

MAISI: Medical AI for Synthetic Imaging

Pengfei Guo^{1*} Can Zhao^{1*} Dong Yang^{1*} Ziyue Xu¹
 Vishwesh Nath¹ Yucheng Tang¹ Benjamin Simon² Mason Belue³
 Stephanie Harmon² Baris Turkbey² Daguang Xu¹
¹NVIDIA ²National Institutes of Health ³University of Arkansas for Medical Sciences

Abstract

Medical imaging analysis faces challenges such as data scarcity, high annotation costs, and privacy concerns. This paper introduces the Medical AI for Synthetic Imaging (MAISI), an innovative approach using the diffusion model to generate synthetic 3D computed tomography (CT) images to address those challenges. MAISI leverages the foundation volume compression network and the latent diffusion model to produce high-resolution CT images (up to a landmark volume dimension of $512 \times 512 \times 768$) with flexible volume dimensions and voxel spacing. By incorporating ControlNet, MAISI can process organ segmentation, including 127 anatomical structures, as additional conditions and enables the generation of accurately annotated synthetic images that can be used for various downstream tasks. Our experiment results show that MAISI's capabilities in generating realistic, anatomically accurate images for diverse regions and conditions reveal its promising potential to mitigate challenges using synthetic data.

1. Introduction

Medical imaging analysis has been integral to modern healthcare, providing critical insights into patient diagnosis, treatment planning, and monitoring. The rapid advancement of machine learning (ML) approaches has revolutionized diagnostic and therapeutic practices in modern healthcare. However, the development of effective ML models in this domain continues to face the following significant challenges [34, 40, 78]: (1) **data scarcity**: the rarity of certain medical conditions (e.g., certain types of cancer, and rare diseases) complicates the data acquisition process, which leads to the limited acquired data that might not adequately represent the diversity of real-world cases. (2) **high human-annotation costs**: annotating medical images, such as MRI

*Equal contribution. The code is available at [MONAI Tutorial](#). The online demo is available at [NVIDIA NIM](#).




Figure 1. (a) A generated high-resolution CT volume (with volume dimensions of $512 \times 512 \times 768$ and voxel spacing of $0.86 \times 0.86 \times 0.92 \text{ mm}^3$) by the proposed method and its corresponding segmentation condition overlaid on generated volume. We show the axial, sagittal, and coronal views from top to bottom, respectively. (b) 3D volume rendering of generated CT by MAISI. The rendering setting is tuned to highlight bone structures and demonstrate the realism of the generated CT volume.

and CT scans, is inherently more expertise-demanding than annotating objects in general images. Medical images often contain subtle features that are critical for accurate diagnosis and treatment. Expert knowledge is usually required to accurately identify and annotate these conditions. (3) **privacy concerns**: conventional data acquisition and process-

ing of medical images often require access to large volumes of patient data, which raises ethical concerns and poses significant logistical challenges due to the sensitive nature of patient information.

To address these limitations, generating synthetic data has emerged as a promising direction. By creating artificial yet realistic medical images, synthetic data can augment existing datasets, reduce the dependency on real patient data, and provide a cost-effective alternative to manual data annotation. With the recent advancement of the generative model, many novel approaches, such as generative adversarial networks (GAN) [21] and Diffusion Models (DM) [29], have been extensively studied for their capacity to generate photo-realistic images in various tasks in general computer vision society. In the context of medical image generation, several generative models have been successfully applied for medical image synthesis, such as multi-contrast MR/CT image synthesis [22, 33, 62], cross-modality image translation [9, 15, 57, 66, 76], and image reconstruction [14, 46, 65, 77].

However, several key challenges are not fully explored in previous studies. **First**, realistic high-resolution (larger than the volume dimension of 512^3) 3D volume generation is still a challenging task due to the huge memory consumption imposed by unified 3D frameworks, which must handle the vast amount of data involved in such high-dimensional representations [59]. Overcoming this memory bottleneck is essential for advancing the realism and applicability of 3D volume generation in clinical contexts. **Second**, the constraint of fixed output volume dimensions and voxel spacing poses substantial limitations in real-world applications [11]. These parameter presets are often incompatible with the diverse requirements of different tasks, such as the analysis of varying anatomical structures. The ability to dynamically adjust both the volume dimensions and the voxel spacing is crucial for enhancing the flexibility and utility of 3D generative models. **Third**, another common limitation of current generative models for medical image generation is their specialization to dedicated datasets or particular types of organs. These models, once trained, are typically not generalizable beyond the specific data and target organ they are developed on, which restricts their broader application in diverse settings. Developing more versatile models that can adapt to multiple datasets and organ types and mitigate the need for extensive retraining is a key objective for advancing the field [10].

In this paper, we propose a method, namely Medical AI for Synthetic Imaging (MAISI), a new framework for high-resolution 3D CT volume generation, which consists of three 3D networks including two foundation models (*i.e.*, a volume compression network, a latent diffusion model [50]) and a ControlNet [71] for versatile generation tasks. Volume Compression Network is trained on a large amount of

data (*i.e.*, 39,206 3D CT volumes) and is responsible for compressing the 3D medical images into latent space and mapping the generated latent features back to image space by a visual encoder and a visual decoder, respectively. To reduce the memory footprint, we introduce the tensor splitting parallelism (TSP) inspiring from the tensor parallelism technique [58], originally proposed for linear layers, to the 3D convolutional layers allowing for the encoding and decoding of high-resolution CT volumes in a unified 3D network. The latent diffusion model in MAISI facilitates the creation of realistic latent features of 3D medical images. Benefiting from a compressed latent space with flexible dimensions and taking body region and voxel spacing as conditions, it enables the generation of complex anatomical structures with a high degree of fidelity while maintaining relatively low memory consumption. The latent diffusion model is trained on 10,277 CT volumes from diverse datasets, encompassing various body regions and disease conditions to enhance its generalizability and robustness, which enables the model to capture the knowledge represented in a wide range of clinical scenarios. Further, the integration of ControlNet [71] into the MAISI framework introduces a mechanism for dynamic control over the generated outputs. This component enhances MAISI's versatility and applicability across a wider range of tasks (*e.g.*, conditional generation based on segmentation masks, as illustrated in Fig. 1, image inpainting, *etc.*). Additionally, this capability minimizes the need for extensive retraining of the two underlying foundation models when transitioning between different tasks or clinical objectives, thereby conserving both time and computational resources.

To summarize, this paper makes the following contributions:

- A novel framework, MAISI, for high-resolution 3D CT volume generation is proposed, which enables the versatile generation of synthetic CT images.
- Tensor splitting parallelism (TSP) is introduced to 3D convolutional networks. To the best of our knowledge, MAISI is the first attempt to generate realistic 3D CT images larger than 512^3 voxels.
- MAISI provides dynamic control over outputs, enabling annotated synthetic images to improve downstream task performance.

