



Noise-Controllable Complex-Valued Diffusion Model for k-Space Data of Hyperpolarized ^{129}Xe Lung MRI Generation

Linxuan Han, Sa Xiao, Muhong Li, Jinghan Liu, and Xin Zhou^(✉)

State Key Laboratory of Magnetic Resonance Spectroscopy and Imaging, National Center for Magnetic Resonance in Wuhan, Wuhan Institute of Physics and Mathematics, Innovation Academy for Precision Measurement Science and Technology, Chinese Academy of Sciences–Wuhan National Laboratory for Optoelectronics, Huazhong University of Science and Technology, Wuhan, China
xinzhou@wipm.ac.cn

Abstract. Hyperpolarized ^{129}Xe lung magnetic resonance imaging (MRI) offers a method for visualizing human lung function. However, its long imaging time limits widespread research and clinical adoption. Deep learning has shown significant potential in addressing undersampled MRI reconstruction challenges. Yet, the clinical novelty of hyperpolarized ^{129}Xe lung MRI results in a particular lack of raw k-space data. To address this, we propose a Noise-Controllable Complex-Valued Diffusion Model (NC-CDM) to augment the available data from limited clinical training set. Specifically, complex-valued convolutional kernels replace traditional ones, enhancing feature extraction and data utilization efficiency by learning rich representations from k-space. In addition, a noise-controllable module is introduced to mitigate estimation biases caused by thermal noise during MRI acquisition in the training phases. Experiments compare the proposed NC-CDM with other state-of-the-art models. Fréchet Inception Distance (FID) and Inception Score (IS) metrics show that our method obtains higher image quality. The generated data, mixed with real data, are subsequently applied to downstream MRI reconstruction task using two deep learning-based MRI reconstruction methods: CasNet and KIKI-net. The results show that reconstruction networks trained on our generated data exhibit superior reconstruction performance.

Keywords: Hyperpolarized ^{129}Xe MRI · k-Space Data Generation · Diffusion Model

1 Introduction

Hyperpolarized ^{129}Xe lung MRI provides a non-invasive and non-ionizing method for evaluating pulmonary function, showing potential for early detection

L. Han and S. Xiao—Equal contribution.




Fig. 1. The framework of our work comprises two steps: (a) Using the proposed Noise-Controllable Complex-Valued Diffusion Model (NC-CDM) to generate MRI k-space data. These data are then mixed with real data to create the training set for downstream reconstruction task. (b) The reconstruction task is used as a downstream task to validate the effectiveness of the generation methods.

and treatment of lung diseases [1]. However, due to the decay of hyperpolarized ^{129}Xe magnetization, faster image acquisition is necessary for practical clinical application [2, 3]. The integration of deep learning techniques has proven highly successful in enhancing the process of MRI reconstruction [4, 5]. Nevertheless, its use in medical imaging is limited by the scarcity of data [6]. This limitation is particularly significant in the emerging field of hyperpolarized ^{129}Xe lung MRI, where raw k-space data for reconstruction research is lacking [7]. Recent advancements in denoising diffusion probabilistic models (DDPM) have shown remarkable success in data synthesis tasks [8–15]. Specifically, DDPM have been utilized to address data scarcity in diagnostic tasks by reliably generating images with specific data structures [9] and pathological features [10], offering a novel solution for data augmentation. However, current generative studies mainly focus on generating spatial domain data, with limited research on frequency domain k-space data generation [7]. Additionally, existing studies have predominantly used a training set with high signal-to-noise ratios, overlooking the impact of noise on the generated samples [16, 17].

