International Journal of Research and Technology (IJRT)
International Open-Access, Peer-Reviewed, Refereed, Online Journal
ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529
| An ISO 9001:2015 Certified Journal |

Generative AI in Healthcare: A Systematic Review of Techniques, Clinical Applications, And Ethical Challenges

[1]Smriti Srivastava [2]Aishwarya
[1/2]Assistant Professor
[1] smriti.srivastava@ipec.org.in [2]Aishwarya@ipec.org.in

ABSTRACT
Generative Artificial Intelligence (AI) has emerged as a transformative paradigm in healthcare by enabling the creation of realistic synthetic data, enhancing clinical decision support, and accelerating biomedical research. Unlike traditional discriminative models, generative AI techniques—such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), diffusion models, and large language models (LLMs)—learn underlying data distributions to generate novel and meaningful outputs. This systematic review provides a comprehensive analysis of state-of-the-art generative AI techniques and their applications across key healthcare domains, including medical imaging, disease diagnosis, drug discovery, personalized treatment planning, electronic health record (EHR) synthesis, and clinical documentation. In addition, the review critically examines ethical, legal, and social challenges associated with generative AI in healthcare, such as data privacy, algorithmic bias, model interpretability, reliability, and regulatory compliance. By synthesizing findings from recent studies, this review highlights current research trends, identifies existing gaps, and discusses future research directions necessary for the responsible and effective integration of generative AI technologies into clinical practice. The insights presented aim to guide researchers, clinicians, and policymakers in leveraging generative AI to improve healthcare outcomes while ensuring ethical and trustworthy deployment.

Keywords:-Generative Artificial Intelligence; Healthcare Applications; Generative Adversarial Networks; Diffusion Models; Large Language Models; Medical Imaging; Drug Discovery; Clinical Decision Support; Ethical Challenges; Privacy and Bias

1.INTRODUCTION

The rapid advancement of Artificial Intelligence (AI) has significantly transformed healthcare systems by enabling data-driven decision-making, automation of clinical workflows, and improved diagnostic accuracy. Traditional AI approaches in healthcare have largely relied on discriminative models that focus on classification or prediction tasks. While effective, these methods are limited in their ability to generate new data or model complex underlying data distributions. In contrast, Generative Artificial Intelligence (AI) represents a powerful class of models capable of learning data representations and generating realistic, high-quality synthetic outputs, thereby opening new opportunities for innovation in healthcare. Generative AI encompasses a range of techniques, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), diffusion models, and large language models (LLMs)[1].

These models have demonstrated remarkable performance in generating medical images, synthesizing electronic health records, assisting in clinical documentation, and supporting drug discovery and molecular design. For instance, GAN-based models have been widely adopted for medical image augmentation and enhancement, improving diagnostic performance in data-scarce scenarios. Similarly, diffusion models have shown superior stability and image quality in medical imaging tasks, while LLMs have enabled natural language understanding and generation for clinical reporting and patient interaction. The adoption of generative AI in healthcare has been further accelerated by the growing availability of large-scale medical datasets and advances in computational resources[2]. Generative models offer promising solutions to persistent challenges such as limited labeled data, class imbalance, and privacy concerns by enabling the generation of realistic synthetic data. These capabilities are particularly valuable in sensitive healthcare environments, where data sharing is often restricted due to regulatory and ethical constraints. Despite its potential, the integration of generative AI into healthcare raises critical ethical, legal, and social concerns[3]. Issues related to patient data privacy, algorithmic bias, model transparency, reliability, and regulatory compliance remain significant barriers to clinical adoption. The use of synthetic data and AI-generated clinical content also introduces risks associated with misinformation, accountability, and patient safety. Therefore, a comprehensive understanding of both the technical advancements and ethical implications of generative AI is essential for its responsible deployment in real-world healthcare settings[4]. Several review studies have explored specific aspects of AI in healthcare; however, a systematic and unified review focusing exclusively on generative AI techniques, their clinical applications, and associated ethical challenges remains limited. This review aims to bridge this gap by systematically analysing existing literature on generative AI in healthcare. The primary contributions of this paper are threefold: (i) to provide a structured overview of generative AI techniques used in healthcare, (ii) to examine their applications across major clinical domains, and (iii) to critically assess ethical, legal, and societal challenges while highlighting future research directions.

