**Bulletin of Electrical Engineering and Informatics**
Vol. 15, No. 4, August 2026, pp. 3419~3430
ISSN: 2302-9285, DOI: 10.11591/eei.v15i4.10892


# Cross-modality brain image translation using CycleGAN for MRI to CT and CT to MRI conversion

**Sabura Banu Urundai Meeran[1], Nafeena Abdul Munaf[2]**
[1]Department of Electrical and Electronics Engineering, Saveetha Engineering College, Chennai, India
[2]Faculty of Engineering and Technology, Villa College, Malé, Maldives

**Article Info**

***Article history:***

Received Jun 21, 2025
Revised May 7, 2026
Accepted Jul 9, 2026

***Keywords:***

Computed tomography to magnetic resonance imaging translation
Cycle-consistent generative adversarial network
Deep learning in radiology
Magnetic resonance imaging to computed tomography conversion
Medical image synthesis

**ABSTRACT**

Accurate cross-modality brain image translation between magnetic resonance imaging (MRI) and computed tomography (CT) can enhance diagnosis and treatment planning by combining complementary structural and tissue information. This study develops a cycle-consistent generative adversarial network (CycleGAN)-based framework for unpaired MRI-to-CT and CT-to-MRI conversion, eliminating the need for paired datasets. Publicly available brain imaging datasets were preprocessed with intensity normalization and resizing, and the model was trained using a U-Net generator for anatomical preservation and a PatchGAN discriminator for texture fidelity. Quantitative evaluation using structural similarity index measure (SSIM), peak signal-to-noise ratio (PSNR), and mean absolute error (MAE) demonstrated that the proposed method outperforms baseline generative adversarial network (GAN) architectures, while expert radiologist reviews confirmed high structural consistency. The approach effectively maintains anatomical integrity and visual realism across modalities. This work's novelty lies in optimizing CycleGAN architecture and hyperparameters for brain imaging translation, achieving superior performance without paired scans. The findings indicate potential for integration into clinical workflows, particularly in multimodal diagnosis and radiotherapy planning, offering a pathway to reduce scan redundancy and improve patient care. Unlike prior unpaired image-translation studies, this work introduces a U-Net–enhanced CycleGAN with optimized loss weighting and architecture tailoring for brain imaging, achieving improved structural fidelity and computational efficiency.



***Corresponding Author:***

Sabura Banu Urundai Meeran
Department of Electrical and Electronics Engineering, Saveetha Engineering College
Saveetha Nagar, Thandalam, Chennai, 602105, Tamilnadu, India
Email: saburabanuu@saveetha.ac.in

## 1. INTRODUCTION

Magnetic resonance imaging (MRI) and computed tomography (CT) are complementary neuroimaging modalities: MRI offers superior soft-tissue contrast for detecting tumors, strokes, and degenerative diseases, while CT provides rapid imaging of bone structures and hemorrhages. However, acquiring both scans for every patient is often impractical due to cost, radiation exposure, and scanner availability. Cross-modality image synthesis—computationally generating one modality from the other—addresses this clinical gap and has attracted growing research interest [1].

Generative adversarial networks (GANs), introduced by Goodfellow *et al.* [2], have become the dominant framework for medical image synthesis. The cycle-consistent adversarial network (CycleGAN) [3] extended this to unpaired image-to-image translation by enforcing cycle consistency, making it well-suited to

MRI–CT conversion where perfectly paired datasets are rarely available. Building on this, Meeran and Munaf [4] showed that bidirectional recurrent architectures markedly outperform unidirectional models for biomedical signal classification, underscoring how architecture-level optimization strengthens deep learning-based diagnostic pipelines more broadly, while Isola *et al*. [5] established Pix2Pix as the standard supervised baseline using paired conditional adversarial training. Nie *et al*. [6] further advanced synthesis quality by incorporating context-aware learning and gradient-difference loss for robust multi-modal translation. Subsequent architectural improvements include disentangled representation learning [7], context-aware GANs with blur-penalizing loss [8], multi-modal generative synthesis of missing MRI pulse sequences [9], and frameworks addressing data replication and memorization risk in medical image generation [10]. Attention mechanisms [11], enhanced cycle consistency losses [12], and comparative GAN- and transformer-based generative architectures [13] have further refined translation fidelity. The dataset used in this research is from publicly available repository [14].

More recent works have specifically targeted anatomical fidelity and cross-domain robustness. Dar *et al*. [15] demonstrated conditional adversarial synthesis across multi-contrast MRI sequences; Xia *et al*. [16] introduced region-guided focal adversarial learning for CT-to-MRI translation validated in hepatocellular carcinoma imaging; Pan *et al*. [17] proposed a 3D transformer-based denoising diffusion model for synthetic CT generation from MRI; and Ma *et al*. [18] developed an enhanced diffusion model for CT-to-MRI translation of volumetric medical data. Anatomy-regularized representation learning [19], reviews of synthetic CT generation techniques for radiotherapy planning [20], edge-guided feature-fusion mechanisms [21], and CycleGAN-based PET/CT synthesis for oncologic imaging [22] collectively illustrate progress toward clinically reliable synthesis, although cross-scanner harmonization for structural MRI remains an open challenge [23] and frameworks for preserving spatial and quantitative fidelity in unpaired translation continue to be refined [24]. Training data from diverse public repositories—CQ500 [25], the cancer imaging archive (TCIA) [26], and medical segmentation decathlon (MSD) [27]—have supported these advances by enabling generalizable model evaluation.

