

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 NOISE AND ANATOMY-GUIDED DIFFUSION MODEL FOR REALISTIC CT IMAGE SYNTHESIS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Generative models, specifically Diffusion Models (DMs), have been quite successful in generating high-quality images. However, DMs rely on large-scale training data. In medical imaging, more specifically for computed tomography (CT), these models struggle in accurately reconstructing anatomical structures due to limited training data. This can cause the wrong depiction of organs, which can impact clinical treatment. Some existing models, although guided by anatomical structures, ignore dose-dependent noise, which is critical in real-world scenarios. To tackle this challenge, we propose a novel diffusion model, namely NA-Diff, which is guided by noise from different dose levels and anatomical structures, leveraging a dual conditional diffusion framework. To facilitate large-scale training of DMs on complex structured CT data, we transform natural images emulating realistic CT noises and leverage them for pre-training, followed by fine-tuning on small CT data. Extensive experimental results demonstrate that NA-Diff generates high-fidelity and noise-aware CT images, effectively delineating the organ-of-interest and bridging the gap between synthetic and real CT.<sup>1</sup>

## 1 INTRODUCTION

Computed tomography (CT) serves a crucial role in modern medical imaging, providing detailed anatomical visualization for diagnosis and treatment planning. Each year, more than 80 million CT scans are performed in the United States, and the number keeps growing rapidly Schultz et al. (2020). However, obtaining high-quality CT images is constrained by radiation dose trade-offs, motion artifacts, resolution limitations, and scanner technology. Acquiring large labeled medical imaging datasets can also be challenging due to privacy concerns and expensive annotation processes.

Generative models have shown promise in medical image synthesis, with approaches such as generative adversarial networks (GANs) Goodfellow et al. (2020); Mirza & Osindero (2014) and variational autoencoders (VAEs) Kingma et al. (2019) constructing realistic images. However, these models often suffer from blurring artifacts and difficulties in maintaining fine anatomical structures. In recent years, diffusion models Rombach et al. (2022); Ho et al. (2020) emerged as a strong alternative for image synthesis by modeling a structured denoising process. Unlike GANs, diffusion models (DMs) have more stable training dynamics and produce detailed and diverse outputs. DMs have been successfully adapted to various domains, including medical imaging Bhattacharya et al. (2024); Konz et al. (2024); al Nomaan Nafi et al. (2024); Munia & Imran (2025), specifically, conditional ones over their unconditional counterparts. Conditional DMs rely on a large amount of labeled data for effective training. To mitigate this dependency, alternatives such as self-supervised learning, weak supervision, data augmentation, and synthetic data generation are being explored Liu et al. (2024); Oh & Jeong (2024). In these scenarios, generated outputs may appear anatomically plausible but lack clinical precision, while augmentation can introduce hallucinations or structural distortions, making them inappropriate for medical applications.

For image synthesis tasks, generative models Goodfellow et al. (2020); Mirza & Osindero (2014); Rombach et al. (2022); Ho et al. (2020) portray remarkable performances, leveraging a satisfactory amount of large data for image generation. Yet, such extensive datasets are not accessible to the medical imaging domain or in real-world clinical environments, which limits the scalability of different approaches Song et al. (2023); Chen et al. (2024); al Nomaan Nafi et al. (2024). In response to

<sup>1</sup>The code will be made available upon acceptance.

the limited data issue, pre-training generative models on natural image datasets (e.g., ImageNet) can be a possible solution. Regardless, it leads to domain shift, as the characteristics (e.g., Hounsfield units, noise patterns) of medical images (e.g., CT images) differ significantly from those of natural images. This domain shift can degrade the quality of synthesized CT images and limit their clinical utility Zhang et al. (2018b). Another line of work introduces annotation-guided CT generation Venkatesh et al. (2024); Bhattacharya et al. (2024); Yang et al. (2024); Konz et al. (2024) to infuse shape constraints and organ-specific priors into the synthesis process. Textual prompts, combined with anatomical masks, have been utilized to generate high-quality 3D CT scans Xu et al. (2024). However, the model is built specifically for 3D volumetric CT synthesis and does not support 2D image generation. Furthermore, medical image synthesis with DMs can produce unrealistic artifacts or hallucinations Cho et al. (2025), which can potentially pose severe risks in subsequent image-based decision-making. Hallucination may arise when the model depends only on anatomy-guided conditioning without noise awareness, while focusing solely on noise guidance may lead to the loss of anatomical fidelity. To address the aforementioned challenges, we propose a novel dual conditional diffusion-based realistic CT image synthesis approach with CT noise and anatomy guidance, leveraging a large number of diverse natural images with emulated CT noise at various dose levels. The model is pre-trained using emulated CT noise guidance, and this is followed by fine-tuning with a small number of real CT images guided by dose-aware noise and anatomy maps. **First, we generate CT-like noise-emulated natural images using our proposed strategy, so we can pre-train on abundant natural images that mimic CT noise characteristics. Then, during fine-tuning on real CT data, we introduce a dual-conditioning setup with both a noise map (dose-aware) branch and an anatomy-guidance branch for CT image synthesis. This combination of synthetic noise-emulated pretraining and dual noise-anatomy conditioning, to our knowledge, has not been explored in prior CT diffusion work and goes beyond a simple diffusion architecture modification.**

The contributions of the paper are summarized as follows:

- A novel framework for CT image synthesis guided by noise and anatomy maps using a dual conditioning technique for preserving realistic noise-aware anatomical properties.
- Emulation of CT images and CT noise-emulated natural images at different dose levels with corresponding noise maps. This procedure is effective and adapts to varying radiation exposures.
- A pre-training strategy utilizing CT noise-emulated natural images at different CT dose levels to learn useful low-level and high-level features from the large dataset, enabling better performance with limited CT image data.
- Extensive evaluations, including external validation, across a range of tasks and metrics, demonstrating consistent and superior performance of NA-Diff in generating realistic CT.

## 2 RELATED WORK

### 2.1 DIFFUSION MODELS

Diffusion models have revolutionized generative modeling, demonstrating unprecedented capabilities in image generation by progressively denoising random noise to synthesize realistic visuals Ho et al. (2020); Song et al. (2020); Song & Ermon (2019). Recent conditional variants refined this approach, enhancing synthesis control by integrating conditions like textual prompts, segmentation masks, or style embeddings Dhariwal & Nichol (2021); Ho & Salimans (2022); Lugmayr et al. (2022); Saharia et al. (2022). Despite significant progress, unconditional models offer limited control over content and structure, while conditional models are often dependent on the availability and quality of external conditioning data. Recent works have addressed these challenges by introducing diffusion Ho et al. (2020); Peebles & Xie (2023) with image conditioning Zhu et al. (2025); Konz et al. (2024); Zhang et al. (2023); Shi et al. (2024), achieving more robust semantic guidance and motivating further exploration in domains that demand precise structural fidelity, such as medical image synthesis.

### 2.2 DIFFUSION FOR MEDICAL IMAGE SYNTHESIS

Latent diffusion-based model MediSyn Cho et al. (2025) effectively generates diverse medical images across multiple imaging modalities using textual guidance. However, text-based conditioning

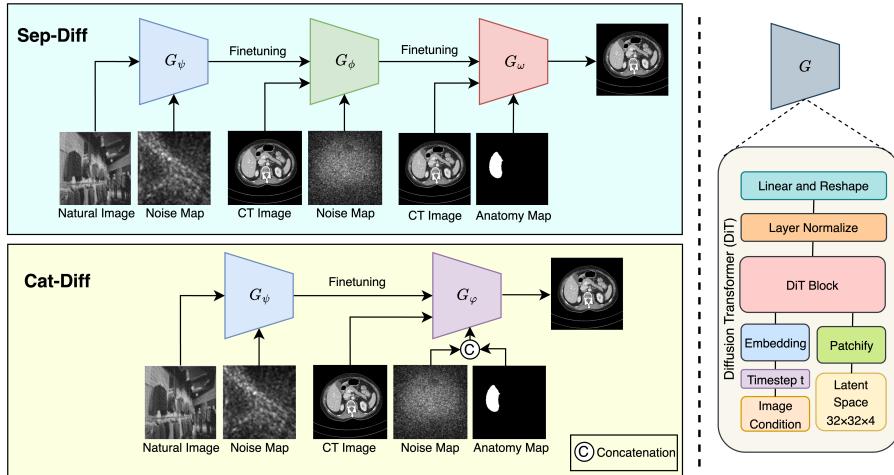


Figure 1: Overview of the proposed NA-Diff framework, illustrating its two variants: Sep-Diff and Cat-Diff. In Sep-Diff variant, the model is trained in three stages leveraging separate conditions while in Cat-Diff, the model utilize dual conditions with two stage training.  $G_\psi$ ,  $G_\phi$ ,  $G_\omega$  and  $G_\varphi$  denotes pre-training model on CT noise-emulated natural images with noise map guidance, fine-tuning model on CT images with noise map guidance, fine-tuning model on CT images with anatomy guidance, and fine-tuning CT images with corresponding noise map and anatomy guidance both, respectively.

is inherently imprecise for fine-grained anatomical control, often resulting in unrealistic anatomical structures Zhang et al. (2024); Chaichuk et al. (2025). While ControlNet Zhang et al. (2023), initially developed for natural image generation, has a mechanism for adding spatial conditioning through structural inputs, making it adaptable for medical domains where spatial fidelity is critical. Yet, direct application to medical domains may suffer from domain shift, limiting anatomical realism and clinical utility without retraining on medical-specific data. Recent anatomy-aware diffusion frameworks embed structural guidance (e.g., segmentation masks, anatomical labels, semantic layouts) directly into the denoising trajectory to ensure radiologically faithful synthesis. Seg-Diff Konz et al. (2024) proposed a segmentation-guided diffusion model that generates medical images by conditioning on segmentation masks at every denoising step. Using cross-attention or dual-stream architectures, some techniques generate image–mask pairs or condition on latent anatomy embeddings, achieving spatial coherence across both 2-D slices and volumetric reconstructions Xing et al. (2024); Zhang et al. (2025); Zheng et al. (2024); Bhat et al. (2025); Jiang et al. (2025); Mao et al. (2025); Bose et al. (2025). Their shared strategy shows that anatomy-guided conditioning markedly boosts structural realism and downstream segmentation performance in CT and related modalities. However, they only depend on precise anatomy inputs at inference and ignore noise-aware conditioning, leaving the problem of converting noisy CT projections into high-quality volumes largely unsolved.

### 2.3 CONTEXT-AWARE MULTI-GUIDANCE DIFFUSION

Relying on a single conditioning process in diffusion models is sometimes flawed, as generating target-like images may require using more than one type of guidance such as structural, and style guidance Konz et al. (2024); Xing et al. (2024); Mao et al. (2025). Incorporating multiple guidance can be positively effective, as it enables the model to capture both stylistic attributes and semantic content with greater fidelity Krishna et al. (2024); Zhu et al. (2023); Shen et al. (2024). For CT image synthesis, both noise conditioning (e.g., dose-aware noise maps) and anatomical guidance are important to enable realistic simulation. The model DCDiff Shen et al. (2024) generates clean CT images with reduced metal artifacts. However, the model lacks awareness of different dose levels and may struggle to generate images with varied noise characteristics corresponding to each dose. To the best of our knowledge, no existing model is able to synthesize CT images that are simultaneously dose-aware and structure-aware, capturing both noise characteristics and anatomical fidelity within a unified framework. The conditioning strategy also varies across different models. For encoding images, most existing models use a VAE-based encoder Kingma et al. (2019) for processing input or

162 conditioning information which is efficient and flexible guidance in the latent space Ho et al. (2020);  
 163 Zhang et al. (2023); Esser et al. (2021); Nie et al. (2024); Peebles & Xie (2023). In contrast, for  
 164 some cases, CLIP (Contrastive Language-Image Pretraining) Radford et al. (2021) can be used for  
 165 image embeddings instead of VAE, as CLIP’s ability to extract semantically rich, global features  
 166 from images or text allows more flexible and effective conditioning for diffusion models Zhu et al.  
 167 (2025).

