

Adaptive Latent Diffusion Model for 3D Medical Image to Image Translation: Multi-modal Magnetic Resonance Imaging Study

Jonghun Kim, Hyunjin Park [†]

Department of Electrical and Computer Engineering
 Sungkyunkwan University, Suwon, Korea

iproj2@g.skku.edu, hyunjinp@skku.edu

Abstract

Multi-modal images play a crucial role in comprehensive evaluations in medical image analysis providing complementary information for identifying clinically important biomarkers. However, in clinical practice, acquiring multiple modalities can be challenging due to reasons such as scan cost, limited scan time, and safety considerations. In this paper, we propose a model based on the latent diffusion model (LDM) that leverages switchable blocks for image-to-image translation in 3D medical images without patch cropping. The 3D LDM combined with conditioning using the target modality allows generating high-quality target modality in 3D overcoming the shortcoming of the missing out-of-slice information in 2D generation methods. The switchable block, noted as multiple switchable spatially adaptive normalization (MS-SPADE), dynamically transforms source latents to the desired style of the target latents to help with the diffusion process. The MS-SPADE block allows us to have one single model to tackle many translation tasks of one source modality to various targets removing the need for many translation models for different scenarios. Our model exhibited successful image synthesis across different source-target modality scenarios and surpassed other models in quantitative evaluations tested on multi-modal brain magnetic resonance imaging datasets of four different modalities and an independent IXI dataset. Our model demonstrated successful image synthesis across various modalities even allowing for one-to-many modality translations. Furthermore, it outperformed other one-to-one translation models in quantitative evaluations. Our code is available at <https://github.com/jongdory/ALDM/>

[†]Corresponding Author.

This study was supported by National Research Foundation (NRF-2020M3E5D2A01084892), Institute for Basic Science (IBS-R015-D1), ITRC support program (IITP-2023-2018-0-01798), AI Graduate School Support Program (2019-0-00421), ICT Creative Consilience program (IITP-2023-2020-0-01821), and the Artificial Intelligence Innovation Hub program (2021-0-02068).




Figure 1. Overview of the proposed image-to-image translation process based on latent diffusion model. Our model utilizes the proposed MS-SPADE to transform the latent representation into the target latents and enables image synthesis in the desired target modality through conditioning.

1. Introduction

Multi-modal images are crucial for a comprehensive evaluation in medical image analysis. To identify clinically important biomarkers, multi-modal 3D images with complementary information are often required. Multi-modal imaging serves as a method to compensate for the limitations of individual imaging techniques and enables precise examination [25, 36, 38]. Medical image analysis, especially diagnosis of brain tumors, is typically performed using magnetic resonance imaging (MRI) and various MRI techniques such as 3D T1-weighted, T1 contrast-enhanced weighted, T2-weighted, and FLAIR are used, because the distinct modalities (T1, FLAIR, etc.) are different imaging sequences that provide varying contrasts and provide complementary information (e.g., edema, enhancing tumor, and necrosis/non-enhancing tumor) about brain tumor that appears over multiple consecutive slices requiring 3D approach. [2, 3, 27]. T1 is sensitive to acute bleeding and clearly displays contrast agents, helping identify vascular structures, tumors, and inflammation areas. T2 highlight conditions like edema or inflammation. FLAIR suppresses cerebrospinal fluid signals and is useful for detecting pathologies such as multiple sclerosis lesions [3, 27, 29, 37]. This is an important application of computer vision in medicine. However, in clinical settings, obtaining multiple modalities can be challenging due to factors such as scan costs, limited scan time, and safety considerations. Consequently, certain modalities may be missing. This absence of

modalities can have a detrimental impact on the quality of diagnosis and treatment. Furthermore, deep learning models that rely on multi-modal also suffer from reduced performance when crucial modalities are unavailable for training.

Many effective deep learning methods have been proposed to address the information gap caused by missing modalities [6, 11, 16, 23, 24, 40, 44]. One approach is to train a single model that learns the shared representation of multi-modalities to handle various missing scenarios [6, 24, 40, 44]. For example, with three out of four MRI modalities missing, we must consider 12 (3×4) missing scenarios. The number of scenarios grows quickly as the number of modalities grows. However, these methods, limited to specific tasks, struggle to generalize across diverse scenarios, not offering a fundamental solution.

Another approach is to generate missing modality images using generative models. Among the existing methods for generating missing modalities, generative adversarial models (GANs) [7, 11, 15, 19, 23, 46] have been used for medical image-to-image translation. However, these methods are 2D and thus do not consider the information from adjacent slices in 3D medical images. They generate slices independently making it difficult to ensure consistent information between adjacent slices. 3D methods for handling medical imaging data are computationally expensive and challenging to apply. To address this issue, a common approach is to train the model using patches during the training stage and generate results using a sliding window method during the inference stage. However, these methods may result in the loss of global information and a potential decrease in image translation performance. Furthermore, existing methods typically generate a given target modality from one or many source(s) (i.e., one-to-one or many-to-one) [11, 23, 45]. To generate multiple targets from one source modality, these methods need a separate model for each target, increasing complexity. Therefore, there is a need for a model that can perform image translation using 3D medical imaging without patch cropping and allow generating flexible targets from a single source modality.

In recent years, the diffusion model (DM) has emerged as a powerful generative model with high-quality generation capabilities, positioning itself as a potential alternative to GANs [9, 12, 17, 31, 34, 35]. The diffusion model leverages cross-attention and flexible conditioning to enable the generation of desired images. Additionally, the latent diffusion model (LDM) defines a denoising task in the compressed latent space of an image significantly reducing computational costs while demonstrating the ability to generate high-quality images at a lower cost [34]. LDM demonstrated its applicability to 3D medical images without the need for patch cropping. This advancement allows for the application of LDM to medical images even in their native 3D form [21, 31].