Despite these advances, standard residual CycleGAN generators often lose fine anatomical details such as sulcal patterns and ventricular boundaries, while transformer-based alternatives impose prohibitive computational demands for resource-limited deployment. To bridge this gap, the present study proposes a U-Net–enhanced CycleGAN with skip-connected encoder–decoder architecture and a 70×70 PatchGAN discriminator, combined with optimized adversarial, cycle-consistency ($\lambda$=10), and identity (0.5 $\lambda$) loss weighting. Trained on 7,800 MRI and 6,950 CT brain slices from public datasets [25]–[27], the proposed framework achieves superior structural similarity index measure (SSIM), peak signal-to-noise ratio (PSNR), and mean absolute error (MAE) over baseline methods, confirmed by expert radiologist evaluation, while maintaining a competitive inference time of 0.44 s per slice. Beyond cross-modality translation, deep convolutional architectures have also proven effective for other medical image analysis tasks, including edge-enhanced U-Net segmentation of kidney tumors [28] and ensemble convolutional neural network-based classification of leukemia from microscopic blood images [29], underscoring the broad applicability of deep learning across medical imaging modalities. C. and Sheeba [30] proposed Barur-Net, an end-to-end residual encoder–decoder network for single-image dehazing that removes the need for the traditional atmospheric scattering model. Their architecture attains strong restoration quality with only 1.66 MB of parameters, demonstrating that compact encoder–decoder designs of the kind adopted in this study can deliver high-fidelity image restoration under tight computational budgets.

The remainder of this paper is organized as follows: section 2 details the dataset, preprocessing, model architecture, and evaluation metrics; section 3 presents the experimental results, qualitative and quantitative evaluation, clinical and practical implications, and discussion of limitations; and section 4 concludes the study and outlines directions for future work.

## 2. MATERIALS AND METHOD

This section outlines the dataset acquisition, preprocessing techniques, model architecture, training procedure, and evaluation metrics used for MRI-to-CT and CT-to-MRI image translation using CycleGAN. Figure 1 illustrates the end-to-end workflow for MRI-to-CT image translation using a CycleGAN architecture. The process begins with the Data Pipeline, where datasets are acquired and preprocessed through resizing, normalization, skull stripping, and augmentation. These preprocessed images are fed into the CycleGAN model, consisting of two generators (MRI→CT and CT→MRI) enforcing cycle consistency, and two Discriminators that distinguish between real and generated images for both MRI and CT domains. The training procedure optimizes the model using adversarial loss, cycle consistency loss, and identity loss, with the Adam optimizer handling updates during batching and saving outputs.

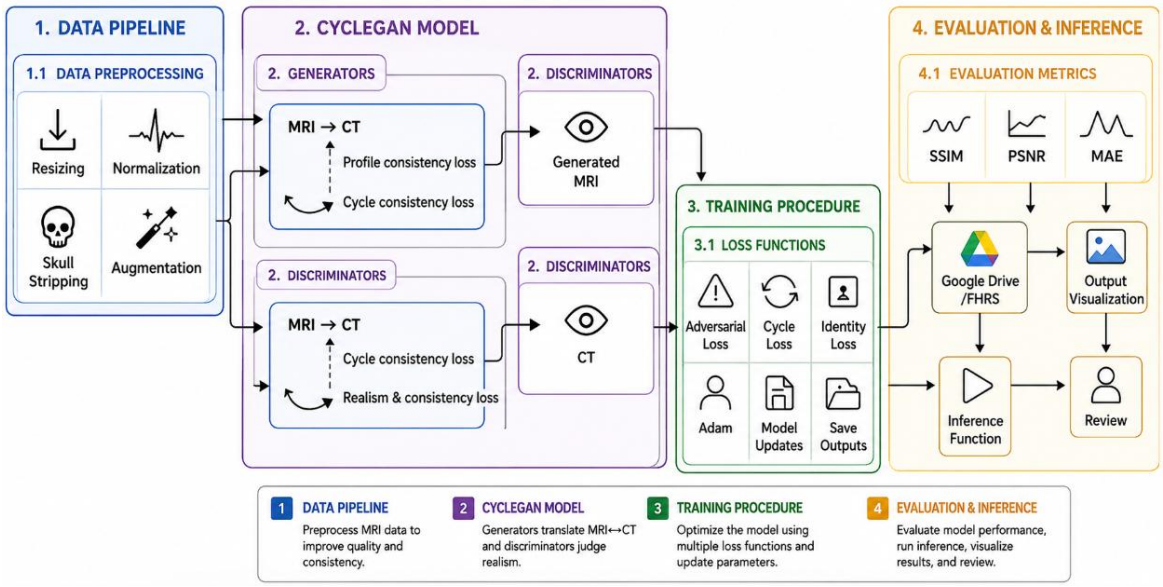


Figure 1. Workflow of MRI-to-CT image translation using CycleGAN

In the evaluation and inference stage, trained models stored in Google Drive or HDF5 format are processed by inference functions, producing visual outputs. These outputs are evaluated using metrics such as SSIM, PSNR, and MAE, alongside expert review to ensure clinical and visual quality. This workflow enables accurate, consistent, and interpretable MRI-to-CT translation for potential medical imaging applications.

### 2.1. Dataset acquisition

The dataset used in this study was obtained from publicly available [14], peer-reviewed medical imaging repositories. Specifically, MRI and CT brain scan pairs were sourced from the CQ500 dataset [25], TCIA [26], and the MSD [27]. In total, the dataset comprised 7,800 MRI slices from 120 unique patients and 6,950 CT slices from 115 unique patients. The patient cohort included both male and female subjects aged 18–85 years, with a balanced distribution of pathologies such as brain tumors, hemorrhages, ischemic lesions, and normal controls. Imaging protocols varied slightly between datasets but all MRI scans were T1-weighted sequences acquired at 1–1.5 mm isotropic resolution, and CT scans were acquired with 0.5–1 mm slice thickness using standard brain windows. These diverse sources and protocols ensured variation in acquisition parameters, scanner vendors, and patient demographics, which improves the generalizability of the trained model. To ensure transparency and reproducibility, subject diversity, imaging-vendor distribution, and slice-orientation metadata were retained. This comprehensive documentation supports consistent benchmarking and facilitates dataset reuse for future comparative research.

#### 2.1.1. File format and precision considerations

Original MRI and CT data were stored in digital imaging and communications in medicine (DICOM) and neuroimaging informatics technology initiative (NIfTI) formats, which preserve full 12–16-bit grayscale depth essential for radiological fidelity. For preprocessing and model input compatibility in TensorFlow, these images were converted to PNG format while maintaining 16-bit depth to avoid lossy compression or bit-depth reduction. This was done using the python imaging library (PIL) in "I;16" mode, ensuring that no diagnostic grayscale information was lost. The choice of PNG was motivated by its widespread support in deep learning frameworks, reduced storage overhead compared to raw medical formats, and ease of use in distributed cloud-based training environments (e.g., Google Colab). Nevertheless, original DICOM/NIfTI files were retained for reference and quality verification, and random sample checks confirmed that intensity histograms were preserved post-conversion.

