

Score-Based Generative Models for Medical Image Segmentation using Signed Distance Functions

Lea Bogensperger¹, Dominik Narnhofer¹, Filip Ilic¹, and Thomas Pock¹

Institute of Computer Graphics and Vision, Graz University of Technology
{lea.bogensperger,dominik.narnhofer,filip.ilic,pock}@icg.tugraz.at

Abstract. Medical image segmentation is a crucial task that relies on the ability to accurately identify and isolate regions of interest in medical images. Thereby, generative approaches allow to capture the statistical properties of segmentation masks that are dependent on the respective structures. In this work we propose a conditional score-based generative modeling framework to represent the signed distance function (SDF) leading to an implicit distribution of segmentation masks. The advantage of leveraging the SDF is a more natural distortion when compared to that of binary masks. By learning the score function of the conditional distribution of SDFs we can accurately sample from the distribution of segmentation masks, allowing for the evaluation of statistical quantities. Thus, this probabilistic representation allows for the generation of uncertainty maps represented by the variance, which can aid in further analysis and enhance the predictive robustness. We qualitatively and quantitatively illustrate competitive performance of the proposed method on a public nuclei and gland segmentation data set, highlighting its potential utility in medical image segmentation applications.

Keywords: Score-based generative models · image segmentation · conditional diffusion models · signed distance function.

1 Introduction

Medical image segmentation approaches are often trained end-to-end in a discriminative manner using deep neural networks [18, 3, 29]. However, also generative models have emerged for image segmentation with the advantage of learning the underlying statistics of segmentation masks conditioned on input images [9, 34, 1]. Apart from generative adversarial networks (GANs), promising candidates in this field are score-based generative models [6, 23, 26], which learn the score of a data distribution to sample from the distribution in the framework of a stochastic differential equation (SDE). Herein, noise is gradually injected to smooth the data distribution until it resembles a simple, tractable prior distribution – a process which can be reversed with the corresponding time-reverse SDE relying on the learned score function.

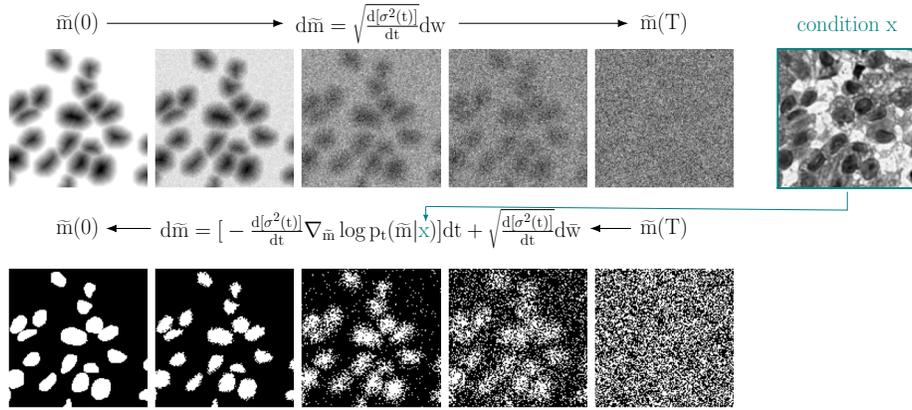


Fig. 1. Schematic of the corruption process (top row from left to right for different $t \in [0, T]$) on an SDF segmentation mask \tilde{m} for a given image x . The forward and reverse processes are governed by the variance-exploding SDE and its time-reverse SDE, respectively, where the latter uses the conditioned score of the distribution of SDF masks \tilde{m} given images x . The corresponding thresholded segmentation masks m are shown in the second row.

Diffusion models can be naturally incorporated to solving inverse problems in medical imaging [25], but they can also be extended to learn a conditional distribution, which makes them well suited for image segmentation. Their potential applicability has already been shown in works [31, 1, 30] using conditional denoising diffusion probabilistic modeling (DDPM). A significant drawback when directly injecting noise on segmentation masks is given by the fact that the distortion process is unnatural with respect to the underlying distribution. One could argue that the statistics of segmentation masks, which are bimodal or contain very few modes depending on the number semantic classes, are not easy to learn as there is no transition between class modes. A remedy is provided by recalling the SDF, a classic image segmentation technique [16], which has also regained attention within newer works on discriminative image segmentation using convolutional neural networks (CNNs) [15, 33, 4]. It is based on the idea that an implicit segmentation map is computed using the SDF for which at any given point in the resulting segmentation map the orthogonal distance to the closest boundary point is computed. Additionally, the distance is denoted with a negative sign for interior regions and a positive sign for the background regions. Thus, the SDF acts as a shape prior in some sense and it represents a smoother distribution of segmentation masks and thus a smooth transition in class modes. Moreover, it naturally promotes smoothness within the transformed binary segmentation map – which in return is obtained by thresholding the SDF map at the object boundaries.

