

Noise Level Adaptive Diffusion Model for Robust Reconstruction of Accelerated MRI

Shoujin Huang¹ , Guanxiong Luo² , Xi Wang³ , Ziran Chen¹ ,
Yuwan Wang¹ , Huaishui Yang¹ , Pheng-Ann Heng³ ,
Lingyan Zhang⁴, and Mengye Lyu^{1*} 

¹ Shenzhen Technology University, Shenzhen, China

² Medical Center Göttingen University, Göttingen, Germany

³ The Chinese University of Hong Kong, Hong Kong, China

⁴ Longgang Central Hospital of Shenzhen, Shenzhen, China

*Corresponding authors: lvmengye@sztu.edu.cn

Abstract. In general, diffusion model-based MRI reconstruction methods incrementally remove artificially added noise while imposing data consistency to reconstruct the underlying images. However, real-world MRI acquisitions already contain inherent noise due to thermal fluctuations. This phenomenon is particularly notable when using ultra-fast, high-resolution imaging sequences for advanced research, or using low-field systems favored by low- and middle-income countries. These common scenarios can lead to sub-optimal performance or complete failure of existing diffusion model-based reconstruction techniques. Specifically, as the artificially added noise is gradually removed, the inherent MRI noise becomes increasingly pronounced, making the actual noise level inconsistent with the predefined denoising schedule and consequently inaccurate image reconstruction. To tackle this problem, we propose a posterior sampling strategy with a novel Noise Level Adaptive Data Consistency (Nila-DC) operation. Extensive experiments are conducted on two public datasets and an in-house clinical dataset with field strength ranging from 0.3T to 3T, showing that our method surpasses the state-of-the-art MRI reconstruction methods, and is highly robust against various noise levels. The code for Nila is available at <https://github.com/Solor-pikachu/Nila>.

Keywords: Image reconstruction · Diffusion model · Parallel imaging

1 Introduction

Magnetic Resonance Imaging (MRI) is a crucial means for medical diagnostics, yet with major drawbacks of long acquisition times and high operational costs. To accelerate MRI scans, a common approach is using multiple coils to acquire the k-space data without obeying the Nyquist sampling theorem, and then to

Shoujin Huang and Guanxiong Luo contribute equally to this work.




Fig. 1: General illustration of diffusion model-based MRI reconstruction. Note that the MRI data (measurements) not only contain signals but also noise due to hardware and subject-related thermal fluctuations. This can interfere with the reverse diffusion process, where artificial Gaussian noise is used for image initialization and gradually removed with a pre-defined schedule.

reconstruct the image with the incorporation of prior knowledge of it[1,2,3], such as sparsity and total variation[4,5]. Deep learning models have been utilized to learn and apply prior knowledge inherently within existing databases, including both supervised[6,7,8,9] and generative paradigms[10].

Recently, diffusion models have been applied to improve reconstruction, where the data consistency term of k-space is integrated as the guidance for the generative process[11,12,13,14,15], i.e., learned reverse process. In diffusion models[16,17], there are: 1) the forward process created with Markov chains where the noise at different levels is added to images that represents the distribution of data; 2) the learned reverse process where the images of data distribution are generated starting from a known distribution, i.e., Gaussian noise.

However, hardware and subject-related thermal fluctuations cause the presence of noise in measured k-space. As illustrated in Fig. 1, this inherent noise propagates into the diffusion model-based reconstruction process through the data consistency term and interferes with the pre-defined noise schedule used for training the reverse process. This is an issue in the late stage of the reverse process when approaching the noise-free data distribution. Consequently, diffusion model-based reconstruction methods may experience sub-optimal performance or even failure when the MRI measurement noise is not negligible, such as in the case of low-field MRI[18], functional MRI (fMRI), and diffusion-weighted imaging (DWI)[19].

In this work, a noise level adaptive data consistency (Nila-DC) operation is proposed for image reconstruction with a lambda rescale function (c.f. Fig. 2 (b)) that ensures robust guidance of k-space regardless of noise when using diffusion models. This approach is validated on multiple datasets of different field strengths, imaging sequences, and noise levels, including challenging scenarios such as accelerated DWI and low-field MRI. In the following sections, the disturbance of noise to the learned reverse process is formulated and experimental results using Nila-DC are presented.