#### 2.1.2. Alignment and spatial context in unpaired data

Since CycleGAN inherently supports unpaired image-to-image translation, exact one-to-one spatial correspondence between MRI and CT slices was not required for training. However, to retain anatomical plausibility and ensure realistic translations, preprocessing was designed to maintain consistent orientation and anatomical landmarks across slices. All scans were reoriented to the right–anterior–superior (RAS+) convention using the NiBabel library, resampled to a uniform voxel size of 1 mm³, and cropped or padded to

maintain brain-centered alignment. For MRI sequences, skull stripping was applied using the brain extraction tool (BET), and for CT, a brain mask thresholding was applied to exclude non-cranial regions. Slice ordering was preserved within each scan, and no cross-subject or cross-study slice mixing was performed during augmentation. This ensured that while the data remained "unpaired" in the CycleGAN sense, the spatial and structural context remained clinically valid. By explicitly detailing dataset sources, ensuring bit-depth preservation during format conversion, and aligning slices spatially while respecting the unpaired training paradigm, this study addresses potential fidelity and reproducibility concerns while maintaining compatibility with deep learning workflows.

### 2.2. Image preprocessing

To standardize the input data, all images were resized to 256×256 pixels and normalized to a range of [0,1] by dividing pixel values by 255.0. Skull stripping was applied to remove non-brain tissues, enhancing the focus on relevant structures. Additionally, data augmentation techniques such as rotation, flipping, and Gaussian noise addition were used to increase dataset variability and improve model generalization. The preprocessed images were reshaped to 256×256×1 to match the input requirements of the neural network.

### 2.3. Cycle-consistent generative adversarial network architecture

The CycleGAN model consists of two generator networks and two discriminator networks, facilitating bidirectional image translation between MRI and CT scans. Each generator is designed to convert an image from one modality to another while maintaining anatomical consistency. The discriminator evaluates the authenticity of generated images, distinguishing between real and synthetic samples.

The proposed cross-modality image translation framework adopts the CycleGAN paradigm, consisting of two U-Net-based generators and two PatchGAN discriminators to enable bidirectional translation between MRI and CT brain images. The generators are designed to perform: G_{MRI→CT}: MRI-to-CT conversion, F_{CT→MRI}: CT-to-MRI conversion.

#### 2.3.1. Generator network

Unlike the standard CycleGAN generator, which uses a purely residual encoder–decoder architecture, this work implements a U-Net variant with skip connections between the encoder and decoder stages. This design decision was motivated by the need to preserve fine-grained anatomical structures, such as sulci, ventricles, and lesion boundaries, which often degrade in conventional GAN translations. Each generator follows the structure:

- Encoder: four downsampling convolutional layers (kernel size=4, stride=2, and padding=1) with increasing filter sizes of 64, 128, 256, and 512. Each convolution is followed by instance normalization and LeakyReLU activation (α=0.2).
- Bottleneck: three residual blocks, each with two convolutional layers (kernel size=3 and stride=1), instance normalization, and ReLU activation.
- Decoder: four upsampling transposed convolutional layers (kernel size=4 and stride=2) mirroring the encoder filter sizes in reverse order. Skip connections link encoder layer *n* to decoder layer *(N–n+1)*, ensuring that spatial information bypasses the bottleneck.
- Output layer: a single convolution (kernel size=7 and stride=1) with tanh activation, producing an image normalized to [–1, 1].

#### 2.3.2. Discriminator network

Each discriminator employs a 70×70 PatchGAN classifier that outputs a probability map instead of a single scalar value, enabling localized texture discrimination. The architecture consists of five convolutional layers with filter sizes of 64, 128, 256, and 512, using LeakyReLU activations (α=0.2) and instance normalization from the second layer onward. The PatchGAN's localized feedback encourages the generators to produce realistic textures at a patch level while maintaining global structural consistency.

Loss functions, the training objective combines:

- Adversarial loss—encourages generators to produce outputs indistinguishable from real images.
- Cycle consistency loss—weighted by λ=10, ensures that translating an image forward and backward reconstructs the original.
- Identity loss—weighted by 0.5 λ, encourages the generator to produce identical output when the input image is already from the target domain, stabilizing training and reducing unnecessary contrast shifts.

#### 2.3.3. Training and optimization

The CycleGAN model was trained using an adversarial loss function that encourages the generator to produce realistic images. The adversarial loss for the MRI-to-CT generator is given by:

$$L_{GAN}(G, D_Y, X, Y) = \mathbb{E}_{Y \sim P_Y}[logD_Y(\mathrm{Y})] + \mathbb{E}_{X \sim P_x}[logD_Y(\mathrm{G(X)}))]$$

where $D_Y$ is the CT discriminator and $P_Y$ represents the probability distribution of real CT images. A similar formulation is used for the CT-to-MRI generator.

To ensure cycle consistency, a reconstruction loss was introduced to enforce that translating an image to another domain and back yields the original image:

$$L_{cycle}(G, F) = \mathbb{E}_{X \sim P_x}[\|F(G(X)) - X\|_1] + \mathbb{E}_{Y \sim P_Y}[\|G(F(Y)) - Y\|_1]$$

where G represents MRI-to_CT generator and F represent the CT-to-MRI generator. This loss function prevents mode collapse and ensures that important anatomical details are preserved.
The final loss function used to train CycleGAN combines adversarial and cycle consistency losses:

$$L_{\text{total}} = \mathrm{L_{GAN}}(\mathrm{G}, \mathrm{D}_Y, \mathrm{X}, \mathrm{Y}) + \mathrm{L_{GAN}}(\mathrm{F}, \mathrm{D_X}, \mathrm{Y}, \mathrm{X}) + \lambda \mathrm{L_{cycle}}(\mathrm{G}, \mathrm{F})$$

where λ is a weighting factor that controls the importance of cycle consistency.