In this work, we propose to fuse medical image segmentation using score-based generative modeling based on SDEs with a segmentation approach relying

on the Euclidian SDF - thereby learning a smooth implicit representation of segmentation masks conditioned on the respective image. The noise-injecting data perturbation process in diffusion models thereby blends in canonically by gradually smoothing the distribution of the SDF mask. Since binary masks only provide a discrete representation of the object region, the score function can be highly sensitive to small changes in the binary mask, leading to segmentation errors and inconsistencies. On the other hand, the SDF provides a smooth and continuous representation of the object boundary, which can be more natural and better suited to model the shape and boundaries of objects in the image.

With this approach, the object boundaries are obtained from the sampling process, which can then be used to threshold for the binary segmentation masks, implying that the segmented objects will be smooth. Moreover, the segmentation uncertainty can be quantified by acquiring multiple segmentations given an input image due to its generative nature, thus further enhancing the robustness and interpretability of the approach and providing valuable insight into the segmentation process.

2 Method

2.1 Image Segmentation using SDFs

Image segmentation is the task of finding a segmentation mask $m \in \mathcal{M}$ that assigns each pixel in an image $x \in \mathcal{X}$ a class, where $\mathcal{X} := \mathbb{R}^{M \times N}$ and $\mathcal{M} := \mathbb{R}^{M \times N}$. Using SDFs the segmentation problem can be rephrased in the context of (signed) distances with respect to the object boundaries, where we consider each pixel m_{ij} of the domain $\Omega = \{1, \dots, M\} \times \{1, \dots, N\}$ with the object \mathcal{S} to be segmented. The SDF map then contains for each m_{ij} the distance to the closest boundary pixel $\partial\mathcal{S}$, where a negative/positive sign denotes the interior/outside of the object, respectively.

Mathematically, the SDF \tilde{m} of a segmentation mask m can thus be computed using the Euclidian distance function for each pixel m_{ij} , which additionally can be truncated at a threshold δ to consider only a span of pixels around the object boundary $\partial\mathcal{S}$. Moreover, the truncation helps to remain agnostic with respect to higher positive distances belonging to the background, as they should not impact the resulting segmentation. The implicit, truncated SDF $\phi(m_{ij})$ is then obtained as follows (also see [16]):

$$\phi(m_{ij}) = \begin{cases} -\min\{\min_{y \in \partial\mathcal{S}} \|y - m_{ij}\|_2, \delta\} & \text{if } m_{ij} \in \mathcal{S}, \\ \min\{\min_{y \in \partial\mathcal{S}} \|y - m_{ij}\|_2, \delta\} & \text{if } m_{ij} \in \Omega \setminus \mathcal{S}, \\ 0 & \text{if } m_{ij} \in \partial\mathcal{S}. \end{cases} \quad (1)$$

Thus we can obtain the full segmentation map $\tilde{m} = \phi(m_{ij})_{\substack{i=1, \dots, M \\ j=1, \dots, N}}$ by applying $\phi(\cdot)$ on the entire segmentation mask m , which is also depicted in Figure 2. Conversely, given \tilde{m} one can easily retrieve the binary segmentation map m by thresholding at 0 to separate segmented objects from background.

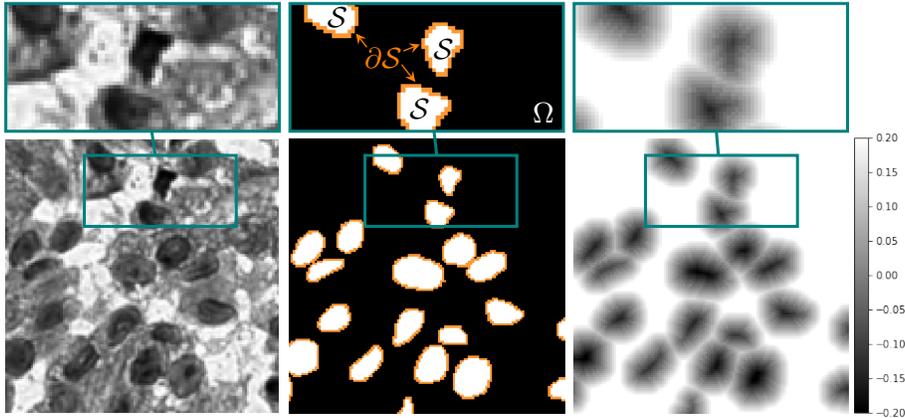


Fig. 2. Given an image x (left), its binary segmentation mask (center) can be transformed into a truncated, normalized SDF mask (right). The zoomed area shows some segmentation objects in detail denoted by S and their boundaries ∂S embedded in the domain Ω .

