

Generative AI in Healthcare: Transforming Medical Imaging, Accelerating Drug Discovery, and Enhancing Clinical Decision-Making

Khursheed Ahmad Ganie¹, Mohit Sharma¹, Priti Panwar¹, Arti²

¹ Allied Health Sciences Department, Chitkara School of Health Sciences, Chitkara University, Punjab India-140401

² Department of Life Sciences Rayat Bahra University, Punjab-140103

DOI : <https://doi.org/10.51583/IJLTEMAS.2025.1412000008>

Received: 13 December 2025; Accepted: 19 December 2025; Published: 26 December 2025

ABSTRACT

Generative Artificial Intelligence (AI) is revolutionizing healthcare with its transformative potential in medical imaging, drug discovery, and clinical decision-making. Generative models such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and large foundation models can synthesize realistic data, emulate biological systems, and accelerate innovation beyond traditional AI methods. In medical imaging, generative AI enhances diagnostic accuracy by enabling high-resolution image reconstruction, noise reduction, anatomical segmentation, and the creation of synthetic datasets to support algorithm training in data-scarce environments. These advancements assist radiologists in early disease detection, treatment planning, and longitudinal patient monitoring. In drug discovery, generative AI accelerates molecule design, lead optimization, and prediction of protein-ligand interactions, reducing time and cost while enabling precision therapeutics and drug repurposing. Clinically, it supports automated report generation, patient-specific treatment simulations, and digital twin development for disease modeling and trial optimization through synthetic patient cohorts. Despite these advances, challenges persist regarding data quality, interpretability, regulatory approval, ethical transparency, and bias mitigation, which are critical for ensuring patient trust and safety. This study explores generative AI's applications across medical imaging, pharmacology, and clinical workflows, highlighting its opportunities, limitations, and future directions toward sustainable, ethical, and patient-centered healthcare integration.

Keywords: Generative AI, Medical Imaging, Drug Discovery, Clinical Applications, Synthetic Data etc.

INTRODUCTION

Generative Artificial Intelligence (Generative AI) has rapidly emerged as one of the most transformative technologies in the field of computer science and biomedical research. It refers to a class of artificial intelligence models that learn to understand and replicate complex data distributions to create new, realistic data samples that resemble real-world inputs. In contrast to classic discriminative models that target to categorize or predict with the help of the already available data, generative models are built with a goal to produce new data instances, which can find numerous applications in healthcare as in Figure 1 [1].

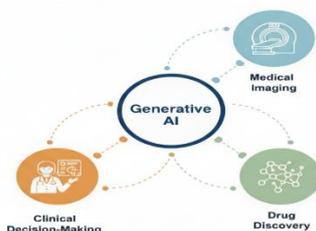


Figure 1: Generative AI application in healthcare

Generative modeling was founded on the development of Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs). The concept of the latent representation of data sampled to give new examples with controlled variability was developed by VAEs [2]. Conversely, GANs are trained on a competitive model where a generator (data generating) and a discriminator (data trying to differentiate between generated and real data) are involved [3]. The adversarial process has found application especially in generating real-world medical images, de-noising scans, and scaling down small datasets. Such models have been enhanced throughout the years to more sophisticated models such as diffusion models and transformer-based models that can be more fidelity and controllable on generative tasks [4]. Recent advancements in the diffusion probabilistic models including models like Denoising Diffusion Probabilistic Model (DDPM) and Latent Diffusion Model (LDM) have demonstrated a high level of performance in terms of high-quality image generation [5]. These models continuously increase the amount of detailing of random noise by successive improvement to high-quality images as a way of giving improved reconstruction of images in medical modalities such as MRI images and CT images. Similarly, transformer-based Large Language Models (LLMs), like GPT inspires, BERT inspires, and others, have surpassed the state of the art in generative text simulation by producing fluent, contextualized and domain specific text [6]. These models can be used in the clinical world to generate synthetic clinical notes, radiology reports and even assist in summarizing the history of patients without losing linguistic fluency and factual integrity [7].

Generative AI data augmentation and privacy-preserving synthetic data generation is among the most important beneficial aspects of AI in the healthcare sector. Privacy laws and the unavailability of expert-labelled samples (particularly rare diseases) can severely restrict medical data. Generative AI may be used to generate realistic but artificial patient data, enables training its models with no exposure to sensitive patient data [8]. This artificial data has the ability to produce balance in classes, pathological variations, and enhance robustness in the models. Moreover, medical images can be enhanced and de-noised with generative models, which allows improving the diagnostic quality of medical images by recovering missing or corrupted information in imaging datasets [9]. The other large-scale use is in molecular design and drug discovery. The process of traditional drug discovery is long and expensive, and in most cases may take years of experimentation and billions of dollars are involved. It can be massively accelerated through the use of generative models that can create new molecular structures that maximize the desired biological and chemical properties [10]. With powerful generative models such as graph neural networks and variational molecular generators, scientists are able to predict binding affinities, molecular stability and pharmacokinetics, speeding up the time between an idea and clinical test by a long way [11]. The use of generative AI can be used to identify promising drug candidates more effectively and faster than the methods of traditional computational chemistry because it searches large chemical spaces through an algorithm.

The clinical decision-making sector may utilize generative AI to contribute to multimodal data generation, clinical records and smart reports. An example of this could be that, using the example of electronic health records (EHRs), the trained LLAMs will be applicable in assisting a medical practitioner, as in summarizing patient encounters, generating discharge plans, and related potential diagnosis suggestions based on their textual and numerical input [12]. Such systems imply a reduced workload of clinicians and decreased documentation errors, which promotes efficiency in the working process and patient outcomes. Moreover, full-scale diagnostic reasoning may be assisted by using generative systems to integrate textual and visual information with genomic information in combination with multimodal AI models, when combined with other systems [13]. Although generative AI use in healthcare systems has a high potential, it also raises important ethical and regulatory concerns. The research concerns surrounding the data bias, model hallucination, interpretability, and the potential of misuse need a narrow scope of validation and control [14]. The quality, safety and interpretability of the generated medical data and recommendations remain a major issue before the large-scale clinical adoption. In order to address such challenges, the interaction between clinicians, AI researchers, and policymakers across diverse spheres is required so that they could come up with guidelines on how to authenticate their models, explain them, and apply them in an ethically sound manner [15]. The future of healthcare is generative AI and the disruptive aspect of such models can be undermined by the fact that they can produce high-quality artificial data, recombine and enhance medical imaging, design new therapeutic molecules and optimize clinical operations. This review will discuss three important fields of medical imaging, drug discovery and clinical decision-making and documentation that generative AI is transforming. All these domains represent the merge

of artificial intelligence and medicine towards a more data-driven, more efficient and more personalized health-related future.

METHODOLOGY OF LITERATURE REVIEW

The present review was based on a systematic and integrative approach that summarized the current research on the use, efficacy and issues of Generative Artificial Intelligence (AI) in healthcare. The methodological framework was formulated so that it would be able to guarantee transparency, reproducibility and scientific rigor in the identification and analysis of the relevant peer content.

Search Strategy

A comprehensive literature search involved the use of various academic databases, such as PubMed, Scopus, IEEE Xplore, ScienceDirect, and Google Scholar, the publications of which were printed in 2013-2025. Searching was done using Boolean operators and Medical Subject Headings (MeSH) where necessary, and it included the following terms: Generative Artificial Intelligence, Generative Adversarial Networks, Diffusion Models, Variational Autoencoders, Large Language Models, Medical Imaging, Drug Discovery and Clinical Decision Support.

Additionally, the reference lists of the major papers and reviews were searched manually to find additional literature not available via database search. Articles that were published in other languages were not taken into account. Articles published in English language were only considered.