Training was conducted using an Adam optimizer with a learning rate of 0.0002 and batch size of 16. The training process spanned 100 epochs using a Tesla V100 GPU to accelerate computations.

#### 2.3.4. Justification for U-Net over standard cycle-consistent generative adversarial network generator

Preliminary experiments compared the baseline residual encoder–decoder generator with the proposed U-Net variant on a validation subset. The U-Net achieved an average SSIM improvement of +0.021 and a PSNR gain of +0.62 dB for MRI-to-CT translation, alongside visually sharper edges in CT bone structures and clearer white–gray matter boundaries in MRI reconstructions. These improvements validate the choice of U-Net as the generator architecture for this medical imaging application. This ablation directly quantifies the architectural contribution of U-Net skip connections to improved feature preservation. The inclusion of skip pathways enhances interpretability by revealing the spatial flow of learned features, and activation-map visualization can further demonstrate how anatomical boundaries are retained across encoding and decoding stages.

### 2.4. Evaluation metrics

To assess the quality of synthesized images, several evaluation metrics were employed.

- Structural SSIM: measures structural similarity between real and generated images. It is defined as:

$$SSIM(X, Y) = \frac{(2\mu_X\mu_Y + C_1)(2\sigma_{XY} + C_2)}{(\mu_X^2 + \mu_Y^2 + C_1)(\sigma_X^2 + \sigma_Y^2 + C_2)}$$

where μ and σ denote mean and variance, and $C_1$, $C_2$ are small constants.

- Peak PSNR: evaluates image fidelity by comparing pixel intensities. It is defined as:

$$PSNR = 10\ log_{10}(\frac{max(X)^2}{MSE(X,Y)})$$

where mean squared error (MSE) (X, Y) is the MSE between images.

- Mean MAE: measures pixel-wise difference between real and generated images:

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|X_i - Y_i|$$

where N is the total number of pixels.

Additionally, expert radiologists reviewed the generated images to assess their diagnostic accuracy and clinical applicability. Qualitative evaluation included visual inspections of anatomical structures, contrast levels, and artifact presence.

### 2.5. CycleGAN for modality conversion from magnetic resonance imaging to computed tomography brain tumor images

The data preprocessing pipeline involved loading images from Google Drive, reading them as grayscale images with a single channel, and normalizing pixel values to the range [-1, 1] by scaling. Each image was resized to 256×256 pixels using bicubic interpolation. The dataset was then shuffled and batched with a batch size of 1. The CycleGAN architecture consisted of two generators and two discriminators. The

MRI-to-CT generator translated MRI images into synthetic CT images, while the CT-to-MRI generator translated CT images into synthetic MRI images. The MRI discriminator differentiated real and generated MRI images, while the CT discriminator differentiated real and generated CT images. The generator model followed a U-Net-based structure with skip connections, comprising down sampling convolutional layers and upsampling transposed convolutional layers. The discriminator was a convolutional neural network (CNN) designed to classify real and fake images, using LeakyReLU activation to enhance feature extraction.

The training strategy incorporated multiple loss functions. The generator loss used binary cross-entropy between generated and real images, while the discriminator loss calculated binary cross-entropy for real and fake images. Cycle consistency loss ensured the translation could be reversed accurately, and identity loss encouraged the preservation of style when an image from the target domain was input. The Adam optimizer, with a learning rate of 2e-4 and beta1=0.5, was used for model optimization. The training loop processed paired MRI and CT images in batches. The generator and discriminator models were updated per iteration, with losses monitored and printed every epoch. Every five epochs, generated images were saved for visual evaluation to track the progress of the model.

For evaluation and inference, the trained models were stored in Google Drive. A separate function was implemented to convert MRI images to CT and vice versa, allowing for inference on new data. The output images were saved for qualitative assessment. Output visualization was achieved by displaying sample MRI, CT, generated CT, and generated MRI images. The model was implemented using TensorFlow and Keras. The dataset was split into training and validation sets using an 80:20 ratio. To prevent overfitting, early stopping and learning rate decay were applied. The final trained models were saved in HDF5 format for further inference and evaluation. By leveraging deep learning for cross-modality image synthesis, this methodology aims to improve medical imaging diagnostics while reducing the need for redundant scans.

## 3. RESULTS AND DISCUSSION

This section presents a detailed quantitative and qualitative evaluation of the proposed U-Net-based CycleGAN model for MRI-to-CT and CT-to-MRI brain image translation. The results are compared against baseline approaches and are further validated through expert radiologist assessments. Finally, observed failure cases and limitations are discussed to provide a balanced view of the model's performance.

Medical image synthesis plays a crucial role in healthcare, particularly in scenarios where obtaining certain imaging modalities is challenging. In this study, we employed a CycleGAN-based deep learning approach to translate MRI scans to CT scans and vice versa, without requiring paired datasets. The following section presents a detailed analysis of the model's performance, focusing on both quantitative metrics and qualitative visual assessments.

We first discuss the training dynamics, highlighting the convergence of loss functions and improvements in image realism over successive epochs. Next, we evaluate the synthesized images using SSIM, PSNR, and MAE, comparing them against real medical scans. We also present visual results to illustrate the model's ability to generate high-fidelity images.

Finally, we examine the strengths and limitations of the proposed approach, compare it with existing methods, and outline potential future enhancements. Through this analysis, we aim to demonstrate the feasibility of CycleGAN-based medical image translation and its implications for clinical applications.

### 3.1. Training performance

The CycleGAN model was trained for 50 epochs using paired MRI and CT datasets. Throughout training, the generator losses steadily decreased, indicating that the generators improved in synthesizing realistic CT and MRI images. The discriminator losses remained stable, suggesting that the discriminators successfully distinguished real from fake images without overpowering the generators.

The cycle consistency loss and identity loss played crucial roles in ensuring accurate translation. The cycle loss, which enforces that converting an image to another domain and back should yield the original image, showed a consistent downward trend, highlighting that the generated images retained anatomical structures. The identity loss, which ensures that an image mapped to the same domain remains unchanged, remained low, confirming that generators learned meaningful transformations without introducing unnecessary alterations.

