Proceedings of CONF-SEML 2026 Symposium: Computational Analysis and Modeling in Complex Intelligent Systems
DOI: 10.54254/2755-2721/2026.GU34155

# *Application of Generative Adversarial Networks in Medical Image Synthesis, Reconstruction and Enhancement*

**Yixiang Gao**

*Manchester City United College, Hubei University, Wuhan, China*
*2229143368@qq.com*

***Abstract.*** Generative adversarial networks (GANs) leverage adversarial training between a generator and a discriminator to effectively model the true distribution of medical images, addressing challenges like limited labeled data, high image noise, and missing multimodal information. As a result, GANs can be applied to tasks including image synthesis, denoising, reconstruction, super-resolution, and cross-modal enhancement. This study reviews recent advances in GANs for medical image synthesis, enhancement, and restoration, highlighting their core mechanisms, typical architectures, and primary applications. In addition, it also discusses common evaluation methods and clinical relevance, analyzes the main technical and practical challenges, and explores potential strategies to facilitate clinical translation of GAN-based approaches. The results indicate that GANs enhance image quality by producing high-fidelity images, reducing noise, aiding reconstruction, and enabling super-resolution, while simultaneously expanding limited training datasets via synthetic image generation.

***Keywords:*** Generative Adversarial Networks, Medical Images, Image Synthesis, Image Inpainting, Image Enhancement

## 1. Introduction

In modern healthcare, medical imaging supports the workflow from disease detection to outcome prediction. In particular, multimodal images like computed tomography (CT), magnetic resonance imaging (MRI), ultrasound, and pathology slides provide noninvasive anatomical and pathological information for clinical decisions. Nevertheless, the translation of precise diagnosis and treatment into clinical practice remains hindered by several challenges: the scarcity of high-quality annotated data, inconsistent imaging quality, and the tendency of conventional image processing approaches to discard critical lesion-related information. Deep generative models such as generative adversarial networks (GANs) offer promising solutions to these issues, owing to their capacity for high-fidelity sample generation, strong data distribution fitting, and flexible semantic manipulation. Since their introduction in 2014, GANs have achieved remarkable progress in image synthesis, enhancement and restoration, cross-modality translation, and lesion segmentation, with partial clinical translation already realized [1]. However, several critical limitations persist: the medical realism of synthesized samples and cross-center generalization capability of GAN-based models remain insufficient; the black-box nature of such models compromises clinical interpretability; the regulatory boundaries for

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.GU34155

data privacy compliance and ethical norms are still ambiguous; and unified standards are lacking for training stability, validation, and quantitative evaluation. This paper examines the applications of GANs in medical imaging, addressing performance, clinical deployment challenges, and technical optimization. By reviewing mechanisms, applications, and limits, and forecasting trends, it maps the GAN framework and its variants, categorizes key applications such as image synthesis, restoration, and enhancement, and highlights research directions. Moreover, it identifies priorities for improving medical realism, model interpretability, privacy protection, and clinical validation, and recommends lightweight models, multi-center training, and use of medical knowledge, supporting standardized translation and improving precision in medical imaging.

## 2. Fundamentals of generative adversarial networks (GANs)

### 2.1. Basic principles and training mechanisms of GAN

Essentially, a GAN consists of a generator $G$ and a discriminator $D$. By leveraging the principle of two-player zero-sum games, it allows both components to enhance each other via adversarial training [2]. In particular, the generator transforms a random noise vector $z$, typically drawn from a Gaussian or uniform distribution, into a synthetic sample $G(z)$ that shares the dimensionality of real data, thus approximating the true data distribution. Meanwhile, the discriminator acts as a binary classifier that outputs D(x) to ascertain whether an input $x$ originates from the real dataset or is generated by $G(z)$. Accordingly, the optimization objective of GAN can be formulated as a minimax problem, with its value function $V(G,D)$ expressed as:

$$\min_G \max_D V(D,G) = E_{x\sim P_{data}}(x)[\log D(x)] + E_{z\sim P_z}(z)[\log(1-D(G(z)))] \tag{1}$$

Where $p_{data}(x)$ denotes the distribution of real data, $p_z(z)$ represents the prior distribution of the input noise, and *E* refers to mathematical expectation. During the training process, the generator and discriminator progressively enhance their respective abilities through adversarial optimization. The discriminator progressively enhances its capability to differentiate between authentic and synthetic samples, whereas the generator iteratively fine-tunes its parameters to produce outputs of increasing realism, thus challenging the discriminator. As both networks undergo co-evolution, the generated data distribution approaches that of the real data. Eventually, the discriminator becomes unable to reliably distinguish real from generated samples, with its output probability converging to 0.5. At this point, the model reaches a Nash equilibrium, indicating the completion of training.

