

Applications of Generative Artificial Intelligence in Medical Imaging: A Review of Opportunities and Risks

Zhouyu Chen*

The First Clinical College, Wenzhou Medical University, Zhejiang 325000, China

**Corresponding author: Zhouyu Chen.*

Abstract

Generative Artificial Intelligence (GAI), with its core capabilities in image synthesis, data augmentation, and cross-modal transformation, is profoundly reshaping the entire workflow of medical imaging, including acquisition, diagnosis, treatment planning, and education and training. This paper systematically reviews the technical pathways and clinical application progress of generative AI in the field of medical imaging. It focuses on analyzing its core value in improving imaging quality, alleviating data scarcity, assisting precision diagnosis and treatment, and empowering primary healthcare. At the same time, it deeply explores potential risks such as “medical hallucination” misleading, black-box effects, data bias, and lagging compliance regulation. On this basis, the paper proposes development strategies emphasizing compliance, standardization, and strong interpretability, aiming to provide theoretical references for the safe implementation and large-scale promotion of generative AI in the field of intelligent medical imaging.

Keywords

generative artificial intelligence, medical imaging, large language models, medical hallucination, smart healthcare

1. Introduction

Medical imaging serves as a core pillar of the modern healthcare system, encompassing radiology, ultrasound, nuclear medicine, interventional therapy, and other departments. It provides critical evidence for clinical diagnosis, treatment efficacy evaluation, and prognosis assessment. With the intensification of population aging and the high prevalence of chronic diseases, the annual growth rate of medical imaging examinations exceeds 10%. Traditional imaging diagnosis faces multiple bottlenecks, including excessive reading workload, insufficient capabilities of primary care physicians, scarcity of data for rare diseases, and the difficulty in balancing radiation dose with imaging quality. Traditional discriminative artificial intelligence (such as CNN networks) mainly focuses on lesion detection and segmentation, making it difficult to address root-cause issues such as data scarcity and missing modalities [1].

With the development of Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and the recent emergence of Diffusion Models and medical multimodal large models (Med-VLM),

generative AI can learn distributional patterns from existing data and generate entirely new, high-fidelity medical images and textual data. This achieves a technological breakthrough characterized by “from nothing to something, from low to high, and from single to multiple” [2, 3].

Based on clinical needs and technological evolution, this paper systematically analyzes the core opportunities and potential risks of generative AI in medical imaging and proposes targeted development strategies, providing theoretical support for the safe implementation and large-scale promotion of generative AI in the field of medical imaging.

2. Core Opportunities of Generative Artificial Intelligence in Medical Imaging

2.1 Improving Imaging Quality and Reducing Radiation Risk

Generative AI resolves the traditional contradiction in imaging of “low dose = low quality,” achieving a win-win outcome for both safety and precision. In low-dose CT (LDCT) examinations, algorithms based on Denoising Diffusion Probabilistic Models (DDPM) can effectively remove artifacts and restore subtle lesion structures while reducing the radiation dose to one-third of conventional examinations, significantly benefiting sensitive populations such as children and pregnant women [4]. Generative AI systems for DSA angiography can reduce intraoperative radiation by 66%, ensuring clear vascular contours while substantially decreasing radiation exposure for both physicians and patients. In addition, in MRI accelerated imaging, generative models not only shorten scanning time by 70%–90%, thereby alleviating equipment shortages and long patient appointment cycles, but also produce images with superior signal-to-noise ratio and resolution compared to traditional accelerated sequences. This significantly improves patient experience and greatly enhances the turnover efficiency of medical resources [5].

Table 1: Comparison of Image Quality Assessment in Low-Dose CT (LDCT) Reconstruction Between Diffusion Models and Traditional Algorithms [16, 17]

Algorithm Model Name	Technical Route Type	Radiation Dose	PSNR (dB) ↑	SSIM ↑	Core Clinical Performance and Advantages/Disadvantages
FBP (Filtered Back Projection)	Traditional physical reconstruction algorithm	10%	23.45 dB	0.612	Extremely fast reconstruction speed, but images exhibit obvious radial artifacts and high noise levels, making it difficult to meet diagnostic needs for subtle lesions.
RED-CNN	Discriminative deep learning model	10%	27.83 dB	0.795	Significantly reduces noise, but suffers from over-smoothing, easily causing blurred tissue boundaries and loss of detailed information such as ground-glass nodules.
Vanilla DDPM	Basic diffusion generative model	10%	29.50 dB	0.842	High texture restoration capability and good visual quality, but inference relies on multi-step iteration, resulting in longer processing time per image.
RE-EDPM / PTDM	Improved diffusion generative model	10%	31.60 dB	0.879	Effectively suppresses noise while maintaining high structural consistency, better preserves subtle anatomical structures, and inference speed is optimized to approximately 0.25 seconds per slice, reaching clinical application level.