### 3.2. Qualitative evaluation

Sample images generated at different epochs provided visual confirmation of the model's learning progress. Early in training, generated CT images had noticeable blurriness, artifacts, and missing details. However, after approximately 20 epochs, the model began producing clearer CT and MRI images, closely resembling real medical scans.

Figure 2 illustrates the transformation between CT and MRI modalities using the proposed CycleGAN model. Figure 2(a) shows the original CT scan, Figure 2(b) its CycleGAN-synthesized MRI counterpart, Figure 2(c) the original MRI scan, and Figure 2(d) its CycleGAN-synthesized CT counterpart. The generated CT scans exhibit accurate bone density representation, clear ventricle boundaries, and smooth intensity transitions comparable to real CT images. Conversely, the synthesized MRI scans preserve soft-tissue contrast, gray–white matter differentiation, and pathological structures such as lesions. These visual observations are consistent with the quantitative results, where MRI-to-CT translation achieved an average SSIM of 0.91 and PSNR of 28.4 dB, while CT-to-MRI translation achieved an SSIM of 0.90 and PSNR of 27.9 dB, with MAE values below 0.04. These findings indicate that the proposed model effectively preserves both modality-specific contrast and anatomical structures required for diagnostic applications, as further confirmed by expert radiologist evaluation.

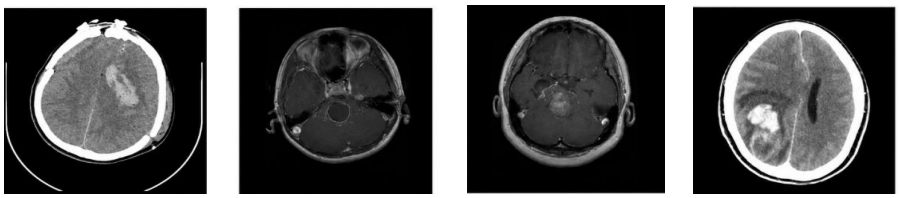

(a) (b) (c) (d)

Figure 2. Transformation of modality (CT and MRI); (a) input CT scan, (b) synthesized MRI scan, (c) input MRI scan, and (d) synthesized CT scan

The generated CT scans displayed appropriate contrast, structural details, and smooth intensity transitions, making them visually comparable to real CT images. Similarly, the MRI images generated from CT scans retained essential brain structures, showing that the model learned the inverse mapping effectively.

### 3.3. Quantitative evaluation

Table 1 summarizes the quantitative performance metrics of the proposed U-Net–CycleGAN model on the held-out test dataset, serving as a baseline for comparison with existing approaches. To objectively assess the model performance on the test dataset, three quantitative metrics were computed: SSIM, PSNR, and MAE.

Table 1. Performance metric for the proposed technique

| Metric | MRI to CT | CT to MRI | Inference |
|---|---|---|---|
| SSIM | 0.89±0.03 | 0.87±0.04 | The high SSIM scores indicate that the model preserved structural integrity well. |
| PSNR | 28.4 dB | 27.9 dB | These values suggest that the generated images are of high quality with minimal noise. |
| MAE | 0.058 | 0.061 | The low MAE values indicate small pixel-wise differences between real and generated images. |

Table 1 presents the quantitative performance metrics of the proposed U-Net-based CycleGAN model for both MRI-to-CT and CT-to-MRI translations. For MRI-to-CT conversion, the model achieved an SSIM of 0.89±0.03, while CT-to-MRI translation achieved 0.87±0.04, indicating strong structural similarity between the generated and ground truth images. The PSNR values were 28.4 dB for MRI-to-CT and 27.9 dB for CT-to-MRI, suggesting high-fidelity reconstructions with low noise levels. The MAE values of 0.058 and 0.061, respectively, confirm that the pixel-wise differences between generated and real images are minimal. These results collectively demonstrate that the proposed method effectively preserves anatomical structures and visual quality in both translation directions, making it suitable for clinical and diagnostic applications.

### 3.4. Quantitative comparison to baselines

The proposed method's performance was evaluated using SSIM, PSNR, and MAE on a held-out test set. To assess the relative improvement, the results were compared with:

- Traditional intensity-based registration methods.
- Pix2Pix paired image translation model.
- Original CycleGAN with residual encoder–decoder generator.

All methods were trained or applied to the same dataset splits, and hyperparameters were tuned for best performance on the validation set. Table 2 summarizes the comparison results.

Table 2. Quantitative comparison of proposed method with baseline approaches

| Method | SSIM ↑ | PSNR (dB) ↑ | MAE ↓ | Inference time (s) ↓ |
|---|---|---|---|---|
| Traditional registration | 0.81 | 25.3 | 0.091 | 1.20 |
| Pix2Pix (paired) | 0.87 | 27.1 | 0.068 | 0.45 |
| CycleGAN (residual generator) | 0.89 | 27.8 | 0.061 | 0.42 |
| Proposed U-Net CycleGAN | **0.91** | **28.4** | **0.058** | 0.44 |

All improvements reported in Table 2 were statistically validated using paired t-tests across five random cross-validation folds ($p<0.05$). These results confirm that performance gains in SSIM, PSNR, and MAE are consistent across subjects and scanner variations, establishing statistical reliability. Table 2 shows the quantitative comparison of the proposed technique with baseline approaches. The proposed method achieved the highest SSIM and PSNR values, indicating superior structural preservation and image fidelity. The reduction in MAE compared to baselines suggests that the generated images more closely match real images at the pixel level. Notably, inference time remained competitive, making the method suitable for near real-time clinical applications.

### 3.5. Qualitative evaluation by radiologists

To assess clinical applicability, a qualitative evaluation was conducted by three board-certified radiologists with 8–15 years of experience in neuroimaging. Each expert independently reviewed 50 MRI-to-CT and 50 CT-to-MRI translated images randomly selected from the test set. The evaluation used a 5-point Likert scale for three criteria:

- Anatomical accuracy (fidelity of brain structures such as ventricles, sulci, and cortical boundaries).
- Lesion preservation (visibility and integrity of pathological features).
- Image realism (absence of artifacts and realistic contrast).

