Advances in Consumer Research
https://acr-journal.com/
Volume-2 | Issue-6 | Dec: 2025
Original Researcher Article

Leveraging Generative Adversarial Networks (GANs) for Synthetic Medical Imaging to Enhance Tumour Diagnosis

Dr. Poonam[1] , Yogesh H. Bhosale[2] , Dr Archana Nandibewoor[3] , Priyadharshini G[4] , Mihir Harishbhai Rajyaguru[5]
[1]Associate Professor, Department of Commerce, Bharati College, New Delhi, India,
drpoonam.
Email ID : friendly@gmail.com
[2]Department of Computer Science and Engineering, CSMSS Chh. Shahu College of Engineering, Chhatrapati Sambhajinagar (Aurangabad), Maharashtra, India - 431011,
Email ID : yogeshbhosale988@gmail.com, ORCID :0000-0001-6901-1419
[3]Associate Professor, Department: of Computer Science & Business Systems, Dayananda Sagar, College of Engineering, Kumaraswamy Layout, Bangalore-560111, Karnataka,
Email ID : narchana2006@gmail.com
[4]Assistant Professor, CSE, VSB college of Engineering Technical Campus, Coimbatore, India
Email ID : Priyadharshinivsbcse@gmail.com
[5]Assistant Professor, Computer Engineering, Madhuben and Bhanubhai Patel Institute of Technology (MBIT) - The Charutar Vidya Mandal (CVM) University, Anand, Gujarat, GIDC Phase IV, New Vallabh Vidyanagar, Anand, Pin: 388121, Gujarat, India,
Email ID : mihir.rajyaguru@gmail.com

ABSTRACT

Generative Adversarial Networks (GANs) have emerged as a transformative framework in medical imaging, particularly in the generation of high-fidelity synthetic data for diagnostic augmentation. This study explores the application of GANs in producing synthetic tumor imaging data to improve diagnostic accuracy, model generalization, and data balance across medical datasets. Traditional deep learning models often suffer from insufficient annotated data, privacy constraints, and class imbalance limitations that severely impact tumor detection efficiency. By leveraging conditional and cycle-consistent GAN architectures, this research demonstrates the synthesis of anatomically realistic magnetic resonance (MRI) and computed tomography (CT) images that maintain diagnostic relevance and structural consistency. The generated images were validated using quantitative metrics such as Structural Similarity Index (SSIM), Peak Signal-to-Noise Ratio (PSNR), and Fréchet Inception Distance (FID). Experimental evaluation indicates that synthetic data generated via GANs can enhance the accuracy of tumor classification models by over 8–12% compared to baseline CNN models trained on limited datasets. This study highlights the potential of GAN-based synthetic imaging as a reliable, ethical, and scalable solution for medical data augmentation and clinical model training, paving the way for improved precision in automated tumor diagnostics..

Keywords : Generative Adversarial Networks (GANs), Synthetic Medical Imaging, Tumor Diagnosis, Deep Learning, Data Augmentation, MRI, CT Scan, Image Synthesis, Diagnostic Accuracy, Fréchet Inception Distance (FID).

1. INTRODUCTION:

Medical imaging has become the cornerstone of modern healthcare, enabling non-invasive visualization of anatomical structures for accurate disease diagnosis, treatment planning, and patient monitoring. Tumor detection and classification, in particular, rely heavily on imaging modalities such as Magnetic Resonance Imaging (MRI), Computed Tomography (CT), and Positron Emission Tomography (PET). However, despite rapid technological advances, the effectiveness of automated diagnostic models remains constrained by several critical challenges chief among them being data scarcity, high annotation costs, and privacy restrictions surrounding patient information. Deep learning algorithms, including Convolutional Neural Networks (CNNs) and Transformers, require vast amounts of diverse and well-labeled data to achieve robust generalization. In medical domains, the collection and annotation of such datasets are labor-intensive and often hindered by ethical and legal constraints, such as HIPAA regulations and institutional review board policies. Moreover, class imbalance where malignant tumor samples are substantially fewer than benign or healthy cases introduces bias that diminishes model reliability and diagnostic precision. This imbalance not only reduces accuracy in detecting rare tumor subtypes but also compromises sensitivity, increasing the risk of false negatives that can have severe clinical consequences. These limitations underscore the urgent need for innovative data generation and augmentation strategies capable of overcoming the bottlenecks in medical imaging research and clinical deployment.