2.2 Conditional Score-Based Segmentation

In generative approaches for image segmentation the goal is to sample from the conditional distribution $p(\tilde{m}|x)$ to obtain a segmentation mask \tilde{m} given an input image x . We frame this task in the setting of SDEs, which gradually corrupt data samples from the data distribution $\tilde{m}(0) \sim p_0$ until a tractable prior distribution p_T is reached. There exists a corresponding time-reverse SDE which can then be leveraged to transform the prior distribution back to the data distribution by using the score of the conditional data distribution at time t , i.e. $\nabla_{\tilde{m}} \log p_t(\tilde{m}|x)$. This score function can be learned using a training set of S paired data samples of images and SDF segmentation masks $\mathcal{D} = \{(\tilde{m}_s, x_s)\}_{s=1}^S$.

In general, SDEs have a drift term and a diffusion coefficient that govern its forward and corresponding reverse evolution. Hereby, we solely focus on so-called variance-exploding SDEs, as presented in [26], although our approach should hold also for variance-preserving SDEs. Thus, for a time process with $t \in [0, T]$, a sequence $\{\tilde{m}(t)\}_{t=0}^T$ is generated by means of additive corruptive Gaussian noise with standard deviation $\sigma(t)$. Using Brownian motion w to denote the noise corruption, the corresponding SDE then reads as

$$d\tilde{m} = \sqrt{\frac{d[\sigma^2(t)]}{dt}} dw. \quad (2)$$

The SDE in (2) can be reversed [2], which requires – when conditioning on images x – the score of the conditional distribution $\nabla_{\tilde{m}} \log p_t(\tilde{m}|x)$. Both the forward and the reverse process are illustratively demonstrated in Figure 1, where the resulting thresholded segmentation masks also show the effect of the corruption

process on the explicit binary segmentation mask (which is obtained by thresholding the SDF segmentation mask). The reverse SDE for \bar{w} , since time is now going backwards such that $t \in [T, 0]$, essentially reads as

$$d\bar{m} = \left[-\frac{d[\sigma^2(t)]}{dt} \nabla_{\bar{m}} \log p_t(\bar{m}|x) \right] dt + \sqrt{\frac{d[\sigma^2(t)]}{dt}} d\bar{w}. \quad (3)$$

The scores of the conditional distribution $\nabla_{\bar{m}} \log p_t(\bar{m}|x)$ can be estimated using either techniques from score matching [8, 23, 24] or from implicit score estimation such as DDPM [21, 6]. Here, we choose to learn a noise-conditional score network $s_{\theta^*}(\bar{m}(t), x, \sigma(t))$ for the reverse-time SDE following the continuous formulation of denoising score matching with $t \in \mathcal{U}(0, T)$ for noise levels from σ_{\min} to σ_{\max} [26]. Thereby, using Tweedie’s formula [5] we minimize the following objective:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_t \left\{ \sigma^2(t) \mathbb{E}_{\bar{m}(0)} \mathbb{E}_{\bar{m}(t)|\bar{m}(0)} \left[\|s_{\theta}(\bar{m}(t), x, \sigma(t)) - \nabla_{\bar{m}} \log p_{0t}(\bar{m}(t)|\bar{m}(0))\|_2^2 \right] \right\}. \quad (4)$$

The perturbation kernel $p_{0t}(\bar{m}(t)|\bar{m}(0))$ has the form of a standard normal distribution using $\sigma(t) = \sigma_{\min} \left(\frac{\sigma_{\max}}{\sigma_{\min}} \right)^t$.

Once the learned scores s_{θ^*} are available, they can be used to sample from the conditional distribution, where there exist several numerical solvers based on the time-reverse SDE. Here, we employ a predictor-corrector sampler as proposed by [26], which alternates between time-reverse SDE steps (the predictor) and Langevin Markov chain Monte Carlo (MCMC) sampling steps (the corrector), see Algorithm 1.

Algorithm 1: Predictor-corrector algorithm to sample from $p(\bar{m}|x)$.

```

1 Choose conditioning image  $x$ , set number of iterations  $K, J$ , set  $r \in \mathbb{R}^+$ 
2  $\bar{m}_K \sim \mathcal{N}(0, \sigma_{\max}^2 I)$ 
3 for  $k = K - 1, \dots, 0$  do
4    $z \sim \mathcal{N}(0, I)$ ;
5    $\bar{m}_k = \bar{m}_{k+1} + (\sigma_{k+1}^2 - \sigma_k^2) s_{\theta^*}(\bar{m}_{k+1}, x, \sigma_{k+1}) + \sqrt{\sigma_{k+1}^2 - \sigma_k^2} z$ ;
6   for  $j = 1, \dots, J$  do
7      $z \sim \mathcal{N}(0, I)$ ;
8      $g = s_{\theta^*}(\bar{m}_k^{j-1}, x, \sigma_k)$ ;
9      $\varepsilon = 2(r\|z\|_2 / \|g\|_2)^2$ ;
10     $\bar{m}_k^j = \bar{m}_k^{j-1} + \varepsilon g + \sqrt{2\varepsilon} z$ ;
11  end
12   $\bar{m}_k = \bar{m}_k^J$ ;
13 end

