

Stable Diffusion Segmentation for Biomedical Images with Single-step Reverse Process

Tianyu Lin¹, Zhiguang Chen², Zhonghao Yan³,
Weijiang Yu² ✉, and Fudan Zheng² ✉

¹ School of Biomedical Engineering, Sun Yat-sen University, China

² School of Computer Science and Engineering, Sun Yat-sen University, China
weijiangyu8@gmail.com, zhengfd3@mail2.sysu.edu.cn

³ International School, Beijing University of Posts and Telecommunications, China

Abstract. Diffusion models have demonstrated their effectiveness across various generative tasks. However, when applied to medical image segmentation, these models encounter several challenges, including significant resource and time requirements. They also necessitate a multi-step reverse process and multiple samples to produce reliable predictions. To address these challenges, we introduce the first latent diffusion segmentation model, named SDSeg, built upon stable diffusion (SD). SDSeg incorporates a straightforward latent estimation strategy to facilitate a single-step reverse process and utilizes latent fusion concatenation to remove the necessity for multiple samples. Extensive experiments indicate that SDSeg surpasses existing state-of-the-art methods on five benchmark datasets featuring diverse imaging modalities. Remarkably, SDSeg is capable of generating stable predictions with a solitary reverse step and sample, epitomizing the model’s stability as implied by its name. The code is available at <https://github.com/lin-tianyu/Stable-Diffusion-Seg>.

Keywords: Biomedical Image Segmentation · Latent Diffusion Model · Stable Diffusion · Reverse Process

1 Introduction

Image segmentation is a crucial task in medical image analysis. To alleviate the workload on medical professionals, numerous automated algorithms for medical image segmentation have been developed. The effectiveness of various neural network architectures, such as Convolutional Neural Networks (CNNs)[19,5] and Vision Transformers (ViT)[22,9], has underscored deep learning as a promising approach to medical image segmentation.

The recent interest in Diffusion Probabilistic Models (DPM)[10,15] among researchers has led to a focus on image-level diffusion models in DPM-based segmentation methods[25,24,26,1]. Image-level diffusion models introduce noise to an image through *forward process*, and generate new images by learning how to decode this noise addition step by step in *reverse process*. DPM-based segmentation methods utilize image conditioning to generate segmentation predictions.

However, these approaches face limitations: (1) generating segmentation maps in pixel space is unnecessary and may lead to inefficient optimization and high computational costs since binary semantic maps have sparse semantic information compared to ordinary images; (2) diffusion models usually require multiple reverse steps to achieve detailed and varied generations, with prior diffusion segmentation models needing several samples to average for stable predictions.

To overcome these challenges, we propose a simple yet efficient segmentation framework called SDSeg, with the following contributions:

- SDSeg is built on Stable Diffusion (SD)[18], a latent diffusion model (LDM)[18,17] that conducts diffusion process on a perceptually equivalent latent space with lower resolution, making the diffusion process computationally friendly.
- A simple latent estimation loss is introduced to empower SDSeg to generate segmentation results on a single-step reverse process, and a concatenate latent fusion technique is proposed to eliminate the need for multiple samples.
- The conditioning vision encoder is set trainable to learn images’ features for segmentation and adapt to multiple medical imaging domains.
- SDSeg performs state-of-the-art on five benchmark datasets and significantly improves diffusion-based segmentation models by reducing training resources, increasing inference speed, and enhancing generation stability.




Fig. 1. The overview of SDSeg. We condition SDSeg via concatenation. In the training stage, we only train the denoising U-Net and vision encoder.

2 Methods

The framework of SDSeg is shown in Figure. 1. For medical images, we introduce a trainable vision encoder τ_θ to encode an image $C \in \mathbb{R}^{H \times W \times 3}$ to its latent representation $z_c = \tau_\theta(C)$. For segmentation maps, we utilize an autoencoder for perceptual compression. As Figure. 1 shows, given a segmentation map $X \in \mathbb{R}^{H \times W \times 3}$ in pixel space, the encoder $\mathcal{E}$ encodes X into a latent representation $z = \mathcal{E}(X)$, and the decoder $\mathcal{D}$ recovers the segmentation map from the latent, giving reconstructions $\hat{X} = \mathcal{D}(z) = \mathcal{D}(\mathcal{E}(X))$, where $z \in \mathbb{R}^{h \times w \times c}$. In practice,

we notice that the autoencoder provided by SD performs well enough for binary segmentation maps, as shown in Figure. 2. Thus, we keep the autoencoder frozen in the training stage, which makes SDSeg an end-to-end method. The diffusion process of SDSeg is conducted on the latent space.