In this study, we propose a model that utilizes the conditioning mechanism of LDM to perform image-to-image translation in 3D medical images without patch cropping maintaining the original size. Additionally, our proposed model enables the generation of multiple target modalities from a single source modality. We achieve this by the proposed multiple-switchable block, noted as MS-SPADE block, which dynamically transforms source latents to target-like latents according to style. By utilizing conditioning in LDM, we can generate the desired target modality more faithfully. The pipeline of our model is illustrated in Figure 1.

Our main contributions are summarized as follows:

- We propose a model based on the LDM that translates a single source modality to various target modalities (one-to-many) in 3D medical images.
- We introduce a switchable block that transfers the source latents to target-like latents using style transfer to enhance the performance of image-to-image translation.
- We validated our method on the BraTS2021 and IXI datasets and despite our model's ability to generate multiple modalities, it demonstrated the highest image translation performance on both datasets.

2. Related Works

Generative adversarial networks have achieved great success in image-to-image translation by performing continuous adversarial training between the generator and discriminator [7, 15, 19, 23, 46]. The generator aims to generate as realistic images as possible, while the discriminator learns to distinguish between the generated images and real images [15]. For instance, Pix2Pix [19] focuses on pixel-to-pixel image synthesis based on paired data enhancing the pixel-to-pixel similarity between real and synthesized images. CycleGAN [46] is a generalized conditional GAN enabling image synthesis even with unpaired data. NICEGAN [7] proposes a more compact architecture by sharing the encoder parts of the generator and discriminator. compared to conventional GAN. RegGAN [23] adds a registration network to perform image-to-image translation while maintaining anatomical structures in medical images. Ea-GAN [5] is a 3D-based network designed to enhance cross-modality MR image synthesis by integrating edge information, capturing both voxel-wise intensity and image structure details. ResViT [11] is an adversarial network that incorporates a vision transformer architecture to capture contextual features in medical image translation tasks. However, these methods all work well in situations where there is a one-to-one mapping between the source and target modalities. To generate multiple target modalities from a single source modality, a separate model is needed for each target modality resulting in an increased number of models.

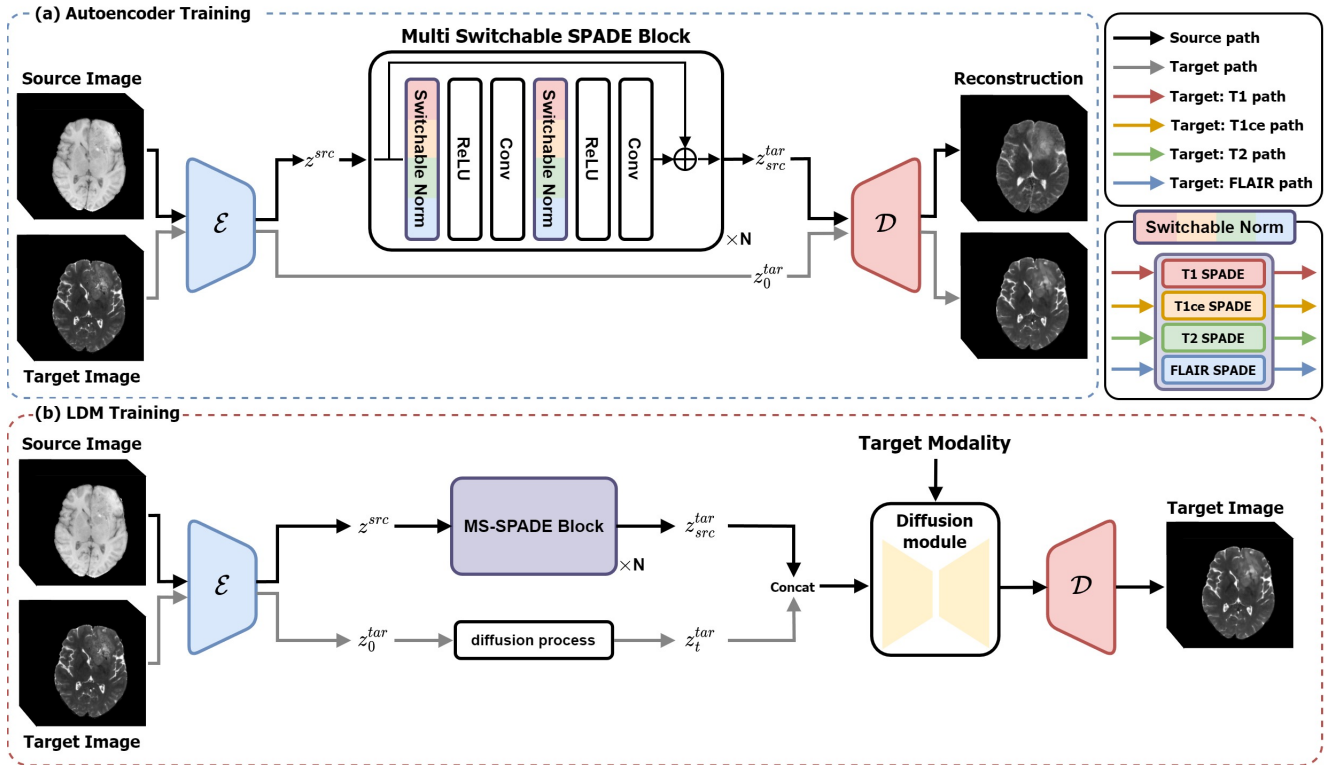


Figure 2. The source and target images are 3D volume images and our method is applied in a 3D manner. In (a), we depict the autoencoder to compute source latents (z^{src}) and the SPADE block to dynamically convert source latents to target-like latents ($z^{src} \rightarrow z_{src}^{tar}$). The switchable block includes normalization layers that are applied differently depending on the target modality allowing translation to target multiple modalities within a single model. The output of the SPADE block is noted as target-like to emphasize the subsequent role of LDM to fully predict the target latents. The block is stacked N times. (b) illustrates the training process of the LDM model to predict the target latents z^{tar} from the target-like latents z_{src}^{tar} obtained from the SPADE block.