```

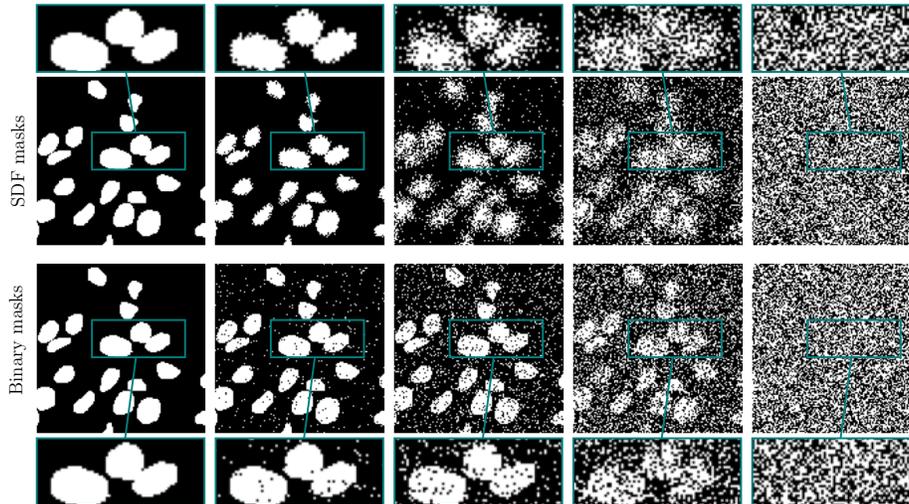


Fig. 3. Comparison of the effect of the corruption process on the resulting thresholded segmentation masks when using an SDF mask \tilde{m} and a binary mask m for a given image x . Note that the SDF representation allows for a more natural distortion process in its thresholded masks, which evolves along the object boundaries instead of directly introducing “hole”-like structures at random pixel positions as it is the case in the thresholded masks when directly using the binary segmentation mask.

2.3 Motivation of the SDF in Conditional Score-Based Segmentation

The motivation of using the SDF in conditional score-based segmentation is due to the nature of the perturbation process in standard diffusion models, which consists of gradually adding noise to the sought segmentation mask in its SDF representation such that its distribution gets smoothed. The destructive process thus implicitly incorporates the boundary information of the segmented objects.

In contrast, if the binary masks are used directly in the perturbation process, the corruption yields “hole”-like structures in the resulting thresholded segmentation masks and there is no possibility to integrate the structure of the segmentation objects within the forward process. A comparison of effect of the destructive process on both variants of segmentation mask representations for varied time steps is shown in Figure 3.

3 Experiments

3.1 Data Sets

For the experimental setting, we utilize two publicly available data sets. The first data set is MoNuSeg [12], which consists of 30 training images and 14 test images. Each image is of size 1000×1000 and overall they contain more than

21,000 annotated nuclei in Haematoxylin and Eosin (H&E) stained microscopic images. For data preprocessing we resort to a structure-preserving color normalization [27] as the different organ sites yield considerable intensity variations in the data. This is followed by a gray scale conversion and all images are subsequently resized to 500×500 . During training, overlapping crops of 128×128 were used with random horizontal and vertical flips for data augmentation.

As a second medical data set, the Gland Segmentation (GlaS) data set [20] was used. It consists of 85 training and 80 test H&E stained microscopic images with annotated glands from colorectal cancer tissue. Again, a structure-preserving color normalization [27] is used also here for data preprocessing, which is followed by resizing all training and test images to 128×128 inspired by [28]. Data augmentation in training is done as with the MoNuSeg data set.

3.2 Architecture & Training

The architecture to learn the noise-conditional score function is adapted from [6, 26]. Further, a conditioning on the image x is required, for which we roughly follow recent works related to conditional generative modeling by concatenating the encoded conditioning image to the network input [17, 7, 19].

For the diffusion parameters we set $\sigma_{\max} = 5$ and $\sigma_{\min} = 1e-3$. Note that the latter is slightly lower than usually proposed in literature, which is due to the SDF data distribution, as denoising score matching requires a perturbation kernel at the lowest noise scale such that the input distribution remains more or less unchanged. For the learning setting we employ Adam’s optimizer [11] with default coefficient values and a learning rate of $1e-4$.