### 168 3 METHOD

170 Fig. 1 illustrates our proposed diffusion-based CT image generation framework NA-Diff. NA-Diff  
 171 includes two variants: (i) separate conditioning with noise map and anatomy mask (Sep-Diff), and  
 172 (ii) concatenate conditioning with noise map and anatomy mask (Cat-Diff). Sep-Diff consists of a  
 173 three-stage training process that systematically enables the synthesis of new CT images. The objec-  
 174 tive is to synthesize CT images conditioned on noise maps or anatomy maps. However, we cannot  
 175 simultaneously condition the model on both the noise map and the corresponding anatomy mask,  
 176 as each stage of the model supports only a single conditioning input. In light of this limitation, we  
 177 introduce Cat-Diff, a two-stage training process that leverages both noise maps and corresponding  
 178 anatomy masks to synthesize CT images.

#### 179 3.1 DiT

180 DiT Peebles & Xie (2023)  
 181 is a transformer-based dif-  
 182 fusion model that operates  
 183 on latent patches for image  
 184 generation task, achieving  
 185 state-of-the-art image  
 186 quality on benchmarks.  
 187 It uses a latent diffusion  
 188 framework and processes  
 189 image representations as  
 190 patches following Vision  
 191 Transformers (ViTs) Dosov-  
 192 vitskiy et al. (2020). Given  
 193 an input image  $x_0$  and a  
 194 corresponding condition  
 195 y (e.g., text-guided, noise  
 196 map, or anatomy-guided

197 image), the model learns to reconstruct  $x_0$  through iterative denoising process:  $p_\theta(\hat{x}_0|x_0, y)$   
 198 where  $\theta$  represents the model parameters. Like other diffusion models Rombach et al. (2022);  
 199 Ho et al. (2020), DiT follows a forward process that gradually adds Gaussian noise to the input  
 200 image and a reverse process that aims to denoise and reconstruct the original image conditioned  
 201 on a noise map or anatomy map. The model uses self-attention mechanisms and multi-head  
 202 attention layers to model long-range dependencies within the image. This allows the network to  
 203 process complex textures and structures efficiently. During DiT training, a noise prediction loss  
 204 is minimized:  $L(\theta) = \mathbb{E}_{x_0, y, t, \epsilon} [\|\epsilon_\theta(x_t, y) - \epsilon\|_2^2]$ , where  $\epsilon_\theta(x_t, y)$  is the predicted noise from the  
 205 diffusion transformer,  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  is the true noise sampled from a Gaussian distribution, and  $\mathbb{E}$   
 206 represents taking the mean over these random variables, ensuring generalization across varying  
 207 noise conditions.

#### 208 3.2 NATURAL IMAGE PRETRAINING

209 **Emulation of CT noise in natural images:** Natural images are by default of high quality. We can  
 210 consider them equivalent to full-dose CT images. An RGB image is first converted to grayscale  
 211 and we map it to CT water density. For simplicity, we assume  $1 \times 1$  pixel size and projection  
 212 angles  $0\text{--}180^\circ$ . The raw projection data (sinogram)  $r$  is generated by performing parallel beam  
 213 projections (Radon transform). The generated sinogram is fed to inverse Radon (iRadon) transform  
 214 to reconstruct the original grayscale CT-like full-dose image  $x_{fd}^{nat}$ . To reduce spectral leakage and  
 215 enhance interpretation, we select ‘Hann’ filter and linear interpolation for iRadon. To match with our  
 CT dataset, we generate lower-dose image ( $x_{ld}$ ) at the same dose level ( $ld$ ). Assuming the number

Table 1: Details of the NA-Diff model variants.

Type	Model	Description
Sep-Diff	NA-Diff-A	Training $G_\phi$ with CT images conditioned on noise map data.
	NA-Diff-B	Training $G_\psi$ with natural image data, followed by fine-tuning ( $G_\phi$ ) with CT images conditioned on noise map data.
	NA-Diff-C	Fine-tuning ( $G_\omega$ ) using NA-Diff-A, with CT images conditioned on anatomy map data.
Cat-Diff	NA-Diff-D	Fine-tuning ( $G_\omega$ ) using NA-Diff-B, with CT images conditioned on anatomy map data.
	NA-Diff-E	Training $G_\varphi$ with CT images conditioned on noise map and corresponding anatomy map.
	NA-Diff-F	Training $G_\psi$ with natural image data, followed by fine-tuning ( $G_\varphi$ ) with CT images conditioned on noise map and corresponding anatomy map.

of photons  $I_0 = 1e^5$  emitted at full dose and water attenuation of  $\mu = 0.02$ , we can then calculate the number of mean transmitted photons for the raw data as  $I_{raw} = ld \cdot I_0 \cdot \exp(-\mu r)$ . Poisson quantum noise is inserted into the raw project data at the target dose level. Therefore, the noisy projection we obtain as:

$$I_{noisy} = -\log(Poisson(I_{raw})/(ld \cdot I_0))/\mu. \quad (1)$$

Using iRadon, the noisy image  $x_{ld}^{nat}$  is then reconstructed from the noisy projection data at the target dose level. Following the approach described in Sec. 3.3 for CT images, natural images are also simulated at additional dose levels, and corresponding noise maps are constructed (Fig. 2).

In the pretraining phase, input natural images are denoted  $x_0^{nat}$ , and corresponding noise maps are denoted by  $y_{noise}^{nat}$ . To synthesize a refined natural image  $\hat{x}_0^{nat}$ , we use the DiT model conditioned on  $y_{noise}^{nat}$ , replacing the text conditioning with an image conditioning strategy. For the pre-training stage of both Sep-Diff and Cat-Diff variants,  $x_0$  and  $y$  are mapped into the latent space using a pre-trained VAE Kingma et al. (2019). Those are patchified and linearly embedded to create token representations:  $z_{x_0} = W_x x_0 + E_{pos}$  and  $z_y = W_y y + E_{pos}$ , where  $W_x, W_y$  are learned embedding matrices,  $E_{pos}$  is the positional encoding, and  $z_{x_0}, z_y$  are the patch embeddings. The timestep embedding is combined with the conditional noise embedding  $c = f_\theta(t) + z_y$ ;  $f_\theta(t)$  is an MLP-based mapping timestep  $t$  to an embedding, and  $c$  is the conditioning vector that modulates the model.

Each Transformer block uses AdaLN-Zero, where the modulation parameters  $(\gamma_t, \beta_t, \alpha_t)$  are predicted from the conditioning vector  $c$ :

$$(\gamma_t, \beta_t, \alpha_t) = \text{MLP}(\text{LayerNorm}(c)), \quad (2)$$

where  $\gamma_t$  and  $\beta_t$  are feature-wise scaling and shifting parameters, and  $\alpha_t$  is a learned residual gating parameter that controls the strength of each residual branch. Given these parameters, a DiT block updates the hidden representation  $x$  as:

$$x = x + \alpha_t \cdot \text{MSA}(\gamma_t \cdot \text{LayerNorm}(x) + \beta_t), \quad (3)$$

$$x = x + \alpha_t \cdot \text{MLP}(\gamma_t \cdot \text{LayerNorm}(x) + \beta_t), \quad (4)$$

where the initial hidden state is the image token sequence  $x = z_{x_0}$ ; that is,  $z_{x_0}$  is the tokenized representation of the input image, and  $x$  is refined iteratively by the stacked Transformer blocks. The model predicts the noise  $\hat{\vartheta}(x_t, t, y)$  and the output AdaLN-Zero parameters are computed as:

$$(\gamma'_t, \beta'_t) = \text{MLP}(\text{LayerNorm}(c)), \quad (4)$$

and the noise prediction is generated by:

$$\hat{\vartheta} = \text{MLP}(\gamma'_t \cdot \text{LayerNorm}(x) + \beta'_t). \quad (5)$$

The denoised image  $\hat{x}_0$  is computed during sampling using the standard DM reconstruction:

$$\hat{x}_0 = \frac{x_t - \sqrt{1 - \bar{\alpha}_t} \hat{\vartheta}}{\sqrt{\bar{\alpha}_t}}, \quad (6)$$

where  $\bar{\alpha}_t$  is the cumulative product of the variance schedule in the forward diffusion process and  $\hat{x}_0$  is the final denoised image. Also,  $\gamma'_t$  and  $\beta'_t$  are dynamically learned based on  $y$  and  $t$ .

### 3.3 NOISE-GUIDED CT FINE-TUNING

**Low dose simulation:** Following Imran et al. (2021), we use the quantum noise properties of low-dose images, assuming a linear relationship through reconstructions with the relative dose levels.

Given real full-dose and lower-dose scans, CT scans can be simulated at any arbitrary dose level ( $d$ ). With the full-dose noise variance  $\sigma_{fd}^2$ , an image  $x_d$  can be formed by adding zero-mean independent noise to the full-dose image  $x_{fd}$ .

$$x_d = x_{fd} + x_{noise}, \quad (7)$$

where,  $x_{noise} \sim \mathcal{N}(0, (1/d - 1)\sigma_{fd}^2)$ .

In order to estimate  $\sigma_{fd}^2$ , we use the difference between  $x_{fd}$  and the available lower-dose image  $x_{ld}$ . We are then able to synthesize images at any arbitrary dose level by determining  $a$  from  $1 + (1/d - 1)a^2 = 1/d$ .

**Noise estimation from single slice/image:** After obtaining noise realization from a single slice of the given  $x_{fd}$  and  $x_{ld}$ , it can be scaled to derive the noise of the desired dose level, i.e.,  $n = b \cdot (x_{ld} - x_{fd})$ . And we can find  $b$  such that

$$b^2 \cdot \left(\frac{1}{d} - 1\right) = \frac{1}{d}. \quad (8)$$

Ideally, noise realizations should be uncorrelated with noise at the routine dose, but correlated with noise at all other dose levels. It can be verified that noise standard deviation (std) in uniform regions of an image matches the expected noise levels in HU. It is, therefore, viable to use noise std map as a means of assessing the quality of CT images at arbitrary dose levels. Using a sliding window with window size  $k \times k$  over the noise realization  $n$  in the image space at  $m \times m$  resolution, the std map is calculated as:

$$y_n^{ct} = \sqrt{\frac{1}{k^2} \sum_{m \times m} (n(k, k))^2}. \quad (9)$$

With the increase in the window size, the noise std scale (HU) is shrunk. We choose  $k = 5$  in order to keep maximum scale gaps in different dose levels.