2. REVIEW METHODOLOGY FOR GENERATIVE AI IN HEALTHCARE

Methodology Component	Description
Review Objective	This systematic review aims to comprehensively analyze existing research on generative artificial intelligence techniques, their clinical applications in healthcare, and the associated ethical, legal, and social challenges, following standard systematic review guidelines to ensure transparency, reproducibility, and comprehensive literature coverage[5].
Research Questions	The review addresses four key questions: (1) identification of generative AI techniques currently used in healthcare applications; (2) exploration of clinical domains where generative AI models are most effectively applied;

	(3) examination of ethical, legal, and social challenges arising from the use of generative AI in healthcare; and (4) identification of research gaps and future research directions in this domain[7].
Literature Search Strategy	A comprehensive literature search was conducted using IEEE Xplore, SpringerLink, Elsevier ScienceDirect, PubMed, and Google Scholar. Studies published between 2018 and 2024 were considered to capture recent advancements. Keyword combinations included "Generative AI," "Healthcare," "Medical Imaging," "GAN," "Diffusion Models," "Large Language Models," "Synthetic Medical Data," and "Ethical Challenges[8]."
Inclusion Criteria	Peer-reviewed journal articles and conference papers; studies focusing on generative AI techniques in healthcare; publications discussing clinical applications or ethical implications; and articles written in the English language[9].
Exclusion Criteria	Non-peer-reviewed articles, editorials, and opinion papers; studies unrelated to healthcare applications; papers focusing exclusively on discriminative AI models; and duplicate or incomplete studies[10].
Study Selection Process	The initial database search yielded a large number of publications. Duplicate records were removed, followed by title and abstract screening for relevance. Full-text reviews were conducted for shortlisted studies, and only those meeting all inclusion criteria were selected, ensuring methodological rigor and minimizing selection bias[11].
Data Extraction	From each selected study, data related to generative AI techniques, healthcare domain, dataset characteristics, evaluation metrics, reported outcomes, and discussed ethical or regulatory concerns were systematically extracted[12].
Data Analysis Method	The extracted data were analyzed using qualitative and thematic analysis to identify common trends, challenges, performance patterns, ethical issues, and research gaps across different healthcare applications [13].
Review Framework	Selected studies were categorized into three major themes: (i) Generative AI Techniques (GANs, VAEs, diffusion models, and large language models); (ii) Clinical Applications (medical imaging, disease diagnosis, drug discovery, electronic health record synthesis, and clinical decision support); and (iii) Ethical and Regulatory Challenges (privacy, bias, transparency, reliability, and compliance)[14].

3. GENERATIVE AI TECHNIQUES IN HEALTHCARE

3.1 GENERATIVE ADVERSARIAL NETWORKS (GANS)

Generative Adversarial Networks (GANs) consist of two competing neural networks: a generator and a discriminator. The generator creates synthetic data samples, while the discriminator attempts to distinguish between real and generated samples. Through this adversarial process, GANs learn to generate highly realistic outputs. In healthcare, GANs have been extensively used for medical image synthesis, image enhancement, data augmentation, and modality translation, such as converting MRI scans to CT images. GAN-based augmentation has been shown to improve diagnostic accuracy, particularly in scenarios involving rare diseases or limited datasets. Despite their effectiveness, GANs suffer from challenges such as training instability, mode collapse, and lack of interpretability. These limitations have motivated the exploration of alternative generative models with improved stability and reliability for clinical use [15].

3.2 Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs) are probabilistic generative models that encode input data into a latent space and reconstruct the data through a decoder network. Unlike GANs, VAEs provide a structured latent representation, making them suitable for tasks requiring data interpolation and uncertainty estimation. In healthcare, VAEs have been applied to medical image reconstruction, anomaly detection, and generation of synthetic patient records[16].

VAEs offer better training stability and interpretability compared to GANs; however, they often produce blurrier outputs, which may limit their effectiveness in high-resolution medical imaging tasks. Hybrid approaches combining VAEs and GANs have been proposed to address these shortcomings[17].

3.3 Diffusion Models

Diffusion models represent a recent advancement in generative AI, achieving state-of-the-art performance in image generation tasks. These models work by gradually adding noise to data and learning to reverse the process to generate high-quality samples. In healthcare, diffusion models have demonstrated superior performance in medical image synthesis, segmentation, and enhancement, particularly in producing high-resolution and anatomically consistent images[18].