3.3 Sampling

To obtain segmentation masks for test images x , we use the predictor-corrector sampler in Algorithm 1. The test images of the MoNuSeg data set are each divided into four evenly sized patches per image, whereas for the GlaS data set the entire test images are processed. In all sampling experiments we set $r = 0.35/0.15$ for the corrector step size scaling for the MoNuSeg and GlaS data set, respectively, as we empirically found this to work best in terms of evaluation metrics. Moreover, we use $K = 500/200$ predictor steps and $J = 2/1$ corrector steps for both data sets, respectively, as this setting revealed to yield best results, despite general low numeric fluctuations amongst different settings.

3.4 Evaluation

The resulting samples are SDF predictions which have yet to be converted to valid segmentation maps. Thus, they have to be thresholded at 0 which represents object boundaries to separate segmented objects (which have 0 at the boundary and negative distances inside assigned) and background (consisting of positive distances). However, due to the employed approach of denoising score matching

we set the threshold to $3\sigma_{\min}$, since we have to assume that there is still remaining noise present at the scale of the smallest noise level. This is crucial to consider since we are interested in the exact boundary.

By leveraging the generative nature of our approach, we are further investigating the effect of averaging over 128 sampling runs. Thereby we obtain the minimum mean square error (MMSE) which gives a more robust prediction, which is shown in increased quantitative scores.

The obtained segmentation masks are then evaluated using the standard metrics F1 score and Intersection over Union (IoU). We compare our method to commonly referred benchmark models, ensuring that both U-Net variants and attention/transformer mechanisms are considered. Moreover, we also consider [30] to obtain a comparison with a generative, conditional DDPM that predicts standard binary segmentation masks in the sampling process. The approach in [30] can also be viewed in the form of an SDE and its time-reverse SDE using a variance-preserving scheme.

3.5 Results

Table 1 shows quantitative results for both data sets with our method and comparison methods. To enable a fair comparison the benchmark results were taken from [28] where possible. The results indicate that our method outperforms the comparison methods on the GlaS data set, but also for the MoNuSeg data set competitive results can be obtained, although slightly worse than some of the comparison methods. For both data sets, using the MMSE by averaging over multiple sampling runs clearly gives a significant boost in quantitative performance. In comparison, the DDPM delivers slightly worse results, but we want to emphasize that they could also be increased by computing the MMSE over several runs before thresholding as it is a generative model – this was also shown in [30] where segmentation ensembles are computed to improve the results.

Table 1. Quantitative segmentation results on the MoNuSeg and GlaS data set.

Method \ Metric	MoNuSeg		GlaS	
	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑
FCN [3]	28.84	28.71	66.61	50.84
U-Net [18]	79.43	65.99	77.78	65.34
U-Net++ [35]	79.49	66.04	78.03	65.55
Res-UNet [32]	79.49	66.07	78.83	65.95
Axial Attention U-Net [29]	76.83	62.49	76.26	63.03
MedT [28]	79.55	66.17	81.02	69.61
DDPM [30]	76.03	61.42	76.81	64.15
Ours	78.13	64.19	82.03	71.36
Ours – MMSE	78.64	64.87	82.77	72.07

Figure 4 shows exemplary qualitative results for both data sets to further highlight the potential applicability of our proposed method. Note that although the evaluation metrics for our approach that can be found in Table 1 are obtained by evaluating the entire images, we show smaller crops here for the sake of a more detailed visual inspection of the segmented objects. One can clearly observe the smooth segmented objects m obtained from thresholding the SDF predictions $\tilde{m}(0)$, which are in good agreement with the groundtruth segmentation masks m_{gt} for both data set samples.

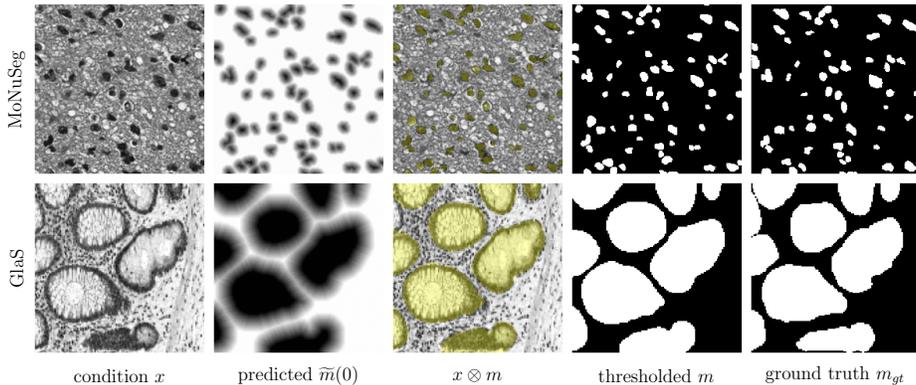


Fig. 4. Exemplary sampled segmentation masks for both data sets. The predicted SDF masks $\tilde{m}(0)$ are directly obtained from the sampling procedure, whereas the thresholded masks m are shown to additionally enable a visual comparison with the depicted ground truth m_{gt} . Furthermore, we provide the condition image with the overlaid thresholded mask $x \otimes m$.

