

No-New-Denoiser: A Critical Analysis of Diffusion Models for Medical Image Denoising

Laura Pfaff^{1,2}, Fabian Wagner², Nastassia Vysotskaya¹, Mareike Thies¹,
Noah Maul¹, Siyuan Mei¹, Tobias Wuerfl², and Andreas Maier¹

¹ Pattern Recognition Lab, FAU Erlangen-Nürnberg, Erlangen, Germany

² Siemens Healthineers AG, Forchheim, Germany

laura.pfaff@fau.de

Abstract. Diffusion models, originally introduced for image generation, have recently gained attention as a promising image denoising approach. In this work, we perform comprehensive experiments to investigate the challenges posed by diffusion models when applied to medical image denoising. In medical imaging, retaining the original image content, and refraining from adding or removing potentially pathologic details is of utmost importance. Through empirical analysis and discussions, we highlight the trade-off between image perception and distortion in the context of diffusion-based denoising. In particular, we demonstrate that standard diffusion model sampling schemes yield a reduction in PSNR by up to 14% compared to one-step denoising. Additionally, we provide visual evidence indicating that diffusion models, in combination with stochastic sampling, have a tendency to generate synthetic structures during the denoising process, consequently compromising the clinical validity of the denoised images. Our thorough investigation raises questions about the suitability of diffusion models for medical image denoising, underscoring potential limitations that warrant careful consideration for future applications.

Keywords: Denoising · Diffusion models · Data fidelity.

1 Introduction

Various medical imaging applications, such as low-field magnetic resonance imaging (MRI) [19, 13, 21] or low-dose computed tomography (CT) [12, 26] inherently suffer from a low signal-to-noise-ratio (SNR). Thus, denoising techniques are commonly employed during or after reconstruction to achieve diagnostic image quality [17, 27]. Conventional denoising strategies include non-local means [18, 12], wavelet-based denoising [11, 2], or iterative techniques [10, 5].

In recent years, deep learning (DL)-based approaches have emerged as leading approaches for image denoising, demonstrating state-of-the-art performance through the extraction of intricate features from extensive datasets [24]. Supervised learning techniques necessitate pairs of noisy and clean images for training. However, the acquisition of noise-free target images can be challenging or even

infeasible in the context of medical imaging. In such scenarios, self-supervised learning methods like Noise2Void [14], Noise2Noise [15], or Noise2Self [4] can be applied to exploit the inherent properties or structure of the data, circumventing the need for noise-free ground-truth images. Alternative approaches use Stein’s unbiased risk estimator (SURE) to obtain outcomes closely resembling those achieved through supervised training [21].

Diffusion models, initially introduced as generative models, have gained widespread popularity in diverse scientific applications. Originally proposed for generating meaningful data samples from random noise [25, 9], diffusion models are now being extensively employed for image denoising [20, 30, 16], also in the context of MRI [29, 8, 7].

In medical imaging, preserving the original image content and avoiding the addition or removal of potentially pathological details is crucial. Thus, the trade-off between the perceptual quality provided by diffusion models and image fidelity has to be carefully considered. In this work, we investigate the application of diffusion models for medical imaging with a focus on MRI denoising. Based on comprehensive experiments using both synthetic and real MRI data, we present two main findings:

1. The iterative sampling, a component of the reverse process in diffusion models, does not contribute to improved denoising quality. On the contrary, it degrades the results.
2. The inherent stochastic nature of the sampling process introduces alterations in image content, yielding a loss in data fidelity.

2 Diffusion Models

A diffusion model characterizes a parameterized Markov chain comprising T discretized states $\mathbf{x}_0, \dots, \mathbf{x}_T$, that can be trained to generate images aligning with a learned data distribution [23, 9, 25]. The forward diffusion process, represented by $q(\mathbf{x}_t | \mathbf{x}_{t-1})$, incrementally adds Gaussian noise to the samples at each state, progressively corrupting the signal. In contrast, the reverse diffusion process aims to recover the prior signal $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$, where θ denotes the parameters of the neural network optimized to estimate p . The transition between each state is determined by a pre-defined noise schedule $\beta_1, \dots, \beta_T$, which is commonly modelled to increase linearly between two given values.