2.1 Latent Estimation

For the training stage, the latent of segmentation map in the first timestep z_0 is added with t time steps of Gaussian noise to get z_t . The forward process of the diffusion can be represented as:

$$z_t = \sqrt{\alpha_t}z_0 + \sqrt{1 - \alpha_t}n \quad (1)$$

where n is random Gaussian noise, and $\bar{\alpha}$ is a hyperparameter for controlling the forward process. The goal of denoising U-Net in every training step is to estimate the distribution of the random Gaussian noise n , formulated as $\tilde{n} = f(z_t; z_c)$, where $f(\cdot)$ denotes the denoising U-Net. The noise prediction loss can be represented as $\mathcal{L}_{noise} = \mathcal{L}(\tilde{n}, n)$.

In tasks aimed at generating varied and semantically rich images, the gradual application of noise estimation in the reverse process can refine the outcomes progressively. However, we believe that the inherently simpler segmentation maps do not substantially benefit from an extensive reverse process. Instead, a proficiently trained denoising U-Net is capable of restoring the latent features containing all necessary structural and spatial characteristics for a segmentation map. Therefore, after obtaining the estimated noise $\tilde{n}$, we can straightforwardly derive the corresponding latent estimation through a simple transformation of Eq. 1:

$$\tilde{z}_0 = \frac{1}{\sqrt{\alpha_t}}(z_t - \sqrt{1 - \alpha_t}\tilde{n}) \quad (2)$$

This technique facilitates the addition of a supervision branch by setting the optimization goal to minimize the difference between the predicted $\tilde{z}_0$ and the true z_0 , with the latent loss function defined as $\mathcal{L}_{latent} = \mathcal{L}(\tilde{z}_0, z_0)$. Thus, the final loss function can be expressed as:

$$\mathcal{L} = \mathcal{L}_{noise} + \lambda\mathcal{L}_{latent} \quad (3)$$

where λ denotes the weight of the latent loss function. In practice, λ is set to 1, and the $\mathcal{L}_{noise}$ and $\mathcal{L}_{latent}$ are mean absolute error.

It is noteworthy that utilizing $\mathcal{L}_{noise}$ along with multiple iterations of DDIM[21] sampling can generate impressive segmentation results. The greatest contribution of introducing $\mathcal{L}_{latent}$ lies in its ability to bypass the unnecessary reverse processes, thereby notably enhancing speed during the inference phase.

2.2 Concatenate Latent Fusion

Stable Diffusion incorporates a cross-attention mechanism to facilitate multi-modal training and generation. Nonetheless, for an image-to-image segmentation




Fig. 2. Visualization of reconstructions and latent representations on BTCV, STS, REF, and CVC. Reconstructions denotes $\tilde{X} = \mathcal{D}(z)$ where latent $z = \mathcal{E}(X)$.

model, prioritizing the extraction of semantic features and structural information from images is essential, whereas multi-modal capabilities might not offer additional advantages. Furthermore, adding cross-attention across several blocks of the denoising U-Net incurs additional computational costs. Thus, it becomes imperative to explore a more efficient method for latent fusion within SDSeg.

Moreover, our observation in Figure. 2 reveals that the segmentation maps exhibit a pronounced spatial correlation with their corresponding latent representations, which might contain the necessary structural and feature information that can benefit segmentation tasks. Consequently, inspired by conventional semantic segmentation methods such as U-Net[19] and DeepLabV3+[5], etc., we employ concatenation, the prevalent and validated effective strategy for integrating an image’s semantic features, to merge the latent representations of segmentation maps with those of image slices.

2.3 Trainable Vision Encoder

In semantic segmentation, a valid vision encoder can extract the necessary structural and semantic features from images, thereby enhancing segmentation results. As an image-conditioned generative model, SDSeg employs a trainable vision encoder to capture the abundant semantic features across images.

The vision encoder τ_θ has the same architecture as the encoder $\mathcal{E}$ and is initialized with its pre-trained weights. Although we discover that simply using a frozen image encoder that is pre-trained on natural images can bring considerable results, we make the vision encoder trainable, thus allowing SDSeg to

adjust to various medical image dataset modalities, enhancing its versatility and effectiveness.

3 Experimental Results

3.1 Datasets and Evaluation Metrics

To comprehensively evaluate the effectiveness and generalization ability of SDSeg, we conduct experiments on three RGB datasets on 2D segmentation task, and two CT datasets on 3D segmentation, as shown in Table. 1.

Table 1. Dataset settings.