2. Related Work

Medical image synthesis has become an increasingly prominent research area, particularly in response to the challenges discussed in Sec. 1. Early approaches [7, 43, 52] to medical image synthesis were predominantly based on traditional image processing techniques, such as

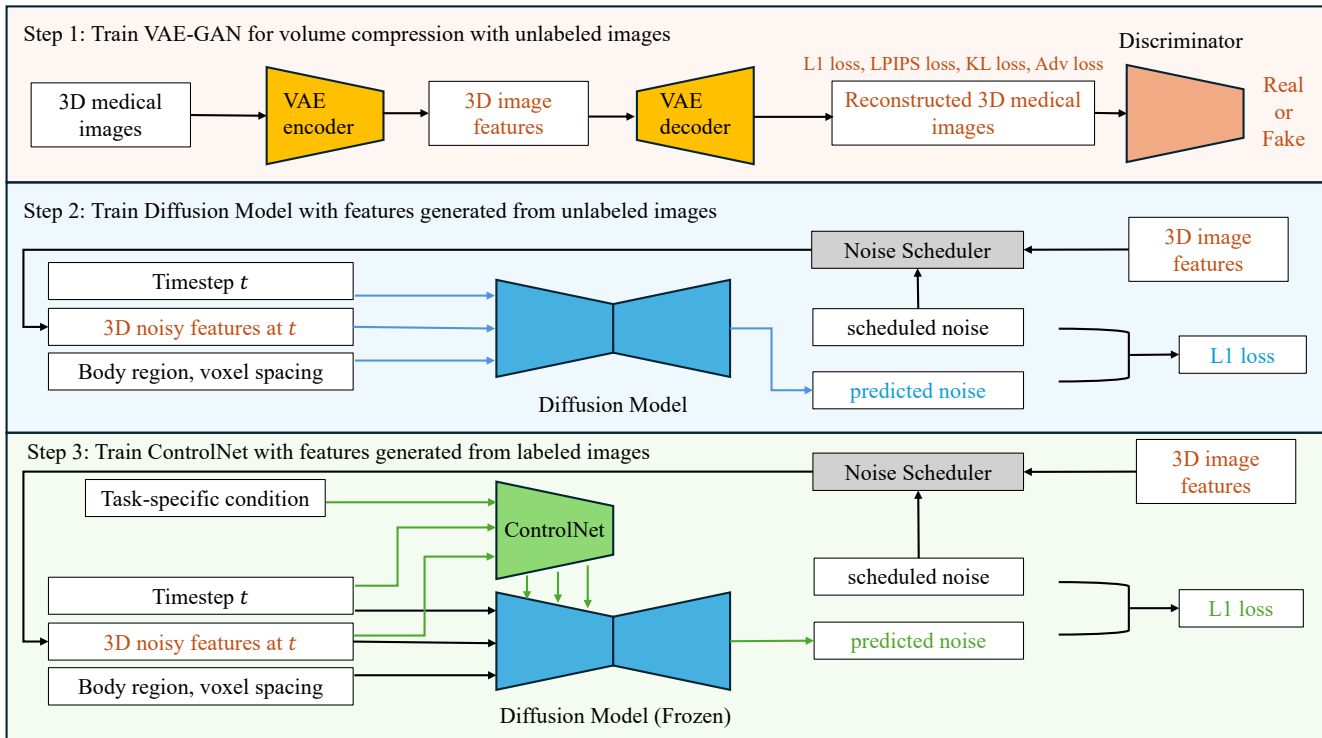


Figure 2. The overview of three development stages of MAISI.

the example-based approach [53] and geometry-regularized dictionary learning [31], which, while effective to some extent, are limited in their ability to generate realistic and diverse medical images. The advent of machine learning, particularly deep learning, has significantly advanced the field, enabling more sophisticated and accurate models for image synthesis [60].

GAN in medical image synthesis. GAN [21], one of generative models [12, 17, 37, 45], has been widely adopted for various tasks, such as MRI/CT image synthesis [22, 33, 62], cross-modality image translation [9, 57, 66, 70], image reconstruction [46, 65] and super-resolution [1, 47, 48, 68], in medical imaging synthesis due to its promising ability to generate high-quality images. One of the most critical applications of GAN in medical imaging is data augmentation by generating annotated images. Several studies [8, 22, 32, 75], have employed GAN to generate lesion images to augment training data for improving downstream tasks to overcome data scarcity issues. However, those methods focus on 2D medical imaging or small volumetric patch synthesis, which is fundamentally limited due to neglecting the inherent complexity and the 3D nature of medical data. In this work, we focus on generating full CT volume in realistic dimensions (up to $512 \times 512 \times 768$) to model complex volumetric features in a unified framework.

DM in medical image synthesis. Diffusion models [29, 50] have recently emerged as a powerful generative model

that has shown great potential in medical imaging synthesis due to its capabilities in high-quality image synthesis, stable training process, and flexibility in conditioning [5, 13, 67]. [35, 38, 56] demonstrate the effectiveness of DM-based methods in generating high-quality 2D medical images that capture intricate details with minimal artifacts, making them suitable for clinical use. GenerateCT [23] is designed to synthesize 3D CT volumes from free-form medical text prompts and accomplishes arbitrary-size CT volume generation by decomposing the process into a sequential generation of individual slices using DM. However, due to the nature of 2D approaches, the issue of 3D structural inconsistencies across slices is noticeable and problematic in the generated images. Application-wise, many recent studies [30, 39, 42, 64, 73, 74] are focusing on tumor synthesizing and improving models' performance in downstream tasks. DiffTumor [10] seeks to enhance the robustness and generalizability of tumor segmentation models across various organs, such as the liver, pancreas, and kidney, by leveraging high-quality synthetic tumors generated through specialized diffusion models. In this work, we focus on achieving conditional generation tailored to versatile tasks by leveraging robust foundation models, which significantly minimizes the need for extensive retraining across different applications, thereby conserving both time and computational resources while maintaining adaptability and efficiency in diverse clinical scenarios.

3. Methodology

As shown in Fig. 2, the development of MAISI involves three stages. In the first stage, the volume compression network (*i.e.*, VAE-GAN [50]) is trained on a substantial dataset comprising 39,206 3D CT volumes and 18,827 3D MRI volumes. This network effectively compresses high-resolution 3D medical images into a latent space that is perceptually equivalent to the image space, reducing memory usage and computational complexity for later stages. In the second stage, a latent diffusion model is trained on 10,277 CT volumes sourced from diverse datasets. This model operates within the compressed latent space, conditioned on specific body regions and voxel spacing, to generate features of realistic and complex 3D anatomical structures in flexible dimensions. Training on a broad range of data enhances the model’s generalizability and adaptability in different tasks. The final stage involves the integration of ControlNet [71] into the MAISI framework. This component allows for dynamic control over the generated outputs by injecting additional conditions into the trained latent DM in the second stage, potentially supporting a wide range of tasks. The integration reduces the need for extensive re-training when the model is adapted to different tasks. In what follows, we provide detailed descriptions of each key component of the MAISI framework.

