ACCELERATING DIFFUSION MODELS VIA PRE-SEGMENTATION DIFFUSION SAMPLING FOR MEDICAL IMAGE SEGMENTATION

Xutao Guo[1,3], Yanwu Yang[1,3], Chenfei Ye[2], Shang Lu [1], Yang Xiang[3,], Ting Ma[1,3,*]*

School of Electronics and Information Engineering, Harbin Institute of Technology at Shenzhen, Shenzhen, China
International Research Institute for Artifcial Intelligence, Harbin Institute of Technology at Shenzhen, Shenzhen, China
Peng Cheng Laboratory, Shenzhen, China

ABSTRACT

Based on the Denoising Diffusion Probabilistic Model (DDPM), medical image segmentation can be described as a conditional image generation task, which allows to compute pixel-wise uncertainty maps of the segmentation and allows an implicit ensemble of segmentations to boost the segmentation performance. However, DDPM requires many iterative denoising steps to generate segmentations from Gaussian noise, resulting in extremely inefficient inference. To mitigate the issue, we propose a principled acceleration strategy, called pre-segmentation diffusion sampling DDPM (PD-DDPM), which is specially used for medical image segmentation. The key idea is to obtain pre-segmentation results based on a separately trained segmentation network, and construct noise predictions (non-Gaussian distribution) according to the forward diffusion rule. We can then start with noisy predictions and use fewer reverse steps to generate segmentation results. Experiments show that PD-DDPM yields better segmentation results over representative baseline methods even if the number of reverse steps is significantly reduced. Moreover, PD-DDPM is orthogonal to existing advanced segmentation models, which can be combined to further improve the segmentation performance.

Index Terms— Diffusion models, medical image segmentation, uncertainty, ensemble

1. INTRODUCTION

Denoising Diffusion Probabilistic Models (DDPM) [1, 2] form a category of deep generative models which has recently become one of the hottest topic in computer vision, due to their promising results in both unconditional and conditional generation tasks [2–4]. Based on the Denoising Diffusion Probabilistic Model (DDPM), medical image segmentation can be described as a conditional image generation task, which allows to compute pixel-wise uncertainty maps of the segmentation and allows an implicit ensemble of segmentations to boost the segmentation performance [5]. Especially in medical applications where subsequent diagnosis or treatment relies on segmentation, algorithms that provide only the most likely hypotheses can lead to misdiagnosis and suboptimal treatment. If multiple consistent hypotheses are provided, they can be used to suggest further diagnostic tests to resolve ambiguity, or experts with access to additional information can select appropriate hypotheses for subsequent steps [5, 6].

The Denoising Diffusion Probabilistic Model (DDPM) consists of two Markov chains. In the forward diffusion process, the clean images are gradually disturbed by Gaussian noise until they are approximated to Gaussian distribution [1]. In the process of inverse diffusion, from the sampled Gaussian noise, the trained denoising deep neural networks is used to iteratively denoise the data to obtain the clean images. Therefore, synthesizing samples from DDPM is achieved by iteratively denoising the sampled Gaussian noise. For the medical image segmentation, we can train the DDPM on the ground truth segmentation, and use the image as a prior during training and in every step during the sampling process [5]. However, a major problem with vanilla DDPM is the inefficiency of inference. Because obtaining clean segmentations from DDPM typically requires hundreds or even thousands of denoising steps, each of which involves forward prediction by the denoising neural network.

To accelerate DDPMs, a few methods have been proposed. Song et al. [7] attempted to reduce the number of diffusion steps by using non-Markovian reverse processes. Unlike DDPM, which defines noise scale as a constant, San-Roman et al. [8] present a adaptive noise scheduling to estimate the noise parameters given the current input at inference time, which requiring less steps. In [9], the authors propose to distill the full sampling process into a faster sampler that requires only half as many steps. Some methods [10, 11] shifted the diffusion process to the latent space using pretrained autoencoders. However, the above methods cannot achieve significant acceleration without sacrificing the quality