Fig. 2: Overview of the proposed method. (a) The proposed data consistency (Nila-DC) operation. The computed gradient ($A^H A x_t - A^H y$) can be noisy due to MRI noise in y (c.f. Eqs. 7 and 8), and is therefore adjusted by an attenuation function (lambda). (b) The attenuation function (c.f. Eq. 10) used to rescale the DC gradient. t is the index of the reverse step. (c) The image reconstruction process, where Gaussian noise initialized x_t undergoes multi-step reverse diffusion process with the guidance from Nila-DC.

2 Methodology

Alike many diffusion model-based techniques, reconstruction is one of the samples from the posterior distribution which is formulated with two components: 1) the likelihood term to describe the data consistency; 2) the diffusion prior trained on image dataset[11,15,14,20]. We observed that the different noise level in many acquisition scenarios affects the pre-defined noise schedule for training a diffusion model and consequently leads to poor performance. In the below, we present a noise level adaptive method to address this issue in the context of using the diffusion denoised probabilistic model (DDPM) as a prior[21,22].

Reconstruction as Bayesian Inversion. Image reconstruction is approached as a Bayesian problem, where the posterior $p(\mathbf{x} | \mathbf{y})$ is determined by the measured k-space $\mathbf{y}$ and a diffusion prior $p(\mathbf{x})$ is trained on an image database[14,20]. Assuming the noise η normally distributed with zero mean and covariance matrix $\sigma_y^2 \mathbf{I}$, the likelihood $p(\mathbf{y} | \mathbf{x})$ for observing the $\mathbf{y}$ determined by $\mathbf{y} = \mathcal{A}\mathbf{x} + \eta$

and given the image $\mathbf{x}$ is given by a complex normal distributions

$$\begin{aligned} p(\mathbf{y} \mid \mathbf{x}) &= \mathcal{CN}(\mathbf{y}; \mathcal{A}\mathbf{x}, \sigma_y^2 \mathbf{I}) \\ &= (\sigma_y^2 \pi)^{-N_p} \exp(-\|\sigma_y^{-1} \cdot (\mathbf{y} - \mathcal{A}\mathbf{x})\|_2^2) , \end{aligned} \quad (1)$$

The forward operator $\mathcal{A}$ is composed of coil sensitivity maps $\mathcal{S}$, a two-dimensional Fourier transform $\mathcal{F}$, and a k-space sampling operator $\mathcal{P}$. Using Bayes' formula, we obtain each intermediate distribution in Fig. 2.

$$p(\mathbf{x}_t \mid \mathbf{y}) \propto p(\mathbf{x}_t) p(\mathbf{y} \mid \mathbf{x}_t) \quad (2)$$

Starting with the initial Gaussian noise $q(\mathbf{x}_N) \sim \mathcal{CN}(0, \mathbf{I})$ at $t = N$, $p(\mathbf{x}_t)$ is obtained with transition kernels $\{p(\mathbf{x}_t \mid \mathbf{x}_{t+1})\}_{t \leq N}$

$$p(\mathbf{x}_t) \propto p(\mathbf{x}_t \mid \mathbf{x}_{t+1}) \cdot \dots \cdot p(\mathbf{x}_{N-1} \mid \mathbf{x}_N) \cdot q(\mathbf{x}_N). \quad (3)$$

With the learned transition kernel $p_\theta(\mathbf{x}_t \mid \mathbf{x}_{t+1})$ in denoising diffusion models, samples are simulated from $p_\theta(\mathbf{x}_t \mid \mathbf{y})$ using the reverse diffusion[17]:

$$\mathbf{x}_{t-1} = \sqrt{\bar{a}_{t-1}} \nabla_{\mathbf{x}_t} \log p_\theta(\mathbf{x}_t \mid \mathbf{y}) + \sqrt{1 - \bar{a}_{t-1}} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{CN}(0, \mathbf{I}), \quad (4)$$

where $\bar{a}_{t-1} := \prod_{s=1}^{t-1} a_s$, $a_{t-1} := 1 - \beta_{t-1}$, β_t following a pre-defined schedule $\{\beta_0, \beta_1, \dots, \beta_T\}$ used in the forward process[16], and

$$\nabla_{\mathbf{x}_t} \log p_\theta(\mathbf{x}_t \mid \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t \mid \mathbf{x}_{t-1}) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} \mid \mathbf{x}_t). \quad (5)$$