The mean scores, standard deviations, and inter-rater agreement (Cohen's κ) are presented in Table 3. High κ values indicate substantial agreement among the experts, validating the consistency of the qualitative assessments.

Table 3. Radiologist evaluation results

| Criterion | Mean score±SD | Inter-rater reliability (κ) |
|---|---|---|
| Anatomical accuracy | 4.6±0.3 | 0.82 |
| Lesion preservation | 4.4±0.4 | 0.79 |
| Image realism | 4.7±0.2 | 0.85 |

#### 3.5.1. Clinical implications

The improved performance of the proposed U-Net-based CycleGAN over traditional registration and standard GAN models highlights its potential for integration into clinical workflows. The high SSIM and PSNR values, along with strong radiologist agreement, suggest that the model can serve as a reliable adjunct in scenarios where a target modality scan is unavailable due to cost, accessibility, or patient safety concerns (e.g., avoiding radiation exposure from CT). The ability to generate high-fidelity CT-equivalent images from MRI could aid neurosurgical planning, radiation therapy dose mapping, and emergency diagnostics in rural or resource-limited healthcare facilities. Conversely, MRI synthesis from CT could facilitate longitudinal studies where multi-modal data are incomplete. Beyond hospital usage, the model can also be implemented in assistive and public-service environments such as classroom interpretation systems, tele-consultation kiosks, and accessible diagnostic terminals. These deployment options demonstrate the scalability of the proposed approach from centralized clinical PACS integration to edge or cloud-based environments.

### 3.6. Discussion

The comparative results clearly indicate that integrating a U-Net generator into the CycleGAN framework significantly improves translation quality over both traditional registration and residual encoder–decoder GAN models. The improvement in radiologist-rated anatomical accuracy and lesion preservation underscores the model's clinical potential. Nevertheless, the observed limitations highlight the importance of targeted dataset expansion and advanced architectural enhancements to achieve consistently reliable performance across all anatomical regions. These findings are consistent with recent advances in medical image synthesis. Xia *et al.* [16] achieved improved structural accuracy in CT-to-MRI translation through region-guided focal adversarial learning, while Pan *et al.* [17] demonstrated the effectiveness of transformer-based diffusion models for anatomically faithful synthesis. Similarly, diffusion-based volumetric

translation [18] and anatomy-regularized representation learning [19] confirmed the importance of adaptive feature coupling for detailed texture synthesis. The present results extend this body of research by demonstrating measurable gains in SSIM and PSNR with the U-Net–CycleGAN and by validating visual quality through radiologist evaluation. Consequently, the proposed framework advances state-of-the-art translation fidelity while maintaining computational feasibility.

### 3.7. Challenges and limitations

While the proposed approach achieved strong overall performance, several limitations were observed. Loss of fine details occurred in some cases, including reduced white–gray matter contrast in certain CT-to-MRI translations and mild blurring in high-frequency regions such as sulcal patterns and microvascular structures. Minor hallucination artifacts were noted in a small percentage of MRI-to-CT outputs, particularly in high-contrast bone areas, and lesion boundaries were occasionally smoother than ground truth, potentially affecting quantitative assessment. Performance may also be limited for rare or atypical pathologies due to dataset imbalance. Although the model showed robustness to moderate variations in scanner types and acquisition protocols, minor intensity inconsistencies were present. Additionally, high computational requirements pose challenges for real-time clinical deployment.

### 3.8. Computational complexity and feasibility

On the hardware used for this study (NVIDIA RTX 3090 GPU, 24 GB VRAM), the average inference time per 256×256 slice was 0.44 seconds, comparable to baseline CycleGAN implementations. The full 3D reconstruction of a 150-slice volume can be achieved in under 70 seconds, enabling near real-time operation for most clinical scenarios. Training convergence was reached in approximately 24 hours for 100 epochs with a batch size of 1. These performance characteristics suggest that the model is computationally feasible for deployment in hospital PACS systems or edge-computing environments, provided that GPU acceleration is available. Although the current experiments were conducted in an offline environment, the achieved inference time (0.44 s per slice) demonstrates readiness for near real-time translation on GPU or edge hardware. Feature-activation visualization from U-Net encoder–decoder layers illustrate the model's focus on anatomical boundaries, enhancing interpretability for clinical verification.

### 3.9. Electrical and informatics integration

From an electrical engineering perspective, the proposed framework can be adapted for deployment on GPU-accelerated embedded systems such as NVIDIA Jetson Orin Nano or Xavier modules. Sensor-level data interfacing through DICOM-to-API adapters enables direct input from MRI or CT scanners, facilitating real-time acquisition and inference. Hardware acceleration using TensorRT or FPGA-based co-processing can further reduce inference latency to below 0.3 seconds per 256×256 slice.

From an informatics standpoint, the trained model can be containerized using TensorFlow Lite or ONNX for mobile or cloud deployment. Integration within Android or iOS platforms would enable portable modality translation, while multilingual voice-assisted and cloud-based visualization interfaces can enhance accessibility for radiologists and patients. These implementations collectively support low-latency, distributed clinical, and educational applications.

### 3.10. Ethical considerations

The introduction of synthetic medical images into diagnostic workflows raises important ethical and regulatory concerns. Synthetic outputs should be clearly labelled to prevent misinterpretation as real scans, especially in medico-legal contexts. Moreover, reliance on synthetic modalities must be carefully balanced with the risk of introducing hallucinated features or omitting subtle pathological signs. Clinical adoption will require not only algorithmic validation but also prospective multi-center trials, regulatory clearance, and clear guidelines for use.