---
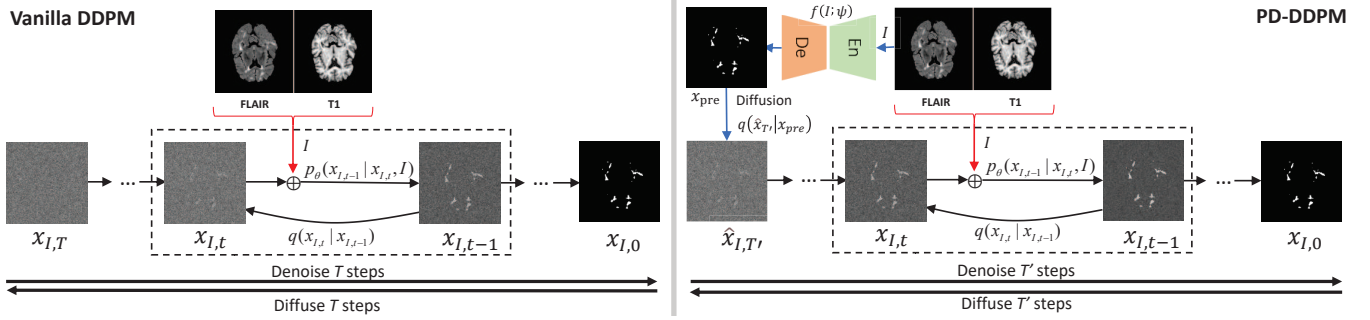
Fig. 1: Compare PD-DDPM with the vanilla DDPM. The training and sampling procedure of the method. In every step t, the conditional information is induced by concatenating the medical images I to the noisy segmentation mask $x_{I,t}$.

of power generation. Some other methods [12, 13] improve sampling efficiency by truncating the forward and reverse diffusion processes, boosting the performance at the same time. But this method needs to combine GAN [14] or VAE [15] models that are difficult to train. And it will break the characteristic that the image dimension remains unchanged during diffusion and sampling in the vanilla DDPM. To sum up, all of the above methods do not implement accelerated sampling specifically for segmentation tasks.

To mitigate the issue, we propose a principled acceleration strategy, called pre-segmentation diffusion sampling DDPM (PD-DDPM), which is specially used for medical image segmentation. The key idea is to obtain pre-segmentation results based on a separately trained segmentation network, and construct noise predictions (non-Gaussian distribution) according to the forward diffusion rule. We can then start with noisy predictions and use fewer reverse steps to generate segmentation results. PD-DDPM not only improves the efficiency of vanilla DDPM without breaking any assumptions, but as an additional benefit, improves the segmentation performance of vanilla DDPM. Experiments show that PD-DDPM yields better segmentation results over representative baseline methods even if the number of reverse steps is significantly reduced. Moreover, PD-DDPM is orthogonal to existing advanced segmentation models, which can be combined to further improve the segmentation performance. Since both pre-segmentation and diffusion operations can be easily implemented in one step, the pre-trained segmentation network brings only minor computational overhead to PD-DDPM.

2. METHOD

2.1. Background on DDPMs

In DDPM [2], the forward diffusion process is a first-order Markov chain that perturbs the data distribution $q(x^0)$ by gradually adding Gaussian noise, with variance $\beta_t \in (0,1)$ at time t, until the data distribution converges to a standard Gaussian distribution. The form of the forward process can

be summarized as follows:

$$
\begin{aligned}
q(x^{1:T}|x^0) &= \prod_{t=1}^{T} q(x^t|x^{t-1}), \\
q(x^t|x^{t-1}) &= \mathcal{N}(x^t; \sqrt{1-\beta_t}x^{t-1}, \beta_t \boldsymbol{I})
\end{aligned}
\tag{1}
$$

where $x^{1:T}$ denotes the set of variables $x^1, x^2, ..., x^T$. $T = 1000 \frown 4000$ is a typical choice for most works. With the limit of small diffusion rate (i.e., β_t is kept sufficiently small), the reverse distribution $q(x^{t-1}|x^t)$ also follows a Gaussian distribution. And thus the reverse process can be approximated using a neural network parameterized Gaussian distribution p_θ, starting at $p(x^T) = \mathcal{N}(x^T; 0, I)$:

$$
\begin{aligned}
p_\theta(x^{0:(T-1)}|x^T) &= \prod_{t=1}^{T} p_\theta(x^{t-1}|x^t), \\
p_\theta(x^{t-1}|x^t) &= \mathcal{N}(x^{t-1}; \mu_\theta(x^t, t), \sigma_t^2 \boldsymbol{I})
\end{aligned}
\tag{2}
$$

