

Open World MRI Reconstruction with Bias-Calibrated Adaptation

Jiyao Liu¹, Shangqi Gao^{2✉}, Lihao Liu³, Junzhi Ning³, Jinjie Wei¹,
Junjun He³, Xiahai Zhuang^{1✉}, and Ningsheng Xu¹

¹ Fudan University, Shanghai, China

² University of Cambridge, Cambridge, United Kingdom

³ Shanghai Artificial Intelligence Laboratory, Shanghai, China

Abstract. Real-world MRI reconstruction systems face the open-world challenge: test data from unseen imaging centers, anatomical structures, or acquisition protocols can differ drastically from training data, causing severe performance degradation. Existing methods struggle with this challenge. To address this, we propose **BiasRecon**, a bias-calibrated adaptation framework grounded in the *minimal intervention principle: preserve what transfers, calibrate what does not*. Concretely, BiasRecon formulates open-world adaptation as an alternating optimization framework that jointly optimizes three components: (1) *frequency-guided prior calibration* that introduces layer-wise calibration variables to selectively modulate frequency-specific features of the pre-trained score network via self-supervised k-space signals, (2) *score-based denoising* that leverages the calibrated generative prior for high-fidelity image reconstruction, and (3) *adaptive regularization* that employs Stein’s Unbiased Risk Estimator to dynamically balance the prior-measurement trade-off, matching test-time noise characteristics without requiring ground truth. By intervening minimally and precisely through this alternating scheme, BiasRecon achieves robust adaptation with fewer than 100 tunable parameters. Extensive experiments across four datasets demonstrate state-of-the-art performance on open-world reconstruction tasks.

Keywords: Open-World Learning · MRI Reconstruction · Test-Time Adaptation

1 Introduction

MRI reconstruction aims to recover high-quality images from undersampled k-space measurements. Traditional approaches rely on hand-crafted priors such as wavelets [11], total variation [2], and low-rank constraints [20]. Deep learning has significantly advanced the field through supervised methods [1, 21], unsupervised approaches [16, 19], and generative prior-based methods [3, 4, 9, 18].

However, clinical deployment presents a fundamentally different challenge. In practice, reconstruction models must operate under **open-world conditions**, where the test data distribution is unknown and can differ drastically from




Fig. 1. Transferability analysis. (a) Open-world distribution shifts from anatomy, center, and modality changes. Models trained on FastMRI knee (inside dashed box) are tested on three OOD scenarios. (b) t-SNE visualization of features from bottom (deeper) and top (shallow) layers across different domains. (c) Frequency decomposition analysis of low-frequency and high-frequency components across different domains.

the training distribution. Such distribution shifts may arise from variations in anatomical structures, imaging protocols, scanner characteristics, or their combinations, as Fig. 1(a) illustrates. In clinical practice, such open-world scenarios are the norm rather than the exception.

Existing reconstruction methods experience significant performance degradation under such distribution shifts, as Fig. 5 demonstrates. We hypothesize that this degradation arises from two sources: (1) the learned image prior becomes misaligned with the target distribution, and (2) hyperparameters tuned on training data become suboptimal for new domains.

A natural question arises: how can we adapt a pre-trained model to unknown distributions with only single-image supervision? We argue that the key insight lies in recognizing that *not all learned knowledge degrades equally under distribution shifts*. Through transferability analysis (Sec. 2.2), we find that deeper layers and low-frequency components are domain-specific, while shallow layers and high-frequency components remain domain-agnostic, as visualized in Fig. 1(b,c). Therefore, efficient adaptation must jointly address *what* to adapt and *how much* to adapt.

