

MAISI-v2: Accelerated 3D High-Resolution Medical Image Synthesis with Rectified Flow and Region-specific Contrastive Loss

Can Zhao¹, Pengfei Guo¹, Dong Yang¹, Yufan He¹, Yucheng Tang¹, Benjamin Simon^{2,3}, Mason Belue⁴, Stephanie Harmon², Baris Turkbey², Daguang Xu¹

¹NVIDIA

²National Institutes of Health

³University of Oxford

⁴University of Arkansas for Medical Sciences
canz@nvidia.com

Abstract

Medical image synthesis is an important topic for both clinical and research applications. Recently, diffusion models have become a leading approach in this area. Despite their strengths, many existing methods struggle with (1) limited generalizability that only work for specific body regions or voxel spacings, (2) slow inference, which is a common issue for diffusion models, and (3) weak alignment with input conditions, which is a critical issue for medical imaging. MAISI, a previously proposed framework, addresses generalizability issues but still suffers from slow inference and limited condition consistency. In this work, we present **MAISI-v2**, the **first** accelerated 3D medical image synthesis framework that integrates rectified flow to enable fast and high quality generation. To further enhance condition fidelity, we introduce a **novel** region-specific contrastive loss to enhance the sensitivity to region of interest. Our experiments show that MAISI-v2 can achieve SOTA image quality with **33× acceleration** for latent diffusion model. We also conducted a downstream segmentation experiment to show that the synthetic images can be used for data augmentation. We release our code, training details, model weights, and a GUI demo to facilitate reproducibility and promote further development within the community.

Code — <https://github.com/NVIDIA-Medtech/NV-Generate-CTMR/tree/main>

GUI demo — <https://build.nvidia.com/nvidia/maisi>

Introduction

Medical image synthesis is a foundational task in medical AI, supporting applications such as data augmentation, modality translation, anomaly simulation, and privacy-preserving data sharing. By generating anatomically plausible and condition-aware 3D images, synthesis models help train robust models in low-data regimes, support cross-modality translation, and enable reproducible benchmarking on synthetic datasets. As medical imaging advances toward more personalized, heterogeneous, and multimodal data, the need for generalizable and controllable synthesis frameworks becomes increasingly critical.

Most existing medical image synthesis methods generally fall into three categories: (1) *direct regression*, which

directly predicts voxel intensities from condition inputs using supervised loss functions (Zhao et al. 2017; Dewey et al. 2019); (2) *GAN-based models*, which generate realistic images by adversarial training (Yang et al. 2020; Nie et al. 2017; Chartsias et al. 2017; Sun et al. 2022); and (3) *diffusion-based models*, which learn to reverse a forward noising process through a series of denoising steps (Dorjsembe et al. 2024; Pinaya et al. 2022; Guo et al. 2025; Zhu et al. 2023; Friedrich et al. 2024; Chen et al. 2024; Wang et al. 2024; Zhu et al. 2024; Wang et al. 2025).

Among these three categories, diffusion models have emerged as the state of the art in image generation, due to their stability and ability to model complex distributions. Their application to 3D medical imaging is particularly promising, yet their clinical deployment remains limited due to three primary challenges: (1) large variation for real-world medical images, (2) slow inference for diffusion models like Denoising Diffusion Probabilistic Models (DDPM) (Ho, Jain, and Abbeel 2020), and (3) weak condition fidelity.

First, many diffusion-based medical synthesis models are trained on narrow datasets—specific organs, modalities, or voxel spacings—limiting their ability to generalize to real-world clinical data, which varies widely in resolution, anatomy, and acquisition protocol (Dorjsembe et al. 2024; Pinaya et al. 2022; Zhu et al. 2023; Friedrich et al. 2024). Second, diffusion models typically require hundreds of iterative denoising steps during inference, which becomes computationally prohibitive in 3D. Existing methods either can only generate relatively small volumes like $256 \times 256 \times 256$ (Dorjsembe et al. 2024; Pinaya et al. 2022; Zhu et al. 2023; Friedrich et al. 2024; Chen et al. 2024; Xu et al. 2024) or take a long time (over 10 minutes) to generate high-resolution volumes such as $512 \times 512 \times 512$ (Guo et al. 2025; Wang et al. 2024). If the cost to generate a synthetic image is too high, the practical value of medical image synthesis will be limited. Third, while recent work in 2D natural image synthesis like ControlNet++ (Li et al. 2024) has improved condition fidelity over ControlNet (Zhang, Rao, and Agrawala 2023), such strategies are underexplored in 3D medical synthesis, often leading to unsatisfying alignment between the conditioning input and the output. This is a particularly critical issue in medical image applications.

To address generalization, the MAISI framework (Guo et al. 2025) and 3D MedDiffusion (Chen et al. 2024) were introduced as unified 3D medical image synthesis systems that leverage latent diffusion models (LDMs) and ControlNet to handle diverse anatomical structures and voxel spacings. However, they both still inherit two major limitations of standard LDM and ControlNet: slow inference and lack of robust condition control.

To tackle efficiency, rectified flow models (Liu, Gong, and Liu 2022) have recently emerged as a powerful alternative to traditional diffusion models. By replacing the stochastic denoising process with a deterministic ordinary differential equation (ODE), rectified flow enables faster sampling while preserving generation quality. This approach has shown promising results in natural image generation and video generation, including Open Sora (Peng et al. 2025) and Stable Diffusion 3.0 (Esser et al. 2024), but has not yet been extended to 3D medical imaging.

To improve condition fidelity, cycle-consistency loss like ControlNet++ (Li et al. 2024) has been proposed in 2D generative models. But it requires to train an additional reverse network using ground-truth conditions, which makes the procedure complicated. Also the error in the reverse network propagates.

In this work, we present **MAISI-v2**, a next-generation framework for accelerated and condition-consistent 3D medical image synthesis. Our contributions are as follows:

- **First Rectified Flow Model for 3D Medical Image Synthesis:** We adapt rectified flow to the 3D medical domain and integrate it into a latent diffusion framework, achieving substantial $33\times$ **acceleration** in inference while preserving image quality.
- **Novel Region-Specific Contrastive Loss:** We propose a region-specific contrastive loss to enhance the model’s sensitivity to small, localized conditions, such as small tumors, while maintaining background consistency.
- **Open-Source Framework and Pretrained Models:** We release code, training details, model weights, and a GUI demo to support reproducibility and provide a strong baseline for future research in controllable and efficient 3D image synthesis.

Related Work

MAISI: Medical AI for Synthetic Imaging

MAISI (Guo et al. 2025) is a modular 3D medical image synthesis framework designed to handle diverse anatomical structures and acquisition protocols. It consists of three key components:

1. **Variational Autoencoder (VAE):** MAISI compresses high-resolution 3D volumes into a compact latent space using a VAE trained with a reconstruction loss. The 1-channel images are spatially compressed by $4 \times 4 \times 4$ and generate 4 channel latent features. The overall compression rate is 16.
2. **DDPM-based Latent Diffusion Model (LDM):** MAISI uses a DDPM-based LDM to generate synthetic latent volumes. Given a latent z , the model learns to denoise

noisy versions z_t over time steps t . The training objective minimizes the standard denoising loss:

$$\mathcal{L}_{\text{ddpm}} = \mathbb{E}_{z, \epsilon, t} \left[\|\epsilon - \epsilon_\theta(z_t, c, t)\|^2 \right] \quad (1)$$

where c includes body region index and voxel spacing conditioning inputs.