Further, we have

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t \mid \mathbf{x}_{t-1}) = \frac{1}{\sqrt{\bar{a}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{a}_t} (\epsilon_\theta(\mathbf{x}_t, t))) \quad (6)$$

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y} \mid \mathbf{x}_t) = -(\mathcal{A}^H \mathcal{A} \mathbf{x}_t - \mathcal{A}^H \mathbf{y}) / \sigma_y^2 \quad (7)$$

Eq. 6 represents the diffusion prior term, and Eq. 7 the data consistency term. Note that Eq. 7 can be derived without approximation errors[23,24] because of the MRI specific assumption in Eq. 1. Initially, when t is large, $\mathbf{x}_t$ has less structural information and higher noise level. When t approaches 0, $\mathbf{x}_t$ is expected to recover structural information and have decreased noise level. However, it is important to note that $\mathbf{y}$ is noisy in nature.

Noise Level Adaptive Data Consistency. During the computation of data consistency, the adjoint of the k-space represents the under-sampled k-space with additive noise. Then we have

$$\mathcal{A}^H \mathbf{y} = \mathcal{A}^H \mathcal{A} \bar{\mathbf{x}} + \mathcal{A}^H \sigma_y, \quad (8)$$

where $\bar{\mathbf{x}}$ is the noise-free image and σ_y is relative to noise level of images in the training dataset (note that σ_y could be estimated via a quick calibration scan or

from image background areas). This σ_y propagates through the diffusion reverse process, adding extra noise σ_y in the image $\mathbf{x}_t$. According to Eqs. 4, 5, 6, 7, 8 and using gradient based sampling method, we have

$$\begin{aligned} \mathbf{x}_{t-1} = & \sqrt{\bar{a}_{t-1}} \left[\frac{1}{\sqrt{\bar{a}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{a}_t} \epsilon_\theta(\mathbf{x}_t, t)) - (\mathcal{A}^H \mathcal{A} \mathbf{x}_t - \mathcal{A}^H \mathcal{A} \bar{\mathbf{x}}) \right] \\ & + \sqrt{\bar{a}_{t-1}} \mathcal{A}^H \sigma_y + \sqrt{1 - \bar{a}_{t-1}} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{CN}(0, \mathbf{I}), \end{aligned} \quad (9)$$

Although the introduction of noisy measurements by data consistency operation may not cause problems during the early stage of the reverse process, as the noise in x_t is gradually removed, σ_y from noisy measurement will become dominant at a certain stage, disturbing the pre-defined denoising schedule. Hence, we propose a new noise level adaptive data consistency operation to robustly utilize of guidance information from MRI data while minimizing the impact of MRI noise. From the pre-defined denoising schedule β_t in $\{\beta_0, \beta_1, \dots, \beta_T\}$, we know that the noise level at $\mathbf{x}_{t-1}$ is $\sqrt{1 - \bar{a}_{t-1}}$; According to Eq. 9, the additive noise propagating from the data consistency operation into the diffusion reverse process is $\sqrt{\bar{a}_{t-1}} \mathcal{A}^H \sigma_y$. To keep denoising schedule β_t unchanged, we introduce a rescale function (Fig. 2(b)) to keep/decay the data consistency term. For simplicity, we use a linear attenuation function $\lambda_t = kt + b$, starting with γ and linearly decaying to 0 at the final moment:

$$\lambda_t = \begin{cases} 1 & , \sqrt{\bar{a}_{t-1}} \mathcal{A}^H \sigma_y < \sqrt{1 - \bar{a}_{t-1}} \\ kt + b & , \sqrt{\bar{a}_{t-1}} \mathcal{A}^H \sigma_y > \sqrt{1 - \bar{a}_{t-1}} \end{cases}, \quad (10)$$

k and b can be computed once σ_y and γ have been defined. In practice, step size is scaled by λ_t . Finally, each reverse step in the proposed MRI reconstruction is defined by:

$$\begin{aligned} \mathbf{x}_{t-1} = & \sqrt{\bar{a}_{t-1}} \left[\frac{1}{\sqrt{\bar{a}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{a}_t} \epsilon_\theta(\mathbf{x}_t, t)) - \lambda_t (\mathcal{A}^H \mathcal{A} \mathbf{x}_t - (\mathcal{A}^H \mathcal{A} \bar{\mathbf{x}} + \mathcal{A}^H \sigma_y)) \right] \\ & + \sqrt{1 - \bar{a}_{t-1}} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{CN}(0, \mathbf{I}). \end{aligned} \quad (11)$$