Building on this insight, we propose **BiasRecon** (Fig. 2), a bias-calibrated adaptation framework grounded in the *minimal intervention principle*: preserve transferable knowledge while calibrating only the components that cause degradation. BiasRecon operationalizes this principle through an alternating optimization framework that directly addresses both hypothesized sources of degradation. To correct the misaligned prior, *frequency-guided prior calibration* adaptively modulates layer-wise and frequency-specific components, achieving effective adaptation with minimal parameters, while *score-based denoising* leverages the calibrated prior for reconstruction. To address the suboptimal hyperparameters, *adaptive regularization* dynamically adjusts the prior-measurement trade-off to match the statistical characteristics of each test sample, using a self-supervised criterion that requires no ground truth. This principled design




Fig. 2. Overview of BiasRecon framework. The top panel illustrates the alternating optimization process of BiasRecon. The framework consists of three key components: a) Frequency-guided Prior Calibration, which adaptively adjusts feature distributions through parameter-efficient modulation; b) Score-based reverse sampling process as the foundation; and c) Regularization Parameter Adaptation, which dynamically balances measurement consistency and prior constraints.

enables BiasRecon to achieve robust open-world adaptation with fewer than 100 tunable parameters.

Contributions. First, we identify through transferability analysis that deeper layers and low-frequency components are domain-specific while shallow layers and high-frequency components remain transferable, motivating the *minimal intervention principle* for open-world adaptation. Second, we propose BiasRecon, an alternating optimization framework with fewer than 100 tunable parameters that operationalizes this principle through frequency-guided prior calibration and adaptive regularization. Third, we demonstrate state-of-the-art performance on open-world MRI reconstruction.

2 Method

2.1 Bias-Calibrated Adaptation Framework

MRI Forward Model. In parallel accelerated MRI, the observation obtained in multi-coil k -space $\mathbf{y}_c \in \mathbb{C}^{N_y}$ is mathematically expressed as $\mathbf{y}_c = \mathbf{A}_c \mathbf{x} + \mathbf{n}$, where $\mathbf{x} \in \mathbb{C}^{N_x}$ is the reconstructed image, $\mathbf{n}$ models Gaussian noise, and $\mathbf{A}_c = \mathcal{M} \mathcal{F} \mathbf{S}_c$ is the forward operator for the c -th coil. Here, $\mathbf{S}_c$ encodes the coil sensitivity map, $\mathcal{F}$ denotes the discrete Fourier transform, and $\mathcal{M}$ is the undersampling mask [6]. For a multi-coil system with C coils, we write $\mathbf{y} = \mathbf{A} \mathbf{x} + \mathbf{n}$.

Bias-Calibrated Adaptation. We adopt score-based generative models [4] as our prior estimator, where a neural network $s_{\theta^*}(\mathbf{x})$ approximates the score function $\nabla_{\mathbf{x}} \log p(\mathbf{x})$. To tackle open-world MRI reconstruction, we introduce a calibration variable δ that modulates the pre-trained score network to accommodate distribution shifts, yielding $s_{\theta^*}(\mathbf{x}; \delta) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x} | \delta)$. We achieve reconstruction by jointly optimizing $\mathbf{x}$ and δ given observation $\mathbf{y}$:

$$(\mathbf{x}, \delta) \in \arg \max_{\mathbf{x}, \delta} p(\mathbf{x}, \delta | \mathbf{y}) \in \arg \max_{\mathbf{x}, \delta} p(\mathbf{y} | \mathbf{x}) p(\mathbf{x} | \delta) p(\delta), \quad (1)$$




Fig. 3. Illustration of frequency-guided prior calibration.




Fig. 4. Dynamic evolution of calibration variables and regularization parameter during the reconstruction process.