### 2.2. Conditional GAN and semantic control

The original GAN generates samples solely from random noise, resulting in uncontrollable content and features, and consequently failing to satisfy directed constraints [3]. To address this issue, Mirza and Osindero proposed the Conditional Generative Adversarial Network (cGAN) in 2014, which first introduced a condition vector into both the generator and discriminator, hence enabling semantic control over the generated content. The key enhancement of cGAN lies in appending a condition vector $y$ to the inputs of both networks. This vector may represent class labels, segmentation masks, reference images, or attribute parameters. The generator processes random noise $z$ while receiving the condition vector $y$, thus generating samples that conform to the semantic features specified by the condition. In contrast, the discriminator judges whether a sample conforms

DOI: 10.54254/2755-2721/2026.GU34155

to the predefined constraints by integrating $y$. Accordingly, the cGAN value function can be formalized as follows:

$$\min_G \max_D V(D,G) = E_{x\sim P_{data}}(x)[\log D(x \mid y)] + E_{z\sim P_z}(z)[\log(1 - D(G(z \mid y)))] \tag{2}$$

Under this framework, the generation process is transformed from unconstrained generation to conditional, guided synthesis, enabling precise semantic control over the outputs. Based on this idea, subsequent models have further enhanced controllability. For instance, Pix2Pix achieves accurate image-to-image translation via paired datasets and pixel-level guidance, while CycleGAN introduces a cycle-consistency loss, thus allowing cross-domain translation even with unpaired data. Moreover, StyleGAN integrates a style control module to enable hierarchical modulation of both global features and local details, affording fine-grained control over texture, contrast, and structural characteristics. These advancements ensure that the generation process not only adheres to conditional constraints but also attains higher fidelity in feature representation and fine-detail synthesis.

## 3. Applications of GAN in medical image data generation and quality optimization

### 3.1. Medical image modal generation and lesion simulation

Through the generation of high-quality images that preserve the features of real scans, medical image synthesis constitutes a core application of GANs in medical imaging, mitigating challenges related to data scarcity, missing modalities, and limited data sharing [4]. In particular, its primary applications encompass cross-modality image synthesis, augmentation and anonymized generation of limited datasets, as well as virtual lesion and pathological feature synthesis. Importantly, cross-modality image synthesis holds the highest clinical value as multimodal medical images offer complementary information. For example, CT clearly visualizes bony structures, MRI excels at distinguishing soft tissues and tumors, and PET reflects metabolic activity. However, some patients cannot receive MRI, and contrast CT or PET carry radiation risks. In addition, some imaging modalities are costly and time-intensive, contributing to widespread modality incompleteness in clinical practice. To address this problem, GANs can synthesize missing modalities from available images without additional examinations. Specifically, CycleGAN and Pix2Pix enable bidirectional translation between CT and MRI, T1-weighted MRI can be used to generate T2-weighted or FLAIR images, and low-dose CT can be synthesized to approximate PET-like images. Consequently, studies have indicated that such synthesized images improve diagnostic accuracy and segmentation performance for brain tumors, prostate cancer, and liver cancer [5, 6].

Beyond cross-modality synthesis, GANs can produce large-scale anonymized medical images, relieving constraints caused by data scarcity and privacy regulations. Due to the high cost of manual annotation, limited samples for rare diseases, and legal restrictions on data sharing, high-quality GAN-generated images have been widely used for model training and rare disease data augmentation. To date, GANs have been applied to pulmonary CT nodules, fundus retinal images, mammograms, and whole-slide pathological images (WSI), effectively mitigating model overfitting in small-sample scenarios [7]. In this context, GANs enable precise regulation of lesion location, size, morphology, and malignancy,generating images that faithfully reproduce realistic disease progression. As a result, these capabilities support the creation of standardized datasets and facilitate model evaluation on rare or extreme cases. This enables the construction of standardized training

DOI: 10.54254/2755-2721/2026.GU34155

datasets and assessment of models on uncommon or extreme cases, providing a reliable foundation for clinical application [8].