The stability and controllability of diffusion models make them promising candidates for clinical applications. However, their high computational cost and longer inference times remain challenges for real-time healthcare systems.

3.4 Large Language Models (LLMs)

Large Language Models (LLMs), such as transformer-based architectures, have revolutionized natural language processing tasks in healthcare. These models are capable of generating clinical notes, summarizing electronic health records, assisting in medical coding, and supporting clinical decision-making. LLMs have also been used for patient interaction, medical question answering, and clinical trial matching. While LLMs offer significant benefits in automating clinical documentation and improving workflow efficiency, concerns related to hallucination, bias, and lack of explain ability pose risks in safety-critical healthcare environments. Ensuring reliable and ethical deployment of LLMs remains an active area of research[19].

3.5 Hybrid and Multimodal Generative Models

Recent research has focused on hybrid and multimodal generative models that integrate multiple data modalities, such as medical images, text, and genomic data. These models aim to provide a holistic understanding of patient health by combining diverse information sources. Multimodal generative AI systems have shown promise in personalized medicine, disease prognosis, and treatment planning. Despite their potential, multimodal models require large, well-annotated datasets and robust integration strategies, which remain challenging in real-world healthcare settings[20].

4. CLINICAL APPLICATIONS OF GENERATIVE AI IN HEALTHCARE

Generative AI has demonstrated significant potential across a wide range of healthcare applications by enhancing data availability, improving diagnostic accuracy, and supporting clinical decision-making. This section reviews the major clinical domains where generative AI techniques have been effectively applied.

4.1 Medical Imaging

Medical imaging is one of the most prominent application areas of generative AI. Generative models are widely used for image synthesis, augmentation, reconstruction, and modality translation in radiology and medical diagnostics. GANs and diffusion models have been applied to generate high-quality synthetic images for modalities such as MRI, CT, X-ray, and ultrasound. These synthetic datasets help address data scarcity and class imbalance, particularly for rare diseases. Generative AI has also been used for image enhancement tasks, including noise reduction, super-resolution, and image completion, leading to improved visualization and diagnostic accuracy. Additionally, cross-modality image translation, such as MRI-to-CT conversion, reduces the need for multiple imaging procedures, thereby minimizing patient exposure to radiation.

4.2 Disease Diagnosis and Prognosis

Generative AI models contribute to improved disease diagnosis and prognosis by learning complex patterns from patient data. Synthetic data generation enables the development of robust diagnostic models even when real-world data is limited. Generative models have been applied to detect diseases such as cancer, cardiovascular disorders, neurological conditions, and infectious diseases. In prognosis and risk prediction, generative models assist in simulating disease progression scenarios, enabling clinicians to evaluate potential outcomes and personalize treatment strategies. These applications are particularly valuable in chronic disease management and early intervention planning.

4.3 Drug Discovery and Molecular Design

Generative AI has revolutionized drug discovery by enabling the generation of novel molecular structures with desired properties. Models such as VAEs, GANs, and diffusion-based architectures are used to design drug candidates, predict molecular interactions, and optimize pharmacological properties. By reducing reliance on trial-and-error methods, generative AI accelerates the drug development pipeline and lowers associated costs. These techniques have been successfully applied in identifying potential treatments for complex diseases, including cancer and rare genetic disorders.

4.4 Electronic Health Record (EHR) Synthesis

The synthesis of electronic health records using generative AI addresses challenges related to data privacy, interoperability, and data sharing. Synthetic EHRs preserve statistical properties of real patient data while minimizing the risk of patient re-identification. These datasets are valuable for research, training, and system testing without compromising sensitive patient information. Generative models have been applied to generate structured clinical data, time-series patient records, and unstructured clinical notes, supporting data-driven healthcare research and development.

4.5 Clinical Documentation and Decision Support

Large language models have significantly improved clinical documentation by automating tasks such as report generation, discharge summaries, and medical coding. These tools reduce administrative burden on healthcare professionals, allowing more time for patient care. Generative AI-based decision support systems assist clinicians by summarizing patient histories, suggesting treatment options, and providing evidence-based recommendations. While promising, such systems must be carefully validated to ensure accuracy, reliability, and patient safety.