3.1. Volume Compression Network

The volume compression model builds upon previous studies [19, 50] and employs a Variational Autoencoder (VAE) trained on combined objectives, which integrates perceptual loss $\mathcal{L}_{\text{lips}}$ [72], adversarial loss [69] \mathcal{L}_{adv} , and L1 reconstruction loss $\mathcal{L}_{\text{recon}}$ on voxel-space. These combined objectives ensure that the volume reconstructions adhere closely to the image manifold and enforce local realism [50]. In addition, we follow [36, 49, 50] adding Kullback-Leibler (KL) regularization \mathcal{L}_{reg} toward a standard normal on the learned latent features for avoiding high-variance latent spaces.

Given a CT volume $x \in \mathbb{R}^{H \times W \times D}$ in grayscale voxel space, where H denotes the height, W the width, and D the depth, the encoder \mathcal{E} of AE downsamples x and generates the latent representation $z = \mathcal{E}(x) \in \mathbb{R}^{h \times w \times d}$ with much smaller spatial dimensions. The decoder \mathcal{D} of AE approximates the reconstructed volume $\tilde{x} = \mathcal{D}(z) = \mathcal{D}(\mathcal{E}(x))$ from the latent features. A 3D discriminator, denoted as \mathcal{C} , is utilized to identify and penalize any unrealistic artifacts in the reconstructed volume \tilde{x} . As shown in Fig. 2 step 1, the overall objective \mathcal{L}_{AE} to train the volume compression network (\mathcal{E}, \mathcal{D}) in MAISI can be defined as follows:

$$\min_{\mathcal{E}, \mathcal{D}} \max_{\mathcal{C}} \left(\mathcal{L}_{\text{recon}}(x, \mathcal{D}(\mathcal{E}(x))) + \mathcal{L}_{\text{lips}}(x, \mathcal{D}(\mathcal{E}(x))) + \mathcal{L}_{\text{reg}}(\mathcal{E}(x)) + \mathcal{L}_{\text{adv}} \right), \quad (1)$$




Figure 3. The schematics of tensor splitting parallelism in MAISI. Feature maps are first partitioned into smaller segments with overlaps and allocated to designated devices. Then, these segments are stitched together to compose the output of the layer.

where

$$\mathcal{L}_{\text{adv}} = \log \mathcal{C}(x) + \log(1 - \mathcal{C}(\mathcal{D}(\mathcal{E}(x))))). \quad (2)$$

Generating high-resolution 3D volumes, particularly those exceeding dimensions of 512^3 voxels, poses a significant challenge due to the substantial memory demands imposed by the 3D convolution networks. In order to address memory bottleneck, previous methods [23, 54] achieve 2D high-resolution image synthesis via an additional super-resolution model. However, in the context of 3D whole-volume generation, the memory consumption can still quickly reach the hardware limitation of modern GPUs (*e.g.*, NVIDIA A100 80G). To overcome GPU memory constraints, sliding window inference [6] is a common technique. It divides the large network input into smaller 3D patches in a sliding-window fashion and then stitches the network’s output of each patch together to form the final results. When used in the 3D medical image segmentation model inference, it can often lead to artifacts/discontinuities along the window boundaries. While overlapping windows can help in segmentation tasks by smoothing over the boundary artifacts of probability maps, we empirically found this issue in transition areas between windows is more pronounced for the synthesis task due to the direct generation of image intensities, and thus the direct adaptation of sliding window inference is not self-sufficient. To minimize the use of the sliding-window approach for image synthesis, we propose a simple yet effective solution by introducing tensor splitting parallelism into convolutional networks. The tensor parallelism [58] is initially developed to distribute the inputs or model weights of matrix multiplication operations in fully connected layers across multiple GPUs. Unlike language models [18] built upon linear layers, the memory bottleneck usually attributes to the large feature maps. As demonstrated in Fig. 3, the proposed TSP is utilized to divide feature maps into smaller segments while preserving necessary overlaps across both convolution and normalization layers of AE. Each segment is assigned to a designated device, and these segments are

subsequently merged to generate the layer’s output. This flexible implementation enables segments to be distributed across multiple devices to accelerate the inference and also allows each segment to be processed sequentially within a single device in a loop to reduce peak memory usage.

3.2. Diffusion Model

The Diffusion Model in MAISI operates on a compressed latent space with flexible dimensions and incorporates body region and voxel spacing as conditional inputs, facilitating the high-fidelity generation of anatomical structures. Diffusion models are probabilistic models that aim to learn a data distribution $p(x)$ by gradually denoising a normally distributed variable. This process is equivalent to learning the reverse dynamics of a fixed Markov Chain over a sequence of T steps. The denoising score-matching [61] is widely adopted in image synthesis tasks [16, 55]. In the context of latent diffusion model [50], the learning model ϵ_θ functions as a uniformly weighted sequence of denoising autoencoders $\epsilon_\theta(z_t, t); t = 1 \dots T$, which are designed to predict a denoised version of the input latent features z_t and z_t represents a noisy variant of the original input at time step t . The neural backbone ϵ_θ is defined as a time-conditional U-Net [51]

As shown in Fig. 2 step 2, the diffusion model in MAISI additionally conditions on both the body region and voxel spacing. The body region is defined by a top-region index i_{top} and a bottom-region index i_{bottom} , indicating the extent of the CT scan coverage. i_{top} and i_{bottom} are defined by 4-dimensional one-hot vectors for head-neck, chest, abdomen, and lower-body regions). We ascertain the body region either through segmentation ground truth or predated segmentation masks from whole-body CT segmentation models, such as TotalSegmentator [63] or VISTA3D [26]. The condition of voxel spacing s is defined by a vector containing three float numbers representing the physical size of each voxel along each of the three dimensions in millimeters. We denote the primary conditions as $c_p := \{i_{\text{top}}, i_{\text{bottom}}, s\}$. Formally, the training objective of MAISI diffusion model is as follows:

$$\mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0,1), t, c_p} \left[\|\epsilon - \epsilon_\theta(z_t, t, c_p)\|_1 \right], \quad (3)$$

where the neural backbone ϵ_θ is configured to condition on time step t and the primary conditions as c_p . Moreover, ϵ_θ undergoes training on the latent variable z_t , which varies in dimensions throughout the training process. This training regimen is designed to facilitate the generation of outputs with flexible volumetric dimensions.