3. **ControlNet:** To provide fine-grained structural guidance, MAISI incorporates ControlNet modules that encode the condition mask and inject it into the LDM through cross-attention or feature modulation. This allows the generation process to be explicitly steered toward the desired anatomy.

Despite its flexibility, MAISI still suffers from slow inference due to DDPM sampling and lacks mechanisms to enforce precise condition fidelity during generation.

Rectified Flow

Rectified Flow is a generative model that learns transport maps between two probability distributions by modeling straight, deterministic paths, referred to as "flows". In theory, there are infinitely many possible paths to map a source distribution π_0 to a target distribution π_1 . Traditional diffusion models rely on stochastic processes, which typically result in curved or noisy trajectories, requiring many iterative steps during inference to follow such a path from noise to data.

In contrast, Rectified Flow formulates the transport as a deterministic ordinary differential equation (ODE), encouraging the flow trajectory to be as straight as possible. This substantially reduces the number of required inference steps and accelerates the generation process. It has been adopted in several state-of-the-art (SOTA) frameworks, including Stable Diffusion 3 (Esser et al. 2024) and Open-Sora (Peng et al. 2025).

Given samples $x_0 \sim \pi_0$ and $x_1 \sim \pi_1$, Rectified Flow aims to learn a time-dependent velocity field $v_t(x)$ such that the solution $\phi_t(x)$ to the ODE:

$$\frac{d\phi_t(x)}{dt} = v_t(\phi_t(x)), \quad \phi_0(x) = x_0, \quad \phi_1(x) = x_1 \quad (2)$$

The model is trained to minimize the deviation of $v_t(x)$ from the ideal straight-line displacement between x_0 and x_1 over time. Specifically, the optimization objective is:

$$\mathcal{L}_{\text{flow}} = \int_0^1 \mathbb{E}_{x_0, x_1, t} \left[\|v_t(x_t, c) - (x_1 - x_0)\|^2 \right] dt \quad (3)$$

where $x_t = (1-t)x_0 + tx_1$ denotes the linear interpolation between the two endpoints. By encouraging $v_t(x)$ to match the constant velocity of a straight path, the model learns to directly connect π_0 and π_1 in a more efficient manner, enabling high-quality image synthesis with significantly fewer steps than traditional diffusion methods.

Methodology

Overview of MAISI-v2

MAISI-v2 builds on the MAISI architecture, by replacing the DDPM-based LDM with a rectified flow model and introducing two loss functions that improve condition fidelity. The training pipeline is similar to MAISI, which is already described in **Related Work**, consists of:

- **VAE:** We reuse the pretrained MAISI VAE to compress 3D volumes without fine-tuning.
- **Rectified Flow-based LDM:** We replace the original DDPM with a rectified flow model to accelerate generation, conditioned on the voxel spacing.
- **ControlNet with Region-specific Contrastive Loss:** The control branch encodes condition masks and guides synthesis through modulation. It is trained with added region-specific contrastive losses, conditioned on the voxel spacing and the segmentation mask.

Training Data

MAISI (Guo et al. 2025), comprising over 10k CT scans. As in MAISI, all images were resampled to dimensions divisible by 128, including 128, 256, 384, 512, and 768 in each axis. There are three differences compared to the original MAISI training data: **(1)** We incorporated the HN-SCC dataset (Grossberg et al. 2020) to improve coverage of the head-and-neck region, as well as an additional in-house dataset containing 1.4k chest and abdomen CT scans. Together, these additions increased the total dataset size from 10k to 12k scans. **(2)** We removed CT scans with fewer than 64 slices to ensure data quality, resulting in 11k usable scans. **(3)** We downsampled high-resolution images to augment the representation of low-resolution cases. The final training set consists of over 107k images generated from the 11k scans.

For ControlNet training, same with MAISI (Guo et al. 2025), we use CT images and the corresponding pseudo segmentation maps with 127 types of human anatomical structures and various lesions, which are obtained from the open-source model VISTA3D (He et al. 2024). For some datasets, manual labels are available for a few organs or tumors. In this case, we replace the corresponding pseudo labels with real labels. This setup allows the ControlNet module to learn spatially accurate synthesis conditioned on semantic labels.

Rectified Flow-based LDM Training

Naively training images with varying shapes all together will cause the batch size to be restricted to 1. And small batch size will make the LDM training process prone to NaN issues for mixed-precision optimization. To improve both stability and convergence speed for training, we developed a three-stage training strategy with mixed-precision optimization. The only conditioning signal is the voxel spacing, which is automatically extracted from the NIFTI header and requires no manual annotation. Training is conducted on 64 NVIDIA A100 GPUs (80GB) over three weeks using the AdamW optimizer.

The three training stages are outlined below.

Pre-training Stage: In the first stage, we use low-resolution images of size $128 \times 128 \times 128$, consisting of both naturally low-resolution volumes and high-resolution images that have been downsampled. Since all images share the same shape, we are able to use a large batch size of 96 on 80GB A100 GPUs. This stage is critical as it not only accelerates training significantly but also helps prevent NaN issues that may arise in subsequent stages. We use a learning rate of $1e-3$ to train this stage for one day.

Main Stage: Once the low-resolution model converges, we move to the second stage and continue training on the full-resolution dataset, where images vary in spatial dimensions. To avoid using batch size as 1, we use bucketed data parallelism. We group images by shape and distribute them to separate GPUs, ensuring that each GPU processes only images of the same shape, and then train the model with distributed data parallel. This allows us to use larger batch sizes where possible. For example, $128 \times 128 \times 128$ images are trained with a batch size of 96, $256 \times 256 \times 128$ with a batch size of 24, and $512 \times 512 \times 768$ with a batch size of 1. We use a learning rate of $1e-3$ and polynomial decay for this stage for 16000 epochs. It took around 10 days. While this strategy greatly accelerates training, it introduces data imbalance issue.

Fine-Tuning Stage: To address the data imbalance issue from the second stage, we add a third fine-tuning stage, where the model is trained using all images mixed together with a fixed batch size of 1 and sampling weights to balance the contribution from different datasets. This ensures the model is exposed to the full data distribution and helps reduce any size-related bias in the final representation. We use a learning rate of $1e-4$ for this stage for around 2000 epochs. It took around 10 days.

ControlNet Training

The ControlNet training was done with 8 A100 GPUs, using AdamW optimizer and a learning rate of $5e-5$ and polynomial decay. We trained the network with 60 epochs and it takes around 2 days.

Let $x_0 \sim \pi_0$ be a latent Gaussian noise sample, and $x_1 \sim \pi_1$ be the latent encoding of a real image I . The conditioning input c (e.g., segmentation mask) is processed through a separate encoder E_θ and fused into the frozen diffusion model and form the generator G_θ which predicts velocity.

$$v_t(x, c) = G_\theta(x_t, c) \quad (4)$$

We train a time-dependent network $G_\theta(x_t, c)$, conditioned on both x_t and c , to minimize the deviation from the linear displacement, as shown in Equation 3.

In our setup, the conditioning input c is a segmentation mask. To emphasize the response in regions of interest (ROIs), specifically tumors, we applied a weighted average loss with a tumor weight of 100. However, this alone was insufficient to ensure clear tumor appearance in the generated outputs. To address this limitation, we introduce a Region-Specific Contrastive Loss.