3 Experiments

Datasets and Baselines. We evaluated our approach on three datasets: 1) The public fastMRI dataset[25]. From the official validation set, we randomly selected 20 samples for each contrast (T1, T1 post contrast, T2, and FLAIR) to form a test set of 80 samples. 2) The public low-field dataset M4Raw[26]. Based on its v1.6 release, 25 individuals from the test set, each with T1, T2, and FLAIR data, were used to form another test set (75 samples in total). 3) An additional in-house clinical dataset. With written informed consent obtained, T1 and FLAIR data were collected from 35 patients with white matter lesions from a local hospital using a 3T MRI Siemens scanner. The proposed method Nila is compared with L1-wavelet SENSE reconstruction[27], diffusion model-based

MRI reconstruction method CSGM[11], Spreco[14], AdaDiff[15]. Following[25], we used equidistant undersampling masks for all experiments because of ease of implementation on MRI machines.

Training and Evaluation Details. The multi-coil images from fastMRI[25] brain training set is coil-combined using ESPIRiT[27] and then used to train our model, using code from guided diffusion[22]. The last four noisy slices were removed from each volume, resulting in a total of $\sim 52k$ slice images for training. We used a learning rate of 0.0001, batch size of 8, and Adam optimizer. Note that this fastMRI trained model is also used to inference on the M4Raw and clinical datasets. For CSGM and Spreco, their official pre-trained models were used[11,14], whereas the AdaDiff model was trained on the same dataset as Nila.

To explore the robustness of different methods, the above models were tested both with and without adding extra Gaussian noise to the fully sampled MRI data. Peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) metrics were computed within the brain region, ignoring background reconstruction errors. Because low-field data is noisy, when evaluating on M4Raw, we set $\sigma_y = 0.05$ for Nila and employed multiple repetition averaged images [26] (6 repetitions for T1 and T2; 4 repetitions for FLAIR) as the reference to calculate the metrics.

Quantitative Results. Even without adding extra noise, our algorithm consistently achieved highest PSNR and SSIM scores on all datasets (Table 1). The differences between our method and others are statistically significant as shown in the supplementary materials. The second best method Spreco performed slightly better than CSGM and Adadiff, likely due to its early stopping mechanism.

When the noise level was respectively controlled (Table 2), all algorithms had reasonable results at low noise level ($\sigma = 0.025$). However, when noise increased, the reconstruction quality of L1-wavelet, CSGM, and Adadiff declined dramatically, sometimes even lower than Zero-filled reconstruction. In contrast, our algorithm remained resilient, achieving the highest scores for scenarios.

Table 1: Quantitative assessment for the three test datasets at different acceleration factors. No extra noise was added. The highest scores are marked in red, second highest blue.

Dataset	fastMRI				Clinical				M4Raw			
Acceleration factor	6 $\times$		8 $\times$		6 $\times$		8 $\times$		3 $\times$		4 $\times$	
Metric	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Zero-filled	26.12	82.27	25.41	80.36	27.07	82.40	26.33	80.34	27.15	87.81	25.97	85.43
L1-wavelet	30.04	89.43	27.31	84.55	31.27	89.60	28.35	84.59	28.86	87.50	28.02	86.57
CSGM	35.04	93.53	32.11	90.78	34.67	92.65	31.57	88.94	26.73	84.53	25.63	82.78
Spreco	32.24	90.70	28.65	85.28	32.27	89.84	28.77	84.07	29.77	90.34	28.88	88.68
AdaDiff	33.25	91.79	29.62	86.66	32.84	90.81	29.60	85.65	29.10	88.29	28.10	86.22
Nila (ours)	37.08	95.74	34.82	94.21	36.10	94.34	33.54	91.55	30.23	91.59	29.86	90.06

Table 2: Quantitative assessment under various noise levels at $6\times$ acceleration factor. Note that the sigma values represent the added noise to full sampled data. The highest scores are marked in red, second highest blue.