Inclusion and Exclusion Criteria

All the articles if they satisfied the following criteria:

1. The article used or presented generative AI models (e.g., GANs, VAEs, diffusion models, or transformer-based LLMs) to a healthcare setting.
2. The research had either one or more of these domains; medical imaging, drug discovery, biomedical data generation, or clinical decision-making systems.
3. The publication in a peer-reviewed journal or significant conference proceedings has to take place between 2013 and 2025.

Articles were excluded if they:

1. Only discriminative AI methods or other irrelevant computational models.
2. They were non-peer-reviewed, editorial commentaries, or not experimental, methodological, or analytical.
3. Were published in non-English languages and without freely available full-text editions.

Data Extraction and Analysis

The initial search identified approximately 250 records. After removing duplicates and applying inclusion criteria, 85 articles were selected for detailed review. Each study was assessed according to the following parameters:

- Generative model type: VAE, GAN, Diffusion, or Transformer-based LLM.
- Application domain: imaging, pharmacological modeling, or clinical AI.
- Outcomes and performance metrics: image fidelity, accuracy, interpretability, or efficiency.

- Reported limitations: data bias, computational constraints, ethical or regulatory barriers.

The thematic groupings of the studies were divided into three areas i.e. medical imaging, pharmaceutical drug discovery, and clinical decision support to guarantee integrative knowledge of cross-disciplinary uses. Theoretical background of generative modeling was laid down by Kingma and Welling [2], Goodfellow et al. [3], and Rombach et al. [4], while applications were analyzed from studies such as Yi et al. [16] in imaging, Zhavoronkov et al. [10] in drug discovery and

Singhal et al. [7] in clinical knowledge modeling. Mittelstadt et al. [14] and Topol [15] had mostly directed the ethical and interpretability perspectives.

The qualitative method of synthesis was used to be able to discover the main trends, pros and cons through the study. The resulting comparative framework is summarized in Table 1.

Table 1. Comparative Summary of Key Literature on Generative AI Applications in Healthcare

Author(s) & Year	Model Type	Application Area	Key Contribution / Findings	Limitations Reported
Kingma & Welling (2013) [2]	Variational Autoencoder (VAE)	Probabilistic modeling and anomaly detection	Introduced latent-space encoding for controlled sampling and synthetic data generation	Generated images often blurred; limited clinical testing
Goodfellow et al. (2014) [3]	Generative Adversarial Network (GAN)	Image synthesis and enhancement	Established adversarial learning enabling realistic medical image generation	Training instability and mode collapse
Rombach et al. (2022) [4]	Latent Diffusion Model (LDM)	Radiological image enhancement	Achieved superior reconstruction of CT/MRI images using latent-space denoising	Requires large datasets and GPU resources
Singhal et al. (2023) [7]	Large Language Model (LLM)	Clinical text generation	Demonstrated transformer-based LLMs capable of encoding clinical reasoning and generating medical narratives	Hallucination and factual inconsistency risks
Chen et al. (2021) [8]	Hybrid VAE–GAN	Synthetic healthcare data generation	Created privacy-preserving synthetic datasets for model training without exposing real patient data	Residual re-identification risk in synthetic data
Zhavoronkov et al. (2019) [10]	Deep Generative Framework	Drug discovery and molecular screening	Accelerated identification of potent DDR1 kinase inhibitors using deep generative modeling	Necessitates biochemical validation of outputs
Sanchez-Lengeling & Aspuru-Guzik (2018) [11]	Variational Molecular Generator	Inverse molecular design	Enabled rapid generation of novel molecules with optimized pharmacological properties	Limited interpretability; experimental validation required

Mittelstadt et al. (2016) [14]	Algorithmic Ethics Framework	Ethical AI and transparency	Conceptualized principles of fairness, accountability, and explainability in healthcare algorithms	Theoretical framework; lacks empirical validation
Topol (2019) [15]	Human-AI Collaboration Model	Clinical decision support	Advocated human-in-the-loop AI systems to enhance diagnostic reliability and clinician trust	Implementation and regulatory barriers
Yi et al. (2019) [16]	GAN	CT and MRI image translation	Applied GANs for cross-modality synthesis and image denoising	Presence of hallucination artifacts; limited interpretability
Dhariwal & Nichol (2021) [17]	Diffusion Model	Image reconstruction	Demonstrated high-fidelity image generation with improved stability compared to GANs	High computational cost and slow sampling
Chung et al. (2022) [18]	Diffusion Model	MRI super-resolution and denoising	Demonstrated enhanced structural detail preservation and improved image quality	Limited clinical generalizability; high training cost

Generative Model Families: A Primer

The contemporary artificial intelligence is based on the generative models enabling machines to create realistic data points of different modalities images, text, audio, and molecular structures. The models have proved essential in the development of artificial medical data in healthcare, enhancement of image quality, as well as clinical decision-making. The significant families of generative models that provide particular capabilities and trade-offs to the medical science are Variational Autoencoders (VAEs), Generative adversarial networks (GANs), Diffusion Models, and Transformer-based Generative Models [19]. Various generative models are in Figure 2 below.

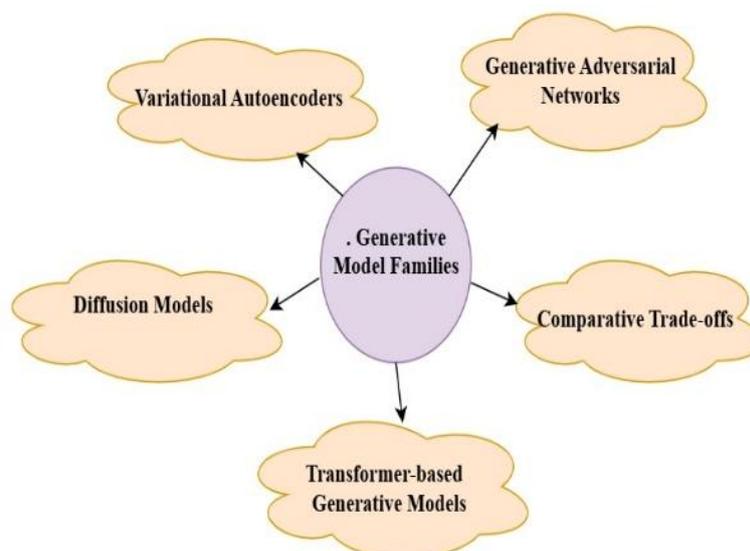


Figure 2: Generative Model Families

A. Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs) are probabilistic generative design frameworks which encode input data in a continuous latent space enabling controlled sampling and interpolating features [2]. The VAE is made up of encoder that encodes the input data into a latent probability distribution and the decoder that recovers data samples that encode the input data. This architecture is efficient in dimensionality reduction as well as the creation of new data points which then obey the same underlying distribution as the training data. VAEs are especially effective in the case of medical anomaly detection since they can be trained on the statistical regularities of normal data, which deviation should be identified as a possible abnormality [20]. They are also used in compressed medical image representations, and are useful in tasks such as telemedicine where bandwidth is crucial. However, VAEs often produce blurrier images compared to adversarial models, due to the trade-off between reconstruction accuracy and latent smoothness [21].

B. Generative Adversarial Networks (GANs)

Goodfellow et al. initiated Generative Adversarial Networks (GANs) to transform the way generative models are trained to adversarially train two models: a generator that generates synthetic samples and a discriminator that differentiates between actual and synthetic samples [3]. Such an adversarial structure pushes the generator to create very realistic examples. Medical imaging is another popular field in which GANs are used to translate images to images, e.g. to convert CT to MRI images, super-resolution reconstruction, and modality synthesis [16]. For instance, GANs have the potential to advance low-dose CT images and produce high-resolution results with little noise, contributing to diagnosis with the lowest radiation dose [22]. In spite of their success, GANs are characterized by the instability in training and mode collapse, in which the generator results in a low diversity of output [23]. Ongoing research addresses these limitations through improved architectures such as Wasserstein GANs and Cycle GANs.