Conditional normalization layers are techniques used for style transfer and image synthesis tasks [13, 18, 28, 30]. Among them, adaptive instance normalization (AdaIN) [18] was initially used for style transfer tasks and has since been adopted in various computer vision tasks. Unlike other normalization techniques, conditional normalization layers involve normalizing layer activations to zero mean and unit variance, followed by denormalization through learned affine transformations from external data. Spatially adaptive normalization (SPADE) [28] has achieved successful image synthesis by applying spatially varying affine transformations to semantic maps. Unlike previous AdaIN, SPADE generates images that contain both style and semantic information allowing the creation of meaningful images using semantic maps, which is well applicable to 3D multi-modal medical imaging. Inspired by previous research, we adopt the SPADE module that transforms the style and semantic information of an image. Additionally, we propose the multi-switchable SPADE (MS-SPADE) module, which allows for flexible translation to various target modalities.

Diffusion probabilistic model has recently gained prominence in image generation delivering state-of-the-art results [8, 9, 12, 17, 31, 35]. Specifically, the denoising diffusion probabilistic model (DDPM) [17] with score matching has demonstrated superior performance in image generation. DDPM leverages a learned Markov chain to transform Gaussian noise distribution into the target distribution. Palette [35] introduces a model that utilizes DM by concatenating source and target images enabling various image-to-image translation tasks such as inpainting, uncropping, image restoration, and colorization without the need for architecture modifications. This model achieves high-fidelity results across different tasks. LDM [34] separates the training process into two stages for efficient processing. First, an autoencoder is trained to provide a perceptually equivalent lower-dimensional representation (e.g., source latents). Subsequently, DM is trained in this latent space reducing complexity and enabling efficient image generation. Importantly, this approach allows for efficient training without the need for patch cropping in 3D images.

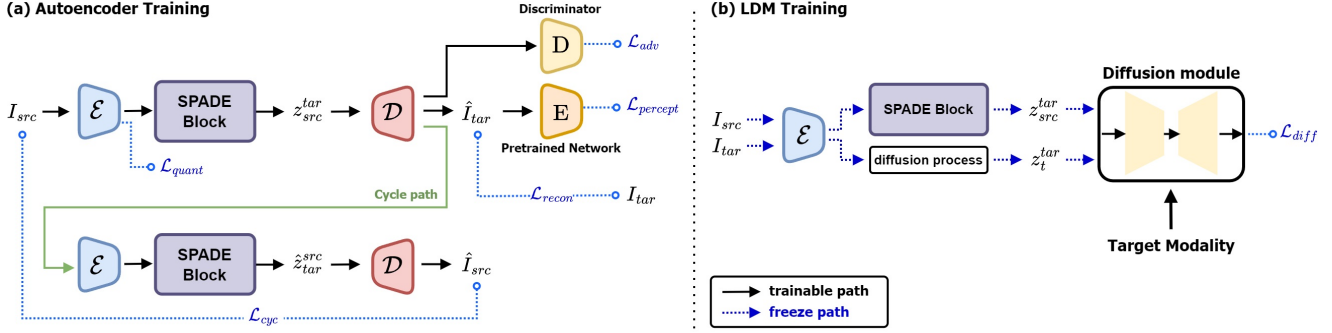


Figure 3. (a) illustrates the training process of the autoencoder during the image compression phase to compute source latents, which involves five different losses: reconstruction, quantization, adversarial objective, perceptual, and cycle consistency. (b) shows the training process of the diffusion model, where the autoencoder is frozen, and only the UNet of the diffusion model is trained. The forward diffusion process is fixed and the input to the UNet consists of the concatenated two latents.

3. Method

To overcome the limitations of single-stage generation approaches in recent generative models, numerous studies [10, 14, 32, 33, 42] have adopted a two-stage approach to combine the strengths of different methods resulting in the development of models with superior performance. Our method trains image-to-image translation in a two-stage process as depicted in Figure 2. In the first stage, we train an autoencoder, inspired by previous research methods such as VQ-VAEs and VQGAN [14, 33, 39], to compress the image into a latent representation that captures its essential features. The latent space, which captures the perceptual representation of the original images, is regularized via vector quantization and allows for style transfer through conditional normalization layers [28]. We also train SPADE block to transform the style of the latents into those of the target modality (as shown in Figure 2 (a)). The proposed MS-SPADE block learns the style parameter of the target and transfers the style from the source to the target. While translation is possible using only the SPADE module, there is a noticeable difference between the true target latents and those obtained through SPADE and thus we note the output of SPADE as target-like latents. To address this discrepancy, we utilize the LDM to reduce the gap via denoising between these two latents, which is the key idea in our model. Subsequently, we combine the target-like latents (out of MS-SPADE) and noisy target latents (out of forward diffusion), similar to the palette approach, to train the DM (as shown in Figure 2 (b)). Our model prioritizes minimizing changes in the source image by first adjusting its latent style and then focusing on generating target-specific latents.

