

Anatomically-Controllable Medical Image Generation with Segmentation-Guided Diffusion Models

Nicholas Konz^{1*}, Yuwen Chen¹, Haoyu Dong¹, and Maciej A. Mazurowski^{1,2,3,4}

¹ Department of Electrical and Computer Engineering, Duke University, NC, USA

² Department of Radiology, Duke University, NC, USA

³ Department of Computer Science, Duke University, NC, USA

⁴ Department of Biostatistics & Bioinformatics, Duke University, NC, USA

*corresponding author nicholas.konz@duke.edu

Abstract. Diffusion models have enabled remarkably high-quality medical image generation, yet it is challenging to enforce anatomical constraints in generated images. To this end, we propose a diffusion model-based method that supports anatomically-controllable medical image generation, by following a multi-class anatomical segmentation mask at each sampling step. We additionally introduce a *random mask ablation* training algorithm to enable conditioning on a selected combination of anatomical constraints while allowing flexibility in other anatomical areas. We compare our method (“**SegGuidedDiff**”) to existing methods on breast MRI and abdominal/neck-to-pelvis CT datasets with a wide range of anatomical objects. Results show that our method reaches a new state-of-the-art in the faithfulness of generated images to input anatomical masks on both datasets, and is on par for general anatomical realism. Finally, our model also enjoys the extra benefit of being able to adjust the anatomical similarity of generated images to real images of choice through interpolation in its latent space. SegGuidedDiff has many applications, including cross-modality translation, and the generation of paired or counterfactual data. Our code is available at <https://github.com/mazurowski-lab/segmentation-guided-diffusion>.

Keywords: diffusion models · image generation · semantic synthesis

Introduction

Denoising diffusion probabilistic models [9] (DDPMs, or just “diffusion models”) have shown extensive applications in medical image analysis [11] due to their ability to generate high-quality, high-resolution images, such as for direct image generation [19,12], image segmentation [27], anomaly detection [18,26], cross-modality image translation [15], and image denoising [6]. However, standard generative models like DDPMs can still fail to create anatomically plausible tissue (Fig 1), and such anatomy is not precisely customizable. Our proposed solution is to incorporate anatomical information as a prior for image generation

via a segmentation mask for different types of tissue, organs, etc., providing the network with a more direct learning signal for anatomical realism.

Generating an image from a mask (a.k.a. *semantic synthesis*) is a type of image-to-image translation task. Existing translation works include GAN-based [28,3,2] and diffusion-based models [26,21,29], yet these do not directly enforce precise pixel-wise anatomical constraints. Recent works introduced fine-tuning large pre-trained text-to-image latent diffusion models (LDMs) for segmentation-conditioning on natural images [29,25], but we found that these adapt poorly to medical images (Sec. 3.1). Instead, we propose the first diffusion model for generating medical images from anatomical segmentations, which is assisted by it being an image-space diffusion model trained from scratch. Image-space-based diffusion models are especially amenable for strict segmentation guidance because a conditioning mask can be used as-is for every small step of the denoising process, with no conversion to some abstract latent space as in certain LDMs [29,25,21] (or GANs) where precise spatial guidance may be lost.




Fig. 1. Standard diffusion models (right) can fail to create realistic tissue even if the overall image appears high-quality, motivating our segmentation-guided model (center).

Segmentation-guided generation would be even more flexible if only certain object classes could be constrained in an input mask, while others are free to be inferred by the model. This opens up further applications such as generating images from incomplete masks ([1]), the generation of anatomically paired/registered data, “counterfactual” analysis of existing annotated data, and others. To solve this, we propose a **mask-ablated training** strategy to provide the model with all possible combinations of missing classes in masks during training, teaching it to make such inferences when generating new images. Notably, we also demonstrate how interpolating within the latent space of our mask ablated-trained model enables generating images with adjustable anatomical similarity to some real image (Sec. 3.2).

In summary, we introduce a diffusion model, “**SegGuidedDiff**” that can flexibly and precisely generate medical images according to anatomical masks. We evaluate our model’s ability to generate images conditioned on a range of

anatomical objects of interest for breast MRI and neck-to-pelvis CT, where it outperforms state-of-the-art mask-conditional generative models in its faithfulness to input anatomical masks, and is on par for general anatomical realism (Sec. 3.1). Our code is publicly released at <https://github.com/mazurowski-lab/segmentation-guided-diffusion> with a focus on usability on any dataset, along with a dataset of paired “pre-registered” generated breast MRIs.