In Sep-Diff variant, with the simulated CT images and the corresponding noise (std) maps, the natural image pre-trained model is fine-tuned on ( $G_\phi$ ). Similar to pre-training of  $G_\psi$ ,  $G_\phi$  is fed by an input CT image  $x_0^{ct}$ , and its corresponding noise map  $y_{map}^{ct}$  for conditioning. The training process is similar to the pre-training phase, and the final layer reconstructs the input CT image  $\hat{x}_0^{ct}$ . However, unlike the DiT approach, which is used in the pre-training phase, the conditioning strategy for CT images corresponding to the map leverages a pre-trained CLIP encoder instead of VAE. Instead of directly using pixel-level features as a condition,  $y_{map}^{ct}$  is first processed using the CLIP processor and mapped into a global semantic embedding using the CLIP encoder. This embedding captures high-level semantic information, which is then used to guide the generative process. Using the diverse dose level images in training, Sep-Diff is enabled to learn the realistic variation in the image quality, retaining textures. **For a single diagnostic task, higher-dose, lower-noise CT images typically improve performance. However, radiation dose cannot be increased freely in practice, especially in screening, pediatrics, and repeated follow-ups. Multi-center datasets also contain diverse protocols and dose levels, leading to substantial variation in noise and texture. Our aim in modeling dose-aware, higher-dose, lower-noise CT images is not to promote increased radiation, but to provide a controllable way to simulate different dose (noise) levels from the same anatomy. This enables systematic study of how downstream tasks behave across dose levels, generation of virtual high-dose images from low-dose scans without extra patient dose, and harmonization of heterogeneous datasets by mapping images to a standardized noise level.**

### 3.4 ANATOMY-GUIDED CT FINE-TUNING

In Sep-Diff, for the second stage of fine-tuning, we introduce anatomy masks as additional guidance. The main reason for incorporating anatomical guidance in CT image synthesis is to enforce anatomical constraints in the generated images, and this helps the model to capture the fundamental structures and distribution of CT images. The anatomy mask is derived from the simulated low-dose CT images using a segmentation model Ward & Imran (2025). The training process is similar to the first stage of fine-tuning in Sep-Diff (i.e.,  $x_0^{ct}$  as CT image and  $y_{mask}^{ct}$  as anatomy guidance). In this phase, we use the DiT layout, modifying the condition strategy, and the final layer reconstructs the denoised image. It maintains anatomical plausibility, followed by the anatomy mask.

324  
 325  
 326  
 327  
 328  
 Table 2: Consistency evaluation of image quality and liver segmentation between synthetic and input  
 329  
 330  
 331  
 332  
 333  
 334  
 335  
 336  
 337  
 338  
 339  
 340  
 341  
 342  
 343  
 344  
 345  
 346  
 347  
 348  
 349  
 350  
 351  
 352  
 353  
 354  
 355  
 356  
 357  
 358  
 359  
 360  
 361  
 362  
 363  
 364  
 365  
 366  
 367  
 368  
 369  
 370  
 371  
 372  
 373  
 374  
 375  
 376  
 377  
 378  
 379  
 380  
 381  
 382  
 383  
 384  
 385  
 386  
 387  
 388  
 389  
 390  
 391  
 392  
 393  
 394  
 395  
 396  
 397  
 398  
 399  
 400  
 401  
 402  
 403  
 404  
 405  
 406  
 407  
 408  
 409  
 410  
 411  
 412  
 413  
 414  
 415  
 416  
 417  
 418  
 419  
 420  
 421  
 422  
 423  
 424  
 425  
 426  
 427  
 428  
 429  
 430  
 431  
 432  
 433  
 434  
 435  
 436  
 437  
 438  
 439  
 440  
 441  
 442  
 443  
 444  
 445  
 446  
 447  
 448  
 449  
 450  
 451  
 452  
 453  
 454  
 455  
 456  
 457  
 458  
 459  
 460  
 461  
 462  
 463  
 464  
 465  
 466  
 467  
 468  
 469  
 470  
 471  
 472  
 473  
 474  
 475  
 476  
 477  
 478  
 479  
 480  
 481  
 482  
 483  
 484  
 485  
 486  
 487  
 488  
 489  
 490  
 491  
 492  
 493  
 494  
 495  
 496  
 497  
 498  
 499  
 500  
 501  
 502  
 503  
 504  
 505  
 506  
 507  
 508  
 509  
 510  
 511  
 512  
 513  
 514  
 515  
 516  
 517  
 518  
 519  
 520  
 521  
 522  
 523  
 524  
 525  
 526  
 527  
 528  
 529  
 530  
 531  
 532  
 533  
 534  
 535  
 536  
 537  
 538  
 539  
 540  
 541  
 542  
 543  
 544  
 545  
 546  
 547  
 548  
 549  
 550  
 551  
 552  
 553  
 554  
 555  
 556  
 557  
 558  
 559  
 560  
 561  
 562  
 563  
 564  
 565  
 566  
 567  
 568  
 569  
 570  
 571  
 572  
 573  
 574  
 575  
 576  
 577  
 578  
 579  
 580  
 581  
 582  
 583  
 584  
 585  
 586  
 587  
 588  
 589  
 590  
 591  
 592  
 593  
 594  
 595  
 596  
 597  
 598  
 599  
 600  
 601  
 602  
 603  
 604  
 605  
 606  
 607  
 608  
 609  
 610  
 611  
 612  
 613  
 614  
 615  
 616  
 617  
 618  
 619  
 620  
 621  
 622  
 623  
 624  
 625  
 626  
 627  
 628  
 629  
 630  
 631  
 632  
 633  
 634  
 635  
 636  
 637  
 638  
 639  
 640  
 641  
 642  
 643  
 644  
 645  
 646  
 647  
 648  
 649  
 650  
 651  
 652  
 653  
 654  
 655  
 656  
 657  
 658  
 659  
 660  
 661  
 662  
 663  
 664  
 665  
 666  
 667  
 668  
 669  
 670  
 671  
 672  
 673  
 674  
 675  
 676  
 677  
 678  
 679  
 680  
 681  
 682  
 683  
 684  
 685  
 686  
 687  
 688  
 689  
 690  
 691  
 692  
 693  
 694  
 695  
 696  
 697  
 698  
 699  
 700  
 701  
 702  
 703  
 704  
 705  
 706  
 707  
 708  
 709  
 710  
 711  
 712  
 713  
 714  
 715  
 716  
 717  
 718  
 719  
 720  
 721  
 722  
 723  
 724  
 725  
 726  
 727  
 728  
 729  
 730  
 731  
 732  
 733  
 734  
 735  
 736  
 737  
 738  
 739  
 740  
 741  
 742  
 743  
 744  
 745  
 746  
 747  
 748  
 749  
 750  
 751  
 752  
 753  
 754  
 755  
 756  
 757  
 758  
 759  
 760  
 761  
 762  
 763  
 764  
 765  
 766  
 767  
 768  
 769  
 770  
 771  
 772  
 773  
 774  
 775  
 776  
 777  
 778  
 779  
 780  
 781  
 782  
 783  
 784  
 785  
 786  
 787  
 788  
 789  
 790  
 791  
 792  
 793  
 794  
 795  
 796  
 797  
 798  
 799  
 800  
 801  
 802  
 803  
 804  
 805  
 806  
 807  
 808  
 809  
 810  
 811  
 812  
 813  
 814  
 815  
 816  
 817  
 818  
 819  
 820  
 821  
 822  
 823  
 824  
 825  
 826  
 827  
 828  
 829  
 830  
 831  
 832  
 833  
 834  
 835  
 836  
 837  
 838  
 839  
 840  
 841  
 842  
 843  
 844  
 845  
 846  
 847  
 848  
 849  
 850  
 851  
 852  
 853  
 854  
 855  
 856  
 857  
 858  
 859  
 860  
 861  
 862  
 863  
 864  
 865  
 866  
 867  
 868  
 869  
 870  
 871  
 872  
 873  
 874  
 875  
 876  
 877  
 878  
 879  
 880  
 881  
 882  
 883  
 884  
 885  
 886  
 887  
 888  
 889  
 890  
 891  
 892  
 893  
 894  
 895  
 896  
 897  
 898  
 899  
 900  
 901  
 902  
 903  
 904  
 905  
 906  
 907  
 908  
 909  
 910  
 911  
 912  
 913  
 914  
 915  
 916  
 917  
 918  
 919  
 920  
 921  
 922  
 923  
 924  
 925  
 926  
 927  
 928  
 929  
 930  
 931  
 932  
 933  
 934  
 935  
 936  
 937  
 938  
 939  
 940  
 941  
 942  
 943  
 944  
 945  
 946  
 947  
 948  
 949  
 950  
 951  
 952  
 953  
 954  
 955  
 956  
 957  
 958  
 959  
 960  
 961  
 962  
 963  
 964  
 965  
 966  
 967  
 968  
 969  
 970  
 971  
 972  
 973  
 974  
 975  
 976  
 977  
 978  
 979  
 980  
 981  
 982  
 983  
 984  
 985  
 986  
 987  
 988  
 989  
 990  
 991  
 992  
 993  
 994  
 995  
 996  
 997  
 998  
 999  
 1000  
 1001  
 1002  
 1003  
 1004  
 1005  
 1006  
 1007  
 1008  
 1009  
 1010  
 1011  
 1012  
 1013  
 1014  
 1015  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025  
 1026  
 1027  
 1028  
 1029  
 1030  
 1031  
 1032  
 1033  
 1034  
 1035  
 1036  
 1037  
 1038  
 1039  
 1040  
 1041  
 1042  
 1043  
 1044  
 1045  
 1046  
 1047  
 1048  
 1049  
 1050  
 1051  
 1052  
 1053  
 1054  
 1055  
 1056  
 1057  
 1058  
 1059  
 1060  
 1061  
 1062  
 1063  
 1064  
 1065  
 1066  
 1067  
 1068  
 1069  
 1070  
 1071  
 1072  
 1073  
 1074  
 1075  
 1076  
 1077  
 1078  
 1079  
 1080  
 1081  
 1082  
 1083  
 1084  
 1085  
 1086  
 1087  
 1088  
 1089  
 1090  
 1091  
 1092  
 1093  
 1094  
 1095  
 1096  
 1097  
 1098  
 1099  
 1100  
 1101  
 1102  
 1103  
 1104  
 1105  
 1106  
 1107  
 1108  
 1109  
 1110  
 1111  
 1112  
 1113  
 1114  
 1115  
 1116  
 1117  
 1118  
 1119  
 1120  
 1121  
 1122  
 1123  
 1124  
 1125  
 1126  
 1127  
 1128  
 1129  
 1130  
 1131  
 1132  
 1133  
 1134  
 1135  
 1136  
 1137  
 1138  
 1139  
 1140  
 1141  
 1142  
 1143  
 1144  
 1145  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153  
 1154  
 1155  
 1156  
 1157  
 1158  
 1159  
 1160  
 1161  
 1162  
 1163  
 1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187  
 1188  
 1189  
 1190  
 1191  
 1192  
 1193  
 1194  
 1195  
 1196  
 1197  
 1198  
 1199  
 1200  
 1201  
 1202  
 1203  
 1204  
 1205  
 1206  
 1207  
 1208  
 1209  
 1210  
 1211  
 1212  
 1213  
 1214  
 1215  
 1216  
 1217  
 1218  
 1219  
 1220  
 1221  
 1222  
 1223  
 1224  
 1225  
 1226  
 1227  
 1228  
 1229  
 1230  
 1231  
 1232  
 1233  
 1234  
 1235  
 1236  
 1237  
 1238  
 1239  
 1240  
 1241  
 1242  
 1243  
 1244  
 1245  
 1246  
 1247  
 1248  
 1249  
 1250  
 1251  
 1252  
 1253  
 1254  
 1255  
 1256  
 1257  
 1258  
 1259  
 1260  
 1261  
 1262  
 1263  
 1264  
 1265  
 1266  
 1267  
 1268  
 1269  
 1270  
 1271  
 1272  
 1273  
 1274  
 1275  
 1276  
 1277  
 1278  
 1279  
 1280  
 1281  
 1282  
 1283  
 1284  
 1285  
 1286  
 1287  
 1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295  
 1296  
 1297  
 1298  
 1299  
 1300  
 1301  
 1302  
 1303  
 1304  
 1305  
 1306  
 1307  
 1308  
 1309  
 1310  
 1311  
 1312  
 1313  
 1314  
 1315  
 1316  
 1317  
 1318  
 1319  
 1320  
 1321  
 1322  
 1323  
 1324  
 1325  
 1326  
 1327  
 1328  
 1329  
 1330  
 1331  
 1332  
 1333  
 1334  
 1335  
 1336  
 1337  
 1338  
 1339  
 1340  
 1341  
 1342  
 1343  
 1344  
 1345  
 1346  
 1347  
 1348  
 1349  
 1350  
 1351  
 1352  
 1353  
 1354  
 1355  
 1356  
 1357  
 1358  
 1359  
 1360  
 1361  
 1362  
 1363  
 1364  
 1365  
 1366  
 1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403  
 1404  
 1405  
 1406  
 1407  
 1408  
 1409  
 1410  
 1411  
 1412  
 1413  
 1414  
 1415  
 1416  
 1417  
 1418  
 1419  
 1420  
 1421  
 1422  
 1423  
 1424  
 1425  
 1426  
 1427  
 1428  
 1429  
 1430  
 1431  
 1432  
 1433  
 1434  
 1435  
 1436  
 1437  
 1438  
 1439  
 1440  
 1441  
 1442  
 1443  
 1444  
 1445  
 1446  
 1447  
 1448  
 1449  
 1450  
 1451  
 1452  
 1453  
 1454  
 1455  
 1456  
 1457  
 1458  
 1459  
 1460  
 1461  
 1462  
 1463  
 1464  
 1465  
 1466  
 1467  
 1468  
 1469  
 1470  
 1471  
 1472  
 1473  
 1474  
 1475  
 1476  
 1477  
 1478  
 1479  
 1480  
 1481  
 1482  
 1483  
 1484  
 1485  
 1486  
 1487  
 1488  
 1489  
 1490  
 1491  
 1492  
 1493  
 1494  
 1495  
 1496  
 1497  
 1498  
 1499  
 1500  
 1501  
 1502  
 1503  
 1504  
 1505  
 1506  
 1507  
 1508  
 1509  
 1510  
 1511  
 1512  
 1513  
 1514  
 1515  
 1516  
 1517  
 1518  
 1519  
 1520  
 1521  
 1522  
 1523  
 1524  
 1525  
 1526  
 1527  
 1528  
 1529  
 1530  
 1531  
 1532  
 1533  
 1534  
 1535  
 1536  
 1537  
 1538  
 1539  
 1540  
 1541  
 1542  
 1543  
 1544  
 1545  
 1546  
 1547  
 1548  
 1549  
 1550  
 1551  
 1552  
 1553  
 1554  
 1555  
 1556  
 1557  
 1558  
 1559  
 1560  
 1561  
 1562  
 1563  
 1564  
 1565  
 1566  
 1567  
 1568  
 1569  
 1570  
 1571  
 1572  
 1573  
 1574  
 1575  
 1576  
 1577  
 1578  
 1579  
 1580  
 1581  
 1582  
 1583  
 1584  
 1585  
 1586  
 1587  
 1588  
 1589  
 1590  
 1591  
 1592  
 1593  
 1594  
 1595  
 1596  
 1597  
 1598  
 1599  
 1600  
 1601  
 1602  
 1603  
 1604  
 1605  
 1606  
 1607  
 1608  
 1609  
 1610  
 1611  
 1612  
 1613  
 1614  
 1615  
 1616  
 1617  
 1618  
 1619  
 1620  
 1621  
 1622  
 1623  
 1624  
 1625  
 1626  
 1627  
 1628  
 1629  
 1630  
 1631  
 1632  
 1633  
 1634  
 1635  
 1636  
 1637  
 1638  
 1639  
 1640  
 1641  
 1642  
 1643  
 1644  
 1645  
 1646  
 1647  
 1648  
 1649  
 1650  
 1651  
 1652  
 1653  
 1654  
 1655  
 1656  
 1657  
 1658  
 1659  
 1660  
 1661  
 1662  
 1663  
 1664  
 1665  
 1666  
 1667  
 1668  
 1669  
 1670  
 1671  
 1672  
 1673  
 1674  
 1675  
 1676  
 1677  
 1678  
 1679  
 1680  
 1681  
 1682  
 1683  
 1684  
 1685  
 1686  
 1687  
 1688  
 1689  
 1690  
 1691  
 1692  
 1693  
 1694  
 1695  
 1696  
 1697  
 1698  
 1699  
 1700  
 1701  
 1702  
 1703  
 1704  
 1705  
 1706  
 1707  
 1708  
 1709  
 1710  
 1711  
 1712  
 1713  
 1714  
 1715  
 1716  
 1717  
 1718  
 1719  
 1720  
 1721  
 1722  
 1723  
 1724  
 1725  
 1726  
 1727  
 1728  
 1729  
 1730  
 1731  
 1732  
 1733  
 1734  
 1735  
 1736  
 1737  
 1738  
 1739  
 1740  
 1741  
 1742  
 1743  
 1744  
 1745  
 1746  
 1747  
 1748  
 1749  
 1750  
 1751  
 1752  
 1753  
 1754  
 1755  
 1756  
 1757  
 1758  
 1759  
 1760  
 1761  
 1762  
 1763  
 1764  
 1765  
 1766  
 1767  
 1768  
 1769  
 1770  
 1771  
 1772  
 1773  
 1774  
 1775  
 1776  
 1777  
 1778  
 1779  
 1780  
 1781  
 1782  
 1783  
 1784  
 1785  
 1786  
 1787  
 1788  
 1789  
 1790  
 1791  
 1792  
 1793  
 1794  
 1795  
 1796  
 1797  
 1798  
 1799  
 1800  
 1801  
 1802  
 1803  
 1804  
 1805  
 1806  
 1807  
 1808  
 1809  
 1810  
 1811  
 1812  
 1813  
 1814  
 1815  
 1816  
 1817  
 1818  
 1819  
 1820  
 18