A visual comparison is additionally depicted in Figure 5, where we compare the thresholded segmentation prediction of our model of a MoNuSeg test image with its DDPM-based (generative) counterpart [30] and two discriminative models, namely U-Net++ [35] and MedT [28]. As can be seen in the provided zoom, the SDF representation of segmentation objects indeed seems to act as a shape prior and thus yields smoother segmentation objects while avoiding artefacts such as single pixels/small structures that are mistakenly classified as foreground objects.

3.6 Segmentation Uncertainty

Since the presented approach is based on a generative scheme, we can sample from the conditional distribution of the SDF given the conditioning image $p(\tilde{m}|x)$. An advantage of this approach is given by the fact that the resulting statistical values allow for the quantification of segmentation uncertainties in the SDF predictions as well as the thresholded masks, which provide additional

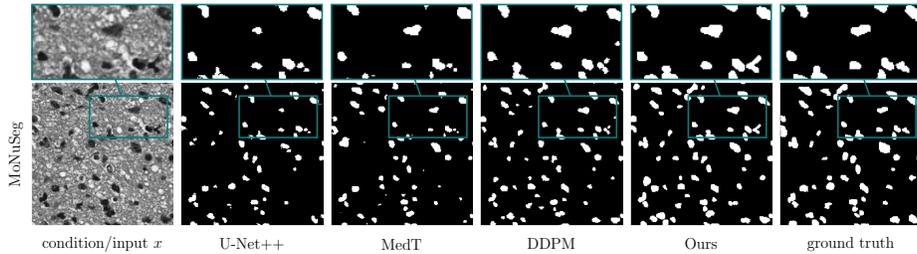


Fig. 5. Qualitative results for a test input image x for discriminative approaches resp. condition image x for the generative approaches. We compare our model with the generative DDPM [30] and two discriminative models, U-Net++ [35] and MedT [28]. Using the SDF in our conditional score-based segmentation approach shows that a shape prior is learned, thus preventing small pixel-wise artefacts in the resulting predictions and yielding smooth segmentation objects.

insights into the segmentation process that are not available with traditional discriminative approaches. An illustration of the aforementioned property on a MoNuSeg data sample can be seen in Figure 6, where the standard deviation maps associated with the SDF predictions and thresholded masks highlight the regions of uncertainty.

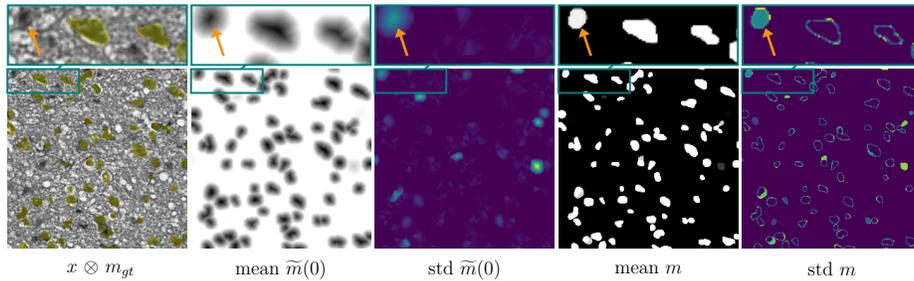


Fig. 6. Segmentation example with according statistical values for the SDF predictions $\tilde{m}(0)$ respectively thresholded masks m . Notably, the image includes a region erroneously segmented as a nuclei, as indicated by the orange arrow. This region is highlighted in the standard deviation maps, which represent the associated uncertainty in the segmentation.

In general, we observe that the standard deviation appears high on transitions from nuclei to background as well as in wrongly detected nuclei or over-segmented parts of nuclei. This encouraging observation leads us to the hypothesis that the uncertainty may be associated directly with segmentation errors similar to what has been shown in [14], see also Figure 7. Here, a visual comparison of the error and standard deviation (uncertainty) maps indicates that the latter very likely has high predictive capability. A detailed analysis including the

mutual information of the two variables, however, is out of scope for this work and subject to future work.