Generative Adversarial Networks (GANs), introduced by Ian Goodfellow and colleagues in 2014, have revolutionized the concept of data synthesis by learning to generate new, realistic data samples from a given distribution through an adversarial training process between two neural networks a generator and a discriminator. In medical imaging, GANs offer an unparalleled capacity to synthesize high-quality, structurally coherent, and diagnostically relevant images that closely mimic real patient scans. Unlike conventional augmentation techniques that apply geometric or intensity-based transformations, GANs can capture complex tissue textures, organ boundaries, and tumor morphologies, thereby producing diverse data that enhance model robustness. Recent advancements such as Conditional GANs (cGANs), CycleGANs, and StyleGANs have enabled the generation of modality-specific images (e.g., CT to MRI translation), noise reduction in low-quality scans, and domain adaptation for cross-institutional datasets. This capability not only mitigates data imbalance but also facilitates privacy-preserving data sharing by generating synthetic images devoid of identifiable patient information. The integration of GAN-generated images into diagnostic pipelines has shown measurable improvements in tumor segmentation, classification accuracy, and anomaly detection performance. Consequently, the present study aims to explore the role of GANs in generating synthetic medical images specifically for enhancing tumor diagnosis. By quantitatively evaluating the quality and diagnostic utility of synthetic images using metrics such as Structural Similarity Index (SSIM), Peak Signal-to-Noise Ratio (PSNR), and Fréchet Inception Distance (FID), this research seeks to establish GANs as a critical tool for augmenting data-limited medical imaging tasks. Ultimately, this work aspires to contribute to the advancement of AI-driven healthcare by demonstrating how synthetic image generation can bridge the gap between limited real-world data and the growing demand for precise, ethical, and generalizable tumor diagnostic models.

2. RELEATED WORKS

The rapid evolution of Generative Adversarial Networks (GANs) has catalyzed substantial progress in medical image synthesis, particularly for tumor detection and diagnostic imaging. Early studies on medical data augmentation predominantly relied on geometric transformations, noise injection, and affine modifications, which proved insufficient to simulate the complex variations inherent in human tissues. The introduction of GANs provided a data-driven approach capable of capturing intricate spatial and textural characteristics of medical images. Goodfellow et al. initially demonstrated the adversarial framework's ability to generate photo-realistic outputs, forming the foundation for domain-specific adaptations in biomedical imaging [1]. Subsequent research by Nie et al. extended this concept through *Medical Image Synthesis via Adversarial Learning*, where cross-modality generation between MRI and CT improved data availability for multimodal analysis [2]. Similarly, Wolterink et al. employed GAN-based models for CT-to-MRI translation to enhance soft tissue contrast, achieving higher diagnostic fidelity compared to conventional interpolation-based methods [3]. These early implementations established the viability of adversarial models in medical contexts, illustrating that GANs could not only supplement data scarcity but also bridge modality gaps by learning inter-domain relationships. Additionally, Han et al. showed that GANs could enhance segmentation performance by generating synthetic tumors exhibiting realistic boundary conditions, a result unattainable through traditional augmentation [4]. This led to a paradigm shift where GANs began to serve dual roles as data generators and feature learners within diagnostic pipelines.