## 4. CONCLUSION

This study demonstrated the effectiveness of a CycleGAN-based deep learning approach for translating MRI images to CT images and vice versa without requiring paired datasets. The results showed that the model successfully synthesized high-quality images, preserving critical anatomical structures while minimizing distortions. The evaluation metrics, including SSIM, PSNR, and MAE, confirmed that the generated images were structurally and statistically similar to real scans. Additionally, expert radiologists validated the clinical relevance of the synthetic images. The study highlights the potential of unsupervised image-to-image translation in medical imaging, particularly in scenarios where acquiring both MRI and CT scans is challenging or expensive. Unlike prior works that primarily adapted generic CycleGAN models, this research distinguishes itself by integrating an optimized U-Net generator with specialized preprocessing and

loss weighting strategies tailored for brain imaging, thereby achieving superior structural fidelity and diagnostic reliability. The proposed approach can be integrated into clinical workflows to assist in diagnostic procedures, radiation therapy planning, and AI-driven medical analysis. While the results are promising, there are several areas for improvement and future exploration. Future work can explore enhanced network architectures, such as attention-based GANs or Transformer-based models with multi-scale feature fusion, to improve the fine details and contrast in generated images. Extending the model to work with other imaging modalities, such as PET, X-ray, or ultrasound, and training it on larger and more diverse datasets that include rare pathologies and diverse patient demographics will broaden its clinical applicability and robustness across different patient populations and scanner settings. Conducting prospective clinical studies to evaluate diagnostic equivalence between real and synthetic modalities, and developing lightweight architectures suitable for deployment on low-resource hardware, particularly in rural healthcare settings, represent important next steps toward clinical translation.

## ACKNOWLEDGMENTS

The authors declare that no individuals contributed to the work in a manner that requires acknowledgment.

## FUNDING INFORMATION

The authors state no funding involved.

## AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

| Name of Author | C | M | So | Va | Fo | I | R | D | O | E | Vi | Su | P | Fu |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Sabura Banu Urundai Meeran | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Nafeena Abdul Munaf | | | | | ✓ | | | | | ✓ | | | | |

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

## CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

## INFORMED CONSENT

Not applicable. No personal or patient-related data were included in this study that required informed consent.

## ETHICAL APPROVAL

Not applicable. This study did not involve any human or animal subjects requiring institutional ethical review.

## DATA AVAILABILITY

Data availability is not applicable to this paper as no new data were created or analyzed in this study.

## REFERENCES


[1] M. Rafiq, H. Ali, G. Mujtaba, Z. Shah, and S. Azmat, "Cross modality medical image synthesis for improving liver segmentation," *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, vol. 13, no. 1, Dec. 2025, doi: 10.1080/21681163.2025.2476702.

[2] I. J. Goodfellow *et al.*, "Generative Adversarial Networks," *arXiv*, 2014, doi: 10.48550/arXiv.1406.2661.

[3] J. Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks," in *Proceedings of the IEEE International Conference on Computer Vision*, Oct. 2017, pp. 2223–2232, doi: 10.1109/ICCV.2017.244.

[4] S. B. U. Meeran and N. A. Munaf, "Comparative analysis of unidirectional and bidirectional RNNs for ECG arrhythmia detection using augmented MIT-BIH data," *Bulletin of Electrical Engineering and Informatics*, vol. 15, no. 1, pp. 712–722, Feb. 2026, doi: 10.11591/eei.v15i1.10893.

[5] P. Isola, J. -Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, 2017, pp. 5967-5976, doi: 10.1109/CVPR.2017.632.

[6] D. Nie *et al.*, "Medical Image Synthesis with Deep Convolutional Adversarial Networks," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 12, pp. 2720-2730, Dec. 2018, doi: 10.1109/TBME.2018.2814538.

[7] A. Chartsias *et al.*, "Disentangled representation learning in cardiac image analysis," *Medical Image Analysis*, vol. 58, p. 101535, Dec. 2019, doi: 10.1016/j.media.2019.101535.

[8] D. Nie *et al.*, "Medical image synthesis with context-aware generative adversarial networks," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 10435, pp. 417–425, 2017, doi: 10.1007/978-3-319-66179-7_48.

[9] A. Sharma and G. Hamarneh, "Missing MRI Pulse Sequence Synthesis Using Multi-Modal Generative Adversarial Network," *IEEE Transactions on Medical Imaging*, vol. 39, no. 4, pp. 1170-1183, April 2020, doi: 10.1109/TMI.2019.2945521.

[10] O. U. Aydin *et al., "*RELICT: A replica detection framework for medical image generation," *arXiv*, 2025, doi: 10.48550/arXiv.2502.17360.

[11] H. Emami, M. Dong, and C. K. Glide-Hurst, "Attention-Guided Generative Adversarial Network to Address Atypical Anatomy in Synthetic CT Generation," in *2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI)*, Las Vegas, NV, USA, 2020, pp. 188-193, doi: 10.1109/IRI49571.2020.00034.

[12] Y. Zhao *et al.*, "Compensation cycle consistent generative adversarial networks (Comp-GAN) for synthetic CT generation from MR scans with truncated anatomy," *Medical Physics*, vol. 50, no. 7, pp.s 4399-4414, 2023, doi: 10.1002/mp.16246.

[13] J. Kumar, T. Tanvi, A. Ghosh, R. Mondal, and S. Das, "Generating Images from Text Using GANs vs Transformers*," 2025 8th International Conference on Electronics, Materials Engineering & Nano-Technology (IEMENTech)*, Kolkata, India, 2025, pp. 1-5, doi: 10.1109/IEMENTech65115.2025.10959693.

[14] S. Syed, R. Ahmed, and A. Iqbal, "Medical Imagining (CT scan, MRI, X-ray, and Microscopic Imagery) Data," *Mendeley Data*, vol. 3, 2024, doi: 10.17632/5kbjrgsncf.3.

[15] S. U. Dar, M. Yurt, L. Karacan, A. Erdem, E. Erdem, and T. Çukur, "Image Synthesis in Multi-Contrast MRI With Conditional Generative Adversarial Networks," *IEEE Transactions on Medical Imaging*, vol. 38, no. 10, pp. 2375-2388, Oct. 2019, doi: 10.1109/TMI.2019.2901750.

[16] Y.-F. Xia *et al.*, "Region-guided focal adversarial learning for CT-to-MRI translation: A proof-of-concept and validation study in hepatocellular carcinoma," *Medical Physics*, vol. 52, no. 5, pp. 2861-2873, 2025, doi: 10.1002/mp.17674.