378 images, we selected 4,000 images with complex-structured and intricate details. Following Sec. 3.2,  
 379 a total of 20,000 natural images and their corresponding noise maps are generated at different dose  
 380 levels (i.e., 5%, 10%, 25%, 50%, and 75%).  
 381

**SDCT Data:** SDCT refers to the six-dose computed tomography, representing our emulated CT dataset. For CT images, we use 40 abdomen scans (25% and 100% doses) from the Low Dose CT Image and Projection Data McCollough et al. (2020) — 20 for training and 20 for testing. After simulation of additional dose levels, a total of 13,314 (e.g., 6696 for training and 6618 for testing) CT images and noise map pairs are obtained. **All CT images in our figures are displayed using a standard soft-tissue window, with a window width (WW) of 400 HU and a window level (WL) of 40 HU.** The anatomy masks are obtained from the simulated low-dose CT images using the liver segmentation model Ward & Imran (2025). All 13,314 simulated low-dose CT images are processed to generate the corresponding liver segmentation masks referred to as anatomy maps. Specifically, liver segmentation models are well-established and generally produce reliable masks. Inaccurate or noisy masks can directly degrade anatomy-guided synthesis and make it difficult to evaluate the method itself. Therefore, we focus on the liver organ for this study. However, the approach is not liver-specific as it can be extended to other organs (e.g., pancreas) with the availability of high-quality organ masks with relative dose levels.

**Training:** We trained the DiT model from scratch, as our task differs from ImageNet-trained DiT models. For training NA-Diff, we start with the DiT-L/4 model ( $256 \times 256$  resolution). With a patch size of 4, L/4 processes a total of 1024 tokens after patchifying the  $32 \times 32 \times 4$  input latent. To process condition images, we use a pre-trained CLIP encoder to obtain global feature embeddings. Natural images are pre-trained for 25 epochs. For CT image synthesis, we train the models for 100 epochs with noise conditioning, 100 epochs with anatomy conditioning, and 150 epochs with dual conditioning. The models are trained on an Intel (R) Xeon (R) w7-2475X, 2600MHz machine with dual NVIDIA A4000 GPUs (16X2=32GB).

**Baseline and Compared Methods:** Considering no existing generative models using CT noise-emulated natural images with noise map conditioning and anatomy guidance, we compare different variant types of the NA-Diff model (Table 1). Direct comparison with our dual condition-guided approach is not entirely applicable; however, we include Seg-Diff Konz et al. (2024) (e.g, the state-of-the-art anatomy-guided diffusion model for CT images) and ControlNet Zhang et al. (2023), as the representative anatomy-guided diffusion baseline and trained with our SDCT dataset, only leveraging anatomy guidance as conditions. For a fairer assessment, we report results from our only anatomy-conditioned variant alongside the dual-conditioned model. In addition, we have shown comparisons with the baseline using the LiTS Bilic et al. (2023) dataset.

**Noise and Anatomy Preservation Evaluation:** NA-Diff model variants are evaluated by predicting diagnostic image quality assessment (IQA). We use a pre-trained IQA model Rifa et al. (2025b) to provide a measure of noise level estimation for the generated images. We also utilize a pretrained segmentation model Ward & Imran (2024) for generating liver masks from the generated images.

**Image Quality:** We use the popular metric CFID Parmar et al. (2022) and KID Bińkowski et al. (2018) for image realism and CMMD Jayasumana et al. (2024) to provide an unbiased and semantically meaningful evaluation of image quality using CLIP features, and LPIPS Zhang et al. (2018a) for perceptual similarity.

Table 3: Quantitative comparison of generative performance using CFID, KID, CMMD, and LPIPS across state-of-the-art baseline and our anatomy-guided NA-Diff variants. Lower scores indicate better fidelity, perceptual quality, and conditional consistency. \* denotes that the model is trained with our CT image data conditioned on the anatomy mask.

Methods	Performance metrics			
	CFID ↓	KID ↓	CMMD ↓	LPIPS ↓
Seg-Diff Konz et al. (2024)	249.769	122.358 ± 13.406	0.973	0.703 ± 0.281
Seg-Diff* Konz et al. (2024)	111.938	103.993 ± 11.255	0.670	0.664 ± 0.145
ControlNet* Zhang et al. (2023)	124.418	104.313 ± 8.5501	0.620	0.580 ± 0.117
NA-Diff-C	89.295	57.601 ± 5.823	0.541	0.387 ± 0.102
NA-Diff-D	<b>53.295</b>	<b>19.039 ± 2.853</b>	<b>0.333</b>	<b>0.382 ± 0.103</b>

Table 4: Consistent IQA scores are observed for NA-Diff-D and NA-Diff-F compared to training data across different dose levels.

Source	Dose Level (%)					
	5	10	25	50	75	100
Train set	1.500	1.450	1.870	2.110	2.290	2.410
NA-Diff-D	2.090	2.020	2.040	2.140	2.140	2.140
NA-Diff-F	1.910	1.960	2.120	2.140	2.160	2.190

432  
433  
434  
Table 5: Ablation study demonstrating the importance of pre-training and dual conditioning. The  
best and second-best results are **bolded** and underlined respectively. Nat. Pre-Train denotes training  
 $G_\psi$  with natural image data.