Forward process: Given a sample $\mathbf{x}_t$ at state t and a noise schedule $\beta_1, \dots, \beta_T$, the standard diffusion process of length T is defined by:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}). \quad (1)$$

Using the reparameterization trick [9], a sample at state t can be directly calculated from clean image $\mathbf{x}_0$:

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}), \quad \text{where } \alpha_t = 1 - \beta_t, \bar{\alpha}_t = \prod_{s=1}^t \alpha_s. \quad (2)$$

Thus, we can express $\mathbf{x}_t$ as:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{(1 - \bar{\alpha}_t)} \epsilon_t, \quad \text{where } \epsilon_t \sim \mathcal{N}(0, 1). \quad (3)$$

Here, $\sqrt{1 - \bar{\alpha}_t}$ defines the standard deviation of the noise at each state $\mathbf{x}_t$. During training, a neural network learns to predict either the clean image $\mathbf{x}_0$ or the noise ϵ_t , given a noisy input image $\mathbf{x}_t$ corresponding to a randomly sampled state t . Temporal information is incorporated at each network layer to convey the position in the diffusion process. Consequently, the model input comprises both the input image $\mathbf{x}_t$ and the corresponding timestep t . During inference, the noise is progressively removed from a given noisy image following a defined reverse sampling process.

Reverse process: In the reverse process, the network prediction is used to model the mean of state $\mathbf{x}_{t-1}$, given $\mathbf{x}_t$ and t as inputs. For *diffusion denoising probabilistic models (DDPMs)* [9], the sampling can be performed as follows:

$$\begin{aligned} p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) &= \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \sigma_t^2 \mathbf{I}), \\ \boldsymbol{\mu}_\theta(\mathbf{x}_t, t) &= \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta^{(t)} \right), \\ \sigma_t^2 &= \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t, \end{aligned} \quad (4)$$

where $\epsilon_\theta^{(t)}$ denotes the model prediction for ϵ_t . Hence, we can sample $\mathbf{x}_{t-1}$ as:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta^{(t)} \right) + \sqrt{\frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t}} \beta_t \mathbf{z}, \quad (5)$$

where $\mathbf{z} \sim \mathcal{N}(0, 1)$.

Diffusion denoising implicit models (DDIMs) [25] share the same training procedure as DDPMs but use implicit sampling:

$$\mathbf{x}_{t-1} = \underbrace{\sqrt{\bar{\alpha}_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta^{(t)}}{\sqrt{\bar{\alpha}_t}} \right)}_{\text{Predicted } \mathbf{x}_0} + \underbrace{\sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \epsilon_\theta^{(t)}}_{\text{Direction pointing to } \mathbf{x}_t} + \underbrace{\sigma_t \mathbf{z}}_{\text{Random noise}}. \quad (6)$$

Different choices of σ_t result in different generative processes. When $\sigma_t = \sqrt{(1 - \bar{\alpha}_{t-1})(1 - \bar{\alpha}_t)\beta_t}$, the generative process becomes a DDPM. In contrast, if $\sigma_t = 0$ for all t , the sampling process becomes deterministic, since no random noise is added. When training the model to predict the clean image $\mathbf{x}_0$ rather than the noise ϵ_t , $\epsilon_\theta^{(t)}$ in Equations 5 and 6 can be derived from Equation 3.

Diffusion models for denoising: In traditional diffusion models, the generative process originates from a random sample derived from a Gaussian noise




Fig. 1. Conditional sampling strategy for denoising applications. Instead of starting from the random noise $\mathbf{x}_T$, the modified reverse process starts from the noisy observation at the corresponding state $\mathbf{x}_t$ to generate the clean image $\mathbf{x}_0$.

distribution, progressively removing noise to synthesize new images. This process can be guided with a conditional image. However, in the context of image denoising, it was demonstrated that initiating the reverse process directly from the noisy image yields beneficial results [20]. In the forward process of diffusion models, samples are corrupted with Gaussian noise. The noise in complex-valued MRI can be modeled as a Gaussian noise distribution [1]. Exploiting this characteristic, we can match the image with the corresponding state in the diffusion process, determined via the estimated noise standard deviation, and initiate the reverse process directly from that state, as illustrated in Figure 1.