1 Method

1.1 A Brief Review of Diffusion Models

Denoising diffusion probabilistic models [9] (DDPMs, or diffusion models for short) are a type of generative latent variable model that learns to sample from some data distribution $p(x_0)$ ($x_0 \in \mathbb{R}^n$) by defining a stochastic process that gradually converts the data to noise—the *forward* process $q(x_t|x_{t-1})$ —and learning to reverse this process via a learned denoising process $p_\theta(x_{t-1}|x_t)$, where θ is the model parameters. Data is generated by iteratively sampling from $p_\theta(x_{t-1}|x_t)$, beginning with a Gaussian noise sample $x_T \sim p(x_T)$, for $t = T - 1, \dots, 0$ (we use $T = 1000$) until an image x_0 is recovered.

Any forward process step can be written explicitly as $x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon$ where $\epsilon \sim \mathcal{N}(0, I_n)$, and $\alpha_t := 1 - \beta_t$ given the variance of the additive pre-scheduled noise β_t , and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. DDPMs can be trained by the usual evidence lower bound (ELBO) maximization, which can be approximately optimized in a relatively simple form by training a network $\epsilon_\theta(x_t, t)$ to predict the noise ϵ added to each datapoint x_0 for various time steps t , with the loss $L = \mathbb{E}_{x_0, t, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2] = \mathbb{E}_{x_0, t, \epsilon} [\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, t)\|^2]$, which has proven to be the typically superior DDPM loss in practice [16].

1.2 Adding Segmentation Guidance to Diffusion Models

Rather than sampling from the unconditional distribution $p(x_0)$, our goal is to condition the generation of some c -channel image $x_0 \in \mathbb{R}^{c \times h \times w}$ to follow some multi-class anatomical mask $m \in \{0, \dots, C - 1\}^{h \times w}$, where C is the number of classes (including background), or in other words, sample from $p(x_0|m)$. While modifying the data likelihood $p(x_0|m)$ to be mask-conditional does not alter the noising process $q(x_t|x_{t-1})$, it does modify the reverse process $p_\theta(x_{t-1}|x_t, m)$ and noise-predicting network ϵ_θ . Propagating this to the ELBO results in a loss of

$$L_m = \mathbb{E}_{(x_0, m), t, \epsilon} [\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, t|m)\|^2] \quad (1)$$

for training our model, where each training image x_0 has some paired mask m . We propose to implement this simply by concatenating m channel-wise to the network input at every denoising step, *i.e.*, modifying the network to have an additional input channel as $\epsilon_\theta(x_t, t|m) : \mathbb{R}^{(c+1) \times h \times w} \rightarrow \mathbb{R}^{c \times h \times w}$, which can be any image-to-image model (see Sec. 3). We use the DDIM algorithm [24] for fast, yet high-quality sampling.

This simple method generates images that are very faithful to input masks (Fig. 2, Table 1), because the denoising process is conditioned on the mask at each of its many gradual steps, allowing the network to follow the masks because they provide helpful spatial information that is directly correlated with the optimal denoised model output that minimizes the loss.

1.3 Mask-Ablated Training and Sampling

Given that our model is mask-guided, the quality of these masks is important; a generated image may be misleading if the input mask is not fully annotated, known as the *partial label problem* [1] in medical image analysis. This is because the model may assume that un-annotated objects should not be present in the output image whatsoever (associating the missing/zero pixel labels as background or some other object), when in reality we may desire for the model to simply “fill in”/infer the unprovided objects.

To alleviate this problem, we propose a **mask-ablated training** (MAT) strategy (Algorithm 1), which provides examples of masks with various numbers and combinations of classes removed for the model to learn to generate images from during training. This can be thought of as a form of self-supervised learning of anatomical object representations (somewhat analogous to MAE [7]). We set all 2^{C-1} of these possible combinations of classes being removed from a given mask in training to occur with equal probability so that the model can handle each equally, although any other balancing of these probabilities for different object classes could be used. Finally, we note that our MAT algorithm is immediately applicable to *any* mask-conditional generative model.