3.1. Image compression to compute latents

Similar to previous work [14, 31, 34], we employ a perceptual compression model to facilitate the application of the diffusion model in the latent space. We also introduce

the MS-SPADE block in the latent space to convert the source latents to the target-like latents. The switchable normalization is illustrated in Figure 2 (a). Our SPADE block learns the mean and standard deviation parameters of the target during training. Thus, during inference, it does not require the target as input but instead uses the target’s style parameters for normalization.

The MS-SPADE block performs different normalizations according to the target modality enabling the transformation of the transformed latents to follow the desired distribution. Let $h \in \mathbb{R}^{N \times C \times H \times W \times D}$ be the input of SPADE Block, where N, C, H, W and D are the size of batch, channel, height, width and depth respectively. The SPADE can be applied as follows:

$$x_{n,c,h,w,d} = \gamma_{c,h,w,d}^{tar}(h) \frac{h_{n,c,h,w,d} - \mu_c}{\sigma_c} + \beta_{c,h,w,d}^{tar}(h) \quad (1)$$

where $x_{n,c,h,w,d}$ denotes the n, c, h, w, d -th element of the feature tensor h , μ_c and σ_c are the mean and standard deviation of the latent features in channel c , and $\gamma_{c,h,w,d}^{tar}$ and $\beta_{c,h,w,d}^{tar}$ are modulation parameters that are learned and applied differently according to the target modality during training.

Furthermore, our compression model is based on VQGAN [14] and is trained using reconstruction loss, quantization loss, perceptual loss based on previous research and patch-based adversarial objectives [14, 22, 39]. Additionally, we incorporate cycle consistency loss to ensure that the original input can be reconstructed back to the original image through the cycle path [46]. The training process is illustrated in Figure 3 (a).

3.2. Diffusion Model

LDM achieves high-resolution image synthesis while reducing computational cost by training the diffusion process in the compressed latent space [34]. Additionally, Palette demonstrates high-quality image synthesis in tasks such as

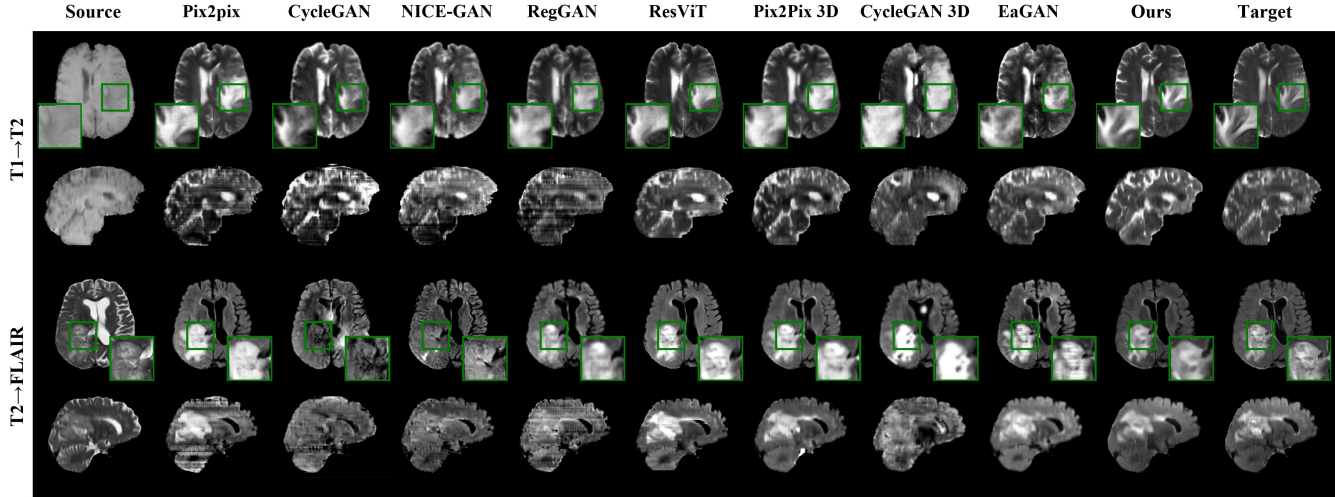


Figure 4. The figures illustrate the results of the proposed model and comparison models on the BraTS2021 dataset for the image-to-image translation tasks (top: T1 → T2, bottom: T2 → FLAIR). These two tasks are representative tasks out of 12 possible tasks. The results are depicted in both the axial (top) and sagittal (bottom) views of the 3D volume. The green bounding boxes highlight the important tumor regions in the synthesized images providing a detailed visual representation.

colorization and inpainting reconstruction without relying on conditioning [35]. Instead, it simply takes the target image as an additional input to the DM resulting in superior results. We employ trained perceptual compression models, which consist of an encoder \mathcal{E} and a decoder \mathcal{D} , to apply the diffusion model in the latent space [34]. Inspired by previous studies [34, 35], we design an LDM that takes the target-like latents besides noisy target latents as inputs. This process is illustrated in Figure 2 (b). The latent z_{src}^{tar} represents the source latents transformed to be similar to the target latents by the SPADE block. It is denoted as target-like latents to emphasize the gap between true target latents and those from SPADE. z_t^{tar} denotes the noisy input latents at time step t . Our diffusion model takes z_{src}^{tar} and z_t^{tar} as inputs at time step t to perform the task of predicting the noise ϵ at that time step. During the training process, the auto-encoder used for image compression is frozen, and only the diffusion model is trained (as shown in Figure 3 (b)). The loss function is defined as follows:

$$\mathcal{L}_{diff} := \mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_{\theta}(z_{src}^{tar}, z_t^{tar}, t)\|_2^2 \right] \quad (2)$$

Our model’s neural backbone, denoted as $\epsilon_{\theta}(z_{src}^{tar}, z_t^{tar}, t)$, is implemented as a time-conditioned UNet [17, 34]. The forward diffusion process is fixed and z_{src}^{tar} and z_t^{tar} can be efficiently obtained from \mathcal{E} during training.