3 Experiments

Diverse techniques for training and sampling processes in diffusion models have been reported in existing literature. To strengthen the fundament of our claims, we trained diffusion models with various configurations. Subsequently, we evaluated all models with different sampling schemes, including stochastic, deterministic, and regularized sampling.

Data: To facilitate the sampling of random states during diffusion model training, access to clean images is essential. However, obtaining completely noise-free MRI data in real-world scenarios is infeasible. Thus, we employed two datasets:

1. Simulated *BrainWeb20* MRI database [3] with 20 anatomical brain models. These were divided into twelve models for training and four each for validation and testing, resulting in 720 image slices for training and 240 image slices each for validation and testing.
2. Real MRI dataset with 23 T2-weighted head scans (MAGNETOM Free.Max, Siemens Healthineers AG, Forchheim, Germany) acquired from healthy volunteers in accordance with all relevant guidelines. This dataset was split into training (19 scans) and validation/test sets (two scans each). To obtain noise-free images, we denoised the dataset using a U-Net architecture [22], trained in a self-supervised manner with SURE [21].

Table 1. The different training configurations. The first column, labeled as DTS, indicates whether the diffusion training scheme was applied, employing a defined noise schedule β and positional encoding.

ID	DTS	Prediction	$\mathbf{x}$ attenuation	β schedule
1	✓	Signal	✓	warm-up (5e-5,1e-2,1000)
2	✓	Noise	✓	warm-up (5e-5,1e-2,1000)
3	✓	Signal	×	warm-up (5e-5,1e-2,1000)
4	✓	Noise	×	warm-up (5e-5,1e-2,1000)
5	✓	Noise	×	constant (1e-4,500)
6	✓	Noise	×	constant (1e-4,200)
7	✓	Signal	×	constant (1e-4,500)
8	✓	Signal	×	constant (1e-4,200)
9	✓	Signal	×	linear (1e-5,1e-3,500)
10	✓	Noise	×	linear (1e-5,1e-3,500)
11	×	Signal	×	$\mathcal{U}[0.05,0.1]$

Training configurations: Table 1 provides an overview of the different training configurations. While some works propose to train the network to predict the noise-free image $\mathbf{x}_0$ [20], others claim that directly predicting the noise leads to more stable results [9]. Usually, in the forward diffusion process the signal $\mathbf{x}_t$ is attenuated with $\sqrt{\alpha_t}$, such that $\mathbf{x}_T$ comprises pure Gaussian noise. However, in denoising scenarios, the reverse process directly commences with a noisy image instead of a random noise sample. Thus, as suggested by Pearl et al. [20], the signal attenuation can be omitted. Importantly, the equations for the reverse sampling schemes have to be adapted accordingly. Furthermore, various functions have been proposed for noise schedule $\beta_{1,\dots,T}$, including linear and constant schedules [6]. Given the rather small noise levels in most MRI applications, Xiang et al. [29] employed a (reverse) warm-up strategy, such that β_t remains constant for the first 300 states. To comprehensively assess diffusion model performance in image denoising, we combined the proposed techniques into diverse training configurations. For comparison, we further trained the diffusion model architecture without the positional encoding on images corrupted with a noise level σ randomly sampled in a fixed interval $[0.05, 0.1]$.

Sampling schemes: During inference, we employed both the standard stochastic DDPM sampling scheme as well as the deterministic DDIM sampling scheme (described in Section 2) for all trained models. To mitigate potential alterations in image content introduced by the sampling process, Chung et al. [7] proposed a low frequency regularization scheme, incorporating a low-pass filtered version of the initial noisy image to keep the low frequency components intact. However, this approach yields attenuated high frequency components defining edges and small structures. Following Yang et al. [30], we regularize the reverse process by including the initial noisy image $\mathbf{x}_T$ in all sampling steps. Further, we report

results for the initial model prediction for $\mathbf{x}_0$ after the first sampling iteration during inference.