2.2 Alleviating Data Scarcity and Privacy Compliance Pressure

Medical data privacy and compliance constraints hinder AI research and development. Generative AI effectively circumvents these restrictions through synthetic data. Synthetic CT, MRI, and ultrasound data can accurately preserve lesion characteristics and anatomical structures without involving personal identity information, thereby complying with regulations such as HIPAA and the Personal Information Protection Law. This enables safe use for model training and algorithm validation. In small-sample scenarios such as rare diseases and pediatric tumors, generative AI can effectively expand datasets, increasing the sensitivity of lesion detection models from 82% to 96% [6, 7]. Furthermore, when combined with Federated Learning, multi-center institutions can utilize generative models to create joint datasets that incorporate the characteristics of each center’s data without sharing raw data, thereby promoting the development of inclusive medical imaging AI [8].

2.3 Assisting Precision Diagnosis and Treatment and Enhancing Efficiency

The synergy between generative AI and discriminative models has constructed a full-process intelligent system of “screening–localization–characterization–reporting.” The chest CT “one-scan multi-task” intelligent agent can automatically complete detection and segmentation of pulmonary nodules, rib fractures, pneumonia, and mediastinal lesions, outputting structured reports within seconds. This reduces physician review time from 30 minutes to less than 5 minutes, greatly liberating clinical productivity [9]. Multimodal GPT for thyroid ultrasound integrates image interpretation, report generation, and risk stratification, reducing unnecessary biopsy rates by 40% and effectively decreasing overdiagnosis and overtreatment. Primary healthcare institutions, with the help of systems such as “Yueyi Zhiying,” can obtain expert-level auxiliary diagnostic capabilities, increasing the detection rate of abnormalities in chest X-rays and CT scans by more than 25% and significantly narrowing the urban-rural healthcare gap. Text-to-image generation models can also convert radiology reports into three-dimensional visualized images, helping non-radiology physicians and patients intuitively understand the condition and improving doctor-patient communication efficiency [10].

Table 2: Multi-Dimensional Performance Comparison of Multimodal Large Models and Traditional AI in Chest X-ray Report Generation Tasks [18, 19]

Model Architecture Category	Representative Technical Solution	F1-CheXbert (Clinical Accuracy) ↑	ROUGE-L (Text Fluency) ↑	Severe Clinical Error Rate (Missed Diagnosis/Hallucination) ↓	Average Physician Review Modification Time ↓
Traditional Baseline Model	CNN (image extraction) + LSTM (text decoding)	0.485	0.254	High (>35%)	~4.5 minutes
Vision-Language Model	Transformer cross-modal attention large model	0.572	0.281	~25%	~3.0 minutes
General Multimodal Large Model	GPT-4V / Gemini Pro	0.615	0.315	~18% (prone to “plausible but fabricated” hallucinations)	~2.5 minutes
Medical-Specific Generative Large Model	FactMM-RAG (with retrieval-augmented generation)	0.689	0.342	Extremely low (<10%) (severe errors reduced by 29%)	~1.5 minutes

2.4 Empowering Precision Treatment Planning

Generative AI demonstrates tremendous clinical value in tumor radiotherapy and surgical planning, providing individualized, visual, and quantitative support for tumor treatment. Cross-modal generation technology can convert MRI into pseudo-CT (pCT), directly generating pseudo-CT images containing electron density information from high soft-tissue-contrast MRI. This eliminates the need for additional CT localization scans for patients, reducing radiation exposure and eliminating spatial errors caused by image registration [11]. 3D medical diffusion models can generate patient-specific organ and tumor models, enabling precise delineation of target volumes and organs at risk, improving the uniformity of radiotherapy dose distribution, and reducing the risk of damage to normal tissues. In orthopedics, neurosurgery, and cardiovascular interventional procedures, generative 3D reconstruction models based on CT/MRI can be used for puncture path planning and vascular intervention simulation, significantly improving the accuracy of preoperative path planning. This has increased the success rate of complex surgeries by 15% and reduced the complication rate by 10% [12].

2.5 Revolutionizing Medical Education and Skills Training

By leveraging generative AI, standardized, repeatable, and highly controllable imaging teaching libraries have been constructed. These can generate simulated images and cases of varying age groups, disease types, and difficulty levels, supporting standardized training and assessment of resident physicians. The models can dynamically adjust lesion size, density, and location to provide targeted training for physicians’ ability to identify and differentially diagnose subtle lesions. Virtual interventional surgery systems can generate real-

time vascular images and anatomical structures, simulating surgical operations and complication management, effectively shortening physicians' learning curves and reducing clinical risks.