Dataset	fastMRI						Clinical					
noisy level	$\sigma = 0.025$		$\sigma = 0.05$		$\sigma = 0.1$		$\sigma = 0.025$		$\sigma = 0.05$		$\sigma = 0.1$	
Metric	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Zero-filled	26.06	81.40	25.90	79.22	25.32	74.08	27.01	81.73	26.84	79.96	26.23	75.31
L1-wavelet	29.07	85.86	26.78	78.15	21.28	64.98	30.15	85.86	27.55	77.51	21.22	63.77
CSGM	31.15	87.64	25.08	76.16	19.47	61.78	31.53	87.16	25.93	75.73	19.64	61.52
Spreco	31.31	89.03	29.62	84.93	25.89	74.23	31.32	88.03	29.58	83.88	26.41	74.10
AdaDiff	29.29	82.93	25.83	73.8	21.99	64.82	29.45	82.74	26.04	73.50	22.12	64.45
Nila (ours)	34.98	93.96	33.56	92.49	31.86	90.38	34.11	91.53	32.80	89.59	31.25	87.15




Fig. 3: Typical reconstructed images. The white numbers indicate the PSNR/SSIM scores. (a) $6\times$ acceleration on fastMRI. (b) $6\times$ acceleration on fastMRI with added Gaussian noise. (c) $6\times$ acceleration on the clinical dataset. Only the proposed method recovered the small white matte lesion as highlighted. (d) $4\times$ acceleration on M4Raw. CSGM and AdaDiff failed to provide usable images.

Qualitative Evaluation. Representative reconstruction results are presented in Fig. 3. Consistent with the qualitative metrics, our algorithm had the highest reconstruction quality regardless of noise levels, clearly revealing the anatomical structures as well as small white matter lesions.

Real-World DWI Reconstruction. Additionally, we experimented with prospectively undersampled 1.5T DWI data ($3\times$ acceleration, $b\text{-value} = 800$). Due to the nature of DWI, ground truth images are not possible to obtain here, but it is obvious that our method again achieved the highest SNR and the least reconstruction artifacts (Fig. 4).




Fig. 4: Reconstruction of prospectively accelerated DWI data. (a) Reconstruction from a single repetition. (b) Averaged from 3 orthogonal diffusion weighting directions each by 3 repetitions. σ_y was set to 0.05 for Nila, as estimated from zero-filled reconstruction.




Fig. 5: Reconstruction of Nila with added noise $\sigma = 0.05$ under different settings. (a) Fully sampled image with added $\sigma = 0.05$ noise. It is for visualization of the noise level. (b - e) Nila reconstruction using σ_y values of 0, 0.025, 0.05, and 0.1, respectively. (d) Fully sampled image without adding extra noise (clean reference). With σ_y set to 0, Nila falls back to conventional DDPM-based MRI reconstruction like Eq. 9. The white numbers indicate the PSNR/SSIM scores of the displayed images.

Ablation Study. The hyperparameter σ_y in Nila affects its reconstruction quality. In particular, when σ_y is set to 0, Eq. 10 becomes an identity function without attenuation effect, and Nila-DC falls back to conventional data consistency operation. To examine the contribution of the proposed Nila-DC operation

and the optimal settings of σ_y , we compared using different σ_y on the clinical dataset at $6\times$ acceleration with added Gaussian noise of $\sigma = 0.05$. As shown in Fig. 5, setting σ_y approximately equal to σ resulted in the best reconstruction. With underestimated σ_y , the reconstruction yield amplified noise, whereas with overestimated σ_y , the reconstruction exhibited over-smoothing. In practice, σ_y can be easily estimated from a quick calibration scan, e.g. using zero flip angle to acquire a few k-space lines, or background areas of zero-filled reconstruction.

4 Conclusion

We identify and address the issue that existing diffusion model-based reconstruction methods are sensitive to the MRI noise level by introducing a noise level adaptive data consistency operation for the reverse diffusion process, which permits robust guidance. The proposed method is comprehensively evaluated to demonstrate outstanding performance under various experimental conditions.

Acknowledgments. This work was supported in part by the National Natural Science Foundation of China under Grant 62101348, the Shenzhen Higher Education Stable Support Program under Grant 20220716111838002, the Natural Science Foundation of Top Talent of Shenzhen Technology University under Grants 20200208, and the Research Grants Council of the Hong Kong Special Administrative Region, China (Project Reference Number: T45-401/22-N).

Disclosure of Interests. The authors declare no competing interests.