Experimental Setup: We trained all networks using PyTorch and the Adam optimizer with default parameters, learning rate $5 \cdot 10^{-5}$, and minibatch size of 8 until the validation loss reached convergence. The training was conducted on an NVIDIA A100 GPU. We utilized the diffusion model network architecture provided by Xiang et al. [29], available on GitHub [28].

Evaluation: We designed our experiments to analyze the influence of the generative sampling process on original image content. Thus, we calculated structural similarity index measure (SSIM) and peak signal-to-noise-ratio (PSNR) w.r.t. to the noise-free target after each iteration. We evaluated the networks trained on simulated BrainWeb data with test images corrupted by two different noise scenarios: 1) $\sigma \in [0.05, 0.1]$, which was used to train the non-diffusion models, and 2) $\sigma = 0.2$. In contrast, the networks trained with real MRI data were tested using the average noise level of each original image.

4 Results and Discussion

Quantitative results for different training configurations and sampling schemes are presented in Table 2 (simulated data) and Table 3 (real data). Despite employing various sampling schemes for the reverse process, none resulted in improvements in PSNR or SSIM for any training configuration compared to the result image derived from the model prediction after the first iteration. In fact, the results even deteriorated. While this effect has never been shown systematically before, it is in line with recent findings [16, 30]. The decline in denoising quality over the course of the reverse process is visually depicted in Figure 2. Particularly, stochastic sampling leads to degraded results by introducing random noise during the sampling process. Although regularizing the sampling process with the initial noisy image flattens the curve, it still does not outperform the denoised image after the first iteration.

In Figure 3, the detrimental effects of the stochastic sampling strategy becomes apparent. After the initial iteration, the model produces an image with details closely resembling the ground truth. However, in the subsequent reverse stochastic sampling process, the image content is altered, posing potential critical implications for diagnosis. By design, diffusion models excel in generating realistic images that align closely with the distribution of training data. In contrast, medical image denoising demands high perceptual quality, with the unyielding requirement of preserving data fidelity. The stochastic sampling contributes to the generation of realistic details, which is advantageous for image synthesis. However, this same stochasticity potentially creates inconsistencies, deviating the generated images from the true acquired information or introducing details that are missing in the measured data. Even the more intuitive, deterministic

Table 2. PSNR and SSIM results (PSNR / SSIM) for different training configurations using the simulated data evaluated on different noise levels. None of the utilized sampling schemes, including deterministic (det), stochastic (stoch), and respective combinations with regularization (reg), demonstrated an improvement across any of the training configurations when compared to the model prediction after the first iteration. Noisy refers to the metrics calculated between noisy input images and ground truth.

	ID	1 iteration	Det	Stoch	Det + Reg	Stoch + Reg
$\sigma \in [0.05, 0.1]$	Noisy	28.8 / 0.633				
	1	40.7 / 0.990	38.4 / 0.985	36.3 / 0.978	39.3 / 0.987	39.9 / 0.988
	2	40.7 / 0.990	39.0 / 0.987	36.6 / 0.980	39.8 / 0.989	40.2 / 0.989
	3	40.4 / 0.990	37.6 / 0.985	35.0 / 0.975	39.1 / 0.988	38.9 / 0.988
	4	40.3 / 0.989	39.1 / 0.986	36.9 / 0.978	39.7 / 0.987	39.6 / 0.987
	5	41.3 / 0.991	40.0 / 0.989	37.7 / 0.983	40.6 / 0.990	40.4 / 0.990
	6	42.0 / 0.993	41.1 / 0.992	38.6 / 0.987	41.6 / 0.992	41.5 / 0.992
	7	41.7 / 0.992	39.8 / 0.989	37.5 / 0.984	40.9 / 0.991	40.8 / 0.991
	8	42.4 / 0.993	41.4 / 0.992	39.0 / 0.988	42.0 / 0.993	41.9 / 0.993
	9	41.1 / 0.991	38.3 / 0.986	36.0 / 0.979	40.2 / 0.989	40.2 / 0.989
	10	40.7 / 0.990	39.1 / 0.986	36.7 / 0.979	40.4 / 0.989	40.4 / 0.989
	11	42.5 / 0.993				
$\sigma = 0.2$	Noisy	20.0 / 0.343				
	9	35.3 / 0.973	34.2 / 0.970	31.9 / 0.956	35.1 / 0.973	35.0 / 0.972
	11	32.0 / 0.904				

Table 3. PSNR and SSIM (PSNR / SSIM) for diffusion models trained with (ID=2) and without (ID=11) the standard training scheme for real MRI data, evaluated on the actual noise level of each image.