Algorithm 1: Segmentation-guided model training with mask ablation.

Input: number of mask classes C , dataset $p(x_0, m)$.

```

repeat
   $x_0, m \sim p(x_0, m)$ 
  for  $c = 1, \dots, C - 1$  do
     $\delta \sim \text{Bernoulli}(0.5)$ 
    if  $\delta = 1$  then
       $m[m = c] = 0$ 
    end
   $\epsilon \sim \mathcal{N}(0, I_n); t \sim \text{Uniform}(\{1, \dots, T\})$ 
   $x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon$ 
  Update  $\theta$  with  $\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(x_t, t|m)\|^2$ 
until converged;

```

2 Datasets

Breast MRI: Our first dataset is a 100-patient subset of the Duke Breast Cancer MRI dataset [23]. We use all 2D image slices from the fat-saturated gradient echo T1-weighted pre-contrast sequence, with a train/test split of 70/15 patients,

resulting in $\sim 12000/2500$ slice images per split. We also keep a held-out training set of 15 patients for additional experiments. All images have full segmentation annotations for (1) breast, (2) blood vessels (BV), and (3) fibroglandular/dense tissue (FGT) provided at [14]. Notably, the FGT and BV have very high variability in shape, size, and other morphological characteristics, posing a challenge for generative models to realistically capture. **CT Organ:** Our second dataset is a 40-patient subset of neck-to-pelvis and abdominal CT scans from [20], with segmentation annotations for liver, bladder, lungs, kidney, and bone. This results in a train/test split of $\sim 11000/2100$ 2D slice images, given a patient-wise split of 24/8, as well as a held-out training set of 8 patients. All generative models are trained on the training sets, and the auxiliary segmentation network, introduced next, is trained on the held-out training sets.

3 Experiments

Training, Architecture and Implementational Details. All images are resized to 256×256 and normalized to $[0, 255]$. We use a UNet architecture [22] for the denoising model ϵ_θ , modified to take two channels (image and mask) as input; see Appendix A for additional training and architecture details.

3.1 Comparison to Existing Image Generation Models

We next compare our segmentation-guided diffusion model (“**SegGuidedDiff**” for short) to state-of-the-art segmentation-conditional image generation models. These are SPADE [17], a GAN-based model that uses spatially-adaptive normalization layers, and ControlNet [29], a recent method for adding spatial conditioning to large pretrained text-to-image diffusion models. Training/implementation details for each are in Appendix A.1; note that we use the standard implementation of SPADE rather than the recent brain MRI SPADE model [5,4] because its modifications to SPADE do not apply to our setting (see Appendix A.1), and that we follow ControlNet’s guidelines for adapting it to medical images. We show example generated images from all models in Fig. 2 (using masks randomly sampled from the test set); more are provided in Appendix B.

Evaluating Faithfulness of Generated Images to Input Masks. To measure how well our model follows an input mask for image generation, we use an auxiliary segmentation network trained on the real training set (a standard UNet; training details in Appendix A.2), to predict segmentations m_{gen}^{pred} on images that were generated from the masks m in the test set, and measure their overlap with (a) m and (b) the model’s predicted segmentations m_{real}^{pred} for the input masks’ original corresponding real images, similar to metrics used in [17]. Our model’s generated images have high overlap for both metrics (> 0.85 Dice coeff., Table 1), showing that our model closely followed the input masks when generating the images, and also outperformed the competing methods.