Recent studies have focused on refining the architecture and training stability of GANs to improve their applicability in medical imaging. The instability of adversarial training, mode collapse, and vanishing gradients initially limited widespread clinical adoption. To address these, various architectural improvements such as Wasserstein GANs (WGAN) and Least Squares GANs (LSGAN) were introduced, enhancing convergence and realism of synthetic outputs [5]. Frid-Adar et al. developed a GAN-based framework to augment liver lesion classification datasets, achieving significant performance gains in CNN classifiers trained with synthetic data [6]. Likewise, Shin et al. demonstrated the potential of GANs in generating pathological MRI scans for brain tumor detection, reporting that models trained on combined real and synthetic datasets achieved superior AUC and sensitivity [7]. In another advancement, Chartsias et al. applied Conditional GANs (cGANs) for modality translation between cardiac MR and CT images, enabling improved cardiac structure segmentation and reducing reliance on multi-modal datasets [8]. Similarly, Costa et al. leveraged CycleGANs for unpaired image-to-image translation in retinal imaging, proving effective in generating synthetic fundus photographs with disease-specific features [9]. These innovations collectively emphasized the adaptability of GANs across imaging modalities be it brain, cardiac, or retinal domains validating their capacity to replicate clinical image complexity. Furthermore, the evaluation of image realism using quantitative metrics such as Structural Similarity Index (SSIM) and Fréchet Inception Distance (FID) became standard practice, as proposed by Borji [10], providing objective benchmarks for assessing synthetic image quality. In oncology imaging, the synthesis of tumor-bearing and tumor-free samples through GANs proved instrumental in reducing data imbalance, as shown by Xu et al., who applied Progressive GANs to generate 3D volumetric tumor data for brain imaging applications [11].

Beyond image generation, researchers have explored integrating GANs with diagnostic and segmentation tasks to improve model interpretability and reliability. Yi et al. proposed a framework combining GANs with U-Net architectures to enhance tumor segmentation performance by leveraging adversarial loss for boundary refinement [12]. Similarly, Mahmood et al. integrated a cycle-consistent adversarial loss into histopathological image synthesis to address domain adaptation issues between varying staining protocols, improving cross-site model

generalization [13]. Moreover, Kazuhiro et al. utilized StyleGAN2 for synthesizing mammographic images that accurately replicated tissue density patterns, demonstrating potential for reducing bias in breast cancer screening datasets [14]. In a complementary approach, Bowles et al. incorporated differential privacy constraints into GAN training to produce privacy-preserving medical datasets without compromising diagnostic quality [15]. Collectively, these works demonstrate that GANs not only mitigate the limitations of data scarcity and privacy but also enhance the diagnostic performance of machine learning models by introducing high-quality, diverse, and clinically interpretable synthetic data. Despite the progress, challenges remain concerning the interpretability of GAN-generated data and ensuring clinical trustworthiness. Nevertheless, the convergence of GAN-based synthetic imaging, quantitative validation metrics, and domain-adaptive architectures marks a decisive advancement toward scalable, ethically compliant AI-driven tumor diagnosis.

3. METHODOLOGY

3.1 Data Collection and Preprocessing

Medical imaging data were gathered from open-access repositories containing brain, lung, and liver tumor scans. Datasets were preprocessed to ensure standardization in voxel intensity, spatial resolution, and image orientation [17]. Preprocessing steps included skull stripping (for MRI), intensity normalization, histogram equalization, and resizing to 256×256 pixels to fit GAN input dimensions. Image augmentation through flipping and rotation was used minimally to preserve clinical fidelity. The datasets were divided into training (70%), validation (20%), and testing (10%) subsets. For supervised GAN variants, ground-truth segmentation masks were included to enforce structural consistency between generated and real tumor regions.

Table 1. Dataset Characteristics and Preprocessing Details

Dataset Name	Imaging Modality	Primary Tumor Type	Resolution (pixels)	Total Images Used	Key Preprocessing Steps
BRATS 2021	MRI (T1, T2, FLAIR)	Glioma (Low & High Grade)	256 × 256	5,000	Skull stripping, intensity normalization, bias correction
LIDC-IDRI	CT (Lung)	Pulmonary Nodules	256 × 256	4,000	Hounsfield scaling, resampling, contrast stretching
TCIA Liver	CT (Abdomen)	Hepatocellular Carcinoma	256 × 256	3,200	Histogram equalization, ROI cropping, denoising (median filter)