3.3. Additional Conditioning Mechanisms

In addition to the primary conditioning on body region and voxel spacing described in the Sec. 3.2, MAISI incorporates an additional mechanism for enhancing the control and flexibility of the generated outputs through the

integration of ControlNet [71]. It is seamlessly embedded into the MAISI architecture with the latent diffusion model, to provide additional conditioning paths that allow for task-specific adaptations. ControlNet [71] is designed to inject auxiliary conditions into the diffusion process, enabling more precise control over the generated anatomical structures. It operates by creating two copies of the neural network blocks: a locked copy that preserves the original model’s knowledge, and a trainable copy that learns to respond to specific conditions. These copies are connected using zero convolution layers, which gradually evolve from zero weights to optimal settings during training. These additional conditions can include a variety of inputs such as segmentation masks for conditional generation based on masks, or masked images and tumor masks for the tumor inpainting [10]. Similar to [38, 44, 71], we employ a compact encoder network to transform the additional condition from its original resolution into latent features, which are denoted by the task-specific condition c_f . This transformation process effectively aligns the additional condition with the spatial dimensions of the latent space. The integration of ControlNet [71] occurs during the third stage (Fig. 2 step 3) of MAISI’s development, where it is trained with the frozen latent diffusion model. The overall learning objective of the entire diffusion algorithm, which incorporates the ControlNet [71], is formulated as follows:

$$\mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0,1), t, c_p, c_f} \left[\|\epsilon - \epsilon_\theta(z_t, t, c_p, c_f)\|_1 \right]. \quad (4)$$

This integration adds a flexible mechanism to MAISI for controlling the generation of 3D anatomical structures. By injecting task-specific conditions, MAISI can be fine-tuned to meet the specific needs of various medical imaging tasks without retraining the two foundation models, making it a versatile tool for various medical image synthesis tasks.

4. Experiments

4.1. Datasets and Implementation Details

To develop and evaluate the proposed MAISI framework, we curate a large-scale medical imaging dataset from publicly available datasets to capture a diverse range of anatomical structures, imaging conditions, and disease states. These datasets are integral to training the three networks within the MAISI framework. The **Volume Compression Network** (MAISI VAE) is trained on a dataset comprising 37,243 CT volumes for training and 1,963 CT volumes for validation, covering the chest, abdomen, and head and neck regions. Additionally, we include 17,887 MRI volumes for training and 940 MRI volumes for validation, spanning the brain, skull-stripped brain, chest, and below-abdomen regions to potentially support MRI modality in future work. The **Latent DM** (MAISI Diffusion Model) was trained using 10,277 CT volumes sourced from multiple public datasets. These datasets are chosen to repre-

sent various clinical scenarios, including different body regions and pathological conditions. Including diverse voxel spacings and anatomical regions as conditional inputs during training is essential to ensure the model’s ability to generate high-fidelity anatomical structures with flexible dimensions. For compatibility with the shape requirement of U-Net [51], we resample the dimensions of volumes to the multiples of 128 in this stage. Supplementary Fig. S1 visualizes the characteristics and spatial complexity of the data involved in training the diffusion model. The **ControlNet** part was further trained using subsets of the datasets used for the diffusion model based on different downstream tasks, with additional annotations such as segmentation masks and tumor labels. For example, segmentation masks with 127 anatomical structures are derived from annotated ground truth or pre-trained models, such as TotalSegmentator [63] and VISTA3D [26]. These additional annotations allow ControlNet to provide fine-grained control over the generation process, enabling tasks such as conditional generation from segmentation masks and tumor inpainting. More details about dataset creation for three development stages can be found in Supplementary Sec. A. We implement all networks using PyTorch [2] and MONAI [6]. The models are trained using the NVIDIA V100 and A100 GPUs. We utilize a quality check function to evaluate the generated images used in downstream tasks, which is designed to verify that the median Hounsfield Units (HU) intensity values for major organs in the CT images are within the established normal range from training data. More details about model training are provided in Supplementary Sec. B.

4.2. Evaluation of MAISI VAE

Dataset	Model	LPIPS ↓	SSIM ↑	PSNR ↑	GPU ↓
MSD Task07	MAISI VAE	0.038	0.978	37.266	0h
	Dedicated VAE	0.047	0.971	34.750	619h
MSD Task08	MAISI VAE	0.046	0.970	36.559	0h
	Dedicated VAE	0.041	0.973	37.110	669h
Brats18	MAISI VAE	0.026	0.977	39.003	0h
	Dedicated VAE	0.030	0.975	38.971	672h

Table 1. Performance comparison of the MAISI VAE model on out-of-distribution datasets versus dedicated VAE models. The “GPU” column shows additional GPU hours for training with one 32G V100 GPU.

To demonstrate the robustness and generalizability of the MAISI VAE model as a foundational model, we test its performance on several out-of-distribution datasets (*i.e.*, unseen during training), including MSD Pancreas Tumor [3] (MSD Task07), MSD Hepatic Vessels [3] (MSD Task08), and BraTS18 [4] (post-contrast T1-weighted MRI). Notably, this application required no additional training, resulting in eliminating any associated training costs of GPU

hours. For comparison, we also train the dedicated VAE models separately on each dataset using 80% of the data, with the same data augmentation techniques and hyperparameters as those employed for the MAISI VAE training, to establish a benchmark for dedicated VAE models.

The results from testing on the remaining 20% of the data, shown in Table 1, revealed that the MAISI VAE model achieved comparable results without additional GPU resource expenditure. This underscores the model’s cost-effectiveness and practicality, suggesting its potential to assist the research community in optimizing resource utilization while maintaining the model’s performance.

4.3. Evaluation of MAISI Diffusion Model

Synthesis quality. We assess the synthesis quality of the standalone MAISI DM by conducting comparisons with several established baseline methods, including DDPM [29], LDM [50], and HA-GAN [62]. The first evaluation focuses on comparing the fidelity of images generated by our model against those produced by the HA-GAN [62], utilizing its publicly available trained weights¹. Given that HA-GAN [62] specifically targets CT images of the chest region, we curate a collection of chest CT datasets for this analysis, including MSD Lung Tumor [3] (MSD Task06), LIDC-IDRI [24], and TCIA COVID-19 [25]. These datasets provide a diverse range of imaging conditions and pathology, enriching the comparative study. We use the Fréchet Inception Distance (FID) [27] as the metric for evaluating the similarity between the distributions of generated images and real counterparts from varied sources. Table 2 presents the average FID for both real and synthesized images across the three datasets. Notably, the MAISI DM significantly outperforms HA-GAN [62] in all datasets, demonstrating its capability to generate images with a much closer appearance to real data.