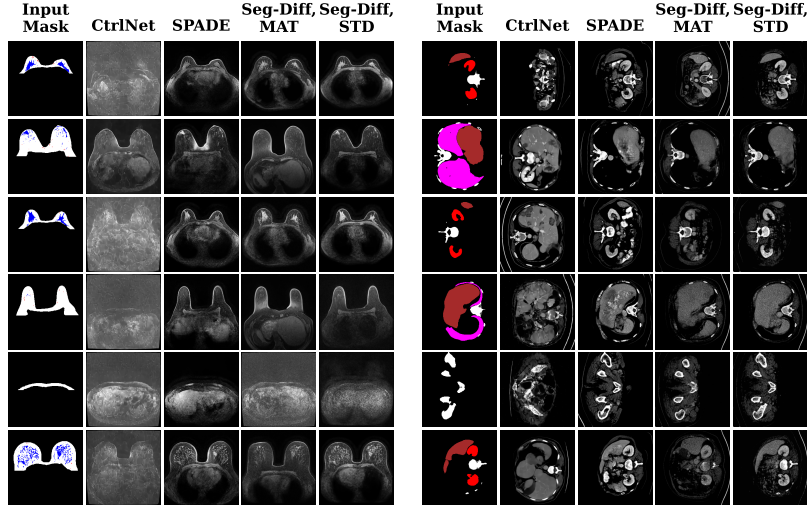


Fig. 2. Visual comparison of our model (SegGuidedDiff, or “Seg-Diff” for short) to existing segmentation-conditional image generation models. For breast MRI, the breast, BV, and FGT segmentations are shown as white, red, and blue, respectively, while for CT, the liver, bladder, lungs, kidneys, and bone are in maroon, orange, pink, red, and white, respectively. “MAT” = max ablated training, “STD” = our standard method.

Evaluating Generated Image Quality. We first attempted to use the common Fréchet Inception Distance (FID) [8] as a metric for quality/realism of generated image features compared to real data, via a CNN image encoder trained on the corresponding dataset. We observed that samples generated from both our segmentation-guided and standard unconditional diffusion models achieved potentially promising results, (*e.g.*, breast MRI feature FID $\simeq 0.5$), yet CNN feature-based metrics like FID fail to capture the global feature of anatomical realism that can differ in images generated by these models (*e.g.*, fibroglandular tissue as shown in Fig. 1), so we caution using such metrics.

Instead, we propose to more precisely measure anatomical realism by determining how well the aforementioned auxiliary segmentation models for the objects of interest can be trained solely on these synthetic images to be able to generalize to real data, using the input masks as targets. We compare the performance of the segmentation models trained on (a) the real held-out training set (Sec. 2) and (b) the set of images generated from all masks corresponding to these images. We split the real test set in half (by patient) into a validation set and a test set to use for these models. The results for this are in Table 2; we see that for both datasets, the segmentation network trained on our model’s synthetic data barely performs worse (by only ≤ 0.04 Dice) than the network trained on real data, implying that our generated images are both highly realistic and faithful to the input masks (especially considering certain objects’

Table 1. Faithfulness of generated images to input masks. m denotes input masks, and m_{gen}^{pred} and m_{real}^{pred} denote the masks predicted for (a) the generated images and (b) the real images corresponding to the input masks, respectively, by an auxiliary segmentation model. Best-performing is shown in bold, and second best is underlined.

	Breast MRI		CT Organ	
Model	$\text{Dice}(m_{gen}^{pred}, m)$	$\text{Dice}(m_{gen}^{pred}, m_{real}^{pred})$	$\text{Dice}(m_{gen}^{pred}, m)$	$\text{Dice}(m_{gen}^{pred}, m_{real}^{pred})$
ControlNet	0.3636	0.3604	0.1132	0.1126
SPADE	<u>0.8473</u>	<u>0.8477</u>	<u>0.8771</u>	<u>0.8603</u>
Ours	0.9027	0.8593	0.8980	0.8797

segmentation difficulty). Our method is on par with SPADE for CT Organ, and slightly worse for breast MRI, while outperforming ControlNet in both cases.

Table 2. Quality of generated images. Real test set performance (Dice coeff.) of a segmentation network for the objects of interest, trained on real data vs. synthetic data generated by different models.

	Real training set	Synthetic training set:		
		ControlNet	SPADE	Ours
Breast MRI	0.8376	0.7570	0.8333	<u>0.7934</u>
CT Organ	0.9075	0.0000	<u>0.8932</u>	0.8981

ControlNet performed poorly for all metrics because it failed to follow input masks closely (if at all) due to its limitations and untested nature of adapting to medical images; more details are in Appendix A.1. We also see that our mask-ablated-trained model (Sec 1.3) follows input masks less strictly than our standard model (Fig. 2), especially for the dataset with more object classes (CT Organ): for breast MRI the model obtained faithfulness metrics (Table 1) of $\text{Dice}(m_{gen}^{pred}, m) = 0.6589$ and $\text{Dice}(m_{gen}^{pred}, m_{real}^{pred}) = 0.6684$, and quality metric (Table 2) of 0.7478, and 0.5952, 0.5963 and 0.7564 for CT Organ, respectively. This is likely because the diversity of object class combinations seen in training scales exponentially with the number of classes in order to prepare the model for all possible combinations (Sec. 1.3), so the model is “spread thin”. However, this relaxation of the mask constraint has its own benefits, detailed next.