4.6 Personalized Medicine

Generative AI enables personalized medicine by integrating patient-specific data, including medical history, imaging, and genomic information. Multimodal generative models support individualized treatment planning and therapy optimization. These approaches have shown promise in oncology, precision diagnostics, and targeted therapies.

5. Ethical, Legal, and Social Challenges

Despite the significant potential of generative AI in healthcare, its adoption raises critical ethical, legal, and social concerns that must be addressed to ensure safe, fair, and responsible use. This section discusses the major challenges associated with the deployment of generative AI systems in healthcare environments.

5.1 Data Privacy and Security

Healthcare data is highly sensitive, and the use of generative AI models requires access to large volumes of patient information. Although synthetic data generation is often promoted as a privacy-preserving solution, generative models may still unintentionally memorize and reproduce sensitive patient details, leading to potential privacy breaches. Ensuring compliance with data protection regulations such as HIPAA and GDPR remains a major challenge, particularly when models are trained on multi-institutional datasets.

5.2 Bias and Fairness

Generative AI models learn from historical data, which may contain inherent biases related to gender, ethnicity, age, or socioeconomic status. If not carefully addressed, these biases can be amplified in generated data and downstream clinical applications, resulting in unfair or inaccurate medical outcomes. Bias in generative models poses serious risks in healthcare decision-making, where fairness and equity are critical.

5.3 Explain ability and Transparency

Many generative AI models operate as black-box systems, making it difficult for clinicians to understand how outputs are generated. Lack of explain ability reduces trust and limits clinical adoption, especially in safety-critical scenarios. Transparent and interpretable generative AI systems are essential for enabling clinicians to validate and justify AI-assisted decisions

5.4 Reliability and Safety

Generative AI models, particularly large language models, may generate incorrect, misleading, or fabricated information, a phenomenon often referred to as hallucination. In healthcare, such errors can have serious consequences for patient safety. Ensuring robustness, reliability, and rigorous validation of generative AI outputs is necessary before deployment in clinical settings.

5.5 Regulatory and Legal Challenges

The regulatory landscape for generative AI in healthcare is still evolving. Existing medical device regulations may not adequately address the dynamic and adaptive nature of generative models. Questions related to accountability, liability, and certification remain unresolved, especially when AI-generated outputs influence clinical decisions. Clear regulatory frameworks are needed to govern the development and deployment of generative AI systems.

5.6 Ethical Use and Human Oversight

The integration of generative AI into healthcare must prioritize ethical principles, including patient autonomy, beneficence, and non-maleficence. Human oversight is essential to ensure that AI-generated outputs support, rather than replace, clinical judgment. Establishing guidelines for responsible AI use and continuous monitoring is critical for maintaining ethical standards.

6. FUTURE RESEARCH DIRECTIONS

Although generative AI has shown considerable promise in transforming healthcare, several research challenges and opportunities remain. Future research should focus on developing privacy-preserving generative models that minimize the risk of sensitive data leakage while maintaining high data utility. Techniques such as federated learning, differential privacy, and secure multi-party computation can play a critical role in enabling collaborative model training across institutions without compromising patient confidentiality. Another important direction involves improving the interpretability and explain ability of generative AI models. Clinicians require transparent systems that provide understandable justifications for AI-generated outputs. Research efforts should aim to integrate explainable AI (XAI) techniques into generative models to enhance trust and facilitate clinical adoption. The robustness and reliability of generative AI systems must also be enhanced to reduce errors, hallucinations, and unpredictable behaviour, particularly in large language models. Rigorous validation protocols, uncertainty quantification, and continuous monitoring are essential for ensuring patient safety in real-world deployments. Additionally, the development of multimodal generative AI systems that integrate medical images, clinical text, genomic data, and wearable sensor data represents a promising area for personalized medicine. Such models can provide a holistic view of patient health and support individualized diagnosis and treatment planning. From a regulatory perspective, future research should explore standardized evaluation benchmarks and ethical frameworks tailored to generative AI in healthcare. Collaboration among researchers, clinicians, policymakers, and industry stakeholders will be critical for establishing responsible guidelines and accelerating the safe integration of generative AI technologies into healthcare systems.