3.3. Modality Conditioning

Inspired by the successful utilization of cross-attention as a conditional image generator in previous studies [8, 34], we incorporate modality conditioning into our model to enhance the image translation to the target modal-

ity. Our approach involves converting the given modality into a one-hot vector y and using it as input for the key and value components of cross-attention during training. The cross-attention in the UNet operates as follows: $\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V$ with

$$Q = W_Q^{(i)} \cdot \varphi_i(z_t), K = W_K^{(i)} \cdot y, V = W_V^{(i)} \cdot y \quad (3)$$

Here, $\varphi_i(z_t) \in \mathbb{R}^{N \times d_e}$ denotes a intermediate representation of the UNet implementing ϵ_{θ} and $W_Q^{(i)} \in \mathbb{R}^{d \times d_e}$, $W_K^{(i)} \in \mathbb{R}^{d \times d_y}$ and $W_V^{(i)} \in \mathbb{R}^{d \times d_y}$ are learnable projection matrices [20, 34, 41]. We redefine the loss function with modality conditioning, expressed as \mathcal{L}_{diff} with y representing the one-hot vector for the given modality class.

$$\mathcal{L}_{diff} := \mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_{\theta}(z_{src}^{tar}, z_t^{tar}, t, y)\|_2^2 \right] \quad (4)$$

For the inference process, we utilize the encoder acquired from the image compression phase. We obtain a target-similar latent by passing the source image through the encoder and the MS-SPADE block. This latent is then concatenated with random noise and used as the conditioning for the diffusion model, facilitating the image translation into the target modality.

4. Experiments

4.1. Datasets

To evaluate the effectiveness of our model, we utilized the multi-modal brain tumor segmentation challenge 2021 (BraTS 2021) dataset [1–3, 27]. We trained our model using the BraTS 2021 training dataset, which consists of 1251




Figure 5. The figures illustrate the results of the proposed model and comparison models on the IXI dataset.

Datasets		BraTS 2021 [1]						IXI [4]		
Tasks		T1 → T2			T2 → FLAIR			T1 → PD		
Dimension	Model	PSNR ↑	NMSE ↓	SSIM ↑	PSNR ↑	NMSE ↓	SSIM ↑	PSNR ↑	NMSE ↓	SSIM ↑
2D	Pix2Pix	24.624±0.962	0.109±0.028	0.874±0.015	24.361±1.061	0.117±0.021	0.846±0.019	25.283±0.861	0.095±0.028	0.882±0.023
	CycleGAN	23.535±1.334	0.155±0.035	0.837±0.028	23.418±0.944	0.164±0.033	0.825±0.035	24.854±1.128	0.106±0.033	0.841±0.027
	NICEGAN	23.721±1.136	0.148±0.029	0.840±0.029	23.643±1.045	0.148±0.022	0.829±0.033	25.330±0.939	0.102±0.027	0.855±0.026
	RegGAN	24.884±0.991	0.094±0.024	0.881±0.017	24.576±1.073	0.112±0.022	0.852±0.028	25.849±0.921	0.087±0.026	0.898±0.018
	ResViT	25.578±0.812	0.088±0.021	0.895±0.018	24.825±1.030	0.108±0.018	0.861±0.021	27.265±0.847	0.076±0.022	0.916±0.021
3D	Pix2Pix	25.181±0.861	0.097±0.031	0.887±0.012	24.602±1.181	0.113±0.021	0.854±0.018	26.387±0.849	0.089±0.032	0.903±0.019
	CycleGAN	23.740±1.198	0.138±0.032	0.835±0.019	23.508±1.301	0.152±0.039	0.822±0.024	25.507±1.133	0.099±0.035	0.882±0.028
	EaGAN	24.884±0.991	0.094±0.024	0.881±0.017	24.576±1.073	0.112±0.022	0.852±0.028	26.849±0.836	0.082±0.021	0.912±0.022
	Ours	25.818±0.857	0.079±0.016	0.904±0.012	25.074±1.085	0.098±0.021	0.867±0.018	27.729±0.885	0.068±0.022	0.921±0.018

Table 1. The values represent the quantitative evaluation results of various comparison models and the proposed model in three quantitative metrics (PSNR, NMSE, and SSIM). The reported format is mean followed by standard deviation. The evaluation metrics were calculated on the 3D volume for both the 2D and 3D methods. Boldface entries indicate the top-performing model for each task.

subjects and includes four MRI modalities (T1, T1ce, T2, FLAIR). For evaluating the image translation capability of our model, we employed the BraTS 2021 validation dataset, which comprises 219 subjects. We also validated our model on the IXI dataset [4], which includes T1, T2, and PD modalities. Out of 574 subjects, 459 were used for training and 115 for testing. For details on the dataset, please refer to the supplementary.

4.2. Implementation Details

The image compression in our model was implemented using VQGAN [14] and the diffusion model was based on LDM [34]. The implementation was done using PyTorch¹ and MONAI² libraries. For detailed information on the model architecture, refer to the supplementary materials. As for the hyperparameters, we used the AdamW [26] optimizer with an learning rate of 2×10^{-6} . we set the number of time steps $T=1000$ with scaled-linearly scheduled noise levels ranging from 0.0015 to 0.0195. The training was conducted on four A100 80GB GPUs with a batch size of 1 per GPU. Additional information on the training hyperparameters can be found in the supplementary materials.