FID ↓ (Avg.)		MSD Task 06	LIDC-IDRI	TCIA COVID-19
Real	MSD Task06	–	3.987	1.858
	LIDC-IDRI	3.987	–	4.744
	TCIA COVID-19	1.858	4.744	–
Synthesis	HA-GAN [62]	98.208	116.260	98.064
	MAISI DM	4.349	6.200	8.346

Table 2. Fréchet Inception Distance of the MAISI model and the baseline method using its released checkpoint with multiple public datasets as the references.

In addition, we retrain all baseline methods using our large-scale datasets, described in Sec. 4.1. For a more comprehensive evaluation of synthesis quality, we utilize an unseen dataset autoPET 2023 [20] as the reference to conduct synthesis quality evaluation. This dataset encompasses whole-body CT scans from patients with various types of

¹<https://github.com/batmanlab/HA-GAN>

Method	FID ↓ (Axial)	FID ↓ (Sagittal)	FID ↓ (Coronal)	FID ↓ (Avg.)
DDPM [29]	18.524	23.696	25.604	22.608
LDM [50]	16.853	10.191	10.093	12.379
HA-GAN [62]	17.432	10.266	13.572	13.757
MAISI DM	3.301	5.838	9.109	6.083

Table 3. Fréchet Inception Distance across three views between MAISI DM and retrained baseline methods using the unseen dataset autoPET 2023 [20] as the reference.

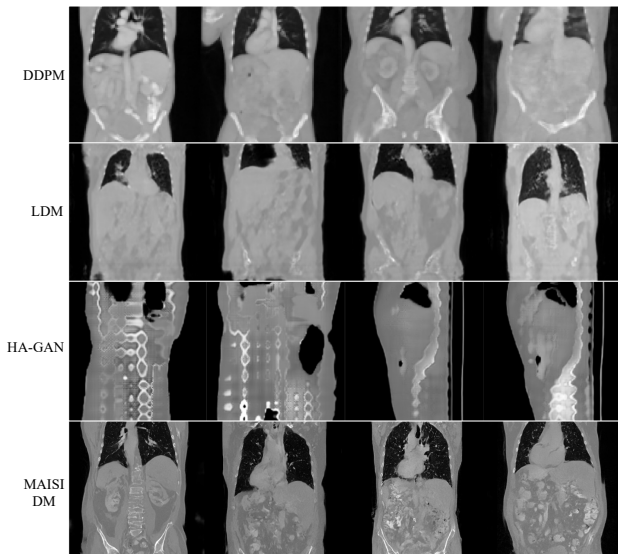


Figure 4. Qualitative comparison of generated images between retrained baseline methods using our large-scale datasets and MAISI DM.

cancer and negative controls. Results in Table 3 demonstrate that our MAISI DM surpasses the retrained baseline models in generating high-quality images in the external evaluation. Fig. 4 presents a visual comparison illustrating that the high-resolution images synthesized by the MAISI DM show improved detail and a more precise representation of global anatomical structures compared to baseline methods.

Response to primary conditions. Fig. 5 illustrates the model’s adaptability to different body regions and voxel spacing conditions. The MAISI model effectively generates anatomically consistent and high-quality images across different primary conditions c_p , demonstrating its flexibility and control over synthesized images.

4.4. Data Augmentation in Downstream Tasks

One of the critical applications of generative models in medical imaging is data augmentation for training deep learning models. To assess the effectiveness of synthetic images in improving model performance, especially for rare medical conditions, we integrate synthetic data generated by MAISI into a standard training pipeline and eval-




Figure 5. The sagittal view of generated CT images from MAISI DM under different primary conditions c_p . From left to right, the voxel spacing is first increased by 50%, followed by a doubling of the output dimensions. The coverage of the generated CT images gradually expands, starting from a local region of the abdomen and extending to the entire chest-abdomen region.

uate it across five tumor types. Specifically, we employ the Auto3Dseg² pipeline—an auto-configuration solution for training medical image segmentation models—to train models on the MSD Task03 [3] (liver tumor), Task06 [3] (lung tumor), Task07 [3] (pancreas tumor), Task10 [3] (colon tumor), and an in-house bone lesion dataset. We conduct experiments by training segmentation models either using only real data (referred to **Real Only** in Fig. 6) or by incorporating synthetic data from different models, thereby demonstrating the impact of synthetic data on data augmentation. To ensure robustness, we performed 5-fold cross-validation and reported the average Dice Similarity Coefficient (DSC) on the testing set across the five folds.

As discussed in Sec.3.3, the integration of ControlNet [71] introduces a flexible mechanism in MAISI, enabling the incorporation of task-specific conditions. To illustrate its versatility, we trained ControlNet [71] for two distinct tasks aimed at generating synthetic data for augmentation purposes. The first task (denoted as **MAISI CT Generation** in Fig. 6) is conditional generation from segmentation masks of 127 anatomical structures, including the five tumor types mentioned earlier. This approach allowed us to generate synthetic data by augmenting real patient tumor masks corresponding to each tumor type. The second task (denoted as **MAISI Inpainting** in Fig. 6) involves training a tumor inpainting model designed to simultaneously support liver, pancreas, and lung tumors, following the setting in [10]. The tumor inpainting model requires a function to simulate tumor masks for adding synthetic tumors into healthy patient data. However, simulating tumors with irregular shapes, such as bone lesions and colon tumors, poses significant challenges. For a comparative analysis, we benchmark against the state-of-the-art tumor synthesis method, DiffTumor [10], using their released model³ which supports liver and pancreas tumors among five tumor