To alleviate the problem of data dependence in undersampled hyperpolarized ^{129}Xe lung MRI reconstruction, this study introduces a Noise-Controllable Complex-Valued diffusion model (NC-CDM). This model is specifically designed to synthesize k-space data, incorporating crucial phase information in the process, as illustrated in Fig. 1(a). Unlike conventional DDPM [18], our approach utilizes complex-valued convolution kernels [19] to better capture complex-valued space information, thereby improving feature extraction and utilization of complex-valued data. Furthermore, Gaussian noise is present in the measured k-space due to thermal fluctuations stemming from hardware and background factors [16]. The presence of thermal noise in measurements introduces deviations in the noise estimation during the reverse process of the diffusion model. These deviations become particularly pronounced as the process approaches the noise-free data distribution. Consequently, when MRI measurements are affected by thermal noise, the DDPM-based data generation methods may produce poor-quality results or even fail [7, 16]. To mitigate this issue, we have developed a noise-controllable module to adjust the predicted noise of the diffusion model during training and facilitate noise controlled data generation during the sampling phase. To validate the effectiveness of the generated data in the downstream

task, we employed a mixed dataset (consisting of both real and generated data) to improve the performance of undersampled hyperpolarized ^{129}Xe lung MRI reconstruction, as depicted in Fig. 1(b). Our contribution are three-fold:

- To the best of our knowledge, this is the first time numerous virtual k-space data have been directly synthesized using a diffusion model for the purpose of training a follow-up model for undersampled hyperpolarized ^{129}Xe lung MRI reconstruction.
- The proposed NC-CDM employs complex-valued convolutions for effective feature extraction and includes a noise-controllable module to mitigate thermal noise.
- Integration of the data generated by the NC-CDM into the training set results in an improvement in the accuracy of reconstruction task on the test set. This demonstrates the efficacy of generated data in downstream task.

2 Method

2.1 Overview of the Proposed Method

The single-coil k-space data, x from hyperpolarized ^{129}Xe lung MRI data, is always corrupted by thermal noise $\eta \sim \mathcal{N}(\mu, \sigma^2)$, where $x, \mu, \sigma \in \mathbb{C}$. Let $x_0 \in \mathbb{C}$ be the clean k-space data, then $x = x_0 + \eta$. Fluctuations in noise levels affect the predetermined noise for training a diffusion model, leading to decreased performance. To address this issue and enhance the quality of generated k-space data, we propose the Noise-Controllable Complex-Valued Diffusion Model (NC-CDM), which is based on a standard UNet architecture within the DDPM. The data is first preprocessed to deal with the thermal noise. Then, traditional convolution is replaced with the complex-valued convolution, and a noise-controllable module is incorporated to address the predicted noise in the diffusion model.

Data Preprocess. To mitigate the impact of thermal noise in the subsequent calculations of NC-CDM, the μ is eliminated from x during preprocessing, as shown in Fig. 2(b). The x is first subjected to inverse Fourier transformation f' to the image domain $y = f'(x)$, followed by the calculation image domain noise from the background in the four image corners. Subsequently, the image domain noise is transformed back into k-space noise $\eta = \eta_R + j\eta_I$, where the real noise is $\eta_R \sim \mathcal{N}(\mu_R, \sigma_R^2)$ and imaginary noise is $\eta_I \sim \mathcal{N}(\mu_I, \sigma_I^2)$, respectively. To finalize preprocessing, we subtract the mean noise (μ_R for real, μ_I for imaginary) from the corresponding parts of the image. The processed data, x'_0 , serves as NC-CDM input at $t = 0$, where $x'_0 = x - \mu = x_0 + \eta', \eta' \sim \mathcal{N}(0, \sigma^2)$.

Noise Predictor of Complex-Valued Diffusion Model. Figure 2(c) presents the noise prediction framework of the NC-CDM at step t . In the NC-CDM, the input is the k-space data. Through the optimization of the loss function (Eq. (3)), the output of the NC-CDM is the predicted noise at step t . After




Fig. 2. Detail of proposed Noise-Controlable Complex-Valued Diffusion Model (NC-CDM). (a) Flowchart of proposed NC-CDM: NC-CDM generates k-space data through diffusion forward and reverse processes. (b) Data preprocess: Performing inverse Fourier transformation on x to obtain y , estimating background noise from the real and imaginary parts of y in all four corners, employing Fourier transform to get η on k-space and subtracting μ from x . (c) Noise predictor of NC-CDM: Utilizing complex-valued convolution and correcting for the estimated noise at time step t . x_t is the complex-valued input at step t .