Region-specific Contrastive Loss To explicitly enforce condition fidelity, we propose a differential conditioning

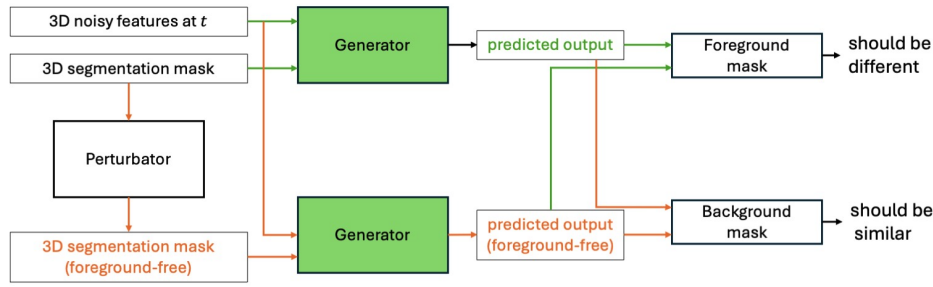


Figure 1: Region-specific Contrastive Loss

strategy that disentangles the model’s response to the region of interest (ROI) versus the background. The key idea is to generate two versions of the output from the same noise input: one conditioned on the original segmentation mask, and another on a perturbed mask where the ROI labels are replaced with the corresponding background labels (e.g., replacing a pancreatic tumor label with a pancreas label). The synthesized images should differ within the ROI (reflecting the conditioning change), while remaining consistent in the background (which is unaltered).

This results in a composite loss with two complementary objectives: 1) **ROI Sensitivity Loss**, which encourages discriminative synthesis in regions where the condition differs; 2) **Background Consistency Loss**, which enforces invariance in regions where the condition remains unchanged. We term this a **Region-specific Contrastive Loss**—not in the classical embedding sense, but as a mechanism that explicitly contrasts outputs under different conditioning inputs within the ROI.

The high-level idea is presented in Figure 1. The detailed procedure is described as follows.

Algorithm 1: Region-Specific Contrastive Loss Computation

- Require:** Generator G_θ , condition mask c_{orig} , cap δ
- 1: Sample noise $x_0 \sim \pi_0$
 - 2: Construct perturbed mask $c_{\text{perturb}} \leftarrow c_{\text{orig}}$: replace ROI labels with corresponding background labels, e.g., replacing pancreatic tumor label with pancreas label.
 - 3: Compute ROI binary mask $m \leftarrow \mathcal{K}(c_{\text{orig}} \neq c_{\text{perturb}})$
 - 4: Generate predictions with c_{orig} and c_{perturb} : $G_\theta(x_t, c_{\text{orig}})$ and $G_\theta(x_t, c_{\text{perturb}})$
 - 5: Compute ROI sensitivity loss \mathcal{L}_{roi} using Eq. (6)
 - 6: Compute background consistency loss \mathcal{L}_{bg} using Eq. (8)
 - 7: **return** $\mathcal{L}_{\text{roi}} + \mathcal{L}_{\text{bg}}$
-

$$\mathcal{D}_{\text{roi}} = \|(G_\theta(x_t, c_{\text{orig}}) - G_\theta(x_t, c_{\text{perturb}})) \odot m\|_{1,m} \quad (5)$$

$$\mathcal{L}_{\text{roi}} = -\min(\mathcal{D}_{\text{roi}}, \delta) \quad (6)$$

$$m^- = 1 - \text{dilate}(m) \quad (7)$$

$$\mathcal{L}_{\text{bg}} = \|(G_\theta(x_t, c_{\text{orig}}) - G_\theta(x_t, c_{\text{perturb}})) \odot m^-\|_{1,m^-} \quad (8)$$

As shown in Eq. (6), we apply an upper bound δ to prevent gradient explosion during training. This is important.

Removing this bound leads to unstable optimization and results in NaN values during training. Intuitively, the goal is not to maximize the difference within the ROI indefinitely, but merely to ensure the model can distinguish them. Since the latents were generated with a trained VAE encoder, we assume the values of feature maps approximately follows Normal Distribution. Therefore, we set $\delta = 2$, corresponding to a \pm standard deviation range.

Memory-Aware Region-specific Contrastive Loss

When GPU memory is a bottleneck, we need to reduce the memory usage of Region-specific Contrastive Loss for large image inputs. To accommodate varying input sizes and GPU memory constraints, we adaptively select the feature source for loss computation.

- For **small to medium-sized inputs**, we compute the loss on the outputs of the ControlNet and frozen diffusion model G_θ . These final output features provide higher spatial fidelity and stronger alignment with the conditioning mask. This is preferred if GPU memory allows.
- For **large inputs**, we apply the loss on the encoded features of the ControlNet encoder E_θ , without passing through the frozen diffusion model. Although these intermediate features are coarser, this strategy enables training on high-resolution volumes within memory limits.

In both cases, the loss terms \mathcal{L}_{roi} and \mathcal{L}_{bg} remain unchanged. The losses are computed either on the final outputs of the ControlNet and diffusion model, or on intermediate features from the ControlNet encoder, depending on input size and memory availability.

Final Objective The full training objective combines all components:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{flow}} + \lambda_{\text{contrast}}(\mathcal{L}_{\text{roi}} + \mathcal{L}_{\text{bg}}) \quad (9)$$

where $\lambda_{\text{contrast}}$ are hyperparameters controlling the weight of the condition fidelity terms.

Quality Check with CT HU Intensities We performed a quality check on the Hounsfield Unit (HU) intensities of generated CT images. Intuitively, the median HU values of major organs in generated images should fall within realistic physiological ranges.

To determine these ranges, we collected the median HU values for the liver, spleen, pancreas, kidneys, lungs, bones, and brain from all the training images. With these values, the