3. Potential Risks of Generative Artificial Intelligence in Medical Imaging

3.1 “Medical Hallucination” Misleading and Clinical Safety Hazards

Generative models may produce non-existent “false lesions” or cause anatomical structure distortions, known as the “hallucination” phenomenon [13]. Due to the highly deceptive visual appearance of the generated images, the absence of strict clinical review mechanisms may lead to misdiagnosis, missed diagnosis, or unnecessary therapeutic interventions. In addition, generated images lack a clear “gold standard,” making it difficult to quantitatively assess their authenticity and clinical validity, which further increases uncertainty in clinical decision-making.

3.2 Lack of Interpretability and Trust Crisis

The decision-making process of most generative models remains in a “black-box” state, making it difficult to clearly explain the basis for generation. Physicians cannot verify the rationality of key decisions, which undermines clinical trust and complicates responsibility attribution. When model outputs contradict physicians' experience, clinicians are unable to trace the logical reasoning process behind the model's generation. This lack of interpretability severely hinders the establishment of trust in AI systems, particularly in critical decision-making scenarios involving severe cases.

3.3 Data Bias and Healthcare Equity Issues

The output of algorithms is highly dependent on the distribution of training data. If the training dataset lacks coverage of remote areas, ethnic minorities, or patients with rare diseases, the model may perform poorly in these populations, thereby exacerbating inequalities in healthcare resource allocation [14]. Data variations caused by different equipment manufacturers and scanning parameters can also reduce the model's generalization ability across institutions and devices, significantly diminishing the effectiveness of AI applications in certain healthcare facilities.

3.4 Lagging Compliance Regulation and Ambiguous Responsibility Attribution

Currently, no systematic approval and regulatory framework for generative medical imaging large models has been established globally. There exists a clear regulatory vacuum regarding the legal status of synthetic data and the attribution of diagnostic responsibility. When generated images are used for clinical decision-making, they may face risks of medical disputes and quality control issues. Furthermore, controversies over healthcare equity arising from inherent data bias in generative models, as well as potential risks of “distortion” or “fabrication” in medical documentation, pose severe challenges to existing medical ethics standards and regulatory systems.

3.5 Challenges in Technology Implementation and Workflow Adaptation

Generative large models have extremely high demands on computing resources, making deployment difficult for primary healthcare institutions. In scenarios requiring real-time imaging or intraoperative guidance with stringent low-latency requirements, the inference speed of existing models remains insufficient. Moreover, most generative AI systems operate independently and lack deep integration with existing PACS, RIS, and electronic medical record systems, easily forming “information silos.” This affects clinical workflow efficiency and constrains the large-scale promotion of the technology.

4. Development Strategies and Future Prospects

4.1 Establishing Compliance Regulation and Approval Systems

Clinical validation standards, quality control specifications, and registration pathways for medical imaging generative models should be formulated, with clear boundaries and legal status defined for the use

of synthetic data. Third-party evaluation institutions for fidelity and safety should be established to promote compliant market entry of products. At the same time, active participation in international standard-setting is encouraged, drawing on advanced experiences to build a regulatory framework that aligns with China's local medical ecosystem.

4.2 Enhancing Model Interpretability and Medical Reliability

To address the “black-box” effect and “hallucination” risks, attention mechanisms, causal inference networks, and anatomical topology constraints should be introduced to map unexplainable latent space features into visualizable generation bases. In addition, an automated independent image quality evaluation system should be developed to automatically detect artifacts and false structures, achieving a closed-loop of “generation–quality control–early warning” to reduce the risk of “hallucination” misleading.

4.3 Advancing Standardized Data Ecosystems

At the national level, standards for medical image acquisition, annotation, and formatting should be established. Through the synergy of federated learning and synthetic data, large-scale, high-quality medical image datasets covering multiple centers, devices, and populations should be constructed. This will enhance model generalization ability while protecting privacy. Algorithmic bias should be eliminated to ensure the generalization performance of models in primary healthcare and other settings, thereby safeguarding healthcare equity.

4.4 Clarifying Ethics and Responsibility Division

The responsibilities of physicians, developers, and healthcare institutions should be clearly defined, adhering to the principle of “human-machine collaboration with physicians making final decisions.” Ethical review should be strengthened to ensure data privacy and healthcare equity. A traceability mechanism should be established to ensure that every step of the generation process is auditable, providing technical support for responsibility attribution in medical disputes [15].

4.5 Technology Implementation and Workflow Adaptation

With the increasing computational demands of generative AI models, primary hospitals need to rely on GPU-accelerated hardware or cloud computing platforms to ensure efficient image generation and real-time inference. It is essential to promote deep integration of generative AI systems with existing PACS, RIS, and electronic medical record (EMR) systems to ensure smooth flow of medical imaging data and prevent the formation of information silos.