### 3.2. Medical image modal generation and lesion simulation

Affected by noise, artifacts, low resolution, and poor contrast, clinical images frequently need quality improvement, yet traditional techniques can impair essential lesion information. By contrast, GANs enable high-quality image enhancement and restoration while preserving key diagnostic information, making them a dominant approach for medical image preprocessing.

Among their applications, medical image enhancement employs GANs to suppress noise, increase resolution, and boost contrast, enabling clinicians to more accurately visualize anatomical structures and lesions. Typical scenarios include low-dose imaging denoising, super-resolution reconstruction, and contrast enhancement. For ionizing radiation modalities such as CT and X-ray, image quality positively correlates with radiation dose. Therefore, low-dose CT (LDCT) images exhibit significant noise, which can obscure small lesions and elevate the likelihood of misdiagnosis. To mitigate this issue, GAN-based denoising models address this by mapping low-dose to routine-dose images, removing noise while keeping lesions and calcifications, thus lowering radiation exposure without reducing diagnostic accuracy [9-11].

In addition to enhancement, medical image restoration seeks to recover true anatomical structures in areas corrupted by artifacts, occlusions, or damage, with common applications like metal artifact removal, motion correction, and defect restoration. Metal artifacts, often caused by dental fillings, orthopedic implants, or cardiac pacemakers, severely obscure surrounding anatomical structures and lesion information in CT scans. And GAN-based metal artifact restoration models learn the structural distribution of artifact-free images and achieve high-quality reconstruction of corrupted regions, outperforming traditional methods. These models have been widely used in dental, orthopedic, and head CT imaging. However, motion artifacts caused by patient breathing, heartbeat, or involuntary movement can result in blurring, ghosting, and distortion. GANs effectively detect and correct these artifacts, restoring sharp anatomical structures without the need for repeated scanning [12]. Besides, GANs also restore wrinkled or damaged regions in pathological slides, glare or occlusions in fundus images, and blurry or out-of-focus areas in endoscopic images, thus providing high-quality data for subsequent diagnosis and analysis.

## 4. Medical image modal generation and lesion simulation

### 4.1. Clinical application challenges

Despite the outstanding performance of GANs in medical image generation, they still face multiple challenges in clinical practice. The realism and generalizability of synthesized images remain limited. Small deviations in lesions or anatomical structures can cause misdiagnosis, and some generated images contain anatomically or pathologically inconsistent pseudo-features that cannot be verified by clinicians. Imaging data vary across devices, centers, and patient populations. Thus, models trained on single-center data often perform poorly when applied across devices or institutions. Furthermore, mode collapse and limited sample diversity restrict data augmentation for prevalent conditions and synthesis of rare cases. The black-box nature of models and the lack of interpretable mechanisms have reduced clinical trust, as physicians cannot directly verify the authenticity of lesions and key anatomical structures. In medical applications, GANs are yet to have a unified clinical validation or evaluation standard. The lack of standardization undermines the

DOI: 10.54254/2755-2721/2026.GU34155

reproducibility of model training and evaluation. It also raises ethical and privacy concerns, as some GAN models may inadvertently expose training data. Besides, synthesized images can retain sensitive patient information, potentially enabling identity reconstruction. Additionally, the inherent instability of GAN training and the lak of standardized protocols exacerbate the challenges of practical implementation. The combination of insufficient standards, technical instability, and regulatory uncertainty significantly impedes the safe and effective application of GANs in healthcare.

### 4.2. Future development strategies

Future research on GANs in medical imaging should be centered on clinical demands. Controllable generative models therefore need to integrate anatomical, pathological, and imaging principles to produce more realistic images, enabling precise control over global structures and local lesions, thereby improving clinical usability. Meanwhile, model interpretability should be enhanced and a systematic clinical validation framework established. Methods such as feature visualization and causal inference can be adopted to reveal the logic of the generation process. In collaboration with clinical experts and regulatory authorities, unified validation criteria and evaluation protocols should be developed to ensure that synthesized images conform to medical priors and strengthen clinical credibility. Privacy preservation and multi-center collaborative learning also represent key directions. By applying federated learning and differential privacy techniques to GANs, a secure multi-center training framework can be established, addressing medical data silos and protecting training data confidentiality. Besides, industry and regulatory agencies should clarify ethical and legal standards to ensure compliant development. In addition, lightweight and multimodal fusion models for clinical deployment hold substantial practical value. The design of low-computation, high-efficiency GAN architectures supports deployment on portable devices and edge computing platforms, thus allowing real-time image enhancement and reconstruction.