C. Diffusion Models

Diffusion models are a more recent type of generative systems that demonstrate impressive results in the creation of high-fidelity images. These architectures are trained to progressively remove random noise in structured data by successive refinement steps effectively undoing a diffusion process that introduces noise to the data with training [5]. Diffusion models are more stable to train, have better mode coverage and state-of-the-art image quality compared to GANs [22]. Diffusion-based models have been utilized to reconstruct MRI, generate histopathological images and de-noise low-quality scans in the medical field [23]. Despite being computationally costly, because of their iterative sampling mechanism, diffusion models are gaining popularity as they are more robust and can produce fine-grained real-world details.

D. Transformer-based Generative Models / Large Language Models (LLMs)

The original transformer-based models, created to process natural language, have become the potent generative models designed to work with different types of data. Transformers are based on self-attention models that learn complex dependencies on sequential data, which can be used to produce coherent and context-sensitive text [6]. Large Language Models (LLM) like GPT, BERT, and Med-PaLM can be used to apply these features to clinical text generation, summary of reports, and generation of structured output based on unstructured clinical data [7]. Besides, multimodal generative tasks have been translated to transformer architectures, including textual, visual, and tabular data - effective in generating medical reports or interpreting radiology images [25]. They are at the heart of the future in AI-based healthcare because of their scalability and flexibility.

E. Comparative Trade-offs

All the generative model families have their own unique advantages and difficulties. VAEs provide interpretable latent space and provide smooth data interpolation but can produce a blurry image. GANs are very much good in generating clear, photorealistic images but at the same time experience instability in their training. Diffusion models are seen to have better quality and diversity of the samples, but need a massive amount of computation power. Transformers and LLMs perform exceptionally well in the text and sequence modeling tasks but require

big data and computational units. The choice of the generative model is, therefore, dependent on the specific healthcare task - be it image synthesis, text generation or multimodal data fusion [15].

Transforming Medical Imaging

One of the most data-intensive and, at the same time, computationally intensive fields of healthcare is medical imaging. Although the imaging modalities, like MRI, CT, PET, and ultrasound, have evolved, several issues still exist, such as the lack of annotated datasets, differences in image quality, and long acquisition time. Generative AI has become a revolution in this area, making it possible to generate synthetic data, cross-modality translation, and image reconstruction not to mention that it enhances segmentation, harmonization, and diagnostic accuracy. The working of generative AI in medical imaging is shown in figure 3. All such innovations contribute to achieving better generalization of the models, cost reductions, and increased accessibility in the clinical imaging processes [26].

A. Data Augmentation and Synthetic Data

Medical imaging data are frequently small-scale, unbalanced, or biased against some population groups because of privacy issues and expert labeling is expensive. Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and diffusion models are generative models that create realistic synthetic images that can be used to augment training datasets to ultimately enhance the performance of downstream classifiers in disease detection and segmentation [27]. As an illustration, GAN-based augmentation has been demonstrated to enhance the accuracy of tumor detection by generating realistic lesion and tissue texture variations in MRI images [28]. The simulation of rare pathologies that are not represented sufficiently in clinical repositories is also supported by synthetic data, which guarantees enhanced generalization of the model [29]. Moreover, artificial amplification is used to adjust the demographic and device-level diversity, reducing the bias in the algorithms. Notably, with a validated synthetic image, it is possible to train models without accessing real patient data, both in terms of ethical and privacy concerns in the context of data protection laws, like HIPAA and GDPR [8].

B. Cross-Modality Synthesis

Another strong tool of generative AI is cross-modality synthesis which improves the translation of imaging modalities including CT and MRI or MRI and PET. Paired datasets can be trained to generate one modality of the other, further simplifying the processes that are both costly and invasive to perform [30]. As an illustration, Cycle GANs have been applied in the synthesis of MRI images given CT scans to enable a physician view soft-tissue contrast without having to perform another MRI scan [31]. This modality completion enables the multi-modal analyses to be complete in a case where some scans are not present, or incomplete and contraindicated. Cross-modality synthesis also enhances the interoperability of data, and as such, it is now possible to combine heterogeneous data sources across imaging centers [32].

C. Image Reconstruction and Enhancement

Image reconstruction and enhancement stimulate the use of generative models that allow high-quality images to be generated based on under sampled or low-dose acquisitions. In MRI, VAEs and GAN reconstruction models are trained on the prior distribution of anatomical structures, which they use to recover missing information in sparse k-space data to generate better and faster reconstructions compared to traditional compressed sensing algorithms [33]. Equally, diffusion-based models can be trained to refine erroneous or under-resolution scans, which give fidelity results maintaining subtle structural information [23]. The use of GANs in CT has been in low-dose reconstruction, which is effective in reducing radiation dose and retaining diagnostic quality [20]. Such innovations help in the acceleration of scanning, reduction of risk to the patient and enhancement of image quality both in the research and clinical practice.

D. Anomaly Detection and Segmentation

Generative models are able to train on the appearance of healthy anatomy and then identify abnormalities by showing differences between normal patterns. Reconstruction-error and latent-space-distance techniques have

the capability to emphasize suspicious areas which can be tumors, lesions, or structural abnormalities [17]. An example is that VAEs which are trained on normal brain MRIs can be used to detect areas of pathological tissue through comparison between input and reconstructed images. Moreover, GAN adversarial training improves segmentation performance, in which the discriminator imposes anatomy realism on segmentation masks [34]. Diffusion models have also been shown to be useful in producing accurate and consistent organ and lesion segmentation maps, and are also more robust on semi-supervised medical segmentation challenges [35]. The use of generative AI in medical imaging is as in the Figure: 3.

E. Image Harmonization and Bias Mitigation

An ongoing difficulty in large-scale medical imaging research is domain shift, which arises due to the change in imaging procedures, scanner vendors, and site factors. This type of heterogeneity impairs the inter-institutional generalization of deep learning models. Generative domain adaptation helps to solve this problem by aligning the images across an image source into a single visual style without anatomical information loss [36]. Standardization of MRI and CT scans across multi-centers studies has been applied to the style transfer frameworks and harmonization frameworks of GAN, improving the reproducibility and decreasing biasing factors because of confounding factors [37]. The strategies will be essential in the achievement of the establishment of robust AI solutions that can be effectively applicable in various populations and imaging contexts.

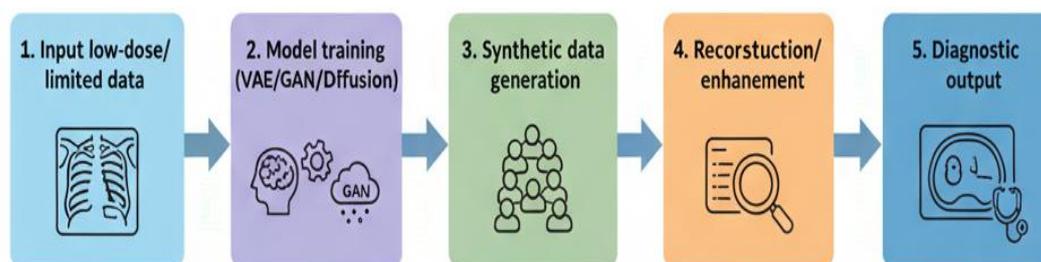


Figure 3: Generative AI in medical Imaging

F. Risks and Failure Modes in Imaging

Regardless of the massive progress in medical imaging, generative AI in medical imaging has its risks. Synthetic images may contain hidden artifacts or hallucinated objects that may be confusing to radiologists or diagnostic codes [38]. Overweight synthetic augmentation without proper validation may end up giving spurious correlation hence rendering the model incompetent. Also, diffusion and adversarial systems may sometimes produce anatomically constrained or clinically misleading images, especially when trained on small and/or biased data sets. It should be ensured that the clinical safety is guaranteed by the strict validation procedures, reader tests, and frameworks of uncertainty quantification [39]. The generative processes must be transparent as well as interpretable to eliminate diagnostic errors and keep confidence among the clinicians.