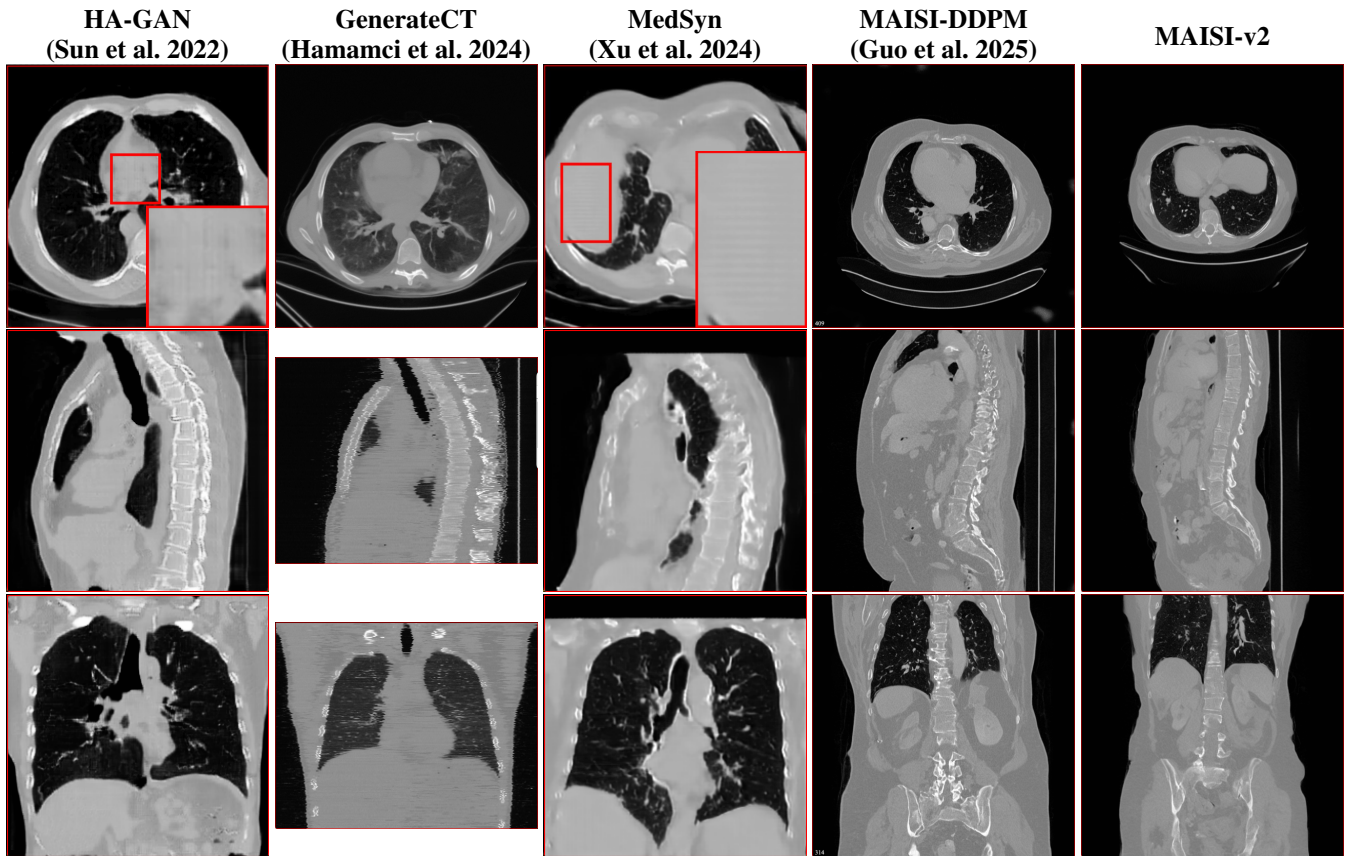


Figure 2: Qualitative comparison across axial (top row), sagittal (middle row), and coronal (bottom row) views. Columns correspond to different methods. MAISI-DDPM and MAISI-v2 in this figure are unconditional synthesis which do not use ControlNet or segmentation maps.

lower threshold for each organ was set as the more extreme of the 5th percentile or the 6-sigma lower bound, and the upper threshold was defined similarly using the 95th percentile and the 6-sigma upper bound.

This quality check is related to choosing $\lambda_{\text{contrast}}$. Large $\lambda_{\text{contrast}}$ leads to results with tumors well aligned with tumor masks. Yet if too large, the rest of the body structures will be wrong and the result will not be able to pass the quality check. Our criterion to choose $\lambda_{\text{contrast}}$ is a large $\lambda_{\text{contrast}}$ that can make nearly 100% of the generated images pass the quality check.

Experiments

We implement all networks using PyTorch (Ansel et al. 2024) and MONAI (Cardoso et al. 2022). All the evaluation experiments were done with 80G H100 GPUs.

Training Data and Implementation Details

The training dataset consists of many public datasets including HNSCC dataset (Grossberg et al. 2020), AbdomenCT-1K (Ma et al. 2021), AeroPath (Støverud et al. 2023), AMOS22 (Ji et al. 2022), COVID-19 (Chowdhury et al. 2020; Rahman et al. 2021), CRLM-CT (Simpson

et al. 2024), CT-ORG (Rister et al. 2019), CTPelvic1K-CLINIC (Liu et al. 2021), LIDC (Armato III et al. 2011), MSD Task03, Task06, Task07, Task08, Task09, Task10 (Antonelli et al. 2022), Multi-organ-Abdominal-CT (BTCV, Pancreas-CT) (Gibson et al. 2018; Landman et al. 2015; Roth et al. 2016), NLST (Team 2011), StonyBrook-CT (Saltz et al. 2021), TCIA Colon (Johnson et al. 2008), TotalSegmentatorV2 (Wasserthal 2023). We also added two in-house datasets, including 237 bone lesion CT and 1386 chest&abdomen CT to further enrich the training dataset, especially for CT images with bone lesions and corresponding masks.

For implementation details in ControlNet, we found that large $\lambda_{\text{contrast}}$ leads to results with tumors well aligned with masks. Yet if too large, the rest of the body structures will be wrong.

In order to tackle this issue, for the 60 training epochs, we empirically choose $\lambda_{\text{contrast}} = 0.01$ for the first 30 epochs to make sure the tumors appear in the generated results, then decrease to $\lambda_{\text{contrast}} = 0.001$ to make the rest of the body structure correct. We experimented with the opposite way, i.e., first use small $\lambda_{\text{contrast}}$ then increase, the tumors did not show up as expected. In the final models, 100% of the results passed the quality check on CT HU intensities.

Model	Steps	Time (s)	FID _{xy}	FID _{yz}	FID _{xz}	FID _{avg}	Device
HA-GAN	1	1	13.813	12.567	14.405	13.595	H100
MedSyn (2-stage DDIM)	50+20	100	18.662	22.171	33.293	24.709	H100
GenerateCT (2D EDM)	25×201	89	7.909	9.256	15.106	10.757	H100
MAISI (DDPM)	1000	198+15	2.199	2.480	2.642	2.441	H100
MAISI (DDIM)	30	6+15	4.855	4.703	4.770	4.776	H100
MAISI-v2 (Rectified Flow)	30	6+15	2.217	2.211	2.538	2.322	H100

Table 1: FID scores and single 80G GPU inference time cost on the OOD dataset AutoPET2023 across orthogonal planes and overall average (lower is better). All synthetic images are $512 \times 512 \times 512$ volumes with voxel spacing of $1 \times 1 \times 1$ mm³. For MAISI and MAISI-v2, the time cost includes 30 steps LDM (6s) or 1000 steps LDM (198s), and VAE decoding (15s) on H100.

Method (50 steps inference)	Output Size	# Voxels	Time (s)	Device
SVD (Blattmann et al. 2023) (Conditioned)	$3 \times 576 \times 1024 \times 25$ (video)	4.4E7	100	A100
Open Sora 2.0 (Peng et al. 2025) (Conditioned)	$3 \times 768 \times 768 \times 128$ (video)	2.3E8	162	H200
MAISI-v2 (Unconditioned)	$512 \times 512 \times 512$ (image)	1.3E8	26	H100
MAISI-v2 (Conditioned)	$512 \times 512 \times 512$ (image)	1.3E8	34	H100

Table 2: Inference time comparison across different methods and image/video sizes on single 80G GPU. MAISI-v2 (Unconditioned) refers to the time cost of LDM and VAE decoding. MAISI-v2 (Conditioned) refers to the time cost of ControlNet, LDM, and VAE decoding. All diffusion models in this table have 50 inference steps. The devices are different since we used the officially reported time cost from the original papers.

Metric	5	10	20	30	40	50	70	100
FID _{xy}	22.774	4.232	2.538	2.217	2.088	1.966	1.818	1.783
FID _{yz}	13.907	3.765	2.446	2.211	2.119	2.001	1.970	1.960
FID _{xz}	24.321	5.266	2.952	2.538	2.405	2.225	2.156	2.159
FID _{avg}	20.334	4.421	2.645	2.322	2.204	2.064	1.981	1.967

Table 3: FID scores across inference steps for MAISI-v2 LDM. Each value represents the score on the OOD AutoPET23 dataset (lower is better).