where $p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(\mathbf{y} | \mathbf{A}\mathbf{x}, \gamma^{-1}\mathbf{I})$ and $p(\mathbf{x} | \boldsymbol{\delta})$ is determined through the score function. Taking the negative log and introducing a proximal term, we split the problem into three alternating subproblems:

$$\begin{cases} P1 : \boldsymbol{\delta}^{t-1} \in \arg \max_{\boldsymbol{\delta}} \underbrace{\log p(\mathbf{x}^t | \boldsymbol{\delta}) + \log p(\boldsymbol{\delta})}_{\text{Prior calibration}}, \\ P2 : \hat{\mathbf{x}}^{t-1} \in \arg \min_{\mathbf{x}^t} \underbrace{\frac{1}{2} \|\hat{\mathbf{x}}^t - \mathbf{x}^t\|_2^2 - \log p(\mathbf{x}^t | \boldsymbol{\delta}^{t-1})}_{\text{Denoising}}, \\ P3 : \hat{\mathbf{x}}^{t-1} = \arg \min_{\hat{\mathbf{x}}} \underbrace{\frac{\gamma}{2} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}\|_2^2}_{\text{Data Fidelity}} + \underbrace{\frac{1}{2} \|\hat{\mathbf{x}}^{t-1} - \hat{\mathbf{x}}\|_2^2}_{\text{Proximity}}. \end{cases} \quad (2)$$

For P2 in Eq. (2), we derive a one-step update:

$$\hat{\mathbf{x}}^{t-1} = \mathbf{x}^t + (\sigma^t)^2 \nabla_{\mathbf{x}} \log p(\mathbf{x} | \boldsymbol{\delta}^{t-1})|_{\mathbf{x}=\mathbf{x}^t}, \quad (3)$$

where the score function is estimated by a bias-calibrated network $\mathbf{s}_{\theta^*}(\mathbf{x}^t, t; \boldsymbol{\delta}^{t-1})$ and $(\sigma^t)^2$ is derived from Tweedie's formula [8]. For P3 in Eq. (2), we use conjugate gradient [3] to efficiently solve the data fidelity term.

2.2 Frequency-guided Prior Calibration

Transferability Analysis. To understand how distribution shifts affect the denoising network in P2 of Eq. (2), we analyze feature representations through t-SNE visualization, as shown in Fig. 1(b,c). We find that *deeper layers and low-frequency components exhibit domain-specific characteristics (encoding anatomical structures like knee vs. brain), while shallow layers and high-frequency components remain domain-agnostic (capturing transferable edges and textures)*. This observation motivates us to design a layer-wise and frequency-specific modulation strategy that selectively calibrates low-frequency components in deeper layers while preserving high-frequency representations in shallow layers.

Parameter-efficient Calibration Strategy. Motivated by the above analysis, we instantiate the calibration variable $\boldsymbol{\delta}$ as a compact parameter vector for the U-Net architecture [13] of our score-based network. Specifically, for each skip connection layer l , we decompose features via Fourier transformation

and introduce two learnable scalars $\alpha_l, \beta_l \in [0, 2]$ to modulate low- and high-frequency components, respectively. The calibration variable is thus defined as $\delta = [\alpha_1, \beta_1, \dots, \alpha_L, \beta_L]^\top \in \mathbb{R}^{2L}$, introducing just $2L$ parameters for an L -layer network, enabling efficient frequency-aware adaptation, as illustrated in Fig. 3.

Self-supervised Optimization. With δ defined above, for P1 in Eq. (2), we optimize δ without ground truth by partitioning the measurement mask into two complementary Gaussian 2D subsampling masks ($\mathcal{M} = \mathcal{M}_A + \mathcal{M}_R$). Through k-space projection and probabilistic modeling, we obtain:

$$\mathcal{L}_{\text{SSL}}(\delta) = \frac{\tau}{2} \|\mathbf{y}_R - \mathbf{A}_R \mathbf{f}_{\theta^*}(\delta; \mathbf{x}_A^t, t)\|_2^2, \quad (4)$$

where $\mathbf{f}_{\theta^*}$ represents the reconstruction from P2 and P3, and τ controls precision. Therefore, P1 in Eq. (2) can be reformulated as:

$$\delta^{t-1} \in \arg \min_{\delta} \mathcal{L}_{\text{SSL}}(\delta) - \log p(\delta), \quad (5)$$

where $p(\delta) = \mathcal{N}(1, 1)$ ensures $\delta = 1$ preserves the original network.