	ID	1 iteration	Det	Stoch	Det + Reg	Stoch + Reg
	Noisy	31.5 / 0.723				
	2	40.4 / 0.981	39.5 / 0.980	37.4 / 0.973	40.0 / 0.981	39.9 / 0.981
	11	41.4 / 0.983				

sampling strategy demonstrates a decreasing quality in denoised images as the sampling process progresses. The reverse sampling process not only compromises the quality of the denoised prediction but also entails a considerable increase in inference time, as the network has to be applied at each iteration.

The diffusion model that was limited to a specific range of noise levels during training (ID=11) demonstrates superior performance compared to diffusion models trained using the standard scheme, when assessed within that specific noise level range. This advantage stems from the model’s specialization in handling the noise level encountered during training. In contrast, when applied to a higher noise level, the diffusion model trained with the standard scheme, exposed to a wider range of noise levels during training, exhibits better performance.




Fig. 2. The evaluation metrics were computed for the denoised network prediction at every step of the reverse diffusion process for training configuration ID = 9, noise level $\sigma = 0.7$, averaged over a randomly selected subset of 50 test images of the BrainWeb data set.




Fig. 3. Exemplary result image demonstrating the negative impact of the stochastic sampling scheme. Although the initial model prediction adeptly captures fine image details, the final prediction reveals an alteration in image content.

Despite the ideal conditions in our experiments, marked by access to noise-free images and precise noise level regulation, our analysis demonstrated the inherent challenges of applying diffusion models to medical image denoising. This highlights a critical concern, as the method’s performance may be further compromised in real-world applications. In practical scenarios, accurately estimating the noise level is a crucial factor in determining the matching state. The spatially variant nature of noise in MRI, especially in parallel imaging contexts, adds an extra layer of complexity. This implies that the method’s limitations could be magnified, casting further doubt on its viability for real-world applications.

Although our dataset was limited due to the necessity of using noise-free data, this constraint mirrors the typical conditions in real-world medical applications. Future work will focus on expanding the analysis to include more extensive datasets, leveraging advanced training schemes and pre-trained models to further validate our results.

5 Conclusion

Our investigation revealed certain advantages in the *training* scheme of diffusion models, contributing to enhanced model robustness. However, we demonstrated that the *reverse* process, employing existing sampling techniques, results in a degradation in denoising quality and image fidelity with each step. The presented findings raise questions regarding the practicality of applying diffusion models for medical image denoising, where data fidelity is of utmost importance.

Disclosure of Interests. L.P. receives PhD funding from Siemens Healthineers AG. F.W. and T.W. are employees of Siemens Healthineers AG. All other authors have no competing interests to declare that are relevant to the content of this article.

