu, "Constructing an effective evaluation system to identify doctors' research capabilities," *Heal. Care Sci.*, vol. 3, no. 1, pp. 67–72, Feb. 2024, doi: 10.1002/hcs2.82.
[101] K. B. Jillahi and A. Iorliam, "A Scoping Literature Review of Artificial Intelligence in Epidemiology: Uses, Applications, Challenges and Future Trends," *J. Comput. Theor. Appl.*, vol. 1, no. 4, pp. 421–445, Apr. 2024, doi: 10.62411/jcta.10350.
[102] O. S. Donaires, L. O. Cezarino, L. B. Liboni, E. M. S. Ribeiro, and F. P. Martins, "Multivariate data analysis of categorical data: taking advantage of the rhetorical power of numbers in qualitative research," *Qual. Quant.*, vol. 57, no. 6, pp. 5283–5312, Dec. 2023, doi: 10.1007/s11135-022-01589-1.
[103] G. Eysenbach, "Infodemiology and Infoveillance: Framework for an Emerging Set of Public Health Informatics Methods to Analyze Search, Communication and Publication Behavior on the Internet," *J. Med. Internet Res.*, vol. 11, no. 1, p. e11, Mar. 2009, doi: 10.2196/jmir.1157.
[104] L. R. Non, A. R. Marra, and D. Ince, "Rise of the Machines - Artificial Intelligence in Healthcare Epidemiology," *Curr. Infect. Dis. Rep.*, vol. 27, no. 1, p. 4, Dec. 2025, doi: 10.1007/s11908-024-00854-8.
[105] W. Samek, T. Wiegand, and K.-R. Müller, "Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models," *arXiv*. Aug. 28, 2017. [Online]. Available: http://arxiv.org/abs/1708.08296
[106] S. Criollo-C, A. Guerrero-Arias, Á. Jaramillo-Alcázar, and S. Luján-Mora, "Mobile Learning Technologies for Education: Benefits and Pending Issues," *Appl. Sci.*, vol. 11, no. 9, p. 4111, Apr. 2021, doi: 10.3390/app11094111.
[107] A. Mashekova, V. Zarikas, Y. Zhao, and E. Y. K. Ng, "The Importance of Explainable Artificial Intelligence Based Medical Diagnosis," *Clin. Exp. Obstet. Gynecol.*, vol. 51, no. 12, Dec. 2024, doi: 10.31083/j.ceog5112268.
[108] A. Gaudio *et al.*, "DeepFixCX: Explainable privacy-preserving image compression for medical image analysis," *WIREs Data Min. Knowl. Discov.*, vol. 13, no. 4, Jul. 2023, doi: 10.1002/widm.1495.

[109] A. Iorliam, A. T. S. Ho, S. Tirunagari, and D. Windridge, "Data Forensic Determination of the Accuracy of International COVID-19 Reporting: Using Zipf's Law for Pandemic Investigation." Apr. 30, 2020. doi: 10.20944/preprints202004.0531.v1.
[110] J. Zhang *et al.*, "Quantifying digital health inequality across a national healthcare system," *BMJ Heal. Care Informatics*, vol. 30, no. 1, p. e100809, Nov. 2023, doi: 10.1136/bmjhci-2023-100809.