¹<https://pytorch.org/>

²<https://monai.io/>

4.3. Evaluation Metrics

We compared the performance of methods with previous approaches commonly used in medical image translation research [43]. The synthesis quality was evaluated using peak signal-to-noise ratio (PSNR), normalized mean squared error (NMSE), and structural similarity index (SSIM) [43]. They were computed between the ground truth images and the synthesized target images. The average and standard deviation of the metrics were reported on an independent test set that did not overlap with the training-validation sets. The evaluation was performed on 3D volumes for all tests. For the 2D methods, the synthesized target images were stacked to form a 3D volume for comparison.

5. Results

5.1. Comparison of Ours to Baselines

Table 1 presents the quantitative evaluations of each method for three tasks: T1 → T2, T2 → FLAIR, and T1 → PD, across two different datasets. These tasks are representative tasks because T1 and T2 are routinely available, while the others are less common. Despite other models performing only one-to-one image synthesis tasks, our model is capable of one-to-many image synthesis and

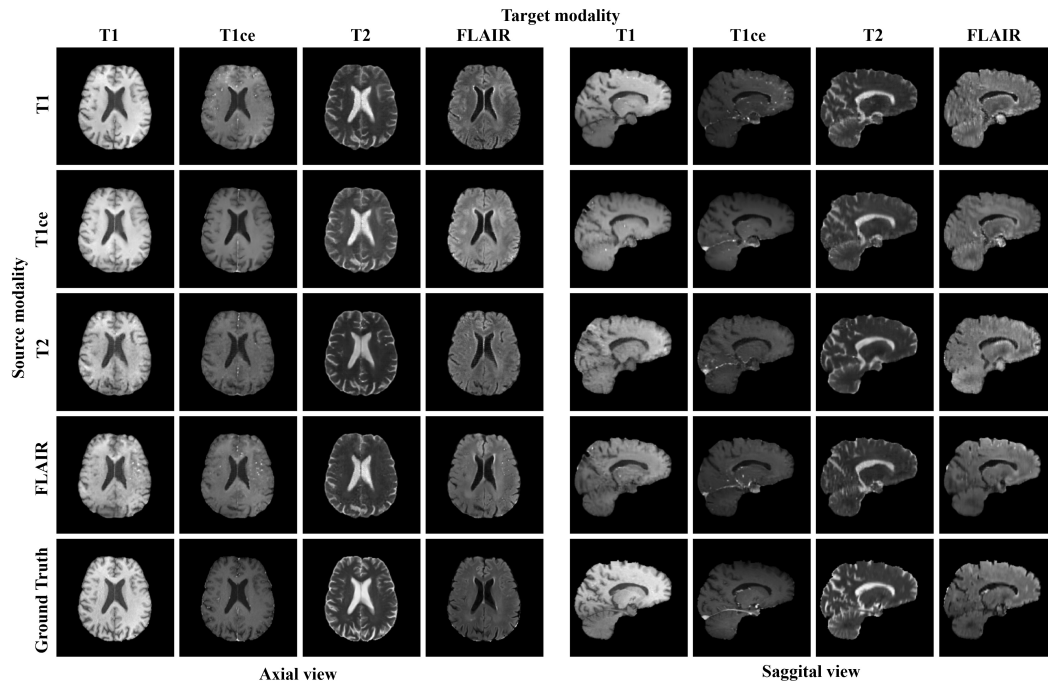


Figure 6. The figures showcase the image translation results from each source modality to the corresponding target modality using our proposed model for all possible combinations. Each row corresponds to a source modality, each column represents a target modality and the last row represents the ground truth.

Source \ Target	T1			T1ce			T2			FLAIR		
	PSNR \uparrow	NMSE \downarrow	SSIM \uparrow	PSNR \uparrow	NMSE \downarrow	SSIM \uparrow	PSNR \uparrow	NMSE \downarrow	SSIM \uparrow	PSNR \uparrow	NMSE \downarrow	SSIM \uparrow
T1	29.001	0.055	0.942	26.119	0.078	0.912	25.818	0.103	0.904	24.842	0.113	0.859
	± 0.643	± 0.025	± 0.022	± 0.816	± 0.022	± 0.015	± 0.857	± 0.030	± 0.014	± 0.728	± 0.034	± 0.019
T1ce	26.228	0.076	0.922	28.759	0.060	0.937	25.990	0.092	0.907	25.204	0.092	0.881
	± 0.794	± 0.027	± 0.033	± 0.885	± 0.019	± 0.015	± 0.859	± 0.032	± 0.908	± 0.811	± 0.050	± 0.037
T2	25.422	0.085	0.908	25.234	0.087	0.895	29.230	0.048	0.942	25.074	0.098	0.867
	± 0.852	± 0.026	± 0.020	± 1.152	± 0.034	± 0.025	± 0.720	± 0.018	± 0.915	± 1.085	± 0.021	± 0.018
FLAIR	25.186	0.090	0.905	25.899	0.094	0.906	26.146	0.086	0.913	28.608	0.058	0.938
	± 0.759	± 0.028	± 0.048	± 1.039	± 0.025	± 0.027	± 0.636	± 0.028	± 0.944	± 0.769	± 0.025	± 0.028

Table 2. The values present the quantitative evaluation of image translation results from source modalities to target modalities using our proposed model. The diagonal entries denote reconstructing from latents and are presented for completeness.

it showed superior performance in quantitative evaluations compared to other models. Moreover, in Figure 4, which depicts qualitative evaluations, our model achieved the closest resemblance to the ground truth and synthesized images that closely matched the challenging tumor regions (green bounding box). Additionally, in the sagittal view, while 2D methods exhibited striped patterns due to their slice-wise inference without considering neighboring slices, our 3D method, including our proposed method, overcame this limitation and demonstrated improved results. These results demonstrate that our proposed model not only allows the generation of multiple modalities from a single modality but also outperforms networks designed for a single task. Furthermore, it shows the successful synthesis of 3D volume images without the need for preprocessing steps such

as patch cropping. Figure 5 shows the evaluations in the IXI dataset.