T steps, by sampling from the noise assigned by the output of the network, we can derive the generated data. In the NC-CDM, the complex-valued denoise network utilizes a UNet-based model. It consists of three modules, including encoder, decoder, and skip connection. In our complex-valued diffusion model, all convolutions are carried out through complex-valued operations to extract detailed feature representations in a complex-valued space. The complex-valued convolution can be formulated as:

$$\begin{bmatrix} \text{Re}(W^{\otimes c}) \\ \text{Im}(W^{\otimes c}) \end{bmatrix} = \begin{bmatrix} \mathbf{X} & -\mathbf{Y} \\ \mathbf{Y} & \mathbf{X} \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} \quad (1)$$

where $\text{Re}()$ and $\text{Im}()$ represent the real and imaginary data, respectively. $\mathbf{a}$ is the real part, and $\mathbf{b}$ is the imaginary part. The convolution kernel denotes as $W = X + jY$, where X represents the real convolution, and Y represents the imaginary convolution. $\otimes$ is the convolutional operator.

Noise-Controlable Module. The noise of hyperpolarized ^{129}Xe in lung MRI is high. This thermal noise infiltrates the diffusion model-based generation process via the data consistency term, disrupting the predefined noise schedule used for training the reverse process. For example, at time step t , $x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\bar{\epsilon}_t$ with clean input x_0 and DDPM parameter α , due to the influence of intrinsic thermal noise η' , the forward process from x'_0 to x_t is rewritten as:

$$\begin{aligned}
x_t &= \sqrt{\alpha_t}x_{t-1} + \sqrt{1 - \alpha_t}\varepsilon_t \\
&= \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\bar{\varepsilon}_t + \sqrt{\bar{\alpha}_t}\eta', \eta' \sim \mathcal{N}(0, \sigma^2)
\end{aligned} \tag{2}$$

The difference between them is $\sqrt{\bar{\alpha}_t}\eta'$. Furthermore, this difference will increase as t decreases. Consequently, diffusion model-based generation methods may experience sub-optimal performance or even failure when the thermal noise of MRI is not negligible. To address this issue, we propose an explicit noise control module during the diffusion process. This module introduces a λ_t into the reverse process to correct errors from the training stage. In simple terms, we use λ_t to attenuate the predicted noise, and the loss function is rewritten as:

$$L_{t-1} = \mathbb{E}_{x_0, \varepsilon \sim \mathcal{N}(0, I)} \left[\|\varepsilon - \lambda_t \varepsilon_\theta(x_t, t)\|^2 \right], \lambda_t = \frac{\sqrt{1 - \bar{\alpha}}}{\sqrt{1 - \bar{\alpha}_t} + \sqrt{\alpha_t}\sigma} \tag{3}$$

The backward process can be written as follows:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \varepsilon_\theta(x_t, t) + \sqrt{\alpha_t} \eta_g \right), \eta_g \sim \mathcal{N}(0, \sigma_g^2) \tag{4}$$

$$x_g = x_0 + \mu_g \tag{5}$$

where x_g is the generated image, μ_g and σ_g are the mean and variance that control the generated image.

2.2 Evaluation

The Fréchet Inception Distance (FID) [20] and Inception Score (IS) [21] are commonly used to assess the visual quality of generative models. However, these metrics can unfairly penalize non-GAN models, and IS may yield overly optimistic results for methods with sampling modifications. To comprehensively evaluate the impact of synthetic samples on downstream task, we trained CasNet [4] and KIKI-Net [5] on a training set comprising a mixture of real and generated images. CasNet is a network specifically designed for reconstructing hyperpolarized ^{129}Xe lung MRI, while KIKI-Net is a general network for reconstructing magnetic resonance images. The networks were trained for image reconstruction with acceleration factors (AF) of 4 and 6, and their performance was evaluated on a test set of real images using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM). Our method was extensively compared with StyleGAN [22], VAE [23], DDPM [18], and DDIM [24] in terms of image generation similarity and reconstruction accuracy against real images. The NC-CDM involves generating in the k-space domain, as opposed to directly generating time-domain images as seen in other approaches. The NC-CDM demonstrates higher reliability and wider acceptance compared to recently proposed networks.