2.3 Regularization Parameter Adaptation

For P3 in Eq. (2), the regularization parameter γ balances measurement consistency and prior constraints. Manual selection is suboptimal for open-world scenarios since γ relates to the noise level, which should be image-specific and updated adaptively. We employ Stein’s Unbiased Risk Estimator (SURE) [24] with Monte Carlo approximation [12] to estimate MSE without ground truth:

$$\mathcal{L}_{\text{Reg}}(\gamma) = \frac{\|\mathbf{x}_{\text{zf}} - \mathbf{h}(\gamma; \mathbf{x}^t)\|^2}{N\epsilon} \boldsymbol{\mu}^T (\mathbf{h}(\gamma; \mathbf{x}^t + \epsilon \boldsymbol{\mu}) - \mathbf{h}(\gamma; \mathbf{x}^t)), \quad (6)$$

where $\boldsymbol{\mu} \sim \mathcal{N}(0, \mathbf{I})$, ϵ is the perturbation scale, and $\mathbf{h}(\gamma; \cdot)$ represents the compositional mapping of P2 and P3 from Eq. (2). To prevent overfitting, we employ early stopping based on sliding window convergence:

$$\mathcal{E}(\mathcal{L}_{\text{Reg}}) = 1 - \frac{\sum_{t'=t-k+1}^t \mathcal{L}_{\text{Reg}}^{t'}}{\sum_{t'=t-2k+1}^{t-k} \mathcal{L}_{\text{Reg}}^{t'}}, \quad (7)$$

where updating stops when $\mathcal{E}(\mathcal{L}_{\text{Reg}}) < \tau_{\text{Reg}}$.

3 Experiments

3.1 Experimental Setup

Datasets. We evaluate on four multi-coil MRI datasets. *FastMRI knee* [22] serves as the training set (973 volumes), discarding the first and last 5 slices of each volume with small objects. We directly evaluate on three open-world test sets without retraining: *FastMRI-Brain* (Cross-Anatomy), *Stanford Knee*

(Cross-Center) [10], and *CMRxRecon* (Combined shift) [17]. For each test set, 100 volumes are randomly selected with 3 slices per volume. Input images are unified to 320×320 through cropping and padding. Coil sensitivity maps are estimated by ESPiRiT [15], implemented with the SigPy Python package. For $4\times$ and $8\times$ undersampling, the ACS regions retain 8% and 4% of central k-space lines, respectively.

Baselines and Evaluation. We compare against TV [2], ConvDecoder [5], MoDL [1], and DDS [3]. Reconstruction quality is evaluated by PSNR, SSIM, and LPIPS [23].

Training and Inference. We train a score-based model on FastMRI knee for 1M iterations with batch size 4 and learning rate $1e-4$. Complex data is processed using MVUE [7]. We test Gaussian1D and Uniform1D undersampling at $4\times$ and $8\times$ acceleration. During inference, calibration variable δ and regularization parameter γ are updated via Adam optimizer with learning rates $1e-4$ and 0.3, respectively. We employ DDIM sampling [14] with 100 timesteps to accelerate the reverse process, and solve each data-consistency update with 5 conjugate-gradient iterations. δ is initialized to identity and optimized for 70 steps, while γ is initialized to 1 and optimized for up to 100 steps, typically stabilizing after about 60 steps. We set $\tau_{\text{Reg}} = 0.001$ with window size $k = 5$.

3.2 Open-World Evaluation

Table 1. Open-world MRI reconstruction performance under $8\times$ acceleration with Gaussian 1D sampling.