Accelerating Drug Discovery

Drug discovery process has been perceived as a long, expensive and unpredictable process that may take more than a decade and billions of dollars to navigate a single compound throughout the entire process before it can be put into the market. The paradigm shift of generative artificial intelligence (AI) is discussion of large chemical spaces, creation of new molecules possessing preferred biological characteristics, predictive and synthetic feedback following which enhances the discovery cycles. Generative AI is transforming the in vivo manner in which chemical and biological design are done through the use of graph-based models, language-inspired molecular generators, and diffusion architectures. The comprehensive pipeline illustrating generative AI process to accelerate drug discovery is in Figure 4.

A. Generative Molecular Design

Generative models are able to effectively search through chemical space by suggesting new molecules based on optimized therapeutic characteristics like solubility, potency, selectivity as well as toxicity [11]. Early approaches

used Variational Autoencoders (VAEs) to encode molecular structures-typically represented as SMILES strings or molecular graphs into latent spaces, allowing smooth interpolation and sampling of new compounds [40]. Graph-based generative models, like Graph VAEs and Graph GANs, model molecules as nodes (atoms) and edges (bonds), and they use structural and topological data that are important to the molecular validity [41].

More recent diffusion-based molecular generators progressively improve random molecular noise into chemically valid compounds, and they are state-of-the-art in terms of diversity and fidelity in de novo molecular design [42]. SMILES strings have also been represented as a series of autoregressive sequence model based on natural language processing to generate new molecules in a similar manner as text in generated grammars [43]. Also, conditional generation has increased control as the model can be requested to acquire properties or biological goals desired such as binding affinity to target receptors over which to wish the output to be formed in the molecular design process [44]. Together, these methods have radically enlarged the chemical design space allowing AI systems to suggest candidate molecules inaccessible to the usual computational chemistry methods.

B. Virtual Screening and Property Prediction

In traditional drug discovery, virtual screening is the process of assessing millions of compounds on a biological target through docking or predictive chemistry. Generative AI enhances this reaction by combining property-predictive networks with generative molecular design, to form candidate libraries which are pre-enriched with desirable pharmacological properties [45]. In essence, this coupling is an effective way of reducing search space in the experiment, reducing the number of compounds that are being subject to expensive wet-lab screening.

For example, generative models based on reinforcement learning and conditional GAN have been applied to bias molecule generation to better drug-likeness or ADMET (Absorption, Distribution, Metabolism, Excretion, and Toxicity) properties [46]. Through real-time drug discovery, AI systems are able to search and analyze novel compounds at the same time, making early drug discovery more informed by data and economical [47].

C. Protein Structure and Design

Generative AI and structural biology have combined to use new opportunities in protein modeling and design. Advances like AlphaFold2 showed that deep learning is capable of predicting protein tertiary structures using amino acid sequences correctly, and this fact made generative protein engineering possible [48]. Continuing on this, generative protein design models have the potential of proposing new protein sequences that can form desired 3D folds or that have particular biochemical properties [49].

Other methods such as Protein GAN and diffusion-based protein models create completely novel proteins or peptides through learning the statistical regularities of protein families in nature [50]. Such systems are used to mediate binding interface design, which allows tailored generation of enzymes, antibodies or therapeutic peptides that are designed to bind a specific ligand or antigen. Additionally, generative models can scan mutations in silico of sequence variations systematically to find stabilizing or functional mutations and experimentally synthesize them [51]. These developments are promising in the design of custom biologics and the next generation protein therapeutics.

D. De Novo Drug Design and Synthesis Planning

De novo drug design is one of the most radically applicable areas of generative AI, which involves models that suggest new synthetically accessible molecules with optimized pharmacological activity [11]. These models can be integrated with retrosynthesis planning tools to produce new compounds in addition to suggesting a possible synthetic pathway to synthesize them in laboratory [52]. As an example, reaction prediction networks in transformer form can predict the results of chemical reactions and proposed multi-step synthesis pathways which may minimize cost and complexity [53].

AI systems are achieving the cycle between design and manufacturing by integrating generative molecular design, reaction prediction, and retrosynthetic analysis. There is a range of end-to-end automated drug discovery workflows, with identification of targets and the generation of compounds, to their synthesis and optimization,

supported by this integration. The possibility of reasoning about chemical feasibility as well as biological functionality has enabled researchers to put a high priority on the compounds which are not only effective but also viable to be generated at scale.

E. Challenges in Molecular Generative Models

In spite of such an impressive advancement, there are some obstacles. A key consideration - most generated molecules need to be chemically valid and synthetically feasible [54]. The second important problem is to balance between novelty versus similarity to known scaffold; too new, the compounds may become unrealistic, too similar, the scaffold is overfitted and does not serve as a source of innovation [55].

In addition to that, generative models must be capable of balancing between exploration and exploitation - exploring new parts of chemical space and enhancing the current promising frameworks. Combinations of active learning schemes and experimental feedback loops help models to perfect themselves over time by trying in the field [56]. Finally, reproducibility and interpretability remain a constant issue with many deep generative systems being a black box that it is difficult to make rationalizations on how certain structures are generated. Nevertheless, using bigger datasets, and incorporating physical, biochemical and experimental restraints, AI-assisted drug discovery is getting much closer to autonomous design systems capable of discovering viable drug targets in weeks, rather than years.

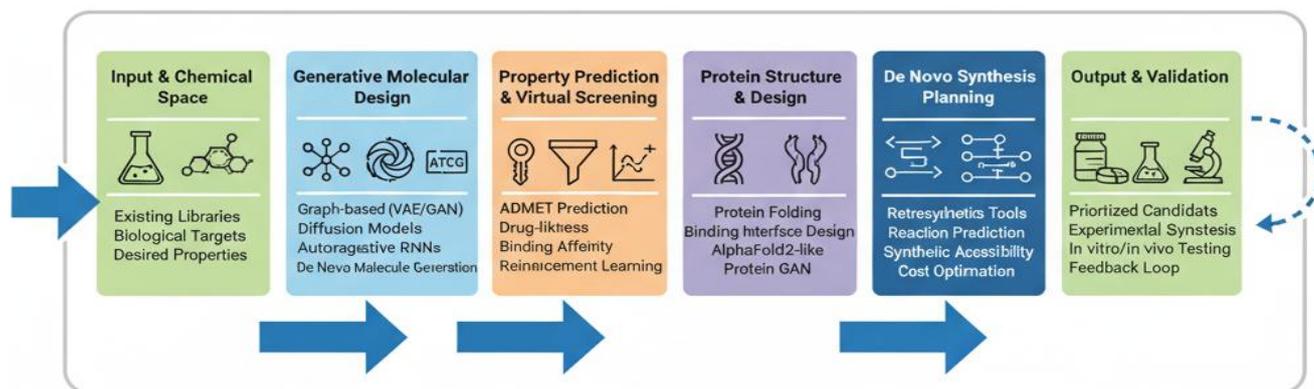


Figure 4: Accelerated Drug Discovery Pipeline Leveraging Generative Artificial Intelligence

Enhancing Clinical Decision-Making

When applied to the clinical processes, the application of the generative artificial intelligence (AI) has the transformative ability to improve the quality of the diagnostic process and patient documentation along with the quality of care in general. AI systems can offer more clinician expertise and simplify decision-making as a result of multimodal clinical data being trained on text, images, laboratory values, genomics, etc. The instruments should reduce the load on the administration, enable more timely and accurate diagnosis and tailor the treatment plan. However, they ought to be successfully implemented with assistance of good validation, interpretability and safety.