The neural network is trained to simulate the reverse process of the diffusion process defined in Equation 1. To generate an image from the reverse process, we first sample x^T from the underlying data distribution by sampling a latent (of the same size as the training data point x^0) from $p(x^T)$ (chosen to be an isotropic Gaussian distribution), and then sequentially draws sample x^{t-1} from $p_\theta(x^{t-1}|x^t)$ for $t = T, T-1, ..., 1$ until we get a new data x^0. The generation process of a DDPM is extremely slow as it need to sample from the transition distribution $p_\theta(x^{t-1}|x^t)$ iteratively, which involves a lot of evaluations of the output of the neural network.

2.2. Pre-segmentation Diffusion Sampling Denoising Diffusion Probabilistic Model

One natural question raised from the vanilla DDPM could be: can we cut the reverse process to T' ($< T$) steps? If we can get the noise samples $x^{T'}$ in advance, it can be achieved. To generate image-specific segmentations, we train DDPM on the ground truth and use images as priors during training

and sampling. To generate image-specific segmentation results, we train DDPM on the ground truth segmentation and use medical images as conditional information during training and sampling process. Therefore, we can train a separate segmentation network f_ψ using medical image conditional information.

$$x_{pre} = f(I; \psi) \tag{3}$$

In the sampling process, we can first obtain the pre-segmentation result through the pre-segmentation network f_ψ as shown in Figure 1. Then, according to Formula 1, the pre-segmentation result is diffused to T' step to obtain the approximate sample $\hat{x}^{T'}$ of $x^{T'}$. In this case the reverse process can be newly defined as:

$$p_\theta(x^{0:(T'-1)}|\hat{x}^{T'}) = \prod_{t=1}^{T'-1} p_\theta(x^{t-1}|x^t) p_\theta(x^{T'-1}|\hat{x}^{T'}),$$
$$p_\theta(x^{t-1}|x^t) = \mathcal{N}(x^{t-1}; \mu_\theta(x^t, t), \sigma_t^2 \boldsymbol{I}) \tag{4}$$

Thus, we can use the denoising neural network to denoise the non Gaussian distribution $\hat{x}^{T'}$ to a clean segmentation x^0 in fewer steps than the vanilla sampling process according to Formula 4. Because the pre-segmentation result is not the real ground truth segmentation, there will be some errors between $\hat{x}^{T'}$ and $x^{T'}$. However, we empirically found through experiments that PD-DDPM boosts the segmentation accuracy of vanilla DDPM. We name the pre-segmentation diffusion sampling DDPM as PD-DDPM. It should be emphasized that PD-DDPM does not break the Gaussian assumption of vanilla DDPM' denoising process. In this way, the number of denoising steps required is greatly reduced, thus effectively alleviating the pain point of high inference cost of DDPM.

3. EXPERIMENTS AND RESULTS

3.1. Datasets

We evaluate our method on the WMH dataset provided by the White Matter Hyperintensities segmentation challenge in MICCAI 2017 [16]. It consists of 60 cases of brain MRI images (3D T1-weighted image and 2D multi-slice FLAIR image) with manual annotations of white matter hyperintensity (binary masks) from three different institutes/scanners. And the manual reference standard is defined on the FLAIR image. So a 2D multi-slice version of the T1 image was generated by re-sampling the 3D T1-weighted image to match with the FLAIR. In this paper, all cases are randomly assigned into five folds. Then we randomly assign these five-folds into a training set (3-fold), a validation set (1-fold), and a test set (1-fold).

3.2. Implementation Details

In this paper, all the networks train using Pytorch using NVIDIA TESLA V-100 (Pascal) GPUs with 32 GB memory.

Fig. 2: Visualization of segmentation and uncertainty maps.