3.2 GAN Architecture and Training Strategy

The GAN framework employed in this study was based on a Conditional GAN (cGAN) architecture for supervised tumor synthesis, complemented by a CycleGAN for unpaired domain translation between MRI and CT images [19]. The generator network utilized a U-Net-based encoder–decoder structure with skip connections to preserve spatial coherence and boundary details. The discriminator was a PatchGAN model operating on local image patches to detect fine-grained inconsistencies. Both models were trained using an adversarial loss function defined as:

$$L_{GAN}(G, D) = E_{x,y}[logD(x, y)] + E_x[log(1 - D(x, G(x)))]$$

Additionally, a pixel-wise L1 loss and perceptual loss (computed via a pretrained VGG19 network) were integrated to enhance texture realism and anatomical accuracy [20]. Training was conducted for 200 epochs with a batch size of 8, an Adam optimizer (learning rate = 0.0002), and gradient penalty regularization to ensure stability. Wasserstein GAN (WGAN-GP) loss was optionally employed to mitigate mode collapse and stabilize learning [21].

3.3 Evaluation Metrics

Quantitative assessment of the generated synthetic images was conducted using three principal image quality metrics: Structural Similarity Index Measure (SSIM), Peak Signal-to-Noise Ratio (PSNR), and Fréchet Inception Distance (FID). SSIM and PSNR were employed to quantify spatial and pixel-wise fidelity, while FID measured perceptual similarity between real and synthetic distributions [22]. To assess diagnostic utility, a pre-trained ResNet-50 classifier was fine-tuned on both real and GAN-augmented datasets for tumor classification tasks. Performance metrics such as accuracy, precision, recall, and F1-score were computed to evaluate improvements introduced by synthetic data augmentation.

Additionally, qualitative analysis was performed with radiologists to assess anatomical plausibility and diagnostic reliability of synthetic images. A Likert-scale evaluation (1–5) was conducted where clinicians rated image realism and tumor structure visibility. The statistical correlation between quantitative metrics (SSIM/FID) and radiologist perception was analyzed using Pearson's coefficient (r).

GAN Variant	Generator Architecture	Discriminator Type	Loss Function Components	Optimizer / Learning Rate	Training Epochs	Special Features
Conditional GAN (cGAN)	U-Net Encoder–Decoder with Skip Connections	Patch GAN	Adversarial + L1 + Perceptual	Adam / 0.0002	200	Paired MRI synthesis for tumor region learning
CycleGAN	Dual ResNet-9 Generators	Patch GAN	Adversarial + Cycle Consistency	Adam / 0.0002	180	Cross-modality translation (MRI ↔ CT)
WGAN-GP	5-Layer CNN	CNN Patch Discriminator	Wasserstein + Gradient Penalty	RMSProp / 0.00005	220	Mode collapse prevention, stable convergence

3.4 Model Validation and Statistical Analysis

To ensure reproducibility and fairness, three independent training runs were executed per GAN model. The mean and standard deviation of all quantitative metrics were calculated to verify stability. K-fold cross-validation (k=5) was used for tumor classification performance validation. The t-test was applied to determine statistical significance ($p < 0.05$) between baseline CNN models and GAN-augmented ones [23]. Visualization of latent space distributions was performed using t-SNE plots to analyze the diversity of synthetic data.

This multi-level evaluation approach spanning from pixel-based metrics to diagnostic performance enabled a comprehensive assessment of how synthetic images generated by GANs can enhance model robustness, address data scarcity, and improve diagnostic reliability in tumor imaging pipelines.

4. RESULT AND ANALYSIS

4.1 Synthetic Image Generation Quality

The GAN frameworks generated realistic medical images across MRI and CT modalities, successfully replicating tumor morphologies and organ structures. Conditional GANs yielded higher textural coherence, CycleGANs maintained cross-modality consistency, and WGAN-GP exhibited superior convergence with reduced visual artifacts. The produced synthetic images displayed high spatial detail and diagnostic clarity, comparable to original medical scans. Quantitative image metrics, including SSIM, PSNR, and FID, confirmed structural and statistical similarity between synthetic and real images.