5.2. Multi Modalities Image Translation

We conducted image translation experiments using our proposed model with different scenarios of translating one source modality to another target modality given four MRI modalities. Figure 6 visualizes the results of image translation tasks from each source modality to the target modality demonstrating successful image synthesis in all cases for the BraTs dataset. The diagonal entries represent tasks where the source and target modalities are the same. The quantitative evaluation of each task in Table 2 reveals the performance of our proposed model for the BraTs dataset. When the target is T1, the best translation performance is




Figure 7. The figures depict the results of the T1 \rightarrow T2 task for each ablation method on the BraTS2021 dataset. Each method represented in the figure corresponds to the methods listed in Table 3

Method	Diff	Palette	SPADE	Reg-type	Metric		
					PSNR \uparrow	NMSE \downarrow	SSIM \uparrow
(a)	✓			VQ	24.902 \pm 1.073	0.097 \pm 0.036	0.876 \pm 0.022
(b)	✓	✓		VQ	25.364 \pm 0.923	0.090 \pm 0.031	0.893 \pm 0.024
(c)			✓	VQ	25.037 \pm 1.019	0.101 \pm 0.026	0.884 \pm 0.032
(d)	✓	✓	✓	KL	25.245 \pm 0.792	0.095 \pm 0.022	0.891 \pm 0.017
ours	✓	✓	✓	VQ	25.818\pm0.857	0.079\pm0.016	0.904\pm0.012

Table 3. The values present the quantitative evaluation metrics of the T1 \rightarrow T2 task for different ablated models: diffusion model, palette, MS-SPADE block, and image compression model for the BraTs dataset. The results are shown and the values in bold indicate the highest performance for each task. Reg-type is the regularization type.

achieved when the source is T1ce. Similarly, for the target of T1ce, the source of F1 yields the highest performance. When the target is T2, using FLAIR as the source produces the best results. Lastly, when the target is FLAIR, the translation performance is highest when the source is T1ce. T1ce, being an MRI with contrast agent administration, contains abundant information about anatomy and tumors compared to other modalities, which explains its superior performance as a source modality in image translation. However, obtaining T1ce, which requires injecting contrast agents can be challenging in many situations. Therefore, our results demonstrate the feasibility of utilizing alternative modalities for image translation when obtaining T1ce is difficult.

5.3. Ablation Study

To demonstrate the incremental value of design components in our model, we conducted several lines of ablation studies. Specifically, we evaluated the performance contribution of the diffusion model, palette module, MS-SPADE block, and image compression model focusing on the T1 \rightarrow T2 task using the BraTS dataset. Table 3 presents the quantitative evaluation results, while Figure 7 illustrates the image synthesis results of each method in the abla-

tion study for the BraTs dataset. In Table 3, (a) represents the performance of the translation when using the diffusion model alone demonstrating that image translation is achievable. However, (b) with the addition of the palette shows superior performance. Additionally, (c) shows the results of conducting style transfer using only the MS-SPADE block without the diffusion model indicating limitations in image translation performance. (d) and "ours" represent the results with different regularization types of the image compression model applied in previous studies [34]. KL-reg, although computationally expensive for 3D training, had an impact on performance due to the adjustment of model size and resulting in slightly blurry compression results. It was observed that the best translation performance was achieved when all the elements were used together.

6. Conclusion

Our method is an important application of computer vision in medicine. We propose a model for multi-modal image translation and conducted comprehensive experiments to evaluate its performance. Our model demonstrated successful image synthesis across different source and target modalities showcasing its versatility. By comparing our method with existing approaches, our method outperformed them in both quantitative and qualitative evaluations. Furthermore, our model showcased the ability to perform one-to-many image synthesis, going beyond the limitations of one-to-one tasks performed by other models. We also showed that our model achieved excellent performance even without pre-processing steps such as patch cropping and could successfully perform image synthesis on 3D medical images. Due to the 3D nature of the proposed approach, our method may be computationally expensive. Future work includes translating from more than two source modalities to a chosen target modality and validating on other medical imaging such as computed tomography.