References

1. Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A.: Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic resonance in medicine* **47**(6), 1202–1210 (2002)
2. Pruessmann, K.P., Weiger, M., Scheidegger, M.B., Boesiger, P.: Sense: sensitivity encoding for fast mri. *Magnetic resonance in medicine* **42**(5), 952–962 (1999)
3. Sodickson, D.K., Manning, W.J.: Simultaneous acquisition of spatial harmonics (smash): fast imaging with radiofrequency coil arrays. *Magnetic resonance in medicine* **38**(4), 591–603 (1997)
4. Lustig, M., Donoho, D.L., Santos, J.M., Pauly, J.M.: Compressed sensing mri. *IEEE signal processing magazine* **25**(2), 72–82 (2008)
5. Block, K.T., Uecker, M., Frahm, J.: Undersampled radial mri with multiple coils. iterative image reconstruction using a total variation constraint. *Magnetic Resonance in Medicine* **57**(6), 1086–1098 (2007)
6. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-end variational networks for accelerated mri reconstruction. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II* 23. pp. 64–73. Springer (2020)

7. Yiasemis, G., Sonke, J.J., Sánchez, C., Teuwen, J.: Recurrent variational network: a deep learning inverse problem solver applied to the task of accelerated mri reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 732–741 (2022)
8. Guo, P., Mei, Y., Zhou, J., Jiang, S., Patel, V.M.: Reconformer: Accelerated mri reconstruction using recurrent transformer. *IEEE transactions on medical imaging* (2023)
9. Huang, J., Fang, Y., Wu, Y., Wu, H., Gao, Z., Li, Y., Del Ser, J., Xia, J., Yang, G.: Swin transformer for fast mri. *Neurocomputing* **493**, 281–304 (2022)
10. Tezcan, K.C., Baumgartner, C.F., Luechinger, R., Pruessmann, K.P., Konukoglu, E.: Mr image reconstruction using deep density priors. *IEEE Transactions on Medical Imaging* **38**(7), 1633–1642 (2019). <https://doi.org/10.1109/TMI.2018.2887072>
11. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.: Robust compressed sensing mri with deep generative priors. *Advances in Neural Information Processing Systems* **34**, 14938–14954 (2021)
12. Fan, Y., Liao, H., Huang, S., Luo, Y., Fu, H., Qi, H.: A survey of emerging applications of diffusion probabilistic models in mri. *arXiv preprint arXiv:2311.11383* (2023)
13. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated mri. *Medical image analysis* **80**, 102479 (2022)
14. Luo, G., Blumenthal, M., Heide, M., Uecker, M.: Bayesian mri reconstruction with joint uncertainty estimation using diffusion models. *Magnetic Resonance in Medicine* **90**(1), 295–311 (2023)
15. Güngör, A., Dar, S.U., Öztürk, Ş., Korkmaz, Y., Bedel, H.A., Elmas, G., Ozbey, M., Çukur, T.: Adaptive diffusion priors for accelerated mri reconstruction. *Medical Image Analysis* p. 102872 (2023)
16. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
17. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
18. Marques, J.P., Simonis, F.F., Webb, A.G.: Low-field mri: An mr physics perspective. *Journal of magnetic resonance imaging* **49**(6), 1528–1542 (2019)
19. Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S.: Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic resonance in medicine* **34**(4), 537–541 (1995)
20. Luo, G., Wang, X., Blumenthal, M., Schilling, M., Rauf, E.H.U., Kotikalapudi, R., Focke, N., Uecker, M.: Generative image priors for mri reconstruction trained from magnitude-only images. *arXiv preprint arXiv:2308.02340* (2023)
21. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021)
22. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
23. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687* (2022)
24. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005* (2021)
25. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., De-fazio, A., Stern, R., Johnson, P., Bruno, M., et al.: fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839* (2018)

26. Lyu, M., Mei, L., Huang, S., Liu, S., Li, Y., Yang, K., Liu, Y., Dong, Y., Dong, L., Wu, E.X.: M4raw: A multi-contrast, multi-repetition, multi-channel mri k-space dataset for low-field mri research. *Scientific Data* **10**(1), 264 (2023)
27. Uecker, M., Tamir, J.I., Ong, F., Lustig, M.: The bart toolbox for computational magnetic resonance imaging. In: *Proc Intl Soc Magn Reson Med.* vol. 24, p. 1 (2016)