Task	Dataset	Target	Training Data	Test Data
2D Binary Segmentation	CVC-ClinicDB[2] (CVC)	Polyp	488 images	62 images
	Kvasir-SEG[12] (KSEG)	Polyp	800 images	100 images
	REFUGE2[14,16] (REF)	Optic Cup	800 images	400 images
3D Binary Segmentation	BTCV ¹ [13]	Abdomen Organ	18 volumes	12 volumes
	STS-3D[7,8] (STS)	Teeth	9 volumes	3 volumes

Our evaluation encompasses three main aspects: Firstly, segmentation results across datasets are assessed using the Dice Coefficient (DC) and Intersection over Union (IoU) metrics. Secondly, we benchmark our model’s efficiency by comparing its computational resource usage and inference speed against other diffusion-based segmentation methods. Thirdly, we evaluate the stability of our generated segmentation results against other diffusion segmentation models using LPIPS[27], PSNR, SSIM, and MS-SSIM. Additionally, we conduct an ablation study to validate the efficacy of our proposed modules.

3.2 Implementation Details

Experimental Settings SDSeg is trained on a single V100 GPU with 16GB RAM. The model is trained for 100,000 steps using AdamW optimizer with a base learning rate of 1×10^{-5} . The batch size is set to 4 by default. We use a KL-regularized autoencoder and LDM model with the downsampling rate $r = \frac{H}{h} = \frac{W}{w} = 8$. SDSeg takes RGB images² as pixel space inputs with $H = W = 256$, and the corresponding latent representation has a shape of $h = w = 32$ with $c = 4$. All model parts are initialized with the pre-trained weights provided by stable diffusion. The additional model parameters of the denoising U-Net for concatenate input are initialized to zeros.

¹ We treat all 13 organs in BTCV as a single target.

² For 1-channel CT slices, we simply repeat 3 times to get 3-channel images.

Inference Stage During the inference stage, we concatenate randomly generated Gaussian noise with the medical image’s latent representation. The denoising U-Net then predicts the estimated noise, allowing SDSeg to derive the latent estimation using Eq. 2. Then, Decoder $\mathcal{D}$ will transfer latent estimation to pixel space to get the final prediction. As shown in Table. 4, SDSeg doesn’t need an external sampler and only needs a single-step reverse to sample one time for a stable prediction.

3.3 Main Results

Comparison with State-of-the-Arts The comparison of our model with several semantic segmentation methods on REF, BTCV, and STS datasets is shown in Table. 2. We also compare our model with state-of-the-art diffusion based segmentation models on CVC, KSEG, and REF datasets, as shown in Table. 3. SDSeg outperforms all other models on the five datasets with various imaging modalities, validating its effectiveness and generalization capability.

Table 2. Comparison with semantic segmentation methods, evaluated by the Dice coefficient metric.

Methods	REF	BTCV	STS
U-Net[19]	80.1	75.9	85.4
U-Net(w/ R50) ³	87.2	90.5	88.4
Swin-UNETR[22]	-	91.3	88.3
nnU-Net[11]	-	91.4	88.9
TransU-Net[4]	85.6	89.1	88.1
SwinU-Net[3]	84.3	86.5	85.8
Ours	89.4	92.8	89.4

Table 3. Comparison with state-of-the-art methods on REF, CVC, and KSEG.

Dataset	Methods	Dice	IoU
CVC	SSFormer[23]	94.4	89.9
	Li-SegPNet[20]	92.5	86.0
	Diff-Trans[6]	95.4	92.0
	Ours	95.8	92.6
KSEG	SSFormer[23]	93.5	89.0
	Li-SegPNet[20]	90.5	82.8
	Diff-Trans[6]	94.6	91.6
	Ours	94.9	92.1
REF	MedSegDiff- V1[25]	86.3	78.2
	MedSegDiff-V2[24]	85.9	79.6
	Diff-Trans[6]	88.7	81.5
	Ours	89.4	81.8

Comparison of computing resource and time efficiency Table. 4 demonstrates the efficiency evaluation results of MedSegDiffs, Diff-U-Net, and SDSeg on BTCV dataset. For a fair comparison, these models are trained on the same server using their source codes. The results highlight SDSeg’s superior efficiency, requiring significantly fewer resources and less time for training. Remarkably, SDSeg’s inference process is about 100 times faster than that of MedSegDiffs and approximately 28 times quicker in generating a single segmentation map.

Table. 4 also compares the reverse process of these models. The latent estimation scheme empowers SDSeg to generate segmentation maps in a single step,

³ Replace the encoder of U-Net to a pre-trained ResNet50.

and the concatenate latent fusion module allows SDDSeg to sample only one time without harming model performance. Moreover, latent estimation makes SDDSeg no longer rely on any external sampler for sampling.