Figure 3: MAISI-v2 segmentation-guided results for small to large volume size and three different regions.

Evaluation of Accelerated LDM

We first evaluate the performance of Rectified Flow LDM, without involving the ControlNet.

For comparison methods, to the best of our knowledge, the only methods that can also generate high-resolution 3D CT images with size equal or larger than $512 \times 512 \times 512$ are MAISI (Guo et al. 2025) and 3D MedDiffusion (Wang et al.

2024), while only MAISI is open-sourced. For MAISI and MAISI-v2, they support various sizes and voxel spacings. We generate synthetic images with size of $512 \times 512 \times 512$ and voxel spacing of $1 \times 1 \times 1 \text{ mm}^3$ for evaluation purpose to show its ability to cover a large body region. Other methods generate images with smaller volume size. HA-GAN (Sun et al. 2022) generates images with a fixed size of $256 \times 256 \times 256$ and unknown voxel spacing as their training data were zoomed to $256 \times 256 \times 256$ without considering the voxel spacing. MedSyn (Xu et al. 2024) generates images with a fixed size of $256 \times 256 \times 256$ and fixed voxel spacing of $1 \times 1 \times 1 \text{ mm}^3$. GenerateCT (Hamamci et al. 2024) generates images with a fixed size of $512 \times 512 \times 201$ and fixed voxel spacing of $0.75 \times 0.75 \times 1.5 \text{ mm}^3$. GenerateCT first generates low-resolution 3D volumes, then applies 2D slice-by-slice upscaling, which causes its low-quality in sagittal and coronal planes.

Qualitative Evaluation: Figure 2 presents representative slices from the axial, sagittal, and coronal planes. GenerateCT (Hamamci et al. 2024) is a 2D model, so it lacks interslice consistency, leading to poor image quality in the sagittal and coronal views. MedSyn (Xu et al. 2024) produces noticeably blurry results with mosaic-like artifacts, such as region inside the red box. HA-GAN (Sun et al. 2022) generates visually sharp images but with mosaic-like artifacts, such as region inside the red box. Also, its voxel spacing is not available, which limits its applicability in real-world medical imaging tasks. Moreover, all three methods are restricted to synthesizing small anatomical regions. In contrast, both MAISI and MAISI-v2 are capable of generating high-quality 3D volumes that span larger body regions while preserving fine anatomical details and realistic structure.

Quantitative Evaluation: We use the AutoPET2023 dataset (Gatidis et al. 2022) as an out-of-distribution (OOD) benchmark to evaluate their generalizability.

For OOD dataset AutoPET2023, it contains a wide range of image resolution. In order to evaluate high-resolution image synthesis, we select original AutoPET2023 images with voxel spacing smaller than $1 \times 1 \times 1 \text{ mm}^3$ and FOV larger than $512 \times 512 \times 512 \text{ mm}^3$, then downsample them to $1 \times 1 \times 1 \text{ mm}^3$ and central crop $512 \times 512 \times 512$ volumes.

For MAISI and MAISI-v2, We generate synthetic images with size of $512 \times 512 \times 512$ and voxel spacing of $1 \times 1 \times 1 \text{ mm}^3$. For other comparison methods, the size of the generated images is fixed. We resampled them to $1 \times 1 \times 1 \text{ mm}^3$ and zero-padded to $512 \times 512 \times 512$. MAISI was trained with 1000 step DDPM. But there are schedulers that can accelerate inference like Denoising Diffusion Implicit Models (DDIM) (Song, Meng, and Ermon 2020) and Elucidated Diffusion Sampler (EDM) (Karras et al. 2022). We tested DDIM scheduler for MAISI.

The metric we use is 2D Fréchet Inception Distance (FID) between synthesized images and the unseen AutoPET2023 dataset for xy , yz , and xz -planes. The FID score measures how closely the distribution of generated samples matches the real data distribution. Lower values indicate higher fidelity and better generalization. To avoid empty slices, we compute FID on the central 50% of the slices. The reason

we use 2D FID in three planes instead of 3D FID is to better reflect human perception.

Table 1 summarizes the results. It shows that MAISI-v2 achieves at least comparable FID with MAISI (DDPM) with much faster speed. For MAISI and MAISI-v2, the time cost in Table 1 is the total of the LDM and VAE decoding time cost. The LDM generates $4 \times 128 \times 128 \times 128$ latent features. They are then decoded by MAISI VAE, which takes 15s to generate a $512 \times 512 \times 512$ volume on H100 using a sliding window patch size of $80 \times 80 \times 80$ on latent features.

While Table 1 compares MAISI-v2’s inference speed with other medical image generation methods, Table 2 extends the comparison to SOTA video generation models, and provides additional perspective on the LDM inference efficiency across domains. Specifically, we include Stable Video Diffusion (SVD) (Blattmann et al. 2023) and OpenSora 2.0 (Peng et al. 2025), both of which use 50 inference steps. To ensure consistency, we also report MAISI-v2’s runtime using 50 inference steps.

Ablation Study: We conduct an ablation study to analyze how the number of inference steps affects the quality of images synthesized by the rectified flow-based MAISI-v2 model. Same as Table 1, FID with the same AutoPET23 reference images is computed along three orthogonal planes (xy , yz , and xz), as well as the averaged score (FID_{avg}).

As shown in Table 3, image quality improves consistently as the number of inference steps increases from 5 to 100, with diminishing returns beyond 30 steps.

Based on this analysis, we choose 30 inference steps as the default setting for most experiments in this paper. This setting offers a favorable balance between synthesis quality and computational efficiency.

Evaluation of ControlNet with Additional Losses

Qualitative Results: Figure 3 shows the generalizability of MAISI-v2 ControlNet for different body regions and voxel sizes. Figure 4 shows qualitative results for MAISI-v2 ControlNet on 5 types of tumors.

Quantitative Evaluation: To assess whether MAISI-generated images can improve downstream tasks, we perform a 5-fold segmentation data augmentation study.

We trained 3D segmentation networks with open-source tool Auto3Dseg¹ on: (1) **Baseline:** Real training images only. (2) **Augmented:** Real training images + same number of segmentation mask-guided synthetic images using different models, including DiffTumor (Chen et al. 2024), MAISI, and the proposed MAISI-v2.

During inference, the input masks used for synthesis are not original tumor masks. Instead, we use those with tumors augmented with erosion, dilation and translation. Note that both MAISI and MAISI-v2 are single models that can generate five types of tumors, rather than five models for five types of tumors. DiffTumor (Chen et al. 2024) is a diffusion model which was optimized specifically for tumor inpainting on the abdomen region for CT images. It takes healthy images and

¹<https://monai.io/apps/auto3dseg>. All the Auto3Dseg model training in this section share the same hyper-parameters.