3.2 Advantages of Mask-Ablated-Training

Sampling from Ablated Masks. The direct benefit of mask ablated training (MAT) is its ability to generate images from masks with classes missing. In Fig. 3 we demonstrate the effect on generated images of ablating certain classes from an input mask for breast MRI (more examples, including for CT Organ, are provided in Appendix B). For example, we see that constraining the BV+FGT

in breast MRI, yet keeping the breast free, results in images that have the latter two classes pre-registered while the breast shape varies, and vice-versa.




Fig. 3. Generating images (even rows) from masks with classes removed (odd rows that are anatomically similar to real rows), shown for breast MRI. **Fig. 4.** Using our model to generate images that are anatomically similar to real images, shown for breast MRI.

Adjustable Anatomical Similarity of Generated Images to Real Images. One application of our model is the adjustable generation of images that are anatomically similar to some real image, which is not immediately possible for existing state-of-the-art GAN-based methods like SPADE. Consider some real image x_0 with anatomical mask m . We can *adjust* the anatomical similarity to x_0 of an image generated from m with our model by interpolating between the synthetic image and the real image in the model’s latent space, as follows.

If the generation/de-noising process given m (starting at timestep $t = T$) is halted at some intermediate step $t = \tilde{t}$ (we use $\tilde{t} = 240$), we obtain a latent representation $x'_{\tilde{t}}$ of the generating image. We can then convert the real image x_0 to the same latent space by applying the *noising* process to x_0 from $t = 0$ to $t = \tilde{t}$ to obtain $x_{\tilde{t}}$. Next, the features of the two images can be mixed via the *interpolated* latent $x_{\tilde{t}}^{\lambda} := (1 - \lambda)x_{\tilde{t}} + \lambda x'_{\tilde{t}}$, where $\lambda \in (0, 1]$ controls the similarity of the mixed features to those of the real image. From here, $x_{\tilde{t}}^{\lambda}$ can be denoised back to image space to obtain the interpolated image x_0^{λ} . An advantage of using our mask-ablated-trained model to generate x_0^{λ} is that certain objects can be constrained while others are free to vary, resulting in higher, yet adjustable, semantic diversity. We demonstrate this with various λ in Figure 4, with only FGT+BV constrained for breast MRI, and only bone constrained for CT Organ.

Conclusion

Our segmentation-guided diffusion model enables superior anatomically-controllable medical image generation, which has many potential applications, including (1) the generation of anatomically rare cases to augment some imbalanced dataset, (2) the synthesis of anatomically-paired/pre-registered data, and (3) *cross-modality* anatomy translation, where our model could be trained on images and masks from one sequence (*e.g.*, T2 MRI), and then supplied with masks from another sequence (*e.g.*, T1 MRI) to create new T2 images from the T1 masks.

However, some limitations of this study are that we did not consider full 3D generation, and we did not compare to ControlNet-like latent diffusion models [21] trained completely from scratch. For future work, we are interested in further improving generated image quality, incorporating image-level class guidance [10] either for pathological or domain-related variables [13], and extending our model to segmentation-guided image *translation*.

Acknowledgments. Research reported in this publication was supported by the National Institute Of Biomedical Imaging And Bioengineering of the National Institutes of Health under Award Number R01EB031575. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Disclosure of Interests. The authors have no competing interests.

References

1. Partial label learning: Taxonomy, analysis and outlook. *Neural Networks* (2023)
2. Cao, S., Konz, N., Duncan, J., Mazurowski, M.A.: Deep learning for breast mri style transfer with limited training data. *Journal of Digital Imaging* **36**(2), 666–678 (2023)
3. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8789–8797 (2018)
4. Fernandez, V., Pinaya, W.H.L., Borges, P., Graham, M.S., Vercauteren, T., Cardoso, M.J.: A 3D Generative Model of Pathological Multi-modal MR Images and Segmentations. In: *Deep Generative Models*. pp. 132–142. Springer Nature Switzerland, Cham (2024)
5. Fernandez, V., Pinaya, W.H.L., Borges, P., Tudosi, P.D., Graham, M.S., Vercauteren, T., Cardoso, M.J.: Can segmentation models be trained with fully synthetically generated data? In: *International Workshop on Simulation and Synthesis in Medical Imaging*. pp. 79–90. Springer (2022)
6. Gong, K., Johnson, K., El Fakhri, G., Li, Q., Pan, T.: Pet image denoising based on denoising diffusion probabilistic model. *European Journal of Nuclear Medicine and Molecular Imaging* pp. 1–11 (2023)
7. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16000–16009 (2022)

8. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
10. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications* (2021)
11. Kazerouni, A., Aghdam, E.K., Heidari, M., Azad, R., Fayyaz, M., Hacıhaliloglu, I., Merhof, D.: Diffusion models in medical imaging: A comprehensive survey. *Medical Image Analysis* p. 102846 (2023)
12. Khader, F., Müller-Franzes, G., Tayebi Arasteh, S., Han, T., Haarbuerger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baekler, B., Foersch, S., et al.: Denoising diffusion probabilistic models for 3d medical image generation. *Scientific Reports* **13**(1), 7303 (2023)
13. Konz, N., Mazurowski, M.A.: Reverse engineering breast mris: Predicting acquisition parameters directly from images. In: *Medical Imaging with Deep Learning* (2023)
14. Lew, C.O., Harouni, M., Kirksey, E.R., Kang, E.J., Dong, H., Gu, H., Grimm, L.J., Walsh, R., Lowell, D.A., Mazurowski, M.A.: A publicly available deep learning model and dataset for segmentation of breast, fibroglandular tissue, and vessels in breast mri. *Scientific reports* **14** **1**, 5383 (2024), <https://api.semanticscholar.org/CorpusID:268251677>
15. Lyu, Q., Wang, G.: Conversion between ct and mri images using diffusion and score-matching models. *arXiv preprint arXiv:2209.12104* (2022)
16. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021)
17. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019)
18. Pinaya, W.H., Graham, M.S., Gray, R., Da Costa, P.F., Tudosiu, P.D., Wright, P., Mah, Y.H., MacKinnon, A.D., Teo, J.T., Jager, R., et al.: Fast unsupervised brain anomaly detection and segmentation with diffusion models. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 705–714. Springer (2022)
19. Pinaya, W.H., Tudosiu, P.D., Dafflon, J., Da Costa, P.F., Fernandez, V., Nachev, P., Ourselin, S., Cardoso, M.J.: Brain imaging generation with latent diffusion models. In: *MICCAI Workshop on Deep Generative Models*. pp. 117–126. Springer (2022)
20. Rister, B., Yi, D., Shivakumar, K., Nobashi, T., Rubin, D.L.: Ct-org, a new dataset for multiple organ segmentation in computed tomography. *Scientific Data* **7**(1), 381 (2020)
21. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10684–10695 (June 2022)
22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* **18**. pp. 234–241. Springer (2015)

23. Saha, A., Harowicz, M.R., Grimm, L.J., Kim, C.E., Ghate, S.V., Walsh, R., Mazurowski, M.A.: A machine learning approach to radiogenomics of breast cancer: a study of 922 subjects and 529 dce-mri features. *British journal of cancer* **119**(4), 508–516 (2018)
24. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=St1giarCHLP>
25. Wang, T., Zhang, T., Zhang, B., Ouyang, H., Chen, D., Chen, Q., Wen, F.: Pretraining is all you need for image-to-image translation. *arXiv preprint arXiv:2205.12952* (2022)
26. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: International Conference on Medical image computing and computer-assisted intervention. pp. 35–45. Springer (2022)
27. Wolleb, J., Sandkühler, R., Bieder, F., Valmaggia, P., Cattin, P.C.: Diffusion models for implicit image segmentation ensembles. In: International Conference on Medical Imaging with Deep Learning. pp. 1336–1348. PMLR (2022)
28. Yang, J., Dvornek, N.C., Zhang, F., Chapiro, J., Lin, M., Duncan, J.S.: Unsupervised domain adaptation via disentangled representations: Application to cross-modality liver segmentation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22. pp. 255–263. Springer (2019)
29. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)