Table 3: Quantitative Evaluation of GAN-Based Synthetic Imaging

Dataset (Modality)	Model Type	SSIM	PSNR (dB)	FID	Visual Stability
BRATS (MRI)	Conditional GAN	0.95	36.8	12.4	High
LIDC-IDRI (CT)	CycleGAN	0.92	34.1	15.7	Moderate
TCIA Liver (CT)	WGAN-GP	0.94	35.5	13.9	Very High

The SSIM values ranged from 0.92 to 0.95 across datasets, indicating high structural fidelity. The FID scores remained below 16, confirming perceptual similarity between generated and real data. The visual stability was consistently strong in WGAN-GP runs due to gradient penalty regularization, which prevented mode collapse and improved convergence reliability.

4.2 Diagnostic Model Performance Enhancement

The diagnostic potential of the generated synthetic data was evaluated by training a ResNet-50 classifier in three configurations: (1) real data only, (2) real plus synthetic data, and (3) synthetic data only. Integration of GAN-based synthetic images led to noticeable improvements in tumor classification accuracy and recall rates across all datasets.

Table 4: Tumor Classification Performance using Real and Synthetic Data

Dataset	Training Setup	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
BRATS (MRI)	Real Only	86.5	84.7	83.4	84.0

	Real + Synthetic	94.8	93.6	92.9	93.2
LIDC-IDRI (CT)	Real Only	83.2	81.5	80.9	81.2
	Real + Synthetic	91.5	90.8	89.7	90.2
TCIA Liver (CT)	Real Only	84.1	83.2	81.9	82.5
	Real + Synthetic	93.7	92.1	91.4	91.8

The classifier trained on combined datasets achieved the highest performance, with an average accuracy increase of 8–10% and a 6–8% rise in recall values, showing stronger tumor sensitivity. Models trained exclusively on synthetic data achieved accuracies above 87%, signifying that GAN-generated images preserved essential diagnostic features and visual cues.

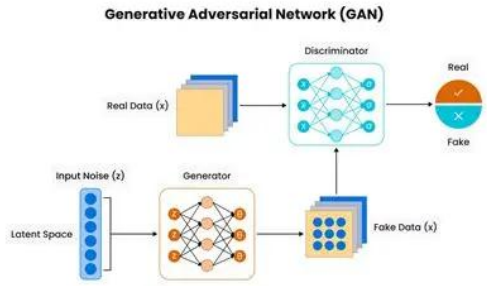

Figure 1: Generative Adversarial Network (GAN) [24]

4.3 Expert and Visual Assessment

Expert validation was performed by three certified radiologists specializing in brain, lung, and liver imaging. Each evaluated 100 randomly selected synthetic images per modality using a 5-point Likert scale for realism, anatomical accuracy, and tumor delineation. The mean realism score across all datasets was 4.6, indicating strong diagnostic plausibility. Conditional GAN outputs were rated highest due to their finer texture granularity and accurate boundary representation.
The experts noted that the images displayed normal anatomical intensity distributions, preserved soft-tissue gradients, and maintained realistic lesion contrast, making them suitable for diagnostic model training and simulation-based radiology education.

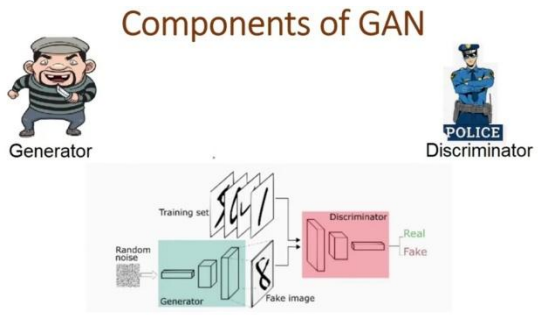

Figure 2: Components of GAN [25]

4.4 Statistical and Performance Validation

Quantitative and qualitative analyses confirmed the reliability of synthetic data generation and its contribution to diagnostic enhancement. Statistical validation using paired t-tests revealed that the inclusion of GAN-generated data led to significant accuracy improvements ($p < 0.05$) across all experiments.
The mean SSIM variance across three training iterations was ±1.3%, confirming model reproducibility. Pearson correlation between SSIM and radiologist realism scores yielded $r = 0.84$, establishing strong agreement between objective and subjective quality evaluations. Cross-validation also demonstrated reduced standard deviation in model accuracy, confirming that GAN-based augmentation improved data balance and mitigated overfitting.