435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 Methods	435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 Nat. Pre-Train	435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 Noise	435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 Anatomy	435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 Performance metrics
				CFID $\downarrow$
				KID $\downarrow$
				CMMID $\downarrow$
				LPIPS $\downarrow$
NA-Diff-A	$\times$	✓	$\times$	245.342
NA-Diff-B	✓	✓	$\times$	152.768
NA-Diff-C	$\times$	$\times$	✓	89.298
NA-Diff-D	✓	$\times$	✓	<b>53.295</b>
NA-Diff-E	$\times$	✓	✓	78.373
NA-Diff-F	✓	✓	✓	<u>64.638</u>
				24.811 $\pm$ 3.228
				<b>0.321</b>
				<b>0.376</b> $\pm$ <b>0.099</b>

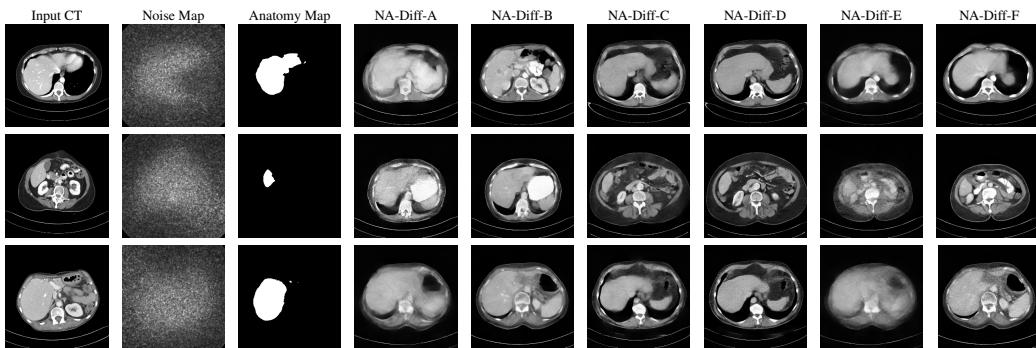


Figure 4: Qualitative comparison among generated images from different NA-Diff variants.

## 4.2 RESULTS AND DISCUSSION

As reported in Table 2, we perform comparisons among the variant types of NA-Diff generated images, leveraging the noise map, the anatomy mask, or both, conditioned. The generated images and the simulated CT test images are assessed by IQA scores (0-4, higher is better) Lee et al. (2025). The pretrained IQA model Rifa et al. (2025b) was trained on a different window-level CT dataset, yet it can approximately estimate the noise level in our generated images. NA-Diff-F has the smallest IQA gap ( $\Delta_{IQA}$ ) among all, demonstrating the effectiveness of the proposed noise-guided training processes. Additionally, the larger  $\Delta_{IQA}$  for the baseline Model A (without natural image pre-training) highlights the importance of CT noise-emulated pre-training in NA-Diff. Moreover, consistent with the dose-IQA relationship, NA-Diff-F generated images are of better quality with the increase in relative dose levels (Table 4), indicating the intrinsic noise aware dose level understanding over NA-Diff-D. In terms of segmentation in Table 2, we show a comparison between the original mask  $m$  and masks  $m_{gen}^{pred}$  predicted for the generated images. The NA-Diff-D and NA-Diff-F model achieves comparatively high scores across all metrics, which indicates good anatomical consistency.

In Table 3, we compare our anatomy-guided model variants (i.e., NA-Diff-C and NA-Diff-D) with the baseline model Seg-Diff Konz et al. (2024). Compared to the anatomy-conditioned baseline Seg-Diff\*, NA-Diff-D achieves substantial improvements across all metrics, where CFID improves by 52.4%, KID by 81.7%, CMMID by 50.4%, and LPIPS by 42.5%. When benchmarked against the ControlNet\* baseline, NA-Diff-D achieves even larger relative gains, with CFID reduced by 57.2%, KID by 81.8%, CMMID by 46.3%, and LPIPS by 34.1%. These results show that our approach can generate anatomy-guided CT images with better fidelity, perceptual quality, and consistency compared to previous methods. The results of the ablation study are shown in Table 5. Across all metrics, natural image pre-training relatively improves performance compared to no pre-training. Comparing

Table 6: Segmentation performance on the downstream task using real data and a combination of real plus synthetic data. The models are trained either on real data only or on a mixture of real and synthetic CT images (paired with masks), and evaluated on the first 1000 cases from the test set (using real images and masks only). Incorporating synthetic data leads to comparable improvements in Dice and IoU across both U-Net and TransUNet architectures.

Methods	Description	Dice	IoU
U-Net Ronneberger et al. (2015)	Real Data	0.874	0.777
	Real + Synthetic Data	0.876	0.779
TransUNet Chen et al. (2021)	Real Data	0.872	0.773
	Real + Synthetic Data	0.877	0.781

486 NA-Diff-D (with pre-training) to NA-Diff-C (without pre-training), we observe substantial improvements  
 487 in CFID by 40.3%, KID by 66.9%, CMMMD by 38.4%, and LPIPS by 1.3%.

488 For dual conditioning, NA-Diff-F’s CFID and KID are higher than NA-Diff-D because CFID is  
 489 conditioned only on anatomy for calculation, favoring NA-Diff-D, and KID reflects greater sample  
 490 diversity. However, NA-Diff-F achieves 3.8% lower CMMMD and 1.5% lower LPIPS, indicating  
 491 better conditional consistency and perceptual similarity with dual conditioning. As shown in Table 6,  
 492 adding synthetic data provides comparable gains in Dice and IoU, which demonstrate the potential  
 493 of our synthetic CT images for downstream segmentation tasks.

494 For completeness, we report the Dice scores  
 495 with 95% confidence intervals: U-Net (Real:  
 496 0.874 [0.857–0.892], Real+Synthetic Data:  
 497 0.876 [0.858–0.893]) and TransUNet (Real:  
 498 0.872 [0.855–0.890], Real+Synthetic Data:  
 499 0.877 [0.860–0.894]), showing that the im-  
 500 provements are consistent and synthetic data  
 501 does not degrade segmentation performance.

502 Furthermore, the qualitative results are pre-  
 503 sented in Fig. 4, where our NA-Diff-F model  
 504 generates the CT images according to the noise-  
 505 aware semantic mask. In addition, as evidenced  
 506 from Table 7, our NA-Diff-D model outper-  
 507 forms the baseline when evaluated on the LiTS  
 508 dataset. We compare the original segmenta-  
 509 tion

510 masks with those predicted from the generated images (i.e., segmentation masks constructed by  
 511 applying a segmentation model Ward & Imran (2025) to the generated images). This evaluation  
 512 demonstrates that our method effectively preserves anatomical structures in the synthesized CT im-  
 513 ages in comparison to the baseline model. Overall, NA-Diff-F performs well while following both  
 514 noise and anatomy guidance. However, if we focus only on anatomy guidance, NA-Diff-D gives fa-  
 515 vorable results. **In addition to the improvement in image generation, NA-Diff can be more resource-  
 516 efficient compared to the baseline Seg-Diff.** Seg-Diff relies on a high-resource setup with  $4 \times$  A6000  
 517 GPUs (48 GB each), whereas our NA-Diff models were trained using only  $2 \times$  A4000 GPUs (32 GB  
 518 total). Despite using far fewer resource requirements, NA-Diff-D still achieves higher Dice, IoU,  
 519 HD, recall, and precision. Regarding data efficiency, Seg-Diff is trained on approximately  $10 \times$   
 520 more real CT images than our NA-Diff model (11000 vs 1116). Considering the scarcity of high-  
 521 quality medical datasets, our NA-Diff model reduces reliance on real CT data by leveraging large  
 522 natural-image collections through our CT-like noise-emulation and pretraining strategies, enabling  
 523 strong synthesis performance with far less real CT data and lower computational cost.

## 5 CONCLUSIONS

524 In this study, we present a novel noise and anatomy-guided diffusion model, NA-Diff, that can  
 525 generate realistic CT images with dose-level aware noise and liver anatomy map guidance. Our  
 526 innovative CT noise-emulated natural image pre-training helps capture complex features from dif-  
 527 ferent CT dose labels. In the fine-tuning phase on our emulated data, the model learns complex  
 528 CT image features with noise and also anatomical maps. Experimental evaluation on IQA and liver  
 529 segmentation demonstrates the realistic quality variation and preservation of anatomical structures,  
 530 along with noise-awareness in the generated CT images. **While higher-dose and low-noise CT scans  
 531 are preferable for diagnosis, noise variation itself is clinically important for downstream tasks such  
 532 as CT IQA, dose optimization, and selecting scans suitable for radiologist review or downstream  
 533 AI models.** These tasks require both high-quality and low-quality examples to accurately identify  
 534 when a scan is too noisy, low-dose, or degraded. Therefore, modeling noise variations is necessary  
 535 to teach the generative model how CT appearance changes under different dose levels, which is  
 536 essential for generating realistic low-quality images and for downstream dose- or quality-aware ap-  
 537 plications. Our noise-aware and anatomically consistent design enables NA-Diff to explicitly learn  
 538 this dose-quality relationship with the focus on the target organ, rather than collapsing toward only  
 539 high-quality reconstructions. Our future work will focus on a more extensive evaluation of NA-Diff  
 across different CT organs and various downstream tasks.

Table 7: The models are fine-tuned (with 3000 random sample images) and evaluated (with 500 random sample images) with segmentation guidance using LiTS Bilic et al. (2023) dataset. Our synthetic datasets show superior performance to the baseline. \* denotes that the model is trained with our CT image data conditioned on the anatomy mask.

Methods	Performance metrics				
	Dice	IoU	HD	Recall	Precision
Seg-Diff*	0.911	0.838	22.629	0.864	0.968
NA-Diff-C	0.910	0.836	21.382	0.866	0.970
NA-Diff-D	<b>0.920</b>	<b>0.853</b>	<b>20.508</b>	<b>0.872</b>	<b>0.974</b>

540 REFERENCES  
541

542 Abdullah al Nomaan Nafi, Md Alamgir Hossain, Rakib Hossain Rifat, Md Mahabub Uz Zaman,  
543 Md Manjurul Ahsan, and Shivakumar Raman. Diffusion-based approaches in medical image  
544 generation and analysis. *arXiv e-prints*, pp. arXiv–2412, 2024.

545 Aditya Bhat, Rupak Bose, Chinedu Innocent Nwoye, and Nicolas Padoy. Simgen: A diffusion-based  
546 framework for simultaneous surgical image and segmentation mask generation. *arXiv preprint*  
547 *arXiv:2501.09008*, 2025.

548 Moinak Bhattacharya, Gagandeep Singh, Shubham Jain, and Prateek Prasanna. Radgazegen: Ra-  
549 diomics and gaze-guided medical image generation using diffusion models. *arXiv preprint*  
550 *arXiv:2410.00307*, 2024.

551 Patrick Bilic, Patrick Christ, Hongwei Bran Li, Eugene Vorontsov, Avi Ben-Cohen, Georgios  
552 Kaassis, Adi Szeskin, Colin Jacobs, Gabriel Efrain Humpire Mamani, Gabriel Chartrand, et al.  
553 The liver tumor segmentation benchmark (lits). *Medical image analysis*, 84:102680, 2023.

554 Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd  
555 gans. *arXiv preprint arXiv:1801.01401*, 2018.

556 Rupak Bose, Chinedu Innocent Nwoye, Aditya Bhat, and Nicolas Padoy. Cosimgen: Controllable  
557 diffusion model for simultaneous image and mask generation. *arXiv preprint arXiv:2503.19661*,  
558 2025.

559 Mikhail Chaichuk, Sushant Gautam, Steven Hicks, and Elena Tutubalina. Prompt to polyp: Medical  
560 text-conditioned image synthesis with diffusion models. *arXiv preprint arXiv:2505.05573*, 2025.

561 Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille,  
562 and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation.  
563 *arXiv preprint arXiv:2102.04306*, 2021.

564 Ying Chen, Hongping Lin, Wei Zhang, Wang Chen, Zonglai Zhou, Ali Asghar Heidari, Huijing  
565 Chen, and Guohui Xu. Icycle-gan: Improved cycle generative adversarial networks for liver  
566 medical image generation. *Biomedical Signal Processing and Control*, 92:106100, 2024.