References

1. Aja-Fernández, S., Vegas-Sánchez-Ferrero, G., Tristán-Vega, A.: Noise estimation in parallel MRI: GRAPPA and SENSE. *Journal of Magnetic Resonance Imaging* **32**(3), 281–290 (2014)
2. Anand, C., Sahambi, J.: MRI denoising using bilateral filter in redundant wavelet domain. In: *TENCON 2008-2008 IEEE Region 10 Conference*. pp. 1–6. IEEE (2008)
3. Aubert-Broche, B., Griffin, M., Pike, G., Evans, A., Collins, D.: Twenty new digital brain phantoms for creation of validation image data bases. *IEEE Transactions on Medical Imaging* **25**, 1410–1416 (2006)
4. Batson, J., Royer, L.: Noise2self: Blind denoising by self-supervision. In: *International Conference on Machine Learning*. pp. 524–533. PMLR (2019)
5. Beister, M., Kolditz, D., Kalender, W.A.: Iterative reconstruction methods in X-ray CT. *Physica medica* **28**(2), 94–108 (2012)
6. Chen, T.: On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972* (2023)
7. Chung, H., Lee, E.S., Ye, J.C.: MR image denoising and super-resolution using regularized reverse diffusion. *IEEE Transactions on Medical Imaging* **42**(4), 922–934 (2022)
8. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated MRI. *Medical image analysis* **80**, 102479 (2022)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
10. Kannengiesser, S., Mailhe, B., Nadar, M., Huber, S., Kiefer, B.: Universal iterative denoising of complex-valued volumetric MR image data using supplementary information. *ISMRM* (2016)
11. Kaur, A., Dong, G.: A complete review on image denoising techniques for medical images. *Neural Processing Letters* pp. 1–44 (2023)
12. Kelm, Z.S., Blezek, D., Bartholmai, B., Erickson, B.J.: Optimizing non-local means for denoising low dose CT. In: *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. pp. 662–665. IEEE (2009)
13. Koonjoo, N., Zhu, B., Bagnall, G.C., Bhutto, D., Rosen, M.: Boosting the signal-to-noise of low-field MRI with deep learning image reconstruction. *Scientific Reports* **11**(1), 8248 (2021)

14. Krull, A., Buchholz, T., Jug, F.: Noise2Void-learning denoising from single noisy images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2129–2137 (2019)
15. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189* (2018)
16. Li, T., Feng, H., Wang, L., Xiong, Z., Huang, H.: Stimulating the diffusion model for image denoising via adaptive embedding and ensembling. *arXiv preprint arXiv:2307.03992* (2023)
17. Lustig, M., Donoho, D., Pauly, J.: Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine* **58**(6), 1182–1195 (2007)
18. Manjón, J.V., Carbonell-Caballero, J., Lull, J.J., García-Martí, G., Martí-Bonmatí, L., Robles, M.: MRI denoising using non-local means. *Medical image analysis* **12**(4), 514–523 (2008)
19. Marques, J., Simonis, F., Webb, A.: Low-field MRI: An MR physics perspective. *Journal of Magnetic Resonance Imaging* **49**(6), 1528–1542 (2019)
20. Pearl, N., Brodsky, Y., Berman, D., Zomet, A., Acha, A.R., Cohen-Or, D., Lischinski, D.: SVNR: Spatially-variant noise removal with denoising diffusion. *arXiv preprint arXiv:2306.16052* (2023)
21. Pfaff, L., Hossbach, J., Preuhs, E., Wagner, F., Arroyo Camejo, S., Kannengiesser, S., Nickel, D., Wuerfl, T., Maier, A.: Self-supervised MRI denoising: leveraging Stein’s unbiased risk estimator and spatially resolved noise maps. *Scientific Reports* **13**(1), 22629 (2023)
22. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 234–241. Springer (2015)
23. 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)
24. Soltanayev, S., Chun, S.: Training and refining deep learning based denoisers without ground truth data. *arXiv preprint arXiv:1803.01314* (2018)
25. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
26. Wagner, F., Thies, M., Gu, M., Huang, Y., Pechmann, S., Patwari, M., Ploner, S., Aust, O., Uderhardt, S., Schett, G., Christiansen, S., Maier, A.: Ultra low-parameter denoising: Trainable bilateral filter layers in computed tomography. *Med Phys* **49**(8), 5107–5120 (2022)
27. Wagner, F., Thies, M., Pfaff, L., Aust, O., Pechmann, S., Weidner, D., Maul, N., Rohleder, M., Gu, M., Utz, J., et al.: On the benefit of dual-domain denoising in a self-supervised low-dose CT setting. In: *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5. IEEE (2023)
28. Xiang, T.: DDM². <https://github.com/StanfordMIMI/DDM2> (2023)
29. Xiang, T., Yurt, M., Syed, A.B., Setsompop, K., Chaudhari, A.: DDM²: Self-supervised diffusion MRI denoising with generative diffusion models. *arXiv preprint arXiv:2302.03018* (2023)
30. Yang, C., Liang, L., Su, Z.: Real-world denoising via diffusion model. *arXiv preprint arXiv:2305.04457* (2023)