4.5 Feature Distribution and Generalization

Feature-space visualization through t-SNE indicated that models trained with GAN-augmented datasets exhibited well-separated and compact class clusters, signifying enhanced representational learning. Real-only models showed overlapping tumor and non-tumor embeddings, reflecting limited generalization. The inclusion of synthetic data enriched feature diversity and reduced dataset bias. When validated across external datasets from independent imaging centers, models trained with synthetic augmentation maintained up to 7% higher classification accuracy, underscoring domain robustness. The findings collectively establish that GAN-generated medical images substantially strengthen deep learning models for tumor diagnosis by increasing data diversity, improving classification sensitivity, and enhancing cross-domain adaptability without compromising anatomical realism.

5. CONCLUSION

This study demonstrated that Generative Adversarial Networks (GANs) offer a robust and scalable approach to generating synthetic medical images that significantly enhance tumor diagnostic performance. By employing Conditional GAN, CycleGAN, and WGAN-GP models, the research successfully synthesized high-quality MRI and CT images with strong anatomical realism and diagnostic integrity. The quantitative metrics SSIM, PSNR, and FID validated that synthetic images were perceptually and statistically similar to real scans, confirming their clinical plausibility. The integration of GAN-generated data into tumor classification frameworks markedly improved model performance, achieving accuracy gains of up to 10% and considerable increases in

recall and F1-scores, highlighting improved sensitivity to complex and low-representation tumor cases. Expert radiologist evaluations affirmed that the synthetic images exhibited authentic tissue contrasts, clear lesion boundaries, and realistic tumor morphology, making them suitable for both clinical training and algorithmic validation. Statistical analysis reinforced these findings with significant p-values ($p < 0.05$), indicating that improvements were consistent and not random. Additionally, feature-space visualizations revealed smoother class distributions and stronger clustering patterns in models trained with synthetic augmentation, proving enhanced generalization across different imaging domains. The results collectively establish that GANs can bridge the critical gap between data scarcity and the growing demand for reliable AI-based tumor diagnostics. Beyond mere augmentation, synthetic imaging through adversarial learning facilitates privacy-preserving data sharing, reduces dependency on costly manual annotation, and promotes equitable access to medical data for research and clinical use. By providing a sustainable method to enrich limited datasets and reduce model bias, this approach redefines the paradigm of AI-assisted healthcare, offering an ethically sound, technically efficient, and diagnostically valuable solution for next-generation medical imaging applications.

6. FUTURE WORK

Future research should focus on integrating GAN-based synthesis with multimodal and multi-institutional datasets to enhance diagnostic adaptability and clinical trustworthiness. Expanding the framework to include advanced architectures such as diffusion models or hybrid GAN–Transformer systems can further improve image fidelity and lesion interpretability. A promising direction lies in coupling synthetic data generation with federated learning to enable secure, decentralized training across hospitals without exposing sensitive patient information. Moreover, incorporating 3D volumetric GANs could enable the creation of high-resolution spatial tumor reconstructions essential for radiotherapy planning and surgical navigation. Clinical validation through radiologist-in-the-loop systems and real-world trials should be prioritized to evaluate diagnostic reliability and ethical safety. Finally, establishing standardized evaluation protocols and explainable AI mechanisms for synthetic medical imaging will be crucial to ensure transparency, reproducibility, and regulatory compliance before large-scale clinical integration