567 Joseph Cho, Mrudang Mathur, Cyril Zakka, Dhamanpreet Kaur, Matthew Leipzig, Alex Dalal, Ar-  
568 avind Krishnan, Eubee Koo, Karen Wai, Cindy S. Zhao, Rohan Shad, Robyn Fong, Ross Wight-  
569 man, Akshay Chaudhari, and William Hiesinger. Medisyn: A generalist text-guided latent diffu-  
570 sion model for diverse medical image synthesis, 2025.

571 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*  
572 *in neural information processing systems*, 34:8780–8794, 2021.

573 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
574 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An  
575 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint*  
576 *arXiv:2010.11929*, 2020.

577 Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image  
578 synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recogni-  
579 tion*, pp. 12873–12883, 2021.

580 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
581 Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the*  
582 *ACM*, 63(11):139–144, 2020.

583 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*  
584 *arXiv:2207.12598*, 2022.

585 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
586 *neural information processing systems*, 33:6840–6851, 2020.

594 Abdullah-Al-Zubaer Imran, Debashish Pal, Bhavik Patel, and Adam Wang. Ssiqa: Multi-task learn-  
 595 ing for non-reference ct image quality assessment with self-supervised noise level prediction. In  
 596 *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 1962–1965. IEEE,  
 597 2021.

598 Sadeep Jayasumana, Sri Kumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and  
 599 Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. In *Pro-  
 600 ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9307–  
 601 9315, 2024.

602 Yifan Jiang, Yannick Lemaréchal, Josée Bafaro, Jessica Abi-Rjeile, Philippe Joubert, Philippe De-  
 603 sprés, and Venkata Manem. Lung-ddpm: Semantic layout-guided diffusion models for thoracic  
 604 ct image synthesis. *arXiv preprint arXiv:2502.15204*, 2025.

605 Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations  
 606 and Trends® in Machine Learning*, 12(4):307–392, 2019.

607 Nicholas Konz, Yuwen Chen, Haoyu Dong, and Maciej A Mazurowski. Anatomically-controllable  
 608 medical image generation with segmentation-guided diffusion models. In *International Confer-  
 609 ence on Medical Image Computing and Computer-Assisted Intervention*, pp. 88–98. Springer,  
 610 2024.

611 Arjun Krishna, Ge Wang, and Klaus Mueller. Multi-conditioned denoising diffusion probabilistic  
 612 model (mddpm) for medical image synthesis. *arXiv preprint arXiv:2409.04670*, 2024.

613 Wonkyeong Lee, Fabian Wagner, Adrian Galdran, Yongyi Shi, Wenjun Xia, Ge Wang, Xuanqin  
 614 Mou, Md Atik Ahmed, Abdullah Al Zubaer Imran, Ji Eun Oh, et al. Low-dose computed to-  
 615 mography perceptual image quality assessment. *Medical Image Analysis*, 99:103343, 2025.

616 Wonkyeong Lee et al. Low-dose Computed Tomography Perceptual Image Quality Assessment  
 617 Grand Challenge Dataset. In *Medical Image Computing and Computer Assisted Intervention*,  
 618 2023. doi: 10.5281/zenodo.7833096.

619 Xinyu Liu, Wuyang Li, and Yixuan Yuan. Diffrect: Latent diffusion label rectification for semi-  
 620 supervised medical image segmentation. In *International Conference on Medical Image Comput-  
 621 ing and Computer-Assisted Intervention*, pp. 56–66. Springer, 2024.

622 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool.  
 623 Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the  
 624 IEEE/CVF conference on computer vision and pattern recognition*, pp. 11461–11471, 2022.

625 Jiawei Mao, Yuhang Wang, Yucheng Tang, Daguang Xu, Kang Wang, Yang Yang, Zongwei Zhou,  
 626 and Yuyin Zhou. Medsegfactory: Text-guided generation of medical image-mask pairs. *arXiv  
 627 preprint arXiv:2504.06897*, 2025.

628 Cynthia McCollough, Bin Chen, David R. Holmes III, Xiaohui Duan, Zhiqiang Yu, Lifeng Yu,  
 629 Shuai Leng, and Joel Fletcher. Low dose ct image and projection data (ldct-and-projection-data)  
 630 (version 6), 2020. URL <https://doi.org/10.7937/9NPB-2637>.

631 Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint  
 632 arXiv:1411.1784*, 2014.

633 Nusrat Munia and Abdullah Al Zubaer Imran. Prompting medical vision-language models to miti-  
 634 gate diagnosis bias by generating realistic dermoscopic images. In *2025 IEEE 22nd International  
 635 Symposium on Biomedical Imaging (ISBI)*, pp. 1–4. IEEE, 2025.

636 Xingyang Nie, Su Pan, Xiaoyu Zhai, Shifei Tao, Fengzhong Qu, Biao Wang, Huilin Ge, and Guojie  
 637 Xiao. Image-conditional diffusion transformer for underwater image enhancement. *arXiv preprint  
 638 arXiv:2407.05389*, 2024.

639 Hyun-Jic Oh and Won-Ki Jeong. Controllable and efficient multi-class pathology nuclei data aug-  
 640 mentation using text-conditioned diffusion models. In *International Conference on Medical Im-  
 641 age Computing and Computer-Assisted Intervention*, pp. 36–46. Springer, 2024.

648 Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties  
 649 in gan evaluation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 650 *recognition*, pp. 11410–11420, 2022.

651

652 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*  
 653 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

654 Ariadna Quattoni and Antonio Torralba. Recognizing indoor scenes. In *2009 IEEE conference on*  
 655 *computer vision and pattern recognition*, pp. 413–420. IEEE, 2009.

656

657 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 658 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
 659 models from natural language supervision. In *International conference on machine learning*, pp.  
 660 8748–8763. PMLR, 2021.

661

662 Kazi Ramisa Rifa, Md Atik Ahamed, Jie Zhang, and Abdullah Imran. Tfkt v2: task-focused knowl-  
 663 edge transfer from natural images for computed tomography perceptual image quality assessment.  
 664 *Journal of Medical Imaging*, 12(5):051805–051805, 2025a.

665

666 Kazi Ramisa Rifa, Md. Atik Ahamed, Jie Zhang, and Abdullah-Al-Zubaer Imran. Task-focused  
 667 knowledge transfer from natural images for ct image quality assessment. In *Medical Imaging*  
 668 *2025: Image Processing*. International Society for Optics and Photonics (SPIE), 2025b.

669

670 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 671 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-  
 672 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

673

674 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomed-  
 675 ical image segmentation. In *International Conference on Medical image computing and computer-  
 676 assisted intervention*, pp. 234–241. Springer, 2015.

677

678 Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David  
 679 Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH*  
 680 *2022 conference proceedings*, pp. 1–10, 2022.

681

682 Carl H Schultz, Romeo Fairley, Linda Suk-Ling Murphy, and Mohan Doss. The risk of cancer from  
 683 ct scans and other sources of low-dose radiation: a critical appraisal of methodologic quality.  
 684 *Prehospital and disaster medicine*, 35(1):3–16, 2020.

685

686 Ruochong Shen, Xiaoxu Li, Yuan-Fang Li, Chao Sui, Yu Peng, and QiuHong Ke. Dcdiff: Dual-  
 687 domain conditional diffusion for ct metal artifact reduction. In *International Conference on Med-  
 688 ical Image Computing and Computer-Assisted Intervention*, pp. 223–232. Springer, 2024.

689

690 Yongyi Shi, Wenjun Xia, Ge Wang, and Xuanqin Mou. Blind ct image quality assessment using  
 691 ddpm-derived content and transformer-based evaluator. *IEEE Transactions on Medical Imaging*,  
 692 43(10):3559–3569, 2024.

693

694 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*  
 695 *preprint arXiv:2010.02502*, 2020.

696

697 Wenbo Song, Yan Jiang, Yin Fang, Xinyu Cao, Peiyan Wu, Hanshuo Xing, and Xinglong Wu.  
 698 Medical image generation based on latent diffusion models. In *2023 International Conference on*  
 699 *Artificial Intelligence Innovation (ICAI)*, pp. 89–93. IEEE, 2023.

700

701 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.  
 702 *Advances in neural information processing systems*, 32, 2019.

703

704 Danush Kumar Venkatesh, Dominik Rivoir, Micha Pfeiffer, Fiona Kolbinger, and Stefanie Speidel.  
 705 Data augmentation for surgical scene segmentation with anatomy-aware diffusion models, 2024.

706

707 Tyler Ward and Abdullah-Al-Zubaer Imran. Annotation-efficient task guidance for medical segment  
 708 anything. *arXiv preprint arXiv:2412.08575*, 2024.

702 Tyler Ward and Abdullah Al Zubaer Imran. Annotation-efficient task guidance for medical segment  
 703 anything. In *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, pp. 1–4.  
 704 IEEE, 2025.

705

706 Zhaohu Xing, Sicheng Yang, Sixiang Chen, Tian Ye, Yijun Yang, Jing Qin, and Lei Zhu. Cross-  
 707 conditioned diffusion model for medical image to image translation. In *International Conference*  
 708 *on Medical Image Computing and Computer-Assisted Intervention*, pp. 201–211. Springer, 2024.

709

710 Yanwu Xu, Li Sun, Wei Peng, Shuyue Jia, Katelyn Morrison, Adam Perer, Afroz Zandifar, Shyam  
 711 Visweswaran, Motahhare Eslami, and Kayhan Batmanghelich. Medsyn: Text-guided anatomy-  
 712 aware synthesis of high-fidelity 3-d ct images. *IEEE Transactions on Medical Imaging*, 2024.

713

714 Yulin Yang, Qingqing Chen, Yinhao Li, Fang Wang, Xian-Hua Han, Yutaro Iwamoto, Jing Liu,  
 715 Lanfen Lin, Hongjie Hu, and Yen-Wei Chen. Segmentation guided crossing dual decoding gen-  
 716 erative adversarial network for synthesizing contrast-enhanced computed tomography images.  
*IEEE Journal of Biomedical and Health Informatics*, 2024.

717

718 Li Zhang, Basu Jindal, Ahmed Alaa, Robert Weinreb, David Wilson, Eran Segal, James Zou, and  
 719 Pengtao Xie. Generative ai enables medical image segmentation in ultra low-data regimes. *Nature*  
*Communications*, 16(1):6486, 2025.

720

721 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image  
 722 diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*,  
 pp. 3836–3847, 2023.

723

724 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
 725 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*  
*726 computer vision and pattern recognition*, pp. 586–595, 2018a.

727

728 Zheyuan Zhang, Lanhong Yao, Bin Wang, Debesh Jha, Gorkem Durak, Elif Keles, Alpay Mede-  
 729 talibeyoglu, and Ulas Bagci. Diffboost: Enhancing medical image segmentation via text-guided  
*730 diffusion model. IEEE Transactions on Medical Imaging*, 2024.

731

732 Zizhao Zhang, Lin Yang, and Yefeng Zheng. Translating and segmenting multimodal medical vol-  
 733 umes with cycle-and shape-consistency generative adversarial network. In *Proceedings of the*  
*734 IEEE conference on computer vision and pattern Recognition*, pp. 9242–9251, 2018b.

735

736 Jian-Qing Zheng, Yuanhan Mo, Yang Sun, Jiahua Li, Fuping Wu, Ziyang Wang, Tonia Vincent, and  
 Bartłomiej W Papież. Deformation-recovery diffusion model (drdm): instance deformation for  
 737 image manipulation and synthesis. *arXiv preprint arXiv:2407.07295*, 2024.

738

739 Qiang Zhu, Kuan Lu, Menghao Huo, and Yuxiao Li. Image-to-image translation with diffusion  
 740 transformers and clip-based image conditioning. *arXiv preprint arXiv:2505.16001*, 2025.