## 5. Conclusion

This study demonstrates that GANs offer revolutionary solutions to the core pain points in the field of medical imaging, effectively alleviating issues such as the scarcity of high-quality annotated data, the lack of multimodal information, and poor quality of clinical medical images. As a result, GANs have exhibited exceptional performance in scenarios including cross-modality image translation (e.g., bidirectional synthesis between CT and MRI), anonymized data augmentation for rare diseases, denoising of low-dose CT, and metal artifact reduction. By breaking down medical data silos, these capabilities significantly improve diagnostic accuracy and model generalization, providing a solid practical foundation for intelligent clinical diagnosis and treatment. However, current research still faces critical bottlenecks. Generated samples tend to produce pseudo-features that violate anatomical or pathological principles, making them difficult to pass clinical validation. Meanwhile, the black-box nature of GANs renders the underlying generation logic unintelligible, which hinders the establishment of clinical trust. Additionally, the risks of privacy leakage in synthesized data and ambiguous ethical norms restrict their large-scale deployment. Furthermore, unstable training and inconsistent evaluation criteria impede the reproducibility and horizontal comparison of relevant studies. In the future, integration of medical prior knowledge, including anatomical and pathological principles, into GAN frameworks is essential to enhance the authenticity and controllability of generated samples. This integration facilitates the development of interpretable GAN frameworks and standardized clinical validation systems, strengthening clinical

DOI: 10.54254/2755-2721/2026.GU34155

credibility. The combination of federated learning and differential privacy technologies enables multi-center collaborative training while safeguarding patient privacy. Finally, the development of lightweight, multimodal fusion models for real-time clinical applications drives the advancement of medical imaging toward greater precision and personalization.

## References

[1] Selvarathinam, A. R. X. L. A., Anbalagan, N., Srinivasu, N. P., et al. (2026). Correction: Parameter-optimized generative adversarial network framework for synthetic MRI generation: Fine-tuning critical variables for enhanced image fidelity. Frontiers in Medicine, 13, 1803906.
[2] Ju, Z., Zhou, W., Kong, L., et al. (2025). HAGAN: Hybrid augmented generative adversarial network for medical image synthesis. Machine Intelligence Research, 22(5), 1-14.
[3] Yang, S. (2025). Research on pelvic CT-to-MRI modality conversion method based on generative adversarial networks (M.S. thesis). Hebei University of Science and Technology, Shijiazhuang, China.
[4] Jin, L., Peng, Y., Wang, J., et al. (2025). Hierarchical multi-scale Mamba generative adversarial network for multi-modal medical image synthesis. Expert Systems with Applications, 280, 127451.
[5] Zhang, Z. (2025). Research on virtual staining of pathological images based on generative adversarial networks (M.S. thesis). Shanghai Normal University, Shanghai, China.
[6] McNulty, R. J., Kho, L., Case, L. A., et al. (2024). Synthetic medical imaging generation with generative adversarial networks for plain radiographs. Applied Sciences, 14(15), 6831.
[7] Casado, R. L. J., Cabello, M. A. M., Baena, L. M. R. (2024). Enhancing histopathological image classification performance through synthetic data generation with generative adversarial networks. Sensors (Basel), 24(12), 3777.
[8] Liu, M. (2024). Research on multimodal brain image generation method based on generative adversarial networks (M.S. thesis). Changchun University of Technology, Changchun, China.
[9] Zhu, Q. (2024). Generative adversarial network-based image enhancement method for occupational pneumoconiosis classification (M.S. thesis). China University of Mining and Technology, Xuzhou, China.
[10] Bai, X. (2024). Research on multi-contrast brain MR image generation based on generative adversarial networks. Linyi University, Linyi, China.
[11] Yang, Y., Liu, J., Zhan, G., et al. (2024). OA-GAN: Organ-aware generative adversarial network for synthesizing contrast-enhanced medical images. Biomedical Physics & Engineering Express, 10(3).
[12] Mahesh, G., Reenu, R., Sweta, J. (2023). Histopathological image synthesis with generative adversarial networks for nuclei segmentation. SN Computer Science, 5(1).