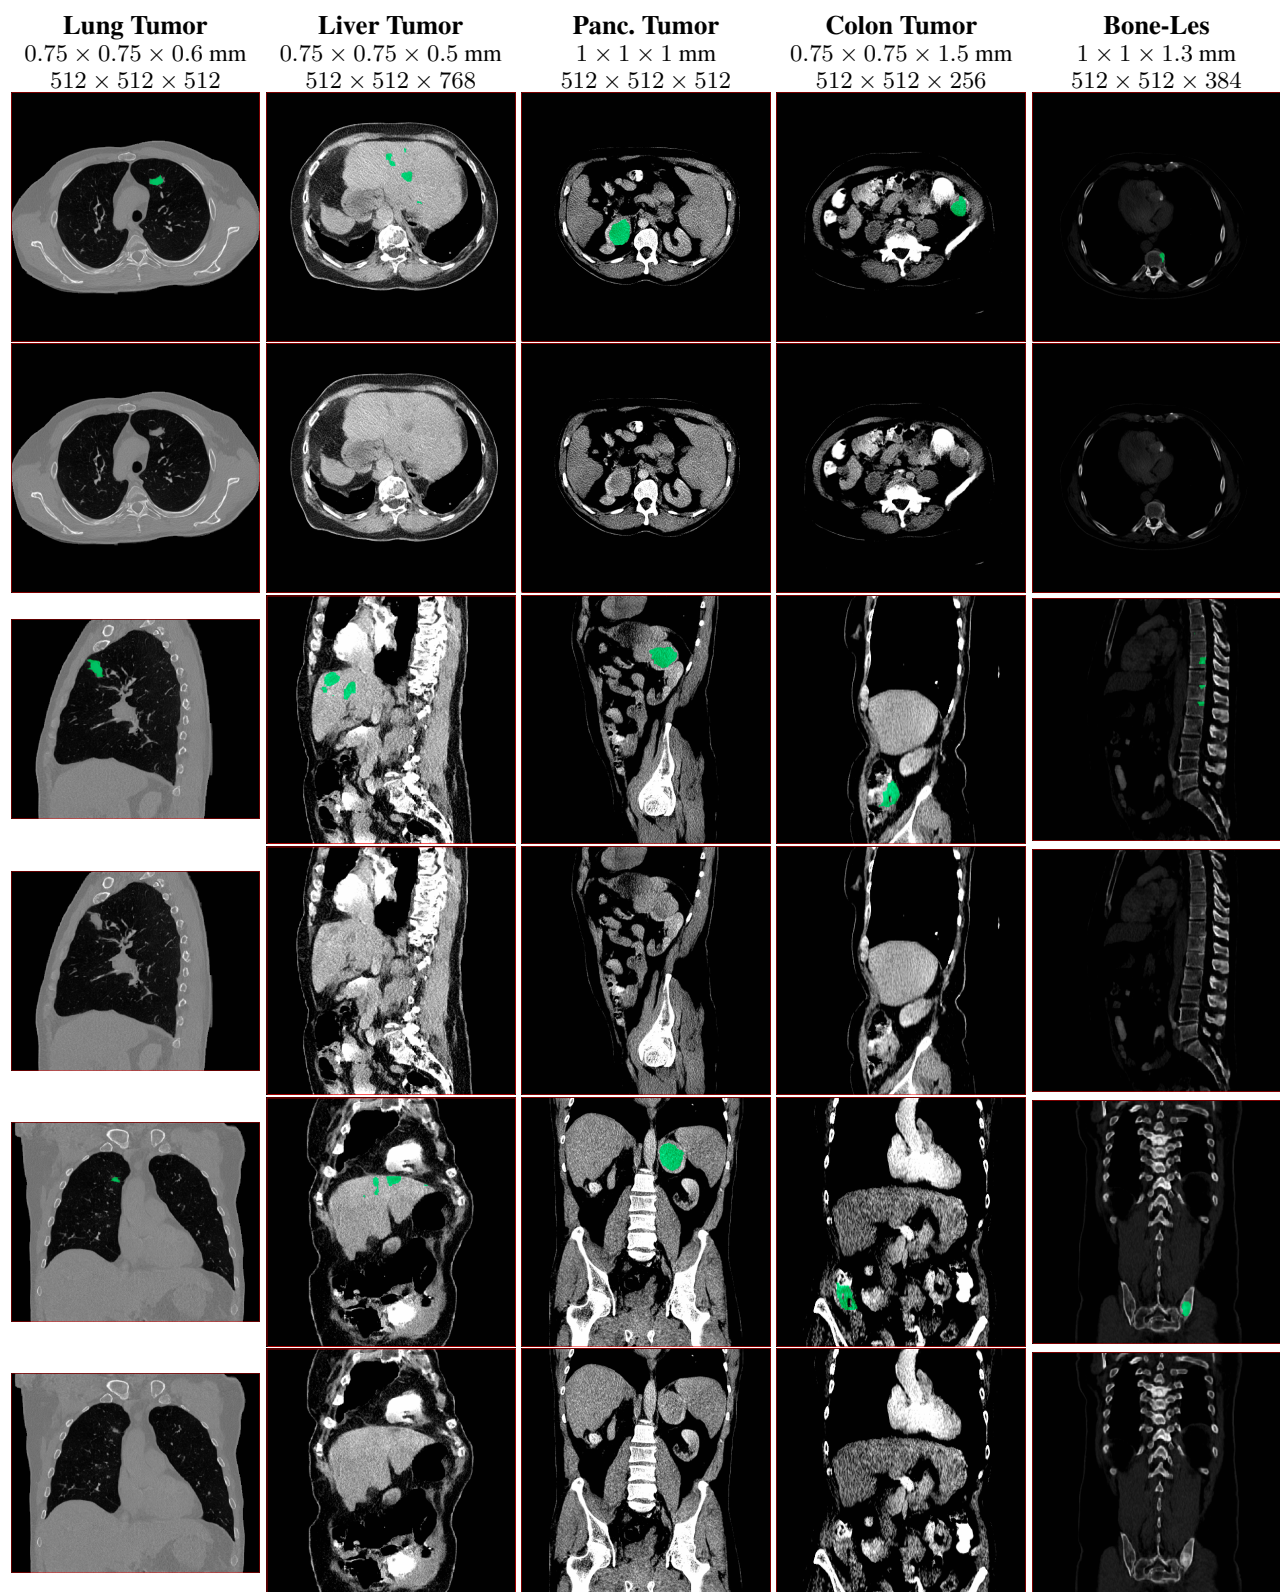


Figure 4: MAISI-v2 segmentation-guided results for five types of tumors. We show results for different voxel spacing and volume size to demonstrate the flexibility of MAISI-v2. Different Hounsfield Unit window is used to better show the contrast between tumor and normal tissues.

Method	Liver Tumor	Lung Tumor	Panc. Tumor	Colon Tumor	Bone-Les
Real Only	0.662	0.581	0.433	0.449	0.504
DiffTumor	0.684 (+2.2%)	– (–%)	0.511 (+7.9%)	– (–%)	– (–%)
MAISI	0.688 (+2.6%)	0.635 (+5.5%)	0.482 (+4.9%)	0.485 (+3.6%)	0.539 (+3.6%)
Ablation study: “MAISI-v2” refers to ControlNet + Contra. loss. † means $p < 0.1$, * means $p < 0.001$					
ControlNet	0.693 (+3.0%)	0.627 (+4.7%)	0.484 (+5.1%)	0.402 (-4.7%)	0.520 (+1.6%)
MAISI-v2	0.695 (+3.3%)	0.655 (+7.5%)†	0.497 (+6.4%)†	0.491 (+4.2%)*	0.537 (+3.3%)*

Table 4: Segmentation Dice scores for five tumor tasks on the test sets. The baseline uses real images only. Others use real + synthetic images generated with different ControlNet training losses. “(X%)” indicates improvement over the “Real Only” baseline. “Panc. Tumor” refers to Pancreatic Tumor.

tumor masks as input and outputs images with inpainted tumors, while MAISI and MAISI-v2 only take segmentation masks as the input.