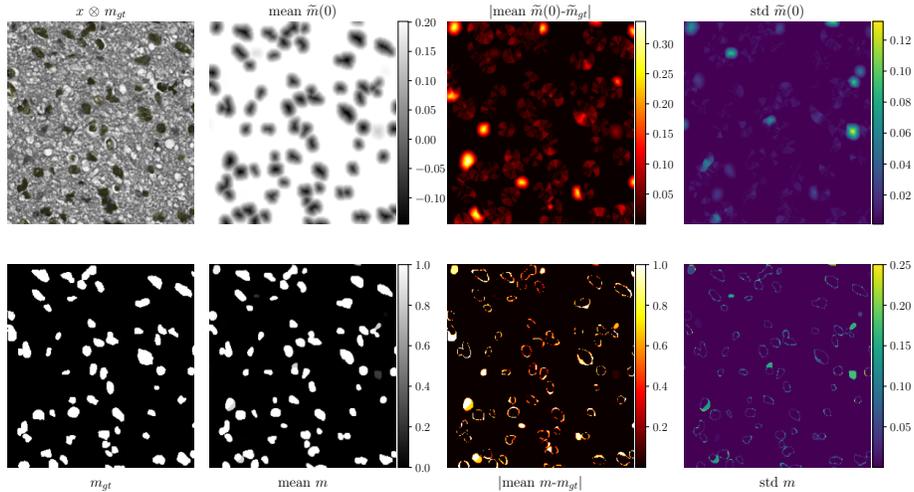


Fig. 7. Segmentation example with statistical values for the predicted $\tilde{m}(0)$ and thresholded masks m , along with the corresponding ground truth images and the absolute error between the predictions and the ground truth. The visual comparison of error and standard deviation maps (uncertainty) suggests that the latter has predictive capability for the error.

3.7 Ablation

As we rely on the standard noise-conditioned score network following [6] for the proposed method, it is more interesting to investigate the influence of the SDF. Therefore, we learn the conditional score function such that segmentation masks m can be sampled given x , where all other settings remain unchanged – including σ_{\min} and σ_{\max} . Hereby, we obtain a mean IoU/mean F1 score of 51.72/67.80 on the MoNuSeg data set, and 61.39/73.97 on the GlaS data set, respectively.

Our approach should work reasonably well with other types of architectures suitable for learning the score of a data distribution. As the primary objective of this work was to introduce a novel concept rather than striving to surpass existing benchmarks, in future work, a more sophisticated network architecture could be considered to further improve the segmentation accuracy.

4 Discussion and Limitations

Generative approaches usually require an increased network complexity and are thus computationally more expensive than their discriminative counterparts. We believe that the advantages of having a generative model outweigh its potential

drawbacks due to the possibility to evaluate statistical quantities, such as averaging multiple predictions and the generation of uncertainty maps, which may give valuable insights and might be crucial especially in medical imaging applications. While the sampling process itself still requires significantly more time than executing a single forward pass of a discriminative network, there is a variety of new methods available to speed up the inference stage of diffusion models [22, 13, 10]. For the sake of simplicity the standard method was chosen, however, there should be nothing to argue against incorporating a different sampling technique.

The proposed approach outperforms the comparison methods only on one of the two used data sets. However, also on the MoNuSeg data set competitive results have been presented that could potentially be improved with a further optimized network architecture. While a state-of-the-art model based on U-Net or transformer architectures might in some cases outperform our quantitative results we still want to emphasize that there is no option of obtaining a segmentation uncertainty as in Figures 6 and 7 with such discriminative approaches. Robustness and reliability should not be traded for minor quantitative increments and improvements.

Moreover, using SDF maps to represent the segmentation masks bears the advantage that a shape prior is learned which favors smooth segmentation objects. The resulting binary segmentation masks thus encompass a different characteristic than when using standard conditional DDPM based approaches [1, 30]. Although the statistics of these underlying segmentation masks are properly learned, the quantitative results do not reflect those findings in the respective metrics.

Further, obtaining SDF predictions can generally offer other potential advantages such as obtaining instance segmentations as a byproduct. By using the watershed transform, the SDF maps could easily be turned into instance segmentation maps and the problem of touching instances is circumvented.

5 Conclusion and Outlook

In this work we proposed a generative approach for medical image segmentation by fusing conditional score-based methods with the concept of representing segmentation maps with the SDF. The potential applicability of the method was demonstrated qualitatively and quantitatively on two public medical data sets. The SDF provides a smoother and more continuous representation of object boundaries compared to binary masks, which makes it a more canonical choice for accurate and robust segmentation.

Furthermore, by leveraging the generative approach, statistical measures such as mean and standard deviation can be calculated to quantify the uncertainty in the segmentation results. This information is especially useful for medical diagnosis and treatment planning, where it is important to know the level of confidence in the segmentation results.

As an outlook for future research we will focus on the extension of the approach to multi-class segmentation by additionally incorporating the exclusivity

of semantic classes per pixel into the predicted SDF maps. Additionally, fuelled by preliminary experimental successes, the possibility of learning the score of the joint distribution of images x and segmentation masks \tilde{m} will be explored. This would provide a powerful framework, where one could directly sample paired training data (\tilde{m}_s, x_s) from the joint distribution or condition on either one of them to sample from the conditional distributions $p(\tilde{m}|x)$ or $p(x|\tilde{m})$.