²<https://monai.io/apps/auto3dseg>

³<https://github.com/MrGiovanni/DiffTumor>

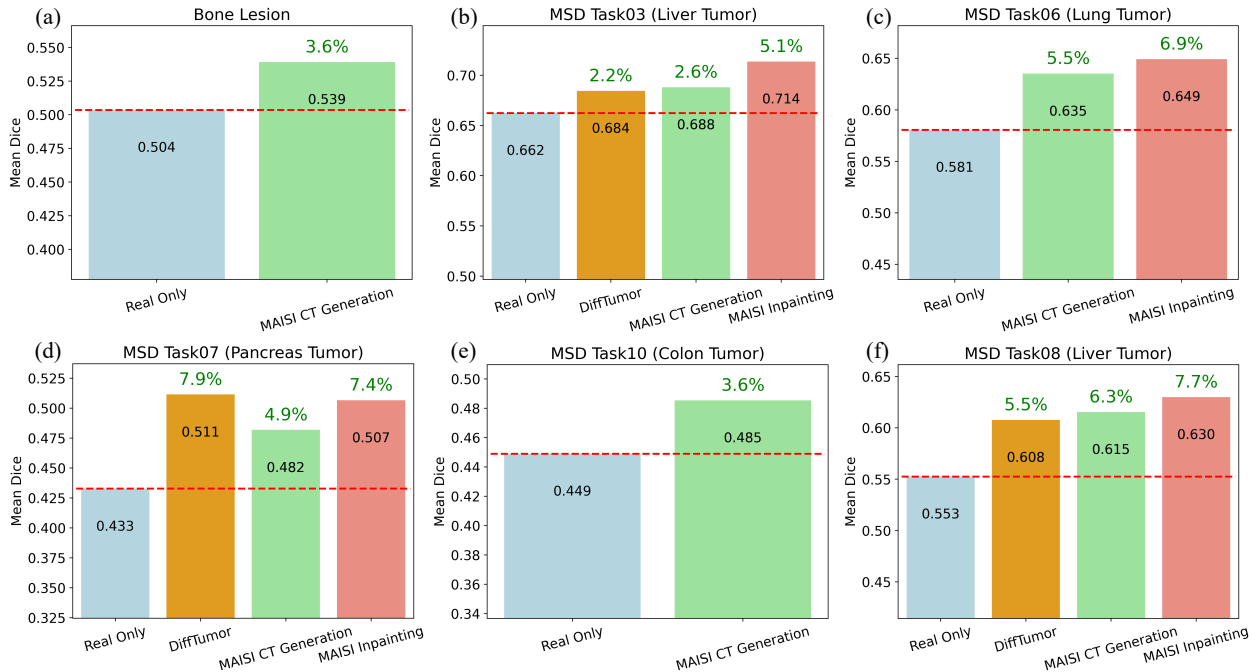


Figure 6. The 5-fold averaged DSC of data augmentation experiments using synthetic data across 5 tumor types. The percentage of relative improvement compared to **Real Only** experiments is shown in green above each bar plot. All reported improvements are significant under the Wilcoxon signed rank test.

types in our experiments.

Results shown in Fig. 6(a)~(e) indicate prominent improvements in DSC scores across all tumor types when incorporating synthetic data from our two augmentation tasks. Specifically, the MAISI CT Generation results in an average DSC improvement of 4% across the five tumor types. The MAISI Inpainting demonstrated a more substantial average improvement of 6.5% in DSC for liver, lung, and pancreas tumors, performing comparably or better than the DiffTumor [10], which trains dedicated synthesis models for each tumor type. Additionally, we conduct an out-of-distribution evaluation by testing tumor segmentation models trained on MSD Task03 [3] on 303 liver tumor samples from MSD Task08 [3]. As shown in Fig. 6(f), models incorporating synthetic data consistently show greater relative performance improvements compared to those evaluated within their original training dataset in Fig. 6(b). These findings underscore the effectiveness of synthetic data as a powerful augmentation strategy to bolster the generalizability of segmentation models. More ablation studies and visualization of synthetic data can be found in Supplementary Sec. C.

5. Discussion and Limitations

While the proposed MAISI demonstrates great potentials in generating high-quality CT images, it is essential to recognize its limitations and potential societal impacts. While

MAISI shows robust performance across various datasets, its ability to accurately represent demographic variations (such as age, ethnicity, and gender differences) in generated anatomy has not been extensively validated. Future studies can focus on ensuring that synthetic data adequately captures this diversity to avoid bias in downstream applications. The capabilities of generating high-resolution images of MAISI, while innovative, still demand substantial computation resources. This could limit accessibility for researchers and institutions with less computational power, potentially widening the gap between high-resource and low-resource entities. Future efforts can focus on improving the accessibility of MAISI, particularly in resource-constrained environments.

6. Conclusion

In this paper, we propose MAISI, a novel framework for generating high-resolution 3D CT volumes using a combination of foundation models and ControlNet [71]. MAISI aims to provide an adaptable and versatile solution for generating anatomically accurate images. Our experiments demonstrate that MAISI can produce realistic CT images with flexible volume dimensions and voxel spacing, offering promising potential to augment medical datasets and improve the performance of downstream tasks.

References

- [1] Waqar Ahmad, Hazrat Ali, Zubair Shah, and Shoaib Azmat. A new generative adversarial network for medical images super resolution. *Scientific Reports*, 12(1):9533, 2022. [3](#)
- [2] Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos, Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Christian Puhersch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo, Phil Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou, Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24)*. ACM, Apr. 2024. [6](#)
- [3] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al. The medical segmentation decathlon. *Nature communications*, 13(1):4128, 2022. [6](#), [7](#), [8](#), [16](#)
- [4] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018. [6](#)
- [5] Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z Li. A survey on generative diffusion models. *IEEE Transactions on Knowledge and Data Engineering*, 2024. [3](#)
- [6] M Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey, Andriy Myronenko, Can Zhao, Dong Yang, et al. Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701*, 2022. [4](#), [6](#), [16](#)
- [7] M Jorge Cardoso, Carole H Sudre, Marc Modat, and Sebastien Ourselin. Template-based multimodal joint generative model of brain data. In *Information Processing in Medical Imaging: 24th International Conference, IPMI 2015, Sabhal Mor Ostaig, Isle of Skye, UK, June 28-July 3, 2015, Proceedings 24*, pages 17–29. Springer, 2015. [2](#)
- [8] Agisilaos Chartsias, Thomas Joyce, Rohan Dharmakumar, and Sotirios A Tsaftaris. Adversarial image synthesis for unpaired multi-modal cardiac data. In *Simulation and Synthesis in Medical Imaging: Second International Workshop, SASHIMI 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 10, 2017, Proceedings 2*, pages 3–13. Springer, 2017. [3](#)
- [9] Agisilaos Chartsias, Thomas Joyce, Mario Valerio Giuffrida, and Sotirios A Tsaftaris. Multimodal mr synthesis via modality-invariant latent representation. *IEEE transactions on medical imaging*, 37(3):803–814, 2017. [2](#), [3](#)
- [10] Qi Chen, Xiaoxi Chen, Haorui Song, Zhiwei Xiong, Alan Yuille, Chen Wei, and Zongwei Zhou. Towards generalizable tumor synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11147–11158, 2024. [2](#), [3](#), [5](#), [7](#), [8](#), [16](#)
- [11] Xiang Chen, Andres Diaz-Pinto, Nishant Ravikumar, and Alejandro F Frangi. Deep learning in medical image registration. *Progress in Biomedical Engineering*, 3(1):012003, 2021. [2](#)
- [12] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018. [3](#)
- [13] 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*, 45(9):10850–10869, 2023. [3](#)
- [14] Mohammad Zalbagi Darestani, Vishwesh Nath, Wenqi Li, Yufan He, Holger R Roth, Ziyue Xu, Daguang Xu, Reinhard Heckel, and Can Zhao. Ir-freestormer: Iterative refinement with fourier-based restormer for accelerated mri reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7655–7664, 2024. [2](#)
- [15] Blake E Dewey, Can Zhao, Jacob C Reinhold, Aaron Carass, Kathryn C Fitzgerald, Elias S Sotirchos, Shiv Saidha, Jiwon Oh, Dzung L Pham, Peter A Calabresi, et al. Deepharmony: A deep learning approach to contrast harmonization across scanner changes. *Magnetic resonance imaging*, 64:160–170, 2019. [2](#)
- [16] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. [5](#)
- [17] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#)
- [18] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. [4](#)
- [19] 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](#)
- [20] Sergios Gatidis, Tobias Hepp, Marcel Früh, Christian La Fougère, Konstantin Nikolaou, Christina Pfannenbergl, Bernhard Schölkopf, Thomas Küstner, Clemens Cyran, and Daniel Rubin. A whole-body fdg-pet/ct dataset with manually annotated tumor lesions. *Scientific Data*, 9(1):601, 2022. [6](#), [7](#)
- [21] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [2](#), [3](#)