741

742 Yuanzhi Zhu, Zhaohai Li, Tianwei Wang, Mengchao He, and Cong Yao. Conditional text image  
 743 generation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer*  
 744 *Vision and Pattern Recognition*, pp. 14235–14245, 2023.

745

746

747

748

749

750

751

752

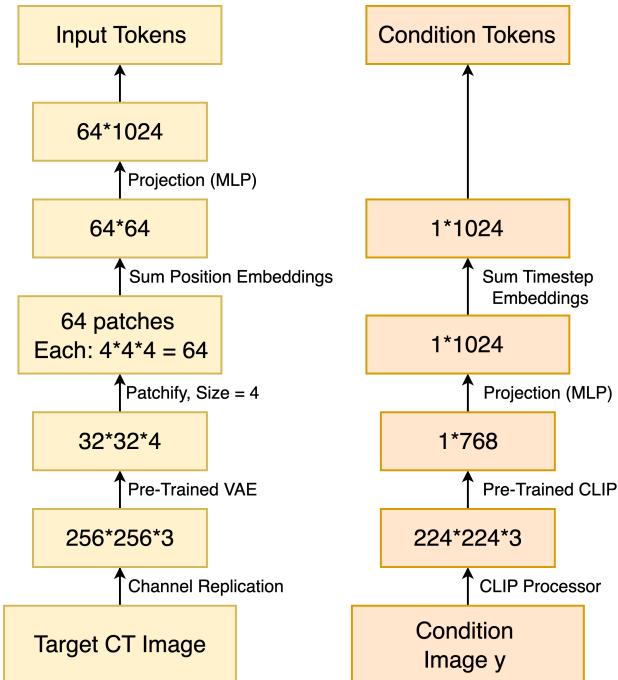
753

754

755

756 A NA-DIFF VARIANTS  
757758 A.1 SEP-DIFF  
759

760 In Sep-Diff, it has three stages of training. In the first stage, the model uses a CT-emulated natural  
761 image and input conditioned on noise maps. This stage is considered a natural image pre-training  
762 phase. Both the input image and the output image are encoded using the VAE Kingma et al. (2019)  
763 encoder. In the second stage, the model is fine-tuned, leveraging CT images and their corresponding  
764 simulated noise maps. In this stage, the input image is also encoded using the VAE encoder; how-  
765 ever, for noise maps, we use the CLIP Radford et al. (2021) model encoder to get a global semantic  
766 embedding as shown in Figure 5. In the final stage, the model fine-tunes utilizing CT images and  
767 their corresponding liver segmentation mask. For the image encoding process, the model follows a  
768 process similar to the second stage.



791 Figure 5: Dimensionality flow of Sep-Diff variants (excluding pre-training). The CT target image is  
792 encoded by a pre-trained VAE into latent patches (e.g., 4x4x4), flattened and projected to the trans-  
793 former’s hidden space. In parallel, the condition image is processed by a pre-trained CLIP model  
794 to produce a global feature vector, which is linearly projected, fused with timestep embeddings, and  
795 combined with input tokens for the transformer.

796 A.2 CAT-DIFF  
797

798 Unlike Sep-Diff, Cat-Diff contains a two-stage training process. In the first stage of training, it  
799 follows the exact same process as Sep-Diff pre-training. However, in the second stage, it combines  
800 the second and third stage training of Sep-Diff into one stage training. Here, the model leverages CT  
801 images with corresponding noise maps and anatomy masks as conditioning. While input images are  
802 encoded using VAE, both conditions use CLIP as an encoder. After getting the global embedding,  
803 we employ concatenation to emphasize dual conditioning. Let  $y_{\text{map}}, y_{\text{mask}} \in \mathbb{R}^{768}$  denote the CLIP  
804 embeddings of the noise map and segmentation mask, respectively. These embeddings are projected  
805 into the transformer’s hidden space using a learnable MLP:

$$\tilde{y}_{\text{map}} = \phi(y_{\text{map}}), \quad \tilde{y}_{\text{mask}} = \phi(y_{\text{mask}}), \quad \rho : \mathbb{R}^{768} \rightarrow \mathbb{R}^{\mathcal{D}}, \quad (10)$$

806 where  $\mathcal{D}$  is the transformer hidden dimension (e.g., 1024), and  $\rho$  is defined as:

$$\rho(y) = \mathcal{W}_2 \cdot \text{SiLU}(\mathcal{W}_1 y + \mathcal{B}_1) + \mathcal{B}_2. \quad (11)$$

810 To preserve the distinct information from both embeddings, we concatenate them along the channel  
 811 dimension:

$$\tilde{y}_{\text{concat}} = [\tilde{y}_{\text{map}} \parallel \tilde{y}_{\text{mask}}] \in \mathbb{R}^{2D}. \quad (12)$$

814 The concatenated vector is then fused and projected back to  $\mathbb{R}^D$  using a second MLP:

$$\tilde{y}_{\text{final}} = \eta(\tilde{y}_{\text{concat}}), \quad \eta : \mathbb{R}^{2D} \rightarrow \mathbb{R}^D. \quad (13)$$

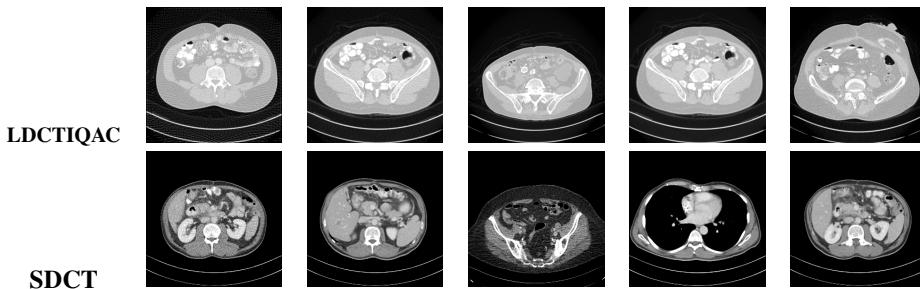
818 Finally, this fused conditioning vector is added to the sinusoidal timestep embedding  $\gamma(t)$  to form  
 819 the complete conditioning representation:

$$c = \tilde{y}_{\text{final}} + f_{\theta}(t), \quad (14)$$

822 where  $f_{\theta}(t) \in \mathbb{R}^D$  is the output of an MLP applied to a sinusoidal timestep embedding. This  
 823 conditioning vector  $c$  is broadcast to the transformer blocks and modulates each residual path using  
 824 adaptive layer normalization (adaLN-Zero).

## 826 B IQA AND SEGMENTATION EVALUATION

828 NA-Diff model variants are evaluated by predicting diagnostic image quality assessment (IQA)  
 829 scores, providing a measure of noise level estimation for the generated CT images. Although the  
 830 IQA score may not be perfectly accurate due to differences in window center (40) and width (400)  
 831 settings between our dataset and the pre-trained IQA model’s Rifa et al. (2025a) training data (as  
 832 shown in Fig. 6), a relative comparison is still feasible. We utilize another pretrained model Ward &  
 833 Imran (2024) for liver segmentation masks comparison; the generated CT images from our model  
 834 are passed through the segmentation model to produce liver masks and show a comparison between  
 835 the model-generated mask and the real mask images. This allows us to assess whether the anatomical  
 836 structures are being effectively preserved in the synthesized images. We report Dice score,  
 837 IoU, Precision, Recall, and Hausdorff Distance (HD) to assess how well the segmentation masks  
 838 predicted from the generated CT images compare to the real ground-truth masks.



840  
 841  
 842  
 843  
 844  
 845  
 846  
 847  
 848 Figure 6: Visualization of the samples of IQA-based model’s Rifa et al. (2025b) trained dataset  
 849 LDCTIQAC Lee et al. (2023) and our SDCT dataset.

## 851 C ADDITIONAL EXPERIMENTS

854 The IQA scores for NA-Diff-D and NA-Diff-F are quite similar across different dose levels as shown  
 855 in Fig. 7. For NA-Diff-F, the IQA scores increase steadily as the dose level rises, indicating that it  
 856 is more effective at handling noise in generated images. Qualitative results comparing the baseline  
 857 model and our anatomy-guided variants are shown in Fig. 8. Our method preserves anatomical  
 858 structures more compared to the baseline.

859 Moreover, as reported in Table 8, our NA-Diff-D model outperforms the baseline when evaluated  
 860 on the LiTS dataset. Compared to the Seg-Diff\* baseline, our NA-Diff-D model achieves the best  
 861 performance across all metrics, improving CFID by 11.98%, KID by 25%, CMMD by 11.4%, and  
 862 LPIPS by 0.4%. This highlights the efficacy of our anatomy-guided approach for generating high-  
 863 fidelity and perceptually consistent CT images. Also, the results in Table 9 highlight the advantage  
 864 of concatenating condition tokens over summing them for dual conditioning.

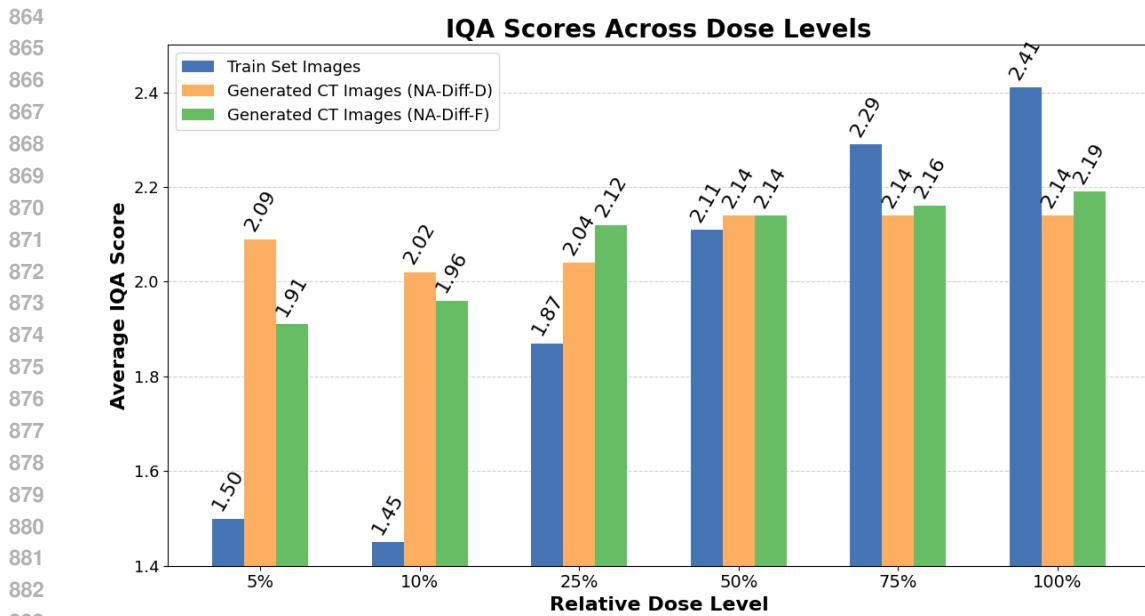


Figure 7: Predicted IQA scores for the NA-Diff-F generated CT images compared to input CT images at different dose levels.

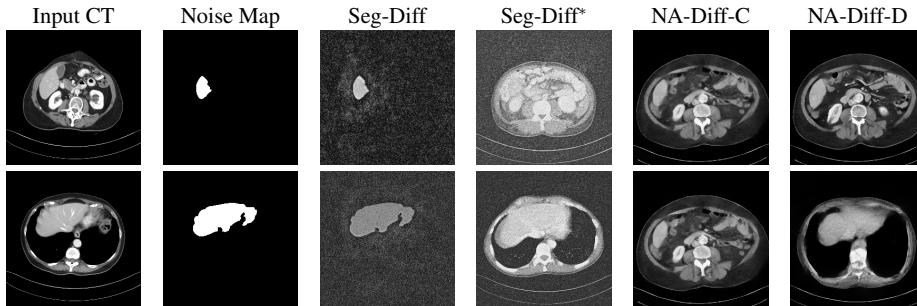


Figure 8: Qualitative results with baseline and antamody guided variants of NA-Diff.