We compare the Dice scores on a held-out test set, reported on four public datasets—hepatic tumor (Task03) (Antonelli et al. 2022), lung tumor (Task06) (Antonelli et al. 2022), pancreatic tumor (Task07) (Antonelli et al. 2022), and colon tumor (Task10) (Antonelli et al. 2022), as well as one in-house dataset for bone lesions (Bone-Les), shown in Table 4.

Results indicate that synthetic data generated by MAISI-v2 improves segmentation accuracy across multiple tumors, demonstrating its potential as a practical data augmentation tool in low-data scenarios.

Ablation Study: We conduct an ablation study on the effect of the two proposed losses. As shown in Table 4, it includes (1) **ControlNet**, with weighted average loss on tumors ($w = 100$) and (2) **ControlNet with contrastive loss**, which is the proposed MAISI-v2. The final column reports the average improvement in Dice score over the real-only baseline. We performed a paired t-test on the values from the two methods. The improvement is significant for 4 out of 5 types of tumors.

Conclusion and Discussion

Tumor alignment for Rectified Flow and DDPM: We found that Rectified Flow ControlNet produces weaker tumor alignment than DDPM ControlNet during inference, especially for colon tumor. Two factors may contribute to this effect. First, Rectified Flow’s deterministic and straight-line transport formulation tends to reduce trajectory diversity (Ma et al. 2025), which may make the model less robust to small conditioning perturbations. Second, because it learns a single deterministic velocity field, prediction errors in small or low-contrast regions can accumulate and amplify during integration. In contrast, DDPM introduces stochastic noise at each reverse step, which helps diffuse such local errors and maintain robustness. Our region-specific contrastive loss mitigates both issues by enforcing local feature consistency and encouraging diverse mappings in the tumor regions. Future work will investigate the underlying mechanisms behind these observations.

Limitations and Future Work: This work has three main limitations. (1) We only trained on CT scans. In the future, we plan to include MRI. (2) We only tested the model on segmentation tasks. Future work will include other applications like image translation, inpainting, etc. (3) Few existing methods support high-resolution 3D synthesis due to GPU memory limitation. Our method is already designed and optimized to train and infer on 24G GPU for small images like $512 \times 512 \times 256$, it still needs 40G for inference and 80G for training for large images like $512 \times 512 \times 768$. In future work, we will further reduce memory use.

Broader Impacts: We share our code, pretrained models, and training details, to make it easier for others to use or fine-tune. Like many other open-sourced large models, we hope to lower the barrier and reduce computing cost for future research in this community. The synthetic images from our model are not meant for clinical diagnosis. They are only used to help train other models, such as for segmentation, as shown in this paper.

Summary: MAISI-v2 is an efficient and flexible method for 3D medical image synthesis. It uses rectified flow and new training losses to improve speed and condition accuracy. Our results show it creates high-quality images and helps improve segmentation models.

Acknowledgments

We thank Weili Nie and Arash Vahdat for valuable discussions.

References

- Ansel, J.; Yang, E.; He, H.; Gimelshein, N.; Jain, A.; Voznesensky, M.; Bao, B.; Bell, P.; Berard, D.; Burovski, E.; et al. 2024. Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, 929–947.
- Antonelli, M.; Reinke, A.; Bakas, S.; Farahani, K.; Kopp-Schneider, A.; Landman, B. A.; Litjens, G.; Menze, B.; Ronneberger, O.; Summers, R. M.; et al. 2022. The medical segmentation decathlon. *Nature communications*, 13(1): 4128.

- Armato III, S. G.; McLennan, G.; Bidaut, L.; McNitt-Gray, M. F.; Meyer, C. R.; Reeves, A. P.; Zhao, B.; Aberle, D. R.; Henschke, C. I.; Hoffman, E. A.; et al. 2011. The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Medical physics*, 38(2): 915–931.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Cardoso, M. J.; Li, W.; Brown, R.; Ma, N.; Kerfoot, E.; Wang, Y.; Murrey, B.; Myronenko, A.; Zhao, C.; Yang, D.; et al. 2022. Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701*.
- Chartsias, A.; Joyce, T.; Dharmakumar, R.; and Tsafaris, S. A. 2017. 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*, 3–13. Springer.
- Chen, Q.; Chen, X.; Song, H.; Xiong, Z.; Yuille, A.; Wei, C.; and Zhou, Z. 2024. Towards generalizable tumor synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11147–11158.
- Chowdhury, M. E.; Rahman, T.; Khandakar, A.; Mazhar, R.; Kadir, M. A.; Mahbub, Z. B.; Islam, K. R.; Khan, M. S.; Iqbal, A.; Al Emadi, N.; et al. 2020. Can AI help in screening viral and COVID-19 pneumonia? *Ieee Access*, 8: 132665–132676.
- Dewey, B. E.; Zhao, C.; Reinhold, J. C.; Carass, A.; Fitzgerald, K. C.; Sotirchos, E. S.; Saidha, S.; Oh, J.; Pham, D. L.; Calabresi, P. A.; et al. 2019. DeepHarmony: A deep learning approach to contrast harmonization across scanner changes. *Magnetic resonance imaging*, 64: 160–170.
- Dorjsembe, Z.; Pao, H.-K.; Odonchimed, S.; and Xiao, F. 2024. Conditional diffusion models for semantic 3d brain mri synthesis. *IEEE Journal of Biomedical and Health Informatics*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Friedrich, P.; Wolleb, J.; Bieder, F.; Durrer, A.; and Cattin, P. C. 2024. Wdm: 3d wavelet diffusion models for high-resolution medical image synthesis. In *MICCAI Workshop on Deep Generative Models*, 11–21. Springer.
- Gatidis, S.; Hepp, T.; Früh, M.; La Fougère, C.; Nikolaou, K.; Pfannenber, C.; Schölkopf, B.; Küstner, T.; Cyran, C.; and Rubin, D. 2022. A whole-body FDG-PET/CT dataset with manually annotated tumor lesions. *Scientific Data*, 9(1): 601.
- Gibson, E.; Giganti, F.; Hu, Y.; Bonmati, E.; Bandula, S.; Gurusamy, K.; Davidson, B.; Pereira, S. P.; Clarkson, M. J.; and Barratt, D. C. 2018. Automatic multi-organ segmentation on abdominal CT with dense V-networks. *IEEE transactions on medical imaging*, 37(8): 1822–1834.
- Grossberg, A.; Elhalawani, H.; Mohamed, A.; Mulder, S.; Williams, B.; White, A. L.; Zafereo, J.; Wong, A. J.; Berends, J. E.; AboHashem, S.; et al. 2020. HNSCC Version 4. *The Cancer Imaging Archive*.
- Guo, P.; Zhao, C.; Yang, D.; Xu, Z.; Nath, V.; Tang, Y.; Simon, B.; Belue, M.; Harmon, S.; Turkbey, B.; et al. 2025. Maisi: Medical ai for synthetic imaging. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 4430–4441. IEEE.
- Hamamci, I. E.; Er, S.; Sekuboyina, A.; Simsar, E.; Tezcan, A.; Simsek, A. G.; Esirgun, S. N.; Almas, F.; Doğan, I.; Dasedelen, M. F.; et al. 2024. Generatect: Text-conditional generation of 3d chest ct volumes. In *European Conference on Computer Vision*, 126–143. Springer.
- He, Y.; Guo, P.; Tang, Y.; Myronenko, A.; Nath, V.; Xu, Z.; Yang, D.; Zhao, C.; Simon, B.; Belue, M.; et al. 2024. Vista3d: Versatile imaging segmentation and annotation model for 3d computed tomography. *arXiv preprint arXiv:2406.05285*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ji, Y.; Bai, H.; Ge, C.; Yang, J.; Zhu, Y.; Zhang, R.; Li, Z.; Zhanng, L.; Ma, W.; Wan, X.; et al. 2022. Amos: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation. *Advances in neural information processing systems*, 35: 36722–36732.
- Johnson, C. D.; Chen, M.-H.; Toledano, A. Y.; Heiken, J. P.; Dachman, A.; Kuo, M. D.; Menias, C. O.; Siewert, B.; Cheema, J. I.; Obregon, R. G.; et al. 2008. Accuracy of CT colonography for detection of large adenomas and cancers. *New England Journal of Medicine*, 359(12): 1207–1217.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577.
- Landman, B.; Xu, Z.; Igelsias, J. E.; Styner, M.; Langerak, T.; and Klein, A. 2015. Segmentation outside the cranial vault challenge. In *MICCAI: multi Atlas labeling beyond cranial vault-workshop challenge*.
- Li, M.; Yang, T.; Kuang, H.; Wu, J.; Wang, Z.; Xiao, X.; and Chen, C. 2024. ControlNet++: Improving Conditional Controls with Efficient Consistency Feedback: Project Page: liming-ai. github. io/ControlNet_Plus.Plus. In *European Conference on Computer Vision*, 129–147. Springer.
- Liu, P.; Han, H.; Du, Y.; Zhu, H.; Li, Y.; Gu, F.; Xiao, H.; Li, J.; Zhao, C.; Xiao, L.; et al. 2021. Deep learning to segment pelvic bones: large-scale CT datasets and baseline models. *International Journal of Computer Assisted Radiology and Surgery*, 16(5): 749–756.
- Liu, X.; Gong, C.; and Liu, Q. 2022. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*.