2.3 Implementation Details

The model is trained using a batch size of 16, with input images resized to 96×96 . The SGD optimizer is employed with an initial learning rate of 0.0001 for training over 3000 epochs. For the training of CasNet and KIKI-net, we utilize the SGD optimizer with an initial learning rate of 0.001 for 200 epochs. The source code is available at https://github.com/TmpAccount25/NC_CDM.

3 Experiments

Dataset. In the study, a total of 94 subjects (including healthy subjects and patients) were enrolled, all providing informed consent. Pulmonary ^{129}Xe imaging parameters included a 3D bSSFP sequence with a matrix size of 96×84 and 24 slices. To address the impact of edge information, only the 16 middle slices from each subject were selected for model training and validation. The dataset was divided by subjects, with 75 patients randomly selected as the training set and the remaining 19 patients as the test set. The generation model was trained on the training set and evaluated on the test set. Subsequently, the models for downstream reconstruction task was trained by mixing the generated data with samples from the training set, and then tested on the test set.

Comparisons with State-of-the-Art Methods. A comparison with state-of-the-art methods is presented in Table 1. Mixing images generated by StyleGAN and VAE with real images did not improve the reconstruction accuracy compared to the baseline (Real Image), due to poor reconstruction quality and significant distribution mismatch with real images. The StyleGAN model exhibited the most significant decline in performance when accelerated by a factor of six, with a decrease of 2.35 in PSNR and 0.1030 in SSIM, while VAE’s performance

Table 1. Quantitative evaluation was conducted on the test set. The first line presents the reconstruction results using only real images as the training set, serving as the baseline. The highest scores are highlighted in bold.

Method	FID	IS	Down-stream task (MR Reconstruction)							
			CasNet				KIKINet			
			$\times 4$		$\times 6$		$\times 4$		$\times 6$	
			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Real Image	—	—	29.71	0.8561	24.05	0.6222	31.8	0.8452	26.48	0.7821
StyleGAN	204.55	1.112	28.89	0.8481	23.61	0.6125	30.29	0.8722	24.13	0.6791
VAE	189.82	1.320	29.32	0.8699	23.72	0.6013	30.73	0.8791	25.72	0.7195
DDPM	160.74	2.091	31.31	0.8724	25.14	0.7094	32.31	0.8940	26.48	0.7871
DDIM	163.70	2.101	30.87	0.8764	25.61	0.7143	32.10	0.8900	26.56	0.7644
Ours	110.12	2.411	33.52	0.9078	27.33	0.8515	34.21	0.9205	30.79	0.8711





Fig. 3. Visual result of the generated data. To better illustrate the quality of the generated images, we present the generated k-space data, which has been transformed into the image domain using inverse Fourier transformation. The final row represents the real data.

decrease was less pronounced. In contrast, DDPM and DDIM showed minor improvements in accuracy under an AF of 4 in CasNet. Specifically, DDPM and DDIM demonstrated increases of 1.6 and 1.16 in PSNR, and 0.0163 and 0.0203 in SSIM, respectively. When an AF of 6 was applied, DDPM and DDIM achieved increases of 1.09 and 1.56 in PSNR, respectively, and improvements of 0.0872 and 0.0921 in SSIM, respectively. However, under the AF of 6, there was no significant improvement in performance for DDPM and DDIM. Our proposed method outperformed other approaches, achieving an FID score of 110.12 and an IS score of 2.411, which suggests a closer match between the generated samples and real sample distribution. Furthermore, our method excelled in the reconstruction task, showing improvements in PSNR by 3.81 and 3.28, and in SSIM by 0.0517 and 0.2293 under the AF of 4 and 6 in CasNet, respectively. Similarly, in KIKI-net under the AF of 4, PSNR improved by 2.41 and 4.31, and SSIM by 7.53 and 8.9 under the AF of 6. These results demonstrate the enhanced adaptability of proposed method to downstream task, addressing the challenge of limited samples. The visual results of images generated using different models are shown in Fig. 3. To visually showcase the quality of the generated images, we display the k-space data converted into the image-domain using inverse Fourier transformation. Our method produces samples that closely resemble real samples in terms of image contrast and structural details.