A. Clinical Text Generation and Documentation

One of the nearest AI generative applications in the health area is the automation of clinical documentation. Transformer-based large language models (LLMs), such as GPT, Med-PaLM, and Bio GPT, can generate, summarize, and structure clinical text with the stunning fluency of transformers [7]. These systems are able to automatically write clinical notes, summarize patient visits and extract structured data of unstructured medical records. By doing so, they have a significant effect of reducing the number of administrative burdens which occupy up to 40 percent of the working time of physicians [57].

The accuracy of coding and recording of bills can also be enhanced because the appropriate ICD or CPT code will be produced because generative models are being trained on electronic health record (EHR) corpora [58].

In addition, they enable enriching and rendering EHR representations more searchable as they assist in transforming fragmented data into coherent and longitudinal narratives that reflect the entire clinical account of a patient. Indicatively, e.g. Google MedLM which is a division of Microsoft Nuance DAX systems are real-time auto-documentation systems that overhear clinician-patient discussions and generate the automated summaries [59]. This ensures efficiency and further allows clinicians to be in a position to focus more on patient care than on data entry.

B. Diagnostic Decision Support

Generative models are also developed to become an important tool of diagnostic decision support. Depending on the synthesis of the patient history, the results of the imaging, and laboratory data, these models might generate the list of the possible diagnoses, suggest the further steps of diagnosis, and even reproduce the possible patient progression [60]. They are able to prioritize use of potential conditions through probabilistic reasoning or Bayesian systems of inference, to inform clinicians towards high-value tests and interventions [61].

For instance, the use of LLMs trained on clinical cases can generate structured chains of reasoning - describing why some diagnoses are more likely to be true - facilitating the educational and clinical decision-making processes [62]. Generative systems such as Med-PaLM 2 have shown near-expert performance on questions in the style of the U.S. medical licensing examinations, which illustrates the possibility of such systems being cognitive collaborators of clinicians [63]. However, there is a need to subject the diagnostic applications to stringent validation with the real-world data to prove reliability and safety before implementation.

C. Multimodal Patient Representations

Another important area of clinical AI advances is the creation of multimodal generative AI systems that integrate various types of patient data - radiological images, lab findings, genetic data, and clinical text into single, context-sensitive representations [64]. The prognostic modeling, risk stratification, and individual patient-specific treatment recommendations based on the integrated patient embeddings are possible [65].

As an example, text note to imaging data can be learned using multimodal transformer networks like Med-BERT and BioViL, in which the networks can generate sensible clinical summaries of both quantitative and qualitative results [66]. Similarly, diffusion as well as generative graph models are also capable of predicting disease progression or even imagining missing modalities (e.g. generate synthetic MRI images with the assistance of CT images) in clinical diagnosis [67]. It is possible to have such confluence of generative modeling and precision medicine, making it possible to transition towards more of a population to patient-specific approach to healthcare.

D. Human-AI Collaboration and Explainability

In healthcare, generative AI is superior as compared to autonomous AI. These systems are an intelligent assistant, which can best be seen as feeding clinicians with synthesized information, describing other available options and helping them to make complex decisions [15]. The implementation of models is valuable to guarantee that clinicians have faith and understanding of model outputs. The aspects of clinical features that best influenced the recommendation or generation of the model can be identified using the visualization tools, such as saliency maps, attention heat map, and counterfactual explanations [68]. Such transparency will help the clinicians examine AI advice critically and this will make them more responsible and capable of being trusted. The models of human - AI collaboration where the clinician will always be the ultimate decision-maker allow avoiding the over-reliance on the automated system without underestimating the complementary nature of human and machine capabilities.

E. Safety, Hallucinations, and Verification

There are high risks associated with generative models in medicine despite the promise. Hallucinations occur in LLMs too, where the outputs consist of statements that are realistic, but not real [69]. Unless included, such errors in clinical practice can lead to an incorrect diagnosis or treatment. In this regard, there is the necessity of having

powerful verification pipelines. These include retrieval-augmented generation (RAG) which is a model where models rely on verifiable information sources e.g. EHRs, clinical guidelines, or peer-reviewed literature [70]. The development of the confidence estimation is another field under the research of the developers and allows AI-based systems to showcase the results of low confident or doubtful results to be reconsidered by a human [71]. Conservative user interface (UI) designs, such as highlighting statements that are not proven, or making clinicians confirm before adding to the medical record help to reduce risk. Also, to monitor model procedures and verify the safety, privacy and fair standards, regulating mechanisms and audit track are being developed [72].

Finally, the implementation of generative AI in clinical practice must be aimed at improving clinician judgment, and not to substitute them. Explainability, safety, and human oversight can be used to convert clinical decision-making through generative models to provide more efficient, data-driven, and equitable health care delivery. The comprehensive framework of generative AI in making clinical decision is in the Figure 5.