- Ma, J.; Zhang, Y.; Gu, S.; Zhu, C.; Ge, C.; Zhang, Y.; An, X.; Wang, C.; Wang, Q.; Liu, X.; et al. 2021. Abdomenct-1k: Is abdominal organ segmentation a solved problem? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6695–6714.
- Ma, Z.; Liu, R.; Liu, S.; Li, J.; and Zhou, B. 2025. Flow Diverse and Efficient: Learning Momentum Flow Matching via Stochastic Velocity Field Sampling. *arXiv preprint arXiv:2506.08796*.
- Nie, D.; Trullo, R.; Lian, J.; Petitjean, C.; Ruan, S.; Wang, Q.; and Shen, D. 2017. Medical image synthesis with context-aware generative adversarial networks. 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, 417–425. Springer.
- Peng, X.; Zheng, Z.; Shen, C.; Young, T.; Guo, X.; Wang, B.; Xu, H.; Liu, H.; Jiang, M.; Li, W.; et al. 2025. Open-sora 2.0: Training a commercial-level video generation model in \$200 k. *arXiv preprint arXiv:2503.09642*.
- Pinaya, W. H.; Tudosiu, P.-D.; Dafflon, J.; Da Costa, P. F.; Fernandez, V.; Nachev, P.; Ourselin, S.; and Cardoso, M. J. 2022. Brain imaging generation with latent diffusion models. In *MICCAI Workshop on Deep Generative Models*, 117–126. Springer.
- Rahman, T.; Khandakar, A.; Qiblawey, Y.; Tahir, A.; Kiranyaz, S.; Kashem, S. B. A.; Islam, M. T.; Al Maadeed, S.; Zughair, S. M.; Khan, M. S.; et al. 2021. Exploring the effect of image enhancement techniques on COVID-19 detection using chest X-ray images. *Computers in biology and medicine*, 132: 104319.
- Rister, B.; Shivakumar, K.; Nobashi, T.; and Rubin, D. L. 2019. CT-ORG: CT volumes with multiple organ segmentations [Dataset]. *The Cancer Imaging Archive*, 21.
- Roth, H.; Farag, A.; Turkbey, E. B.; Lu, L.; Liu, J.; and Summers, R. M. 2016. Data from pancreas-CT. (*No Title*).
- Saltz, J.; Saltz, M.; Prasanna, P.; Moffitt, R.; Hajagos, J.; Bremer, E.; Balsamo, J.; and Kurc, T. 2021. Stony brook university covid-19 positive cases. (*No Title*).
- Simpson, A. L.; Peoples, J.; Creasy, J. M.; Fichtinger, G.; Gangai, N.; Keshavamurthy, K. N.; Lasso, A.; Shia, J.; D’Angelica, M. I.; and Do, R. K. 2024. Preoperative CT and survival data for patients undergoing resection of colorectal liver metastases. *Scientific Data*, 11(1): 172.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Støverud, K.-H.; Bouget, D.; Pedersen, A.; Leira, H. O.; Langø, T.; and Hofstad, E. F. 2023. AeroPath: An airway segmentation benchmark dataset with challenging pathology. *arXiv preprint arXiv:2311.01138*.
- Sun, L.; Chen, J.; Xu, Y.; Gong, M.; Yu, K.; and Batmanghelich, K. 2022. Hierarchical amortized GAN for 3D high resolution medical image synthesis. *IEEE journal of biomedical and health informatics*, 26(8): 3966–3975.
- Team, N. L. S. T. R. 2011. Reduced lung-cancer mortality with low-dose computed tomographic screening. *New England Journal of Medicine*, 365(5): 395–409.
- Wang, H.; Liu, Z.; Sun, K.; Wang, X.; Shen, D.; and Cui, Z. 2024. 3D MedDiffusion: A 3D Medical Diffusion Model for Controllable and High-quality Medical Image Generation. *arXiv preprint arXiv:2412.13059*.
- Wang, H.; Liu, Z.; Sun, K.; Wang, X.; Shen, D.; and Cui, Z. 2025. 3D MedDiffusion: A 3D Medical Latent Diffusion Model for Controllable and High-quality Medical Image Generation. *IEEE Transactions on Medical Imaging*.
- Wasserthal, J. 2023. Dataset with segmentations of 117 important anatomical structures in 1228 CT images. *Zenodo* <https://doi.org/10.5281/ZENODO,6802613>.
- Xu, Y.; Sun, L.; Peng, W.; Jia, S.; Morrison, K.; Perer, A.; Zandifar, A.; Visweswaran, S.; Eslami, M.; and Batmanghelich, K. 2024. Medsyn: Text-guided anatomy-aware synthesis of high-fidelity 3d ct images. *IEEE Transactions on Medical Imaging*.
- Yang, H.; Sun, J.; Carass, A.; Zhao, C.; Lee, J.; Prince, J. L.; and Xu, Z. 2020. Unsupervised MR-to-CT synthesis using structure-constrained CycleGAN. *IEEE transactions on medical imaging*, 39(12): 4249–4261.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhao, C.; Carass, A.; Lee, J.; He, Y.; and Prince, J. L. 2017. 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*, 291–298. Springer.
- Zhu, L.; Codella, N.; Chen, D.; Jin, Z.; Yuan, L.; and Yu, L. 2024. Generative enhancement for 3d medical images. *arXiv preprint arXiv:2403.12852*.
- Zhu, L.; Xue, Z.; Jin, Z.; Liu, X.; He, J.; Liu, Z.; and Yu, L. 2023. Make-a-volume: Leveraging latent diffusion models for cross-modality 3d brain mri synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 592–601. Springer.