Ablation Study: To evaluate the efficacy of the proposed complex-value network with a noise-controllable module in enhancing data quality generation, several ablation experiments were conducted. Table 2 shows all the results of these ablation studies. The initial row represents the baseline DDPM, which is the DM network without complex-valued convolutions and noise-controllable module.

Table 2. Quantitative evaluation across different modules. The highest scores are highlighted in bold.

Method		FID	IS	Down-stream task (MR Reconstruction)							
				CasNet				KIKINet			
				$\times 4$		$\times 6$		$\times 4$		$\times 6$	
Complex	NC			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
$\times$	$\times$	160.74	2.091	31.31	0.8724	25.14	0.7094	32.31	0.8940	26.48	0.7871
$\checkmark$	$\times$	157.81	2.196	32.21	0.8941	26.21	0.7909	32.98	0.9033	27.70	0.7821
$\checkmark$	$\checkmark$	110.12	2.411	33.52	0.9078	27.33	0.8515	34.21	0.9205	30.79	0.8711

**Fig. 4.** Ablation of the number of generated images by proposed Noise-Controllable Complex-Valued Diffusion Model (NC-CDM).

Firstly, replacing traditional convolutions with complex-valued convolutions led to enhancements in all metrics compared to the baseline network, showcasing the capability of complex-valued convolutions in extracting intricate information beyond temporal features. Moreover, the incorporation of the noise control module notably boosted the accuracy of all metrics, with FID of 110.12 and IS of 2.411. When employing CasNet with the AF of 4 and 6, SSIM improved from 0.8941 to 0.9078 and from 0.7909 to 0.8515, respectively. Furthermore, when utilizing KIKI-net with the AF of 4 and 6, SSIM improved significantly, from 0.9033 to 0.9205, and from 0.7821 to 0.8711, respectively. The reconstruction accuracy under different acceleration factors also reached its peak. This underscores the module’s effectiveness in generating low signal-to-noise ratio images. To verify the influence on number of generated image on the downstream task, Fig. 4 demonstrates the correlation between the quantity of artificially generated data integrated into the raw data and the reconstruction accuracy. With an increase in the number of generated images, CasNet and KIKI-net attain PSNR values of 33.52 and 34.21 for the reconstruction task with an AF of 4, and 27.33 and 30.79 for the reconstruction task with an AF of 6.

4 Conclusion

In conclusion, this study introduces an innovative approach to mitigate sample scarcity in medical image analysis. Unlike traditional image-domain generation methods, we utilize a complex-valued diffusion model to directly synthesize k-space data for hyperpolarized ^{129}Xe lung MRI. To mitigate the effects of thermal noise in MRI with low signal-to-noise ratios, a noise-controllable module is integrated, which reduces noise impact during the forward process and enhancing the alignment of generated samples with real data distributions in contrast and tissue detail. We evaluate the proposed method's efficacy by comparing it with state-of-the-art techniques using FID and IS metrics and assessing its impact on downstream task through validation reconstructions in acceleration factors of 4 and 6 with CasNet and KIKI-Net, respectively. Experimental results demonstrate the superior performance of our approach, highlighting its potential for downstream applications and advancing image synthesis techniques. Although the hyperpolarized ^{129}Xe lung MRI data we used were collected using a single coil, the proposed method is not limited to the MRI acquisition method and is applicable to the acquisition of multicoil and single coil MRI.

Acknowledgements. This work was supported by National key Research and Development Program of China (2022YFC2410000), National Natural Science Foundation of China (82127802, 21921004), Key Research Program of Frontier Sciences (ZDBS-Y-JSC004), the Strategic Priority Research Program of the Chinese Academy of Sciences (XDB0540000, XDC0170000), Hubei Provincial Key Technology Foundation of China (2021ACA013, 2023BAA021), and Natural Science Foundation of Hubei Province (2023AFB1061).