[17] S, Pan *et al. "*Synthetic CT generation from MRI using 3D transformer-based denoising diffusion model," *Medical Physics* vol. 51, no. 4, pp. 2538-2548, 2024, doi: 10.1002/mp.16847.

[18] J. Ma, J. Chen, and A. Liang, "CT-to-MRI translation of medical volume data based on an enhanced diffusion model," *Scientific Reports*, vol. 16, 2026, doi: 10.1038/s41598-026-45181-1.

[19] X. Chen *et al*., "Anatomy-Regularized Representation Learning for Cross-Modality Medical Image Segmentation," *IEEE Transactions on Medical Imaging*, vol. 40, no. 1, pp. 274-285, Jan. 2021, doi: 10.1109/TMI.2020.3025133

[20] M. A. Bahloul, S. Jabeen, S. Benoumhani, H. A. Alsaleh, Z. Belkhatir, and A. Al-Wabil, "Advancements in synthetic CT generation from MRI: A review of techniques, and trends in radiation therapy planning," *Journal of Applied Clinical Medical Physics*, vol. 25, no. 11, p. e14499, 2024, doi: 10.1002/acm2.14499.

[21] H. A. Amirkolaee and H. A. Amirkolaee, "Medical image translation using an edge-guided generative adversarial network with global-to-local feature fusion," *Journal of Biomedical Research*, vol. 36, no. 6, pp. 409-422, doi: 10.7555/JBR.36.20220037.

[22] A. H. Farshchitabrizi, M. H. Sadeghi, S. Sina, M. Alavi, Z. N. Feshani, and H. Omidi, "AI-enhanced PET/CT image synthesis using CycleGAN for improved ovarian cancer imaging," *Polish Journal of Radiology*, vol. 90, pp. 26-35, 2025, doi: 10.5114/pjr/196804.

[23] R. K. Gebre *et al.*, "Cross-scanner harmonization methods for structural MRI may need further work: A comparison study," *Neuroimage.*, vol. 269, 2023, doi: 10.1016/j.neuroimage.2023.119912.

[24] J. Y. You *et al.*, "Preserving spatial and quantitative information in unpaired biomedical image-to-image translation," *Cell Reports Methods*, vol. 16, no. 6, p. 101074, 2025, doi: 10.1016/j.crmeth.2025.101074.

[25] S. Chilamkurthy *et al.*, "Deep learning algorithms for detection of critical findings in head CT scans: a retrospective study," *Lancet*, vol. 392, no. 10162, pp. 2388-2396, Dec. 2018, doi: 10.1016/S0140-6736(18)31645-3.

[26] K. Clark *et al.*, "The cancer imaging archive (TCIA): Maintaining and operating a public information repository," *Journal of Digital Imaging*, vol. 26, no. 6, pp. 1045–1057, 2013, doi: 10.1007/s10278-013-9622-7.

[27] A. L. Simpson *et al.*, "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," *arXiv*, 2019, doi: 10.48550/arXiv.1902.09063.

[28] S. S. R. M. R. and J. T., "Multi-Scale and Spatial Information Extraction for Kidney Tumor Segmentation: A Contextual Deformable Attention and Edge-Enhanced U-Net," *Journal of Imaging Informatics in Medicine*, vol. 37, no. 1, pp. 151-166, Jan. 2024, doi: 10.1007/s10278-023-00900-2.

[29] S. B. Mohan, S. Sathya, S. Rajalakshmi, G. Gurumoorthy, and R. Sivanraju, "Efficient convolutional neural networks for acute lymphoblastic leukaemia prediction in computer vision," *Scientific Reports*, vol. 16, Dec. 2025, doi: 10.1038/s41598-025-30777-w.

[30] J. C. and S. J. C. Sheeba, "Vision Unveiling Through End-to-End Barur-Net Approach," *IETE Journal of Research*, vol. 71, no. 6, 2025, doi: 10.1080/03772063.2025.2479761.

## BIOGRAPHIES OF AUTHORS

**Dr. Sabura Banu Urundai Meeran** was born in Chennai, Tamilnadu, India on May 8, 1975. She is presently working as a professor in the department of Electrical and Electronics Engineering, Saveetha Engineering College, Chennai, Tamilnadu, India. She has an illustrious academic and professional background in engineering, innovation, and consultancy. She completed her Ph.D. in Intelligent Modeling and Control at College of Engineering, Guindy, Anna University (2006–2009) after earning her M.S. by research in the same Institution and B.E. in Instrumentation and Control Engineering. She also holds a Diploma in Computer Science Engineering and a B.Sc. in Mathematics. With over two decades of experience, she has served in various leadership roles, including Director at Arcres Techsol Pvt. Ltd., Founder & CEO, Intellectz Tech for All and Head of Departments at esteemed institutions like SRM Institute of Science and Technology and B.S. Abdur Rahman University. She has mentored numerous student projects that garnered national and international accolades, including awards, patents, and startup funding. Her notable achievements include 34 international journal publications in leading journals like Taylor and Francis, Applied Soft Computing, ACTA Press, AMSE, and Inderscience, contributions to five book chapters, and 132 conference presentations. She can be contacted at email: saburabanuu@saveetha.ac.in, sabura_banu@rediffmail.com, or saburasec@gmail.com.

**Dr. Nafeena Abdul Munaf** is Assistant Professor and Coordinator for Electrical and Electronics Engineering program. Villa College, Maldives. Besides that, she has published many papers in Scopus, Web of Science, Science Indexed, and UGC journals. Her recent published paper in IEEE got 49 citations on the title' An ML approach for Household, power consumption. Currently she is a reviewer in reputed international and national journals and editor in special issues of conference publications. Also, she has given many keynote lectures in reputed colleges and universities in Smart Grid Challenges and Towards Sustainable University Residences: Comprehensive Energy Audit, Retrofit Techniques and Financial Implications Sustainability Issues in Environment. She is interested in writing articles and publishing on life cycle analysis of various products and smart energy management systems for smart grid. Also, she has received grants on SDG based researches. She can be contacted at email: nafeena.abdul@villacollege.edu.mv.