**Table 1**: Segmentation scores with different methods.

Methods	Dice	HD95	Jaccard	F1
U-Net [17]	0.787	3.935	0.656	0.738
AttUnet [18]	0.799	4.190	0.673	0.753
U-Net++ [19]	0.798	3.915	0.672	0.751
Bayesian U-Net [20]	0.798	3.803	0.672	0.777
Probabilistic U-Net [21]	0.792	3.777	0.663	0.763
Vanilla DDPM [5]	0.796	4.179	0.669	0.779
TDPM [12]	0.801	3.608	0.676	0.776
ES-DDPM [13]	0.803	3.523	0.678	0.777
PD-DDPM	**0.812**	**3.494**	**0.689**	**0.800**

We optimized all configurations with the Adam optimizer with the learning rate 1e-4 and the weight decay 1e-5. The batch size is set to 12. For the WMH, images and annotation labels were randomly cropped to 128 × 192 patches. We choose a cosine noise schedule for $T = 1000$ steps. The number of channels in the first layer of denoising network is chosen as 128, and we use one attention head at resolution 16. In PD-DDPM, we choose the AttUnet [18] as the pre-segmentation model.

3.3. Comparison of Segmentation Performance

The inference process of DDPM is a stochastic process. So we can implicitly ensemble of segmentation masks to boost the segmentation performance. For every image of the test set, we sample 5 different segmentation masks. Then 5 different segmentation masks are ensembled by averaging and is thresholded at 0.5 to obtain a binary segmentation. In Table 1, the Dice score, the Jaccard index the 95 percentile Hausdorff Distance (HD95), and F1 are presented.

We conduct quantitative experiments to compare our method with a range of representative methods. Here, the

Fig. 3: The Dice and uncertainty on testing set with respect to T'. The horizontal axis represents the size of T'.

Fig. 4: The average and standard deviation of the dice on testing set with respect to ensemble size (when T'=300).

U-Net [17], AttUnet [18] and, U-Net++ [19] are the most representative deep learning model in the segmentation field, but none of them can estimate the uncertainty of segmentation. Bayesian U-Net [20] and Probabilistic U-Net [21] are representative methods that can estimate the uncertainty of segmentation. And we also compare PD-DDPM with other accelerating DDPMs, including TDPM [12] and ES-DDPM [13]. It should be emphasized that the size of ensemble in the comparison methods is also set to 5.

Table 1 shows PD-DDPM (when T'=300) achieve the best results with respect to all four metrics. And PD-DDPM outperforming the vanilla DDPM. For visualization of the segmentation and uncertainty maps, we select three images I_1, I_2, and I_3 from the test set in Figure 2.

3.4. Determining Optimal T'

Here, we analyzed the impact of the hyperparameter T' on PD-DDPM. By varying the T' among $\{50, 100, 200, 300, 400, 500, 600, 700, 800, 1000\}$, we train PD-DDPM for WMH segmentation. As shown in Figure 3, PD-DDPM achieves the best Dice score when $T' = 300$. And we also analyze the effect of T' on uncertainty estimation (ie, the variation of predicting softmax output). Figure 3 shows the uncertainty increases with T' when $T' < 500$. Then the T' was further increased, the uncertainty tended to saturate.

3.5. Effect of the size of ensembles

The optimal size of an ensemble is a task specific parameter that needs to be optimized [22]. Figure 4 shows that (1) the ensemble with multiple masks outperformed the only one mask. (2) when ensemble sizes increased, performance tended to saturate. We set the ensemble size to 5 in our methods. Figure 4 also shows standard deviation of segmentation performance with respect to different ensemble sizes. The variation of segmentation performance was reduced on dice metrics when the ensemble size increased. It demonstrated that the ensemble

Table 2: Segmentation scores of PD-DDPM with different pre-segmentation Dice scores.

Pre-segmentation accuracy	Dice	Hd95	Jaccard	F1
Zero (Dice=0.000)	0.353	9.899	0.235	0.677
Pre-seg (Dice=0.754)	0.790	4.364	0.662	0.776
Pre-seg (Dice=0.770)	0.795	4.121	0.667	0.770
Pre-seg (Dice=0.785)	0.798	3.974	0.671	0.786
Pre-seg (Dice=0.799)	**0.812**	**3.494**	**0.689**	**0.800**

model not only boost the segmentation performance but also guarantee a robust segmentation result.

3.6. Effect of pre-segmentation accuracy

Because the pre-segmentation result is not the real ground truth segmentation, there will be some errors between $\hat{x}^{T'}$ and $x^{T'}$. Here we analyzed the effect of pre-segmentation accuracy on PD-DDPM. As shown in Table 2, the higher the pre-segmentation Dice score, the higher performance of the PD-DDPM (when T'=300). So PD-DDPM can be combined with existing advanced segmentation networks to further improve performance and obtain uncertainty estimates.