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

References

1. Li, H., et al.: Damaged lung gas exchange function of discharged COVID-19 patients detected by hyperpolarized ^{129}Xe MRI. *Sci. Adv.* **7**, eabc8180 (2021)
2. Aviles-Rivero, A., Debroux, N., Williams, G., Graves, M., Schönlieb, C.: Compressed sensing plus motion (CS+M): a new perspective for improving undersampled MR image reconstruction. *Magn. Reson. Med.* **68**, 101933 (2021)
3. Xiao, S., et al.: Considering low-rank, sparse and gas-inflow effects constraints for accelerated pulmonary dynamic hyperpolarized ^{129}Xe MRI. *J. Magn. Reson. Imag.* **290**, 29–37 (2018)
4. Duan, C., et al.: Fast and accurate reconstruction of human lung gas MRI with deep learning. *Magn. Reson. Med.* **82**(6), 2273–2285 (2019)
5. Eo, T., et al.: KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magn. Reson. Med.* **80**(5), 2188–2201 (2018)
6. Luo, Y., Yang, Q., Fan, Y., Qi, H., Xia, M.: Measurement guidance in diffusion models: insight from medical image synthesis. *IEEE Trans. Pattern Anal. Mach. Intell.* **46**(12), 7983–7997 (2024)

7. Li, Z., et al.: Encoding enhanced complex CNN for accurate and highly accelerated MRI. *IEEE Trans. Med. Imaging* **43**(5), 1828–1840 (2024)
8. Fan, Y., et al.: A survey of emerging applications of diffusion probabilistic models in MRI. *Meta-Radiology* **2**(2), 100082 (2024)
9. Li, S., Lin, Y., Chen, H., Cheng, K.: Iterative online image synthesis via diffusion model for imbalanced classification. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 371–381. Springer (2024)
10. Cobb, B., Cook, G., Reader, A.: Improved classification learning from highly imbalanced multi-label datasets of inflamed joints in [99mTc]maraciclaitide imaging of arthritic patients by natural image and diffusion model augmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 339–348. Springer (2024)
11. Pan, S., et al.: 2D medical image synthesis using transformer-based denoising diffusion probabilistic model. *Phys. Med. Biol.* **68**(10), 105004 (2023)
12. Han, K., Xiong, Y., You, C., et al.: Medgen3d: a deep generative framework for paired 3D image and mask generation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 759–769. Springer (2023)
13. Dar, S., et al.: Investigating data memorization in 3D latent diffusion models for medical image synthesis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 56–65. Springer (2023)
14. Zhang, Y., et al.: Texture-preserving diffusion model for CBCT-to-CT synthesis. *Med. Image Anal.* **99**, 103362 (2025)
15. Chen, Y., et al.: Spherical harmonics-based deep learning achieves generalized and accurate diffusion tensor imaging. *IEEE J. Biomed. Health Inform.* **29**(1), 456–467 (2025)
16. Huang, S., et al.: Magnetic resonance image denoising for Rician noise using a novel hybrid transformer-CNN network (HTC-net) and self-supervised pretraining, pp. 498–508. Springer (2024)
17. Shi, S., et al.: Noise level adaptive diffusion model for robust reconstruction of accelerated MRI. *Med. Phys.* **52**(3), 1643–1660 (2025)
18. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Adv. Neural. Inf. Process. Syst.* **33**, 6840–6851 (2020)
19. Xie, Y., Li, Q.: Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 655–664. Springer (2022)
20. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
21. Salimans, T., et al.: Improved techniques for training GANs. In: *Advances in Neural Information Processing Systems*, vol. 29 (2016)
22. Karras, T., et al.: Alias-free generative adversarial networks. *Adv. Neural. Inf. Process. Syst.* **34**, 852–863 (2021)
23. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. *arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114)* (2013)
24. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint [arXiv:2010.02502](https://arxiv.org/abs/2010.02502)* (2020)