7. CONCLUSION

Generative Artificial Intelligence has emerged as a powerful and transformative technology with the potential to significantly enhance healthcare delivery, research, and clinical decision-making. This systematic review has provided a comprehensive overview of generative AI techniques, including GANs, VAEs, diffusion models, and large language models, and examined their applications across key healthcare domains such as medical imaging, disease diagnosis, drug discovery, electronic health record synthesis, and clinical documentation. While generative AI offers substantial benefits in addressing data scarcity, improving diagnostic accuracy, and enabling personalized medicine, its adoption is accompanied by critical ethical, legal, and social challenges. Issues related to data privacy, bias, explain ability, reliability, and regulatory compliance must be carefully addressed to ensure safe and equitable use.

Future advancements in generative AI should prioritize responsible design, robust validation, and human-centered deployment strategies. By addressing current limitations and fostering interdisciplinary collaboration, generative AI can be effectively leveraged to improve healthcare outcomes while maintaining ethical integrity and patient trust.

References

1. Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al., “Generative Adversarial Networks,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
2. Kingma, D. P., and Welling, M., “Auto-Encoding Variational Bayes,” *International Conference on Learning Representations (ICLR)*, 2014.
3. Ho, J., Jain, A., and Abbeel, P., “Denoising Diffusion Probabilistic Models,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
4. Esteva, A., Robicquet, A., Ramsundar, B., et al., “A Guide to Deep Learning in Healthcare,” *Nature Medicine*, vol. 25, no. 1, pp. 24–29, 2019.
5. Topol, E. J., “High-performance medicine: the convergence of human and artificial intelligence,” *Nature Medicine*, vol. 25, pp. 44–56, 2019.
6. Shorten, C., and Khoshgoftaar, T. M., “A survey on Image Data Augmentation for Deep Learning,” *Journal of Big Data*, vol. 6, no. 60, 2019.
7. Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., and Greenspan, H., “GAN-based synthetic medical image augmentation for increased CNN performance,” *Neurocomputing*, vol. 321, pp. 321–331, 2018.
8. Chen, R. J., Lu, M. Y., Chen, T. Y., Williamson, D. F. K., and Mahmood, F., “Synthetic data in machine learning for medicine and healthcare,” *Nature Biomedical Engineering*, vol. 5, pp. 493–497, 2021.

9. Miotto, R., Wang, F., Wang, S., Jiang, X., and Dudley, J. T., "Deep learning for healthcare: review, opportunities and challenges," *Briefings in Bioinformatics*, vol. 19, no. 6, pp. 1236–1246, 2018.
10. Rajpurkar, P., Chen, E., Banerjee, O., and Topol, E. J., "AI in health and medicine," *Nature Medicine*, vol. 28, pp. 31–38, 2022.
11. Beam, A. L., and Kohane, I. S., "Big data and machine learning in health care," *JAMA*, vol. 319, no. 13, pp. 1317–1318, 2018.
12. Chen, T., Kornblith, S., Norouzi, M., and Hinton, G., "A Simple Framework for Contrastive Learning of Visual Representations," *ICML*, 2020.
13. Dosovitskiy, A., et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *ICLR*, 2021.
14. Lehman, E. P., et al., "Does BERT Pretrained on Clinical Notes Reveal Sensitive Data?" *Proceedings of the ACM Conference on Health, Inference, and Learning*, 2021.
15. Shickel, B., Tighe, P. J., Bihorac, A., and Rashidi, P., "Deep EHR: A Survey of Recent Advances in Deep Learning Techniques for Electronic Health Record Analysis," *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 5, pp. 1589–1604, 2018.
16. Yu, K. H., Beam, A. L., and Kohane, I. S., "Artificial intelligence in healthcare," *Nature Biomedical Engineering*, vol. 2, pp. 719–731, 2018.
17. Van der Schaar, M., et al., "How artificial intelligence and machine learning can help healthcare systems respond to COVID-19," *Machine Learning*, vol. 110, pp. 1–14, 2021.
18. Wachter, S., Mittelstadt, B., and Floridi, L., "Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation," *International Data Privacy Law*, vol. 7, no. 2, pp. 76–99, 2017.
19. London, A. J., "Artificial intelligence and black-box medical decisions: accuracy versus explainability," *Hastings Center Report*, vol. 49, no. 1, pp. 15–21, 2019.
20. Amodei, D., et al., "Concrete Problems in AI Safety," *arXiv preprint*, arXiv:1606.06565.