References

1. Amit, T., Nachmani, E., Shaharbany, T., Wolf, L.: Segdiff: Image segmentation with diffusion probabilistic models. arXiv preprint arXiv:2112.00390 (2021)
2. Anderson, B.D.: Reverse-time diffusion equation models. *Stochastic Processes and their Applications* **12**(3), 313–326 (1982)
3. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **39**(12), 2481–2495 (2017)
4. Brissman, E., Johnander, J., Felsberg, M.: Predicting signed distance functions for visual instance segmentation. In: 2021 Swedish Artificial Intelligence Society Workshop (SAIS). pp. 1–6. IEEE (2021)
5. Efron, B.: Tweedie’s formula and selection bias. *Journal of the American Statistical Association* **106**(496), 1602–1614 (2011)
6. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
7. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.* **23**(47), 1–33 (2022)
8. Hyvärinen, A., Dayan, P.: Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research* **6**(4) (2005)
9. Iqbal, A., Sharif, M., Yasmin, M., Raza, M., Aftab, S.: Generative adversarial networks and its applications in the biomedical image segmentation: a comprehensive survey. *International Journal of Multimedia Information Retrieval* **11**(3), 333–368 (2022)
10. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* **35**, 26565–26577 (2022)
11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
12. Kumar, N., Verma, R., Sharma, S., Bhargava, S., Vahadane, A., Sethi, A.: A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE transactions on medical imaging* **36**(7), 1550–1560 (2017)
13. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems* **35**, 5775–5787 (2022)
14. Narnhofer, D., Habring, A., Holler, M., Pock, T.: Posterior-variance-based error quantification for inverse problems in imaging. ArXiv [abs/2212.12499](https://arxiv.org/abs/2212.12499) (2022)
15. Naylor, P., Laé, M., Reyat, F., Walter, T.: Segmentation of nuclei in histopathology images by deep regression of the distance map. *IEEE transactions on medical imaging* **38**(2), 448–459 (2018)
16. Osher, S., Fedkiw, R., Piechor, K.: Level set methods and dynamic implicit surfaces. *Appl. Mech. Rev.* **57**(3), B15–B15 (2004)
17. Özdenizci, O., Legenstein, R.: Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
18. 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)

19. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–10 (2022)
20. Sirinukunwattana, K., Pluim, J.P., Chen, H., Qi, X., Heng, P.A., Guo, Y.B., Wang, L.Y., Matuszewski, B.J., Bruni, E., Sanchez, U., et al.: Gland segmentation in colon histology images: The glas challenge contest. *Medical image analysis* **35**, 489–502 (2017)
21. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International Conference on Machine Learning. pp. 2256–2265. PMLR (2015)
22. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
23. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* **32** (2019)
24. Song, Y., Garg, S., Shi, J., Ermon, S.: Sliced score matching: A scalable approach to density and score estimation. In: Uncertainty in Artificial Intelligence. pp. 574–584. PMLR (2020)
25. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. In: International Conference on Learning Representations (2022)
26. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2021)
27. Vahadane, A., Peng, T., Albarqouni, S., Baust, M., Steiger, K., Schlitter, A.M., Sethi, A., Esposito, I., Navab, N.: Structure-preserved color normalization for histological images. In: 2015 IEEE 12th international symposium on biomedical imaging (ISBI). pp. 1012–1015. IEEE (2015)
28. Valanarasu, J.M.J., Oza, P., Hacihaliloglu, I., Patel, V.M.: Medical transformer: Gated axial-attention for medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24. pp. 36–46. Springer (2021)
29. Wang, H., Zhu, Y., Green, B., Adam, H., Yuille, A., Chen, L.C.: Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV. pp. 108–126. Springer (2020)
30. 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)
31. Wu, J., Fang, H., Zhang, Y., Yang, Y., Xu, Y.: Medsegdiff: Medical image segmentation with diffusion probabilistic model. arXiv preprint arXiv:2211.00611 (2022)
32. Xiao, X., Lian, S., Luo, Z., Li, S.: Weighted res-unet for high-quality retina vessel segmentation. In: 2018 9th international conference on information technology in medicine and education (ITME). pp. 327–331. IEEE (2018)
33. Xue, Y., Tang, H., Qiao, Z., Gong, G., Yin, Y., Qian, Z., Huang, C., Fan, W., Huang, X.: Shape-aware organ segmentation by predicting signed distance maps. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12565–12572 (2020)
34. Xun, S., Li, D., Zhu, H., Chen, M., Wang, J., Li, J., Chen, M., Wu, B., Zhang, H., Chai, X., et al.: Generative adversarial networks in medical image segmentation: a review. *Computers in Biology and Medicine* **140**, 105063 (2022)

35. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4. pp. 3–11. Springer (2018)