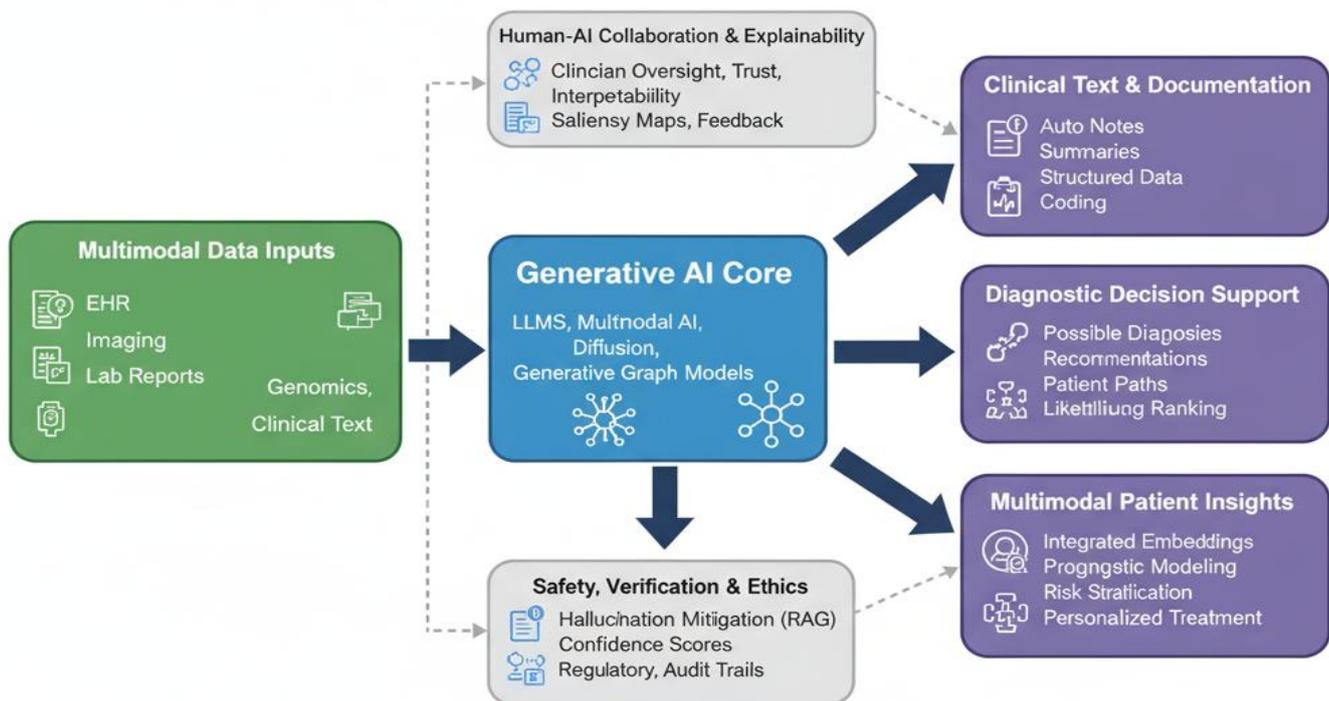


Figure 5: Generative AI in clinical decision-making: A comprehensive framework

Evaluation, Validation and Regulatory Considerations

Generative AI translation into clinical practice needs to be evaluated carefully, validated openly, and follow the existing regulatory guidelines. Conventional performance measures - e.g. image fidelity or text fluency do not suffice in a medical context where clinical safety and diagnosis accuracy take precedence. The evaluation protocols ought to include task-specific measures which entail sensitivity, specificity, false-negatives, and calibration, and downstream clinical utility [73]. Technical quality in the imaging domain can be evaluated based on such metrics as structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR) but, clinical quality measurement based on radiologist rating or diagnostic accuracy testing is still required [74]. The problem is that the real-world effectiveness and safety need to be defined, and prospective clinical trials and randomized controlled experiments are the only possible variants to achieve it. Despite the fact that the retrospective validation is an easy preliminary indicator of the feasibility, it may be challenging to mirror the concerns of the clinical workflow and patient heterogeneity [75]. Prospective assessment will be employed to ensure that the generative models are the same across different populations, different healthcare settings, and different imaging systems and to remove bias and improve generalizability [76].

To aid in maintaining transparency, traceability, and risk management in AI-based medical devices, government regulatory bodies such as the U.S. Food and Drug Administration (FDA), the European Medicines Agency

(EMA), and the Central Drugs Standard Control Organization (CDSCO) in India have demanded transparency, traceability, and risk management [77]. Practitioners: to develop, document datasets they are going to learn and test on, have pre-established use-cases, and, introduce post-market surveillance to note model drift or unexpected harms [78].

The other factor is the privacy of data. The synthetic models which generalize based on medical data may facilitate data sharing without loss of confidentiality, but re-identification threats still exist since they may still contain traces of real patient data [79]. Therefore, in order to verify the compliance with the frameworks such as the HIPAA or GDPR, it is necessary to conduct empirical tests on privacy-protecting tools - such as differential privacy or membership inference audits [80].

Ethical and Social Implications

The prospect of providing healthcare services with the help of generative artificial intelligence (AI) raises serious ethical and social concerns related to the potential of the branch. One of the most significant issues is prejudice and fairness. The systems that are modeled using datasets that lack the presence of a particular demographic, ethnic, or socioeconomic group will potentially perpetuate or intensify the existing healthcare inequality [81]. An example is that when imaging or text data are mostly predominant with data of high-income groups or certain ethnicities, the diagnostic accuracy can be worse with marginalized groups [82]. Although an imbalance can be resolved using synthetic data generation, which creates heterogenous samples, it may also encourage systemic bias in the event that the training data generated to create these samples is imbalanced [8].

Other serious issues are accountability and liability. When generative AI is used to assist in diagnostic reasoning, documentation or treatment recommendations, the issue of responsibility between clinicians, institutions and developers is complicated in a scenario where generative AI is helpful [83]. Regulatory frameworks treat clinicians as ultimate decision-makers, and as AI systems increasingly gain autonomy, there is a need to establish more explicit governance structures that would help place liability in the event of damage or malpractice [84].

It is also necessary that consent and transparency. Patients must be informed of situations when AI systems can assist in their diagnosis or care decisions and when they appreciate their autonomy and trust clinical practice [85]. In addition, provenance and validation of the synthetic datasets training models are also to be reported, as the source of data affects reliability and fairness [86]. The promotion of the open documentation and ethical auditing and the open-ended dataset design will be the determinants to ensure that the positive aspects of generative AI will be equally beneficial to all patients.

Best Practices and Recommendations

Effective deployment of generative AI in healthcare requires adherence to rigorous best practices. Robust data practices involve assembling diverse, representative, and well-annotated datasets, with clear documentation of provenance and preprocessing steps. Model validation should include multi-site and prospective studies, explicitly reporting failure modes and edge cases. Explainability and uncertainty measures, along with evidence retrieval, ensure interpretability and trust. Maintaining a human-in-the-loop approach keeps clinicians central, using AI to augment rather than replace judgment. Compliance is simplified through initial engagement in regulation and through open reports, whereas privacy-sensitive strategies such as differential privacy, federated learning, and synthetic data audits - reduce the risk of patient confidentiality.

DISCUSSION

The reviewed evidence highlights that Generative Artificial Intelligence (AI) represents a transformative advancement in healthcare, particularly within medical imaging, drug discovery, and clinical decision support.

The evidence reviewed underscores the fact that Generative Artificial Intelligence (AI) is a revolutionary development in healthcare, especially in the medical imaging, drug discovery, and clinical decision support. GANs and Diffusion models have proven potential in medical imaging when using them in image reconstruction, denoising and cross-modality synthesis [3], [4], [16], [17]. The diffusion-based models are superior to the

traditional GANs in structural fidelity and noise reduction but use significantly more computing power [18]. These developments serve as part of the minimized radiation dose and increased diagnostic quality, making generative modeling a fundamental aspect of the next-generation imaging processes.

Generative frameworks are used in drug discovery to do de novo molecular design and lead compound optimization which accelerates preclinical research. The works by Zhavoronkov et al. [10] and Sanchez-Lengeling and Aspuru-Guzik [11] demonstrate that deep generative systems are able to recognize bioactive compounds with a high degree of precision and decreases the time needed to develop and the cost of the experiment itself. But, experimental validation, and toxicity testing are still important precedents to clinical translation.

Transformer-based Large Language Models (LLMs) have become the decision-support tool in the field of clinical decision-making that can reason and produce reports with medical accuracy [7]. Although they can process complex clinical data, there are still problems of hallucination, propagation of bias, and not being interpretable. It is crucial to consider using human-in-the-loop systems that have been suggested by Topol [15] to make sure the system is reliable and ethically acceptable.

Ethical and regulatory issues that were revealed by Mittelstadt et al. [14] focus on the need to implement transparency, fairness, and accountability concerning the use of AI in clinical settings. The disadvantages that fail to support the broad adoption of these challenges include data privacy, model bias, and small cross-institutional validation.

Overall, there is a significant potential of generative AI in improving precision medicine with the help of autonomy, simulation, and data generation. However, its introduction in clinical practice needs well-defined validation systems, regulatory controls and ongoing cross-disciplinary cooperation to make it secure, reliable and reproducible.

Future Directions

- A. Multimodal foundation models trained on large-scale, heterogeneous clinical datasets to enable zero/few-shot clinical tasks.
- B. Integration of physics and domain knowledge into generative priors for medically plausible synthesis (e.g., physics-informed reconstruction in imaging).
- C. Closed-loop discovery platforms that tightly integrate generative design, active learning, and automated synthesis/assay systems.
- D. Personalized generative models for individualized treatment planning and longitudinal health trajectory simulation.

CONCLUSION

Generative artificial intelligence (AI) has emerged as a transformative technology with the potential to reshape multiple facets of healthcare. In medical imaging, generative models enable synthetic data augmentation, cross-modality synthesis, high-fidelity reconstruction, and anomaly detection, thereby enhancing diagnostic accuracy, reducing scan times, and mitigating risks from limited or imbalanced datasets. In drug discovery, AI-driven molecular design, virtual screening, protein engineering, and retrosynthesis planning accelerate the identification of novel compounds and biologics, reducing both the time and cost associated with traditional pipelines. Additionally, in clinical decision-making, transformer-based models and multimodal systems support automated documentation, risk stratification, personalized treatment recommendations, and simulation of patient trajectories, complementing clinician expertise and enabling data-driven, precision healthcare.