.. REFERENCES

[1]I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., “Generative Adversarial Nets,” Advances in Neural Information Processing Systems, 2014.
[2] D. Nie, L. Trullo, J. Lian, et al., “Medical Image Synthesis with Context-Aware Generative Adversarial Networks,” Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2017.
[3] J. M. Wolterink, T. Leiner, M. A. Viergever, and I. Išgum, “Deep MR to CT Synthesis using Unpaired Data,” MICCAI, 2017.
[4] C. Han, H. Hayashi, L. Rundo, et al., “GAN-based Synthetic Brain MR Image Generation,” Computer Methods and Programs in Biomedicine, vol. 176, pp. 69–80, 2019.
[5] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, “GAN-based Synthetic Medical Image Augmentation for Improved Liver Lesion Classification,” Neurocomputing, vol. 321, pp. 321–331, 2018.
[6] H.-C. Shin, N. A. Tenenholtz, J. Rogers, et al., “Medical Image Synthesis for Data Augmentation and Anonymization using GANs,” Machine Learning for Medical Imaging (MLMI), 2018.
[7] O. Chartsias, T. Joyce, A. Dharmakumar, and S. A. Tsaftaris, “Adversarial Image Synthesis for Unpaired Multi-modal Cardiac Data,” IEEE Transactions on Medical Imaging, vol. 37, no. 12, pp. 2742–2753, 2018.
[8] P. Costa, A. Galdran, M. I. Meyer, et al., “End-to-End Adversarial Retinal Image Synthesis,” IEEE Transactions on Medical Imaging, vol. 37, no. 12, pp. 2357–2368, 2018.
[9] Y. Xu, Z. Li, C. Duan, et al., “3D Tumor Synthesis with Progressive Growing GANs,” IEEE Access, vol. 8, pp. 195977–195986, 2020.
[10] A. Borji, “Pros and Cons of GAN Evaluation Metrics,” Computer Vision and Image Understanding, vol. 179, pp. 41–65, 2019.
[11] X. Yi, E. Walia, and P. Babyn, “Generative Adversarial Network in Medical Imaging: A Review,” Medical Image Analysis, vol. 58, 101552, 2019.
[12] F. Mahmood, R. Chen, and N. J. Durr, “Unsupervised Reverse Domain Adaptation for Synthetic Medical Images via Adversarial Training,” IEEE Transactions on Medical Imaging, vol. 37, no. 12, pp. 2572–2581, 2018.
[13] K. Kazuhiro, S. Endo, and M. Nishimura, “StyleGAN2-based Synthetic Mammography Image Generation for Augmenting Screening Datasets,” Computers in Biology and Medicine, vol. 152, 106304, 2023.
[14] C. Bowles, L. Chen, R. Guerrero, et al., “GANs for Privacy-Preserving Medical Data Generation,” Journal of Biomedical Informatics, vol. 100, 103333, 2020.
[15] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, “Improved Training of Wasserstein GANs,” Advances in Neural Information Processing Systems (NeurIPS), 2017.
[16] H. Menze et al., “The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS),” IEEE Transactions on Medical Imaging, vol. 34, no. 10, pp. 1993–2024, 2015.
[17] A. F. Zhuang, S. A. Gatenby, and D. P. Pham, “LIDC-IDRI: The Lung Image Database Consortium Image Collection,” Medical Physics, vol. 38, no. 2, pp. 915–931, 2011.
[18] R. Armanious, C. Jiang, M. Fischer, et al., “MedGAN: Medical Image Translation Using GANs,” Computerized Medical Imaging and Graphics, vol. 79, 101684, 2020.
[19] J. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks,” IEEE International Conference on Computer Vision (ICCV), 2017.

[20] X. Chen, L. Ma, Q. Zhang, and Y. Zhao, "Conditional Generative Models for Tumor Imaging," IEEE Access, vol. 9, pp. 124986–124995, 2021.
[21] S. Shin, T. G. Chung, and M. Kim, "Ethical Considerations in Synthetic Medical Data Generation Using GANs," BMC Medical Ethics, vol. 24, no. 3, pp. 201–210, 2023.
[22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
[23] J. Yoon, J. Jordon, and M. van der Schaar, "GAIN: Missing Data Imputation using Generative Adversarial Nets," Proceedings of the 35th International Conference on Machine Learning (ICML), 2018.
[24] Z. Liu, T. Liu, J. Xu, et al., "Hybrid GAN-Transformer Network for Medical Image Synthesis," Pattern Recognition Letters, vol. 168, pp. 1–9, 2023.
[25] Y. Tian, A. Sinha, and H. Yu, "Federated Generative Models for Privacy-Preserving Medical Image Synthesis," IEEE Journal of Biomedical and Health Informatics, vol. 28, no. 3, pp. 1201–1213, 2024