5. Conclusion

Overall, generative artificial intelligence, with its core capabilities in image generation, quality optimization, cross-modal fusion, and data synthesis, provides a systematic solution for the digital-intelligent advancement of medical imaging. It demonstrates tremendous value in improving diagnostic efficiency, reducing radiation risk, overcoming data bottlenecks, and empowering primary healthcare. Current technologies have transitioned from laboratory research to clinical application and have been validated in fields such as radiology, oncology, orthopedics, and interventional medicine. However, challenges remain, including “medical hallucination,” deficiencies in system integration, and lagging regulatory and ethical frameworks.

Looking toward the future, a strategic approach guided by clinical needs, prioritizing compliance, and emphasizing human-machine collaboration should be adopted to promote the deep integration of generative AI with medical imaging. This will help build a safe, efficient, inclusive, and precise intelligent medical imaging system, contributing to the realization of the Healthy China 2030 strategy. As technology matures and regulations improve, generative AI is expected to become the core engine driving the digital-intelligent transformation of medical imaging, reshaping healthcare service models and benefiting hundreds of millions of patients.

References

- [1] Zhang, R. R., Pu, L. H., Zhang, W., Qi, X. Y., & Yang, F. (2019). Current status and prospects of artificial intelligence applications in medical imaging education. *Modern Preventive Medicine*, 46(24), 4527-4609.
- [2] Tang, B. (2024). Generative artificial intelligence empowering new quality productive forces: A dual deconstruction based on Marxist productivity theory and new development concepts. *Academic Forum*, (12), 68-81.
- [3] Peng, C. (2025). Research on generative artificial intelligence in visual art creation [Master's thesis, Donghua University]. <https://doi.org/10.27012/d.cnki.gdhhu.2025.000326>.
- [4] Song, Y., Shen, L., Xing, L., & Ermon, S. (2022). Solving inverse problems in medical imaging with score-based generative models. *IEEE Transactions on Medical Imaging*, 41(4), 818-828.
- [5] Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R., & Rosen, M. S. (2018). Image reconstruction by domain-transform manifold learning. *Nature*, 555(7694): 60-65.
- [6] Sandfort, V., Yan, K., Pickhardt, P. J., et al. (2019). Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Scientific Reports*, 9(1), Article 16884.
- [7] Yue, Y. Y. (2022). Research on medical image enhancement algorithms based on generative adversarial networks [Master's thesis, Nanjing University of Posts and Telecommunications].
- [8] Rieke, N., Hancox, J., Li, W., et al. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3(1), 119.
- [9] Zhang, H. Y., Wang, T. B., Li, M. Z., Zhao, Z., Pu, S. L., & Wu, F. (2022). A review of visual-language multimodal pre-training. *Journal of Image and Graphics*, 27(9), 2652-2682. <https://doi.org/10.11834/jig.220173>.
- [10] Tu, T., Azizi, S., Driess, D., et al. (2024). Towards generalist biomedical AI. *NEJM AI*, 1(3), Article AIoa2300138.
- [11] Maspero, M., Savenije, M. H. F., Dinkla, A. M., et al. (2018). Dose evaluation of fast synthetic-CT generation using a generative adversarial network for general pelvis MR-only radiotherapy. *Physics in Medicine & Biology*, 63(18), 18500.
- [12] Che, L., Zhang, Z. Q., Zhou, J. J., & Li, L. (2024). Current research status and development trends of generative artificial intelligence. *Science and Technology Review*, 42(12), 35-43. <https://doi.org/10.3981/j.issn.1000-7857.2023.03.00366>
- [13] Cohen, J. P., Luck, M., & Honari, S. (2018). Distribution matching losses can hallucinate features in medical image translation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (pp. 529-536). Springer.
- [14] Seyyed-Kalantari, L., Zhang, H., McDermott, M. B. A., et al. (2021). Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine*, 27(12), 2176-2182.
- [15] Wang, S., Liu, T. Y., Wang, C., & Liu, Y. Y. (2024). Exploring the capabilities and risk delineation of generative AI in medical applications. *Medicine & Philosophy*, 45(12), 1-5.
- [16] Li, Q., et al. (2025). Low-dose CT imaging using a regularization-enhanced efficient diffusion probabilistic model. *Biomedical Signal Processing and Control*, 105, 107580.
- [17] Dalla Serra, F., et al. (2025). Multimodal generation of radiology reports using knowledge-grounded extraction of entities and relations. In *Proceedings of the Association for Computational Linguistics (ACL)* (pp. 615-624).
- [18] Alharbi, A. O. (2026). Automated radiology report generation using multimodal foundation models: A systematic review of clinical accuracy and safety. *The Review of Diabetic Studies*, 12(1).

- [19] Chen, S. (2025). A review of multimodal large model based medical image report generation. *Journal of Medical Artificial Intelligence*, 9(2), 248-255.

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Open Access

This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