In Fig. 10, NA-Diff-F demonstrates strong noise map awareness, adapting its generation quality following varied noise map levels. Furthermore, we have the overlay results from NA-Diff-D and NA-Diff-F as shown in Fig. 9, which show good anatomical alignment between the generated images and the segmentation masks.

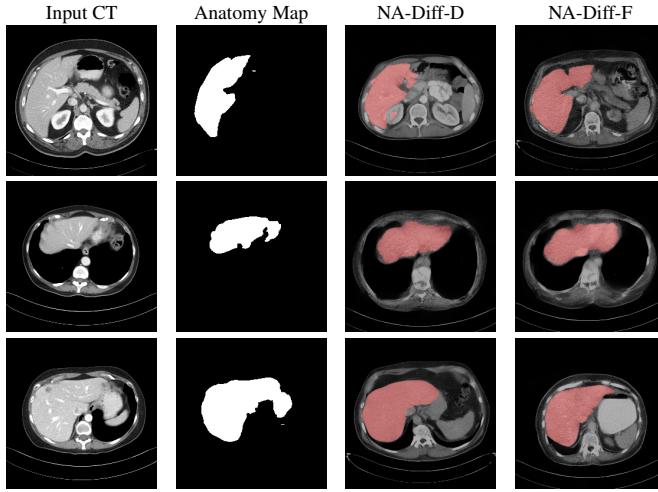
Algorithm 2 shows how the input data is conditioned. The grayscale image  $x$  is first repeated across three channels and encoded into a latent  $z$  using the VAE encoder  $\mathcal{E}$ . The anatomy map  $n$  and mask  $m$  are separately encoded by the CLIP image encoder  $\phi$  to obtain feature vectors  $f_n$  and  $f_m$ . These are projected through  $W_{clip}$  and fused by  $W_{cat}$  to form the condition embedding  $c$ .

Algorithm 1 describes the DiT forward pass with adaLN-Zero conditioning. The latent  $z$  is patch-embedded with fixed sinusoidal positional embedding  $P$ . A timestep  $t$  is mapped into an embedding  $e_t$  (via a sinusoidal MLP) and added to  $c$  to form the global conditioner  $h$ . In each DiT block,  $h$  is used by adaLN to generate scale  $\Gamma$ , shift  $\Delta$ , and gating  $g$  parameters that modulate the LayerNorm outputs before multi-head self-attention (MSA) and MLP layers. After  $L$  such blocks, the final layer produces the predicted noise  $\hat{\epsilon}_\theta$ . This prediction is used in the diffusion loss to train the model.

## D CT NOISE SIMULATION

Diffusion models require large-scale training data. However, using natural images for pre-training can lead to domain shift issues. To address this, we emulate CT-specific noise patterns within natural images, creating CT dose-emulated datasets for the pre-training phase. Fig. 11 illustrates examples

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939



940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Figure 9: Qualitative comparison showing overlay outputs from our NA-Diff-D and NA-Diff-F variants. The visualization highlights segmentation accuracy and generative performance.

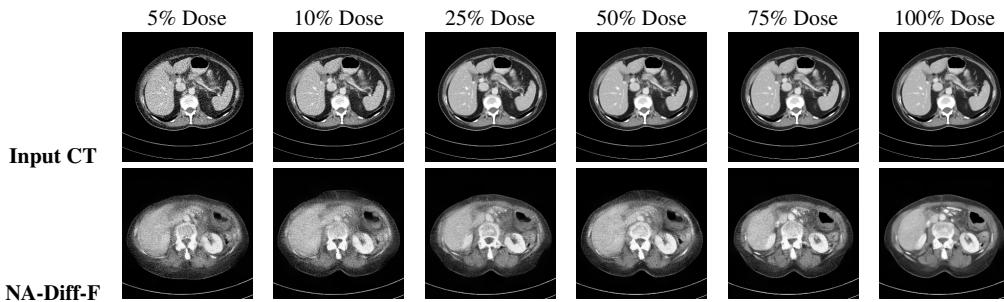


Figure 10: Qualitative comparisons at different noise levels show that our NA-Diff-F model effectively captures and adapts to the provided noise maps.

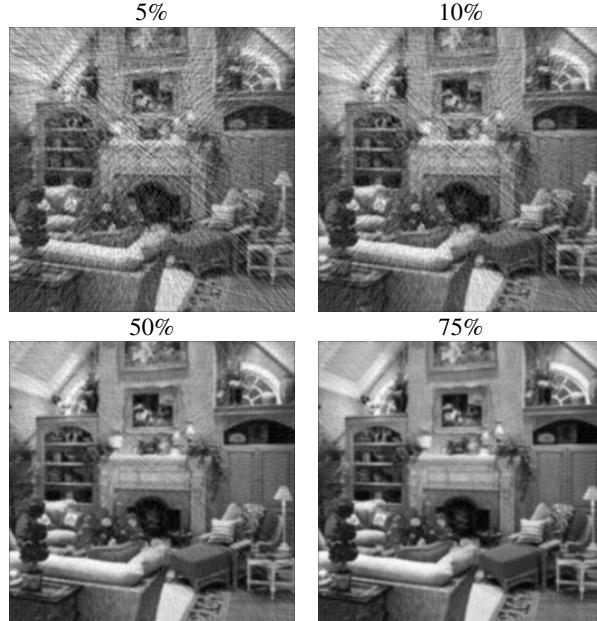
972  
 973  
 974  
 975  
 976  
 Table 8: The models are fine-tuned (with 3000 random sample images) and evaluated (with 500  
 977 random sample images) with segmentation guidance using the LiTS Bilic et al. (2023) dataset. Our  
 978 synthetic datasets show superior performance to the baseline. \* denotes that the model is trained  
 979 with our CT image data conditioned on the anatomy mask. The best and second-best results are  
 980 **bolded** and underlined respectively.  
 981

977 Methods	978 Performance metrics			
	979 CFID ↓	980 KID ↓	981 CMMMD ↓	982 LPIPS ↓
983 Seg-Diff*	60.5636	25.3853 ± 3.065	0.376	0.3833 ± 0.1023
984 NA-Diff-C	89.2949	57.6014 ± 5.8231	0.541	0.3866 ± 0.1021
985 NA-Diff-D	<b>53.2949</b>	<b>19.0386 ± 2.8532</b>	<b>0.333</b>	<b>0.3817 ± 0.1029</b>

982 Table 9: Ablation study highlighting the importance of using concatenation in dual conditioning.  
 983 The sum and concatenation denotes operations on condition token.  
 984

985 Methods	986 Operation	987 Performance metrics			
		988 CFID ↓	989 KID ↓	990 CMMMD ↓	991 LPIPS ↓
992 NA-Diff-F	993 Sum	81.3672	52.7426 ± 5.2481	0.492	0.3824 ± 0.1018
994 NA-Diff-F	995 Concat.	<b>64.6379</b>	<b>24.8106 ± 3.2275</b>	<b>0.3205</b>	<b>0.3760 ± 0.0994</b>

996 of natural images augmented with CT-like noise characteristics. Moreover, Fig. 12 visualizes simulated CT images corresponding to relative dose levels.  
 997  
 998  
 999  
 1000  
 1001  
 1002  
 1003  
 1004  
 1005  
 1006  
 1007  
 1008  
 1009  
 1010  
 1011  
 1012  
 1013  
 1014



1015 Figure 11: Visualization of natural images at different simulated dose levels.  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025

1026  
1027**Algorithm 1** DiT forward pass with AdaLN-Zero

---

1028  
1029  
1030 **Require:** latent image tokens  $z$ , condition embedding  $c$ , timestep embedding  $t$   
1031 1:  $X \leftarrow \text{PatchEmbed}(z) + P$  ▷ initial token sequence  
1032 2:  $h \leftarrow \text{sincosMLP}(t) + c$  ▷ conditioning vector  
1033 3: **for**  $\ell = 1$  to  $L$  **do**  
1034 4:    $(\gamma_t^{(\ell)}, \beta_t^{(\ell)}, \alpha_t^{(\ell)}) \leftarrow \text{MLP}(\text{LN}(h))$   
1035 5:    $X \leftarrow X + \alpha_t^{(\ell)} \cdot \text{MSA}(\gamma_t^{(\ell)} \odot \text{LN}(X) + \beta_t^{(\ell)})$   
1036 6:    $X \leftarrow X + \alpha_t^{(\ell)} \cdot \text{MLP}(\gamma_t^{(\ell)} \odot \text{LN}(X) + \beta_t^{(\ell)})$   
1037 7: **end for**  
1038 8:  $(\gamma'_t, \beta'_t) \leftarrow \text{MLP}(\text{LN}(h))$   
1039 9:  $\hat{\vartheta} \leftarrow \text{unpatchify}(W_{\text{out}}(\gamma'_t \odot \text{LN}(X) + \beta'_t))$   
1040 10: **return**  $\hat{\vartheta}$  ▷ predicted noise  
1041

---

1042  
1043  
1044  
1045  
1046**Algorithm 2** Data conditioning

---

1047  
1048 **Require:**  $x$  (gray),  $n$  (map),  $m$  (mask);  $\mathcal{E}, \phi$   
1049 1:  $z \leftarrow 0.18215 \cdot \mathcal{E}(\text{repeat}(x, 3)).\text{sample}()$   
1050 2:  $f_n \leftarrow \phi(n); f_m \leftarrow \phi(m)$   
1051 3:  $c \leftarrow W_{\text{cat}}([W_{\text{clip}}(f_n) \parallel W_{\text{clip}}(f_m)])$   
1052 4: **return**  $(z, c)$   
1053

---

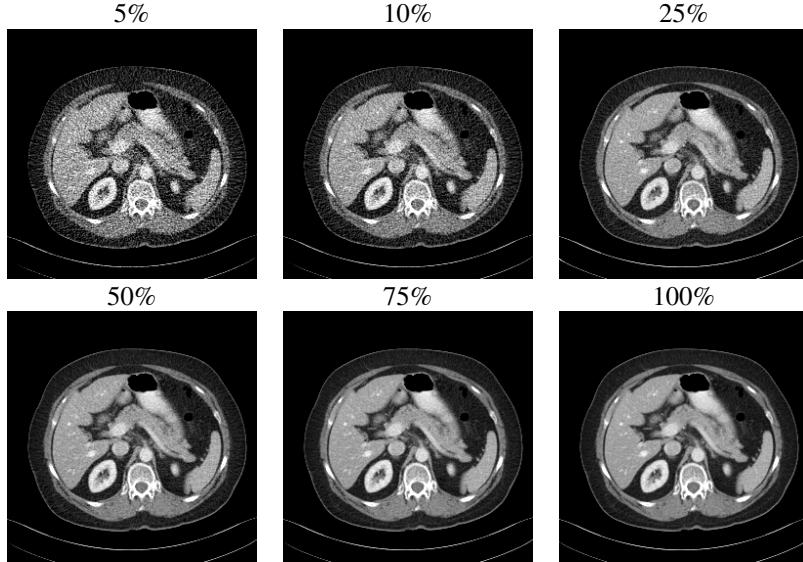
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

Figure 12: Visualization of CT images across different simulated dose levels (e.g., only 25% and 100% are provided; other levels are constructed using our low-dose simulation technique).