4. CONCLUSION

To accelerating DDPM for medical image segmentation, this paper propose a pre-segmentation diffusion sampling DDPM (PD-DDPM), which is specially used for medical image segmentation. We empirically find that PD-DDPM achieves the best results under the parameter settings of this paper when $T' = 300$ (T=1000). Experiments show even with a significantly smaller number of reverse sampling steps, PD-DDPM also outperforming the vanilla DDPM. Compared with some existing acceleration methods, our method also get the best result. Further, PD-DDPM is orthogonal to existing advanced segmentation models, which can be combined to further improve model performance and obtain uncertainty estimates.

5. REFERENCES

[1] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., & Ganguli, S. (2015, June). Deep unsupervised learning using nonequilibrium thermodynamics. In International Conference on Machine Learning (pp. 2256-2265). PMLR.

[2] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems, 33, 6840-6851.

[3] Yang, L., Zhang, Z., & Hong, S. (2022). Diffusion models: A comprehensive survey of methods and applications. arXiv preprint arXiv:2209.00796.

[4] Croitoru, F. A., Hondru, V., Ionescu, R. T., & Shah, M. (2022). Diffusion models in vision: A survey. arXiv preprint arXiv:2209.04747.

[5] Wolleb, J., Sandkuehler, R., Bieder, F., Valmaggia, P., & Cattin, P. C. (2021, December). Diffusion Models for Implicit Image Segmentation Ensembles. In Medical Imaging with Deep Learning.

[6] P. Suetens, "Fundamentals of medical imaging, 3rd edition," 2017.

[7] Song, Jiaming, Chenlin Meng, and Stefano Ermon. "Denoising Diffusion Implicit Models." International Conference on Learning Representations. 2020.

[8] San-Roman R, Nachmani E, Wolf L. Noise estimation for generative diffusion models[J]. arXiv preprint arXiv:2104.02600, 2021.

[9] Salimans T, Ho J. Progressive Distillation for Fast Sampling of Diffusion Models[C]//International Conference on Learning Representations. 2021.

[10] Vahdat A, Kreis K, Kautz J. Score-based generative modeling in latent space[J]. Advances in Neural Information Processing Systems, 2021, 34: 11287-11302.

[11] Rombach R, Blattmann A, Lorenz D, et al. High-resolution image synthesis with latent diffusion models[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2022: 10684-10695.

[12] Zheng, Huangjie, et al. "Truncated Diffusion Probabilistic Models and Diffusion-based Adversarial Auto-Encoders." arXiv preprint arXiv:2202.09671 (2022).

[13] Lyu Z, Xu X, Yang C, et al. Accelerating Diffusion Models via Early Stop of the Diffusion Process[J]. arXiv preprint arXiv:2205.12524, 2022.

[14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In Advances in Neural Information Processing Systems, Vol. 27. 139–144.

[15] Kingma D P, Welling M. Auto-encoding variational bayes[J]. arXiv preprint arXiv:1312.6114, 2013.

[16] Kuijf H J, Biesbroek J M, De Bresser J, et al. Standardized assessment of automatic segmentation of white matter hyperintensities and results of the WMH segmentation challenge[J]. IEEE transactions on medical imaging, 2019, 38(11): 2556-2568.

[17] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation[C]//International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 2015: 234-241.

[18] Oktay O, Schlemper J, Folgoc L L, et al. Attention u-net: Learning where to look for the pancreas[J]. arXiv preprint arXiv:1804.03999, 2018.

[19] Zhou Z, Siddiquee M M R, Tajbakhsh N, et al. Unet++: Redesigning skip connections to exploit multiscale features in image segmentation[J]. IEEE transactions on medical imaging, 2019, 39(6): 1856-1867.

[20] Gal, Yarin, and Zoubin Ghahramani. "Dropout as a bayesian approximation: Representing model uncertainty in deep learning." international conference on machine learning. PMLR, 2016.

[21] Kohl, Simon, et al. "A probabilistic u-net for segmentation of ambiguous images." Advances in neural information processing systems 31 (2018).

[22] Li H, Jiang G, Zhang J, et al. Fully convolutional network ensembles for white matter hyperintensities segmentation in MR images[J]. NeuroImage, 2018, 183: 650-665.