Table 4. Comparison of training resources, inference speed, and reverse process settings on BTCV (1568 slices). The reverse process is assessed by *steps* $\times$ *samples*.

Methods	Training Time (hours)	Training Resources ($\times$ GPUs)	Inference Time (hours)	Inference Speed (samples/s)	Diffusion Sampler	Reverse Process	Dice
MedSegDiff-V1	≈ 48	16GB $\times 4$	≈ 7	0.30	DPM-Solver	50 $\times$ 25	79.24
MedSegDiff-V2	≈ 49	16GB $\times 4$	≈ 7	0.31	DPM-Solver	50 $\times$ 25	83.52
Diff-U-Net[26] ⁴	≈ 16	24 GB $\times 4$	$\approx 1/2$	0.87	DDIM	10 $\times$ 1	91.89
Ours	≈ 12	16GB $\times 1$	$\approx 1/4$	2.01	DDIM	10 $\times$ 1	92.09
Ours	$\approx$ 12	16GB $\times 1$	$\approx$ 1/13	8.36	X	1$\times$1	92.76

Stability Evaluation Since diffusion models are generative models, the samples they generate can exhibit variability. However, diversity is not considered an advantageous trait in the context of medical segmentation models, as medical professionals necessitate the assistance of artificial intelligence to be consistent and reliable. Given a trained model and fixed test data, we evaluate the stability of the diffusion-based segmentation models on the following two tasks:

1. **Dataset-level Stability:** performs repeated inferences on test data to measure variability across different inferences using the LIPIS[27] metric;
2. **Instance-level Stability:** examines the model’s consistency under varying initial noise by conducting repeated inferences under fixed conditions, with PSNR, SSIM, and MS-SSIM as metrics.

Table. 5 showcases SDDSeg’s significant stability across these tests, underscoring its reliability in segmentation despite different initial noises.

Table 5. Comparison of stability evaluation on BTCV. ‘Seg’ denotes segmentation maps; ‘Score’ represents predicted probability scores.

Methods	LPIPS $\downarrow$		PSNR $\uparrow$		SSIM $\uparrow$		MS-SSIM $\uparrow$	
	Seg	Score	Seg	Score	Seg	Score	Seg	Score
MedSegDiff-V2	0.3139	0.2904	11.9271	14.4506	0.5780	0.4662	0.6399	0.6228
Diff-U-Net	0.0633	0.0672	23.7158	24.6675	0.9668	0.9666	0.9442	0.9397
Ours	0.0199	0.0143	27.6348	31.5537	0.9796	0.9764	0.9897	0.9909

⁴ Diff-U-Net uses 3D sliding window infer. Inference speed is estimated as $\frac{\text{slices}}{\text{time}}$.

Ablation Study Our ablation studies assess the contribution of each component within SDSeg, as detailed in Table. 6. The baseline model relies on stable diffusion with cross-attention for generating image-conditioned segmentation maps. The incorporation of latent fusion concatenation notably enhances performance, allowing for efficient learning of spatial information and features. Additionally, the trainable encoder markedly improves performances by extracting relevant semantic features from segmentation targets. While the latent estimation loss function marginally boosts performance, its primary advantage lies in significantly accelerating the reverse process, thus enabling SDSeg to discard traditional samplers for a single-step reverse process, as illustrated in Figure. 3.

Table 6. Ablation study on BTCV and REF.

Latent Estimation	Concatenate Latent Fusion	Trainable Image Encoder	BTCV		REFUGE2	
			Dice	IoU	Dice	IoU
X	X	X	32.67	23.69	28.31	20.36
X	✓	X	80.31	72.27	76.79	69.37
X	✓	✓	91.89	85.41	88.79	80.29
✓	✓	✓	92.76	85.49	89.36	81.68




Fig. 3. Comparison of DDIM convergence speed with and without latent estimation loss on BTCV. $\lambda = 1$ denotes that SDSeg is trained on latent estimation loss $\mathcal{L}_{latent}$.

4 Conclusion

In this paper, we propose SDSeg, a novel and efficient framework for medical image segmentation utilizing stable diffusion. We introduce a latent estimation strategy enabling single-step latent prediction, thereby eliminating the need for

a multi-step reverse process. The model employs concatenate latent fusion for integrating learned image latent that effectively guides the segmentation task. Furthermore, a trainable vision encoder enhances the model’s capability to learn image features and adapt to diverse image modalities. SDSeg achieves state-of-the-art performance across five segmentation datasets, substantially reducing training resource requirements and accelerating the inference process while maintaining remarkable stability.