References

- [1] Ujjwal Baid, Satyam Ghodasara, Suyash Mohan, Michel Bilello, Evan Calabrese, Errol Colak, Keyvan Farahani, Jayashree Kalpathy-Cramer, Felipe C Kitamura, Sarthak Pati, et al. The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314*, 2021. 5, 6
- [2] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin Kirby, John Freymann, Keyvan Farahani, and Christos Davatzikos. Segmentation labels and radiomic features for the pre-operative scans of the tcga-lgg collection. *The cancer imaging archive*, 286, 2017. 1, 5
- [3] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1):1–13, 2017. 1, 5
- [4] Biomedical Image Analysis Group Imperial College London Centre for the Developing Brain King’s College London. Information eXtraction From Images. <https://brain-development.org/ixi-dataset/>, 2018. Accessed: December 15, 2022. 6
- [5] Yu Biting, et al. Ea-gans: edge-aware generative adversarial networks for cross-modality mr image synthesis. *IEEE transactions on medical imaging*, 38(7):1750–1762, 2019. 2
- [6] Cheng Chen, Qi Dou, Yueming Jin, Hao Chen, Jing Qin, and Pheng-Ann Heng. Robust multimodal brain tumor segmentation via feature disentanglement and gated fusion. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III 22*, pages 447–456. Springer, 2019. 2
- [7] Runfa Chen, Wenbing Huang, Binghui Huang, Fuchun Sun, and Bin Fang. Reusing discriminators for encoding: Towards unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8168–8177, 2020. 2
- [8] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021. 3, 5
- [9] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2, 3
- [10] Bin Dai, David Wipf, and Minhao Jiang. Diagnosing and enhancing vae models (iclr 2019). 4
- [11] Onat Dalmaz, Mahmut Yurt, and Tolga Çukur. Resvit: residual vision transformers for multimodal medical image synthesis. *IEEE Transactions on Medical Imaging*, 41(10):2598–2614, 2022. 2
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 2, 3
- [13] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*, 2016. 3
- [14] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 4, 6
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 2
- [16] Mohammad Havaei, Nicolas Guizard, Nicolas Chapados, and Yoshua Bengio. Hemis: Hetero-modal image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, pages 469–477. Springer, 2016. 2
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3
- [18] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. 3
- [19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2
- [20] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021. 5
- [21] Firas Khader, Gustav Müller-Franzes, Soroosh Tayebi Arasteh, Tianyu Han, Christoph Haarbuerger, Maximilian Schulze-Hagen, Philipp Schad, Sandy Engelhardt, Bettina Baeßler, Sebastian Foersch, et al. Denoising diffusion probabilistic models for 3d medical image generation. *Scientific Reports*, 13(1):7303, 2023. 2
- [22] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*, 2014. 4
- [23] Lingke Kong, Chenyu Lian, Detian Huang, zhenjiang li, Yanle Hu, and Qichao Zhou. Breaking the dilemma of medical image-to-image translation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 1964–1978. Curran Associates, Inc., 2021. 2
- [24] Rongjian Li, Wenlu Zhang, Heung-Il Suk, Li Wang, Jiang Li, Dinggang Shen, and Shuiwang Ji. Deep learning based imaging data completion for improved brain disease diagnosis. In *Medical Image Computing and Computer-Assisted*

- Intervention–MICCAI 2014: 17th International Conference, Boston, MA, USA, September 14–18, 2014, Proceedings, Part III 17*, pages 305–312. Springer, 2014. [2](#)
- [25] Cirong Liu, Q Ye Frank, Cecil Chern-Chyi Yen, John D Newman, Daniel Glen, David A Leopold, and Afonso C Silva. A digital 3d atlas of the marmoset brain based on multi-modal mri. *Neuroimage*, 169:106–116, 2018. [1](#)
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. [6](#)
- [27] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014. [1](#), [5](#)
- [28] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019. [3](#), [4](#)
- [29] Sohil H Patel, Laila M Poisson, Daniel J Brat, Yueren Zhou, Lee Cooper, Matija Snuderl, Cheddhi Thomas, Ana M Franceschi, Brent Griffith, Adam E Flanders, et al. T2–flair mismatch, an imaging biomarker for idh and 1p/19q status in lower-grade gliomas: a tcga/tcia project. *Clinical cancer research*, 23(20):6078–6085, 2017. [1](#)
- [30] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [3](#)
- [31] Walter HL Pinaya, Petru-Daniel Tudosiu, Jessica Dafflon, Pedro F Da Costa, Virginia Fernandez, Parashkev Nachev, Sebastien Ourselin, and M Jorge Cardoso. Brain imaging generation with latent diffusion models. In *Deep Generative Models: Second MICCAI Workshop, DGM4MICCAI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 22, 2022, Proceedings*, pages 117–126. Springer, 2022. [2](#), [3](#), [4](#)
- [32] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. [4](#)
- [33] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. [4](#)
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [35] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. [2](#), [3](#), [5](#)
- [36] Yan Shen and Mingchen Gao. Brain tumor segmentation on mri with missing modalities. In *Information Processing in Medical Imaging: 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings 26*, pages 417–428. Springer, 2019. [1](#)
- [37] Gaurav Shukla, Gregory S Alexander, Spyridon Bakas, Rahul Nikam, Kiran Talekar, Joshua D Palmer, and Wenying Shi. Advanced magnetic resonance imaging in glioblastoma: a review. *Chin Clin Oncol*, 6(4):40, 2017. [1](#)
- [38] Tong Tong, Katherine Gray, Qinquan Gao, Liang Chen, Daniel Rueckert, Alzheimer’s Disease Neuroimaging Initiative, et al. Multi-modal classification of alzheimer’s disease using nonlinear graph fusion. *Pattern recognition*, 63:171–181, 2017. [1](#)
- [39] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. [4](#)
- [40] Gijs Van Tulder and Marleen de Bruijne. Why does synthesized data improve multi-sequence classification? In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part I 18*, pages 531–538. Springer, 2015. [2](#)
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [5](#)
- [42] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. : Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. [4](#)
- [43] Xin Yi, Ekta Walia, and Paul Babyn. Generative adversarial network in medical imaging: A review. *Medical image analysis*, 58:101552, 2019. [6](#)
- [44] Tongxue Zhou, Stéphane Canu, Pierre Vera, and Su Ruan. Brain tumor segmentation with missing modalities via latent multi-source correlation representation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part IV 23*, pages 533–541. Springer, 2020. [2](#)
- [45] Tao Zhou, Huazhu Fu, Geng Chen, Jianbing Shen, and Ling Shao. Hi-net: Hybrid-fusion network for multi-modal mr image synthesis. *IEEE Transactions on Medical Imaging*, 39(9):2772–2781, 2020. [2](#)
- [46] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#), [4](#)