- [22] Pengfei Guo, Puyang Wang, Rajeev Yasarla, Jinyuan Zhou, Vishal M Patel, and Shanshan Jiang. Anatomic and molecular mr image synthesis using confidence guided cnns. *IEEE transactions on medical imaging*, 40(10):2832–2844, 2020. [2](#), [3](#)
- [23] Ibrahim Ethem Hamameci, Sezgin Er, Anjany Sekuboyina, Enis Simsar, Alperen Tezcan, Ayse Gulnihhan Simsek, Sevval Nil Esirgun, Furkan Almas, Irem Dogan, Muhammed Furkan Dasdelen, et al. Generatect: Text-conditional generation of 3d chest ct volumes. *arXiv preprint arXiv:2305.16037*, 2023. [3](#), [4](#)
- [24] Matthew C Hancock and Jerry F Magnan. Lung nodule malignancy classification using only radiologist-quantified image features as inputs to statistical learning algorithms: probing the lung image database consortium dataset with two statistical learning methods. *Journal of Medical Imaging*, 3(4):044504–044504, 2016. [6](#)
- [25] Stephanie A Harmon, Thomas H Sanford, Sheng Xu, Evrim B Turkbey, Holger Roth, Ziyue Xu, Dong Yang, Andriy Myronenko, Victoria Anderson, Amel Amalou, et al. Artificial intelligence for the detection of covid-19 pneumonia on chest ct using multinational datasets. *Nature communications*, 11(1):4080, 2020. [6](#)
- [26] Yufan He, Pengfei Guo, Yucheng Tang, Andriy Myronenko, Vishwesh Nath, Ziyue Xu, Dong Yang, Can Zhao, Benjamin Simon, Mason Belue, et al. Vista3d: Versatile imaging segmentation and annotation model for 3d computed tomography. *arXiv preprint arXiv:2406.05285*, 2024. [5](#), [6](#), [16](#), [20](#)
- [27] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. [6](#)
- [28] Irina Higgins, Loic Matthey, Arka Pal, Christopher P Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR (Poster)*, 3, 2017. [16](#)
- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [2](#), [3](#), [6](#), [7](#)
- [30] Qixin Hu, Junfei Xiao, Yixiong Chen, Shuwen Sun, Jie-Neng Chen, Alan Yuille, and Zongwei Zhou. Synthetic tumors make ai segment tumors better. *arXiv preprint arXiv:2210.14845*, 2022. [3](#)
- [31] Yawen Huang, Leandro Beltrachini, Ling Shao, and Alejandro F Frangi. Geometry regularized joint dictionary learning for cross-modality image synthesis in magnetic resonance imaging. In *Simulation and Synthesis in Medical Imaging: First International Workshop, SASHIMI 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 21, 2016, Proceedings 1*, pages 118–126. Springer, 2016. [3](#)
- [32] Yuankai Huo, Zhoubing Xu, Hyeonsoo Moon, Shunxing Bao, Albert Assad, Tamara K Moyo, Michael R Savona, Richard G Abramson, and Bennett A Landman. Synsegnet: Synthetic segmentation without target modality ground truth. *IEEE transactions on medical imaging*, 38(4):1016–1025, 2018. [3](#)
- [33] Thomas Joyce, Agisilaos Chartsias, and Sotirios A Tsaftaris. Robust multi-modal mr image synthesis. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part III 20*, pages 347–355. Springer, 2017. [2](#), [3](#)
- [34] Mintong Kang, Bowen Li, Zengle Zhu, Yongyi Lu, Elliot K Fishman, Alan Yuille, and Zongwei Zhou. Label-assemble: Leveraging multiple datasets with partial labels. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2023. [1](#)
- [35] Bardia Khosravi, Frank Li, Theo Dapamede, Pouria Rouzrokh, Cooper U Gamble, Hari M Trivedi, Cody C Wyles, Andrew B Sellergren, Saptarshi Purkayastha, Bradley J Erickson, et al. Synthetically enhanced: unveiling synthetic data’s potential in medical imaging research. *EBioMedicine*, 104, 2024. [3](#)
- [36] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [4](#)
- [37] Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019. [3](#)
- [38] Nicholas Konz, Yuwen Chen, Haoyu Dong, and Maciej A Mazurowski. Anatomically-controllable medical image generation with segmentation-guided diffusion models. *arXiv preprint arXiv:2402.05210*, 2024. [3](#), [5](#)
- [39] Yuxiang Lai, Xiaoxi Chen, Angtian Wang, Alan Yuille, and Zongwei Zhou. From pixel to cancer: Cellular automata in computed tomography. *arXiv preprint arXiv:2403.06459*, 2024. [3](#)
- [40] Jie Liu, Yixiao Zhang, Jie-Neng Chen, Junfei Xiao, Yongyi Lu, Bennett A Landman, Yixuan Yuan, Alan Yuille, Yucheng Tang, and Zongwei Zhou. Clip-driven universal model for organ segmentation and tumor detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21152–21164, 2023. [1](#)
- [41] Xiangde Luo, Wenjun Liao, Jianghong Xiao, Jieneng Chen, Tao Song, Xiaofan Zhang, Kang Li, Dimitris N Metaxas, Guotai Wang, and Shaoting Zhang. Word: A large scale dataset, benchmark and clinical applicable study for abdominal organ segmentation from ct image. *Medical Image Analysis*, 82:102642, 2022. [20](#)
- [42] Fei Lyu, Mang Ye, Andy J Ma, Terry Cheuk-Fung Yip, Grace Lai-Hung Wong, and Pong C Yuen. Learning from synthetic ct images via test-time training for liver tumor segmentation. *IEEE transactions on medical imaging*, 41(9):2510–2520, 2022. [3](#)
- [43] Michael I Miller, Gary E Christensen, Yali Amit, and Ulf Grenander. Mathematical textbook of deformable neuroanatomies. *Proceedings of the National Academy of Sciences*, 90(24):11944–11948, 1993. [2](#)
- [44] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(5):4296–4304, Mar. 2024. [5](#)
- [45] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan.

- Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. 3
- [46] Cheng Peng, Pengfei Guo, S Kevin Zhou, Vishal M Patel, and Rama Chellappa. Towards performant and reliable undersampled mr reconstruction via diffusion model sampling. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 623–633. Springer, 2022. 2, 3
- [47] Cheng Peng, Wei-An Lin, Haofu Liao, Rama Chellappa, and S Kevin Zhou. Saint: spatially aware interpolation network for medical slice synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7750–7759, 2020. 3
- [48] Chi-Hieu Pham, Carlos Tor-Díez, H el ene Meunier, Nathalie Bednarek, Ronan Fablet, Nicolas Passat, and Fran ois Rousseau. Multiscale brain mri super-resolution using deep 3d convolutional networks. *Computerized Medical Imaging and Graphics*, 77:101647, 2019. 3
- [49] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014. 4
- [50] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bj orn 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, 7
- [51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 5, 6
- [52] Snehashis Roy, Aaron Carass, and Jerry Prince. A compressed sensing approach for mr tissue contrast synthesis. In *Information Processing in Medical Imaging: 22nd International Conference, IPMI 2011, Kloster Irsee, Germany, July 3-8, 2011. Proceedings 22*, pages 371–383. Springer, 2011. 2
- [53] Snehashis Roy, Aaron Carass, and Jerry L Prince. Magnetic resonance image example-based contrast synthesis. *IEEE transactions on medical imaging*, 32(12):2348–2363, 2013. 3
- [54] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 4
- [55] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4713–4726, 2022. 5
- [56] Fangxin Shang, Jie Fu, Yehui Yang, Haifeng Huang, Junwei Liu, and Lei Ma. Synfundus: A synthetic fundus images dataset with millions of samples and multi-disease annotations. *arXiv preprint arXiv:2312.00377*, 2023. 3
- [57] Hoo-Chang Shin, Neil A Tenenholz, Jameson K Rogers, Christopher G Schwarz, Matthew L Senjem, Jeffrey L Gunter, Katherine P Andriole, and Mark Michalski. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In *Simulation and Synthesis in Medical Imaging: Third International Workshop, SASHIMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 3*, pages 1–11. Springer, 2018. 2, 3
- [58] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019. 2, 4
- [59] Satya P Singh, Lipo Wang, Sukrit Gupta, Haveesh Goli, Parasuraman Padmanabhan, and Bal azs Guly as. 3d deep learning on medical images: a review. *Sensors*, 20(18):5097, 2020. 2
- [60] Youssef Skandarani, Pierre-Marc Jodoin, and Alain Lalonde. Gans for medical image synthesis: An empirical study. *Journal of Imaging*, 9(3):69, 2023. 3
- [61] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 5
- [62] Li Sun, Junxiang Chen, Yanwu Xu, Mingming Gong, Ke Yu, and Kayhan Batmanghelich. Hierarchical amortized gan for 3d high resolution medical image synthesis. *IEEE journal of biomedical and health informatics*, 26(8):3966–3975, 2022. 2, 3, 6, 7
- [63] Jakob Wasserthal, Hanns-Christian Breit, Manfred T Meyer, Maurice Pradella, Daniel Hinck, Alexander W Sauter, Tobias Heye, Daniel T Boll, Joshy Cyriac, Shan Yang, et al. Totalsegmentator: robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence*, 5(5), 2023. 5, 6, 16
- [64] Linshan Wu, Jiaxin Zhuang, Xuefeng Ni, and Hao Chen. Freetumor: Advance tumor segmentation via large-scale tumor synthesis. *arXiv preprint arXiv:2406.01264*, 2024. 3
- [65] Yutong Xie and Quanzheng Li. Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 655–664. Springer, 2022. 2, 3
- [66] Heran Yang, Jian Sun, Aaron Carass, Can Zhao, Junhoon Lee, Jerry L Prince, and Zongben Xu. Unsupervised mrt-to-ct synthesis using structure-constrained cyclegan. *IEEE transactions on medical imaging*, 39(12):4249–4261, 2020. 2, 3
- [67] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023. 3
- [68] Chenyu You, Guang Li, Yi Zhang, Xiaoliu Zhang, Hongming Shan, Mengzhou Li, Shenghong Ju, Zhen Zhao,

- Zhuiyang Zhang, Wenxiang Cong, et al. Ct super-resolution gan constrained by the identical, residual, and cycle learning ensemble (gan-circle). *IEEE transactions on medical imaging*, 39(1):188–203, 2019. 3
- [69] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021. 4
- [70] Mahmut Yurt, Salman UH Dar, Aykut Erdem, Erkut Erdem, Kader K Oguz, and Tolga Çukur. mustgan: multi-stream generative adversarial networks for mr image synthesis. *Medical image analysis*, 70:101944, 2021. 3
- [71] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 2, 4, 5, 7, 8
- [72] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 4
- [73] Xiaoman Zhang, Weidi Xie, Chaoqin Huang, Ya Zhang, Xin Chen, Qi Tian, and Yanfeng Wang. Self-supervised tumor segmentation with sim2real adaptation. *IEEE Journal of Biomedical and Health Informatics*, 27(9):4373–4384, 2023. 3
- [74] Zhaoxiang Zhang, Hanqiu Deng, and Xingyu Li. Unsupervised liver tumor segmentation with pseudo anomaly synthesis. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 86–96. Springer, 2023. 3
- [75] Zizhao Zhang, Lin Yang, and Yefeng Zheng. Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 9242–9251, 2018. 3
- [76] Can Zhao, Aaron Carass, Junghoon Lee, Yufan He, and Jerry L Prince. Whole brain segmentation and labeling from ct using synthetic mr images. In *Machine Learning in Medical Imaging: 8th International Workshop, MLMI 2017, Held in Conjunction with MICCAI 2017, Quebec City, QC, Canada, September 10, 2017, Proceedings 8*, pages 291–298. Springer, 2017. 2
- [77] Can Zhao, Blake E Dewey, Dzung L Pham, Peter A Calabresi, Daniel S Reich, and Jerry L Prince. Smore: a self-supervised anti-aliasing and super-resolution algorithm for mri using deep learning. *IEEE transactions on medical imaging*, 40(3):805–817, 2020. 2
- [78] Yi Zhou, Xiaodong He, Shanshan Cui, Fan Zhu, Li Liu, and Ling Shao. High-resolution diabetic retinopathy image synthesis manipulated by grading and lesions. In *International conference on medical image computing and computer-assisted intervention*, pages 505–513. Springer, 2019. 1