Acknowledgments. This study was funded by the Program of Science and Technology of Guangdong (Grant No. 2020B1111170009), the 2022 Industrial Technology Basic Public Service Platform Project of China (Grant No. 2022-228-219), and the Fundamental Research Funds for the Central Universities, Sun Yat-sen University (Grant No. 23xkjc016).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Baranchuk, D., Voynov, A., Rubachev, I., Khrulkov, V., Babenko, A.: Label-efficient semantic segmentation with diffusion models. In: International Conference on Learning Representations (2022)
2. Bernal, J., et al.: Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized Medical Imaging and Graphics* **43**, 99–111 (2015). <https://doi.org/10.1016/j.compmedimag.2015.02.007>
3. Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation. In: European conference on computer vision. pp. 205–218. Springer (2022)
4. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306 (2021)
5. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Computer Vision – ECCV 2018. pp. 833–851. Lecture Notes in Computer Science (2018). https://doi.org/10.1007/978-3-030-01234-2_49
6. Chowdary, G.J., Yin, Z.: Diffusion transformer u-net for medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 622–631. Springer (2023)
7. Cui, W., et al.: Ctooth: A fully annotated 3d dataset and benchmark for tooth volume segmentation on cone beam computed tomography images. In: Intelligent Robotics and Applications. pp. 191–200. Lecture Notes in Computer Science (2022). https://doi.org/10.1007/978-3-031-13841-6_18
8. Cui, W., et al.: Ctooth+: A large-scale dental cone beam computed tomography dataset and benchmark for tooth volume segmentation. In: Data Augmentation, Labelling, and Imperfections. pp. 64–73. Lecture Notes in Computer Science (2022). https://doi.org/10.1007/978-3-031-17027-0_7
9. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In:

- Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 574–584 (2022)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems. vol. 33, pp. 6840–6851 (2020)
 11. Isensee, F., et al.: nnu-net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021). <https://doi.org/10.1038/s41592-020-01008-z>
 12. Jha, D., Smedsrud, P.H., Riegler, M.A., Halvorsen, P., de Lange, T., Johansen, D., Johansen, H.D.: Kvasir-seg: A segmented polyp dataset. In: International Conference on Multimedia Modeling. pp. 451–462. Springer (2020)
 13. Landman, B., et al.: Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge. In: Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge. vol. 5, p. 12 (2015)
 14. Li, F., et al.: Development and clinical deployment of a smartphone-based visual field deep learning system for glaucoma detection. *npj Digital Medicine* **3**(1), 1–8 (2020). <https://doi.org/10.1038/s41746-020-00329-9>
 15. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021)
 16. Orlando, J.I., et al.: Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical Image Analysis* **59**, 101570 (2020). <https://doi.org/10.1016/j.media.2019.101570>
 17. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4195–4205 (2023)
 18. 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. pp. 10684–10695 (2022)
 19. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. pp. 234–241. Lecture Notes in Computer Science (2015). https://doi.org/10.1007/978-3-319-24574-4_28
 20. Sharma, P., Gautam, A., Maji, P., Pachori, R.B., Balabantaray, B.K.: Li-segpnnet: Encoder-decoder mode lightweight segmentation network for colorectal polyps analysis. *IEEE Transactions on Biomedical Engineering* **70**(4), 1330–1339 (2023). <https://doi.org/10.1109/TBME.2022.3216269>
 21. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
 22. Tang, Y., Yang, D., Li, W., Roth, H.R., Landman, B., Xu, D., Nath, V., Hatamizadeh, A.: Self-supervised pre-training of swin transformers for 3d medical image analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20730–20740 (2022)
 23. Wang, J., et al.: Stepwise feature fusion: Local guides global. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. pp. 110–120. Lecture Notes in Computer Science (2022). https://doi.org/10.1007/978-3-031-16437-8_11
 24. Wu, J., Fu, R., Fang, H., Zhang, Y., Xu, Y.: Medsegdiff-v2: Diffusion based medical image segmentation with transformer. arXiv preprint arXiv:2301.11798 (2023)
 25. Wu, J., Fu, R., Fang, H., Zhang, Y., Yang, Y., Xiong, H., Liu, H., Xu, Y.: Medsegdiff: Medical image segmentation with diffusion probabilistic model. In: Medical Imaging with Deep Learning. pp. 1623–1639. PMLR (2024)

26. Xing, Z., Wan, L., Fu, H., Yang, G., Zhu, L.: Diff-unet: A diffusion embedded network for volumetric segmentation. arXiv preprint arXiv:2303.10326 (2023)
27. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)