However, realizing the full potential of generative AI requires rigorous evaluation and validation across retrospective and prospective clinical settings, multi-site studies, and randomized trials. Ethical and social considerations—including bias, equity, consent, transparency, and accountability—must be carefully addressed to ensure fair and responsible deployment. Safety mechanisms such as uncertainty quantification, retrieval augmentation, and human-in-the-loop oversight are critical to prevent hallucinations, misdiagnoses, or

overreliance on AI outputs. Furthermore, regulatory engagement, adherence to standards, and privacy-preserving techniques are essential to maintain trust and compliance in clinical contexts.

With careful stewardship, interdisciplinary collaboration between machine learning researchers, clinicians, regulatory authorities, and patients can ensure that generative AI serves as an augmentative tool rather than a replacement for human expertise. These models have the capacity to shorten drug development cycles, improve diagnostic precision, and enhance personalized care, ultimately contributing to more efficient, equitable, and patient-centered healthcare systems. As technology matures, the integration of generative AI will likely become a cornerstone of modern medicine, bridging computational innovation with tangible clinical impact.

Ethical Statement

None of the authors have conducted any research on humans or animals for this paper.

Conflicts of Interest

The authors declare no conflicts of interest related to this work.

Data Availability Statement

Since this study did not create or analyse any new data, data sharing is not applicable to this article.

FUNDING

The author(s) declare that no financial support was received for the research, authorship, and/or publication of this article.

Author Contribution Statement

Khursheed Ahmad Ganie: Concept, Writing, Methodology and Editing

Mohit Sharma: Writing, Methodology, Formal Analysis and Validation

Priti Panwar: Editing and Review

Arti Sharma: Proofreading

Conflict of Interest

The authors do not have any conflicts of interest.

REFERENCES

1. Cacciamani GE, Siemens DR, Gill I. Generative artificial intelligence in health care. *J Urol.* 2023;210(5):723–5.
2. Kingma DP, Welling M. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114. 2013 Dec 20.
3. Goodfellow IJ, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. *Adv Neural Inf Process Syst.* 2014;27.
4. Rombach R, Blattmann A, Lorenz D, et al. High-resolution image synthesis with latent diffusion models. In: *Proc IEEE/CVF Conf Comput Vis Pattern Recognit*; 2022. p. 10684–95.
5. Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. *Adv Neural Inf Process Syst.* 2020;33:6840–51.
6. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. *Adv Neural Inf Process Syst.* 2017;30.
7. Singhal K, Azizi S, Tu T, et al. Large language models encode clinical knowledge. *Nature.* 2023;620(7972):172–80.

8. Chen RJ, Lu MY, Chen TY, et al. Synthetic data in machine learning for medicine and healthcare. *Nat Biomed Eng.* 2021;5(6):493–7.
9. Armanious K, Jiang C, Fischer M, et al. MedGAN: Medical image translation using GANs. *Comput Med Imaging Graph.* 2020;79:101684.
10. Zhavoronkov A, Ivanenkov YA, Aliper A, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nat Biotechnol.* 2019;37(9):1038–40.
11. Sanchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: generative models for matter engineering. *Science.* 2018;361(6400):360–5.
12. Patel BN, Rosenberg L, Willcox G, et al. Human–machine partnership with artificial intelligence for chest radiograph diagnosis. *NPJ Digit Med.* 2019;2(1):111.
13. Esteva A, Chou K, Yeung S, et al. Deep learning-enabled medical computer vision. *NPJ Digit Med.* 2021;4(1):5.
14. Mittelstadt BD, Allo P, Taddeo M, et al. The ethics of algorithms: mapping the debate. *Big Data Soc.* 2016;3(2):2053951716679679.
15. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med.* 2019;25(1):44–56.
16. Yi X, Walia E, Babyn P. Generative adversarial network in medical imaging: a review. *Med Image Anal.* 2019;58:101552.
17. Dhariwal P, Nichol A. Diffusion models beat GANs on image synthesis. *Adv Neural Inf Process Syst.* 2021;34:8780–94.
18. Chung H, Lee ES, Ye JC. MR image denoising and super-resolution using regularized reverse diffusion. *IEEE Trans Med Imaging.* 2022;42(4):922–34.
19. Bond-Taylor S, Leach A, Long Y, et al. Deep generative modelling: a comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models. *IEEE Trans Pattern Anal Mach Intell.* 2021;44(11):7327–47.
20. Zimmerer D, Kohl SA, Petersen J, et al. Context-encoding variational autoencoder for unsupervised anomaly detection. *arXiv preprint arXiv:1812.05941.* 2018 Dec 14.
21. Burgess CP, Higgins I, Pal A, et al. Understanding disentangling in β -VAE. *arXiv preprint arXiv:1804.03599.* 2018 Apr 10.
22. Wolterink JM, Leiner T, Viergever MA, et al. Generative adversarial networks for noise reduction in low-dose CT. *IEEE Trans Med Imaging.* 2017;36(12):2536–45.
23. Salimans T, Goodfellow I, Zaremba W, et al. Improved techniques for training GANs. *Adv Neural Inf Process Syst.* 2016;29.
24. Chen S. A review of multimodal large model-based medical image report generation. *Front Med Sci Res.* 2025;7(3).
25. Lundervold AS, Lundervold A. An overview of deep learning in medical imaging focusing on MRI. *Z Med Phys.* 2019;29(2):102–27.
26. Frid-Adar M, Klang E, Amitai M, et al. Synthetic data augmentation using GAN for improved liver lesion classification. In: *IEEE Int Symp Biomed Imaging*; 2018. p. 289–93.
27. Han C, Rundo L, Araki R, et al. Infinite brain MR images: PGGAN-based data augmentation for tumor detection. In: *Neural Approaches to Dynamics of Signal Exchanges.* Singapore: Springer; 2019. p. 291–303.
28. Shin HC, Tenenholtz NA, Rogers JK, et al. Medical image synthesis for data augmentation and anonymization using GANs. In: *Simulation and Synthesis in Medical Imaging.* Cham: Springer; 2018. p. 1–11.
29. Chatsias A, Joyce T, Giuffrida MV, et al. Multimodal MR synthesis via modality-invariant latent representation. *IEEE Trans Med Imaging.* 2017;37(3):803–14.
30. Wolterink JM, Dinkla AM, Savenije MH, et al. Deep MR to CT synthesis using unpaired data. In: *Simulation and Synthesis in Medical Imaging.* Cham: Springer; 2017. p. 14–23.
31. Hiasa Y, Otake Y, Takao M, et al. Cross-modality image synthesis from unpaired data using CycleGAN. In: *Simulation and Synthesis in Medical Imaging.* Cham: Springer; 2018. p. 31–41.
32. Hammernik K, Klatzer T, Kobler E, et al. Learning a variational network for reconstruction of accelerated MRI data. *Magn Reson Med.* 2018;79(6):3055–71.

33. Xue Y, Xu T, Zhang H, et al. SEGAN: Adversarial network with multi-scale L1 loss for medical image segmentation. *Neuroinformatics*. 2018;16(3):383–92.
34. Zhang Z, Yao L, Wang B, et al. Diffboost: Enhancing medical image segmentation via text-guided diffusion model. *IEEE Trans Med Imaging*. 2024;17 Dec.
35. Dewey BE, Zhao C, Reinhold JC, et al. DeepHarmony: A deep learning approach to contrast harmonization across scanner changes. *Magn Reson Imaging*. 2019;64:160–70.
36. Pinto MS, Paolella R, Billiet T, et al. Harmonization of brain diffusion MRI: concepts and methods. *Front Neurosci*. 2020;14:396.
37. Cohen JP, Luck M, Honari S. Distribution matching losses can hallucinate features in medical image translation. In: *MICCAI Conf Med Image Comput Comput-Assist Interv*. Cham: Springer; 2018. p. 529–36.
38. D'Antonoli TA. Ethical considerations for artificial intelligence: an overview of the current radiology landscape. *Diagn Interv Radiol*. 2020;26(5):504.
39. Gómez-Bombarelli R, Wei JN, Duvenaud D, et al. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent Sci*. 2018;4(2):268–76.
40. Simonovsky M, Komodakis N. GraphVAE: Towards generation of small graphs using variational autoencoders. In: *Int Conf Artif Neural Netw*. Cham: Springer; 2018. p. 412–22.
41. Hoogeboom E, Satorras VG, Vignac C, et al. Equivariant diffusion for molecule generation in 3D. In: *Int Conf Mach Learn*. 2022. p. 8867–87.
42. Segler MH, Kogej T, Tyrchan C, et al. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Cent Sci*. 2018;4(1):120–31.
43. Nigam A, Pollice R, Krenn M, et al. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (STONED) algorithm for molecules using SELFIES. *Chem Sci*. 2021;12(20):7079–90.
44. Chen H, Engkvist O, Wang Y, et al. The rise of deep learning in drug discovery. *Drug Discov Today*. 2018;23(6):1241–50.
45. Popova M, Isayev O, Tropsha A. Deep reinforcement learning for de novo drug design. *Sci Adv*. 2018;4(7):eaap7885.
46. Gao W, Coley CW. The synthesizability of molecules proposed by generative models. *J Chem Inf Model*. 2020;60(12):5714–23.
47. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583–9.
48. Strokach A, Becerra D, Corbi-Verge C, et al. Fast and flexible design of novel proteins using graph neural networks. *bioRxiv*. 2019 Dec 10:868935.
49. Repecka D, Jauniskis V, Karpus L, et al. Expanding functional protein sequence spaces using generative adversarial networks. *Nat Mach Intell*. 2021;3(4):324–33.
50. Anishchenko I, Pellock SJ, Chidyausiku TM, et al. De novo protein design by deep network hallucination. *Nature*. 2021;600(7889):547–52.
51. Coley CW, Green WH, Jensen KF. Machine learning in computer-aided synthesis planning. *Acc Chem Res*. 2018;51(5):1281–9.
52. Schwaller P, Vaucher AC, Laino T, et al. Prediction of chemical reaction yields using deep learning. *Mach Learn Sci Technol*. 2021;2(1):015016.
53. Elton DC, Boukouvalas Z, Fuge MD, et al. Deep learning for molecular design—a review of the state of the art. *Mol Syst Des Eng*. 2019;4(4):828–49.
54. Arús-Pous J, Blaschke T, Ulander S, et al. Exploring the GDB-13 chemical space using deep generative models. *J Cheminform*. 2019;11(1):20.
55. Griffiths RR, Hernández-Lobato JM. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chem Sci*. 2020;11(2):577–86.
56. Shanafelt TD, Dyrbye LN, Sinsky C, et al. Relationship between clerical burden and characteristics of the electronic environment with physician burnout. *Mayo Clin Proc*. 2016;91(7):836–48.
57. Teng F, Ma Z, Chen J, et al. Automatic medical code assignment via deep learning approach for intelligent healthcare. *IEEE J Biomed Health Inform*. 2020;24(9):2506–15.
58. Baker HP, Dwyer E, Kalidoss S, et al. ChatGPT's ability to assist with clinical documentation: a randomized controlled trial. *JAAOS*. 2024;32(3):123–9.

59. Agrawal A, Gans J, Goldfarb A. *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Cambridge, MA: Harvard Business Press; 2022.
60. Obermeyer Z, Emanuel EJ. Predicting the future—big data, machine learning, and clinical medicine. *N Engl J Med*. 2016;375(13):1216.
61. Jin Q, Yuan Z, Xiong G, et al. Biomedical question answering: a survey of approaches and challenges. *ACM Comput Surv*. 2022;55(2):1–36.
62. Nori H, King N, McKinney SM, et al. Capabilities of GPT-4 on medical challenge problems. *arXiv preprint arXiv:2303.13375*. 2023 Mar 20.
63. Rajpurkar P, Chen E, Banerjee O, et al. AI in health and medicine. *Nat Med*. 2022;28(1):31–8.
64. Miotto R, Li L, Kidd BA, et al. Deep patient: an unsupervised representation to predict the future of patients from EHRs. *Sci Rep*. 2016;6(1):26094.
65. Wang Y, Yin C, Zhang P. Multimodal risk prediction with physiological signals, medical images and clinical notes. *Heliyon*. 2024;10(5).
66. Pinaya WH, Tudosiu PD, Dafflon J, et al. Brain imaging generation with latent diffusion models. In: *MICCAI Workshop on Deep Generative Models*. Cham: Springer; 2022. p. 117–26.
67. Holzinger A, Carrington A, Müller H. Measuring the quality of explanations: the system causability scale (SCS). *KI-Künstl Intell*. 2020;34(2):193–8.
68. Ji Z, Lee N, Frieske R, et al. Survey of hallucination in natural language generation. *ACM Comput Surv*. 2023;55(12):1–38.
69. Lewis P, Perez E, Piktus A, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Adv Neural Inf Process Syst*. 2020;33:9459–74.
70. Macrae C. Governing the safety of artificial intelligence in healthcare. *BMJ Qual Saf*. 2019;28(6):495–8.
71. Amann J, Blasimme A, Vayena E, et al. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Med Inform Decis Mak*. 2020;20(1):310.
72. Varoquaux G, Cheplygina V. Machine learning for medical imaging: methodological failures and recommendations. *NPJ Digit Med*. 2022;5(1):48.
73. Preiksaitis C, Rose C. Opportunities and challenges of generative artificial intelligence in medical education: a scoping review. *JMIR Med Educ*. 2023;9:e48785.
74. Ibrahim H, Liu X, Rivera SC, et al. Reporting guidelines for clinical trials of artificial intelligence interventions: the SPIRIT-AI and CONSORT-AI guidelines. *Trials*. 2021;22(1):11.
75. Nagendran M, Chen Y, Lovejoy CA, et al. Artificial intelligence versus clinicians: systematic review of design and reporting standards. *BMJ*. 2020;368.
76. Kaladharan S, Manayath D, Gopalakrishnan R. Regulatory challenges in AI/ML-enabled medical devices: a scoping review. *J Med Devices*. 2024;18(4):040801.
77. Pannonhalmi Á, Sipos B, Kurucz RI, et al. Advancing regulatory oversight of medical device trials to align with clinical drug standards in the European Union. *Pharmaceuticals*. 2025;18(6):876.
78. Yale AJ. *Privacy-preserving synthetic health data generation and evaluation*. Rensselaer Polytechnic Institute; 2020.
79. Stadler T, Oprisanu B, Troncoso C. Synthetic data—anonymisation groundhog day. In: *31st USENIX Security Symposium (USENIX Security 22)*; 2022. p. 1451–68.
80. Obermeyer Z, Powers B, Vogeli C, et al. Dissecting racial bias in an algorithm used to manage population health. *Science*. 2019;366(6464):447–53.
81. Seyyed-Kalantari L, Zhang H, McDermott MB, et al. Underdiagnosis bias of AI algorithms in chest radiographs. *Nat Med*. 2021;27(12):2176–82.
82. Price WN, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. *JAMA*. 2019;322(18):1765–6.
83. Mennella C, Maniscalco U, De Pietro G, et al. Ethical and regulatory challenges of AI technologies in healthcare: a narrative review. *Heliyon*. 2024;10(4).
84. Grote T, Berens P. On the ethics of algorithmic decision-making in healthcare. *J Med Ethics*. 2020;46(3):205–11.
85. Mittelstadt B. Principles alone cannot guarantee ethical AI. *Nat Mach Intell*. 2019;1(11):501–7.