Method	FastMRI Knee (In-Distribution)			FastMRI-Brain (Cross-Anatomy)		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Zero-filling	19.83 $\pm$ 1.52	0.691 $\pm$ 0.025	0.329 $\pm$ 0.031	23.72 $\pm$ 1.83	0.785 $\pm$ 0.028	0.236 $\pm$ 0.022
<i>Conventional CS method</i>						
TV	26.04 $\pm$ 1.63	0.771 $\pm$ 0.023	0.313 $\pm$ 0.025	28.14 $\pm$ 1.86	0.817 $\pm$ 0.027	0.187 $\pm$ 0.024
<i>Internal deep image prior</i>						
ConvDecoder	29.97 $\pm$ 2.15	0.745 $\pm$ 0.029	0.257 $\pm$ 0.027	31.27 $\pm$ 2.73	0.876 $\pm$ 0.022	0.141 $\pm$ 0.022
<i>Supervised Method</i>						
MoDL	31.67 $\pm$ 1.98	0.773 $\pm$ 0.028	0.215 $\pm$ 0.024	32.53 $\pm$ 0.99	0.807 $\pm$ 0.064	0.137 $\pm$ 0.019
<i>External deep image prior</i>						
DDS	31.64 $\pm$ 1.89	0.780 $\pm$ 0.031	0.200 $\pm$ 0.021	32.46 $\pm$ 1.63	0.864 $\pm$ 0.029	0.140 $\pm$ 0.016
Ours	31.76$\pm$2.01	0.792$\pm$0.032	0.195$\pm$0.018	33.65$\pm$1.42	0.885$\pm$0.019	0.136$\pm$0.025

Method	Stanford Knee (Cross-Center)			CMRxRecon (Combined)		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Zero-filling	19.51 $\pm$ 1.48	0.691 $\pm$ 0.027	0.327 $\pm$ 0.029	22.82 $\pm$ 1.76	0.676 $\pm$ 0.031	0.164 $\pm$ 0.021
<i>Conventional CS method</i>						
TV	24.12 $\pm$ 1.72	0.788 $\pm$ 0.026	0.306 $\pm$ 0.028	31.15 $\pm$ 1.92	0.804 $\pm$ 0.029	0.182 $\pm$ 0.023
<i>Internal deep image prior</i>						
ConvDecoder	30.84 $\pm$ 2.25	0.753 $\pm$ 0.031	0.264 $\pm$ 0.028	34.71 $\pm$ 2.42	0.854 $\pm$ 0.025	0.133 $\pm$ 0.021
<i>Supervised Method</i>						
MoDL	30.66 $\pm$ 1.85	0.803 $\pm$ 0.029	0.203 $\pm$ 0.023	32.98 $\pm$ 1.94	0.877$\pm$0.021	0.143 $\pm$ 0.022
<i>External deep image prior</i>						
DDS	29.23 $\pm$ 1.78	0.822 $\pm$ 0.027	0.198 $\pm$ 0.022	33.72 $\pm$ 1.96	0.865 $\pm$ 0.026	0.138 $\pm$ 0.019
Ours	31.29$\pm$1.92	0.846$\pm$0.028	0.185$\pm$0.021	35.32$\pm$2.12	0.878$\pm$0.024	0.116$\pm$0.020




Fig. 5. Qualitative comparison of MRI reconstruction results under $8\times$ Gaussian 1D undersampling.

Robustness against Distribution Shifts. Table 1 demonstrates BiasRecon’s effectiveness across diverse distribution shifts. On the in-distribution FastMRI knee dataset, our method slightly outperforms baselines by adapting to image-specific noise levels. More importantly, BiasRecon achieves the best performance across all open-world scenarios: +1.19 dB PSNR over DDS on *Cross-Anatomy*, +2.06 dB on *Cross-Center*, and +1.60 dB on *Combined*, which is the most challenging scenario involving cross-anatomy, cross-center, and cross-modality shifts simultaneously. These consistent improvements validate our two hypothesized sources of degradation: frequency-guided prior calibration corrects the misaligned prior, while adaptive regularization adjusts the prior-measurement trade-off for unseen acquisition conditions. Qualitative results in Fig. 5 further confirm that BiasRecon produces reconstructions most consistent with ground truth.

Robustness against undersampling. We evaluated the robustness against different undersampling operators, including Uniform1D and Gaussian1D, with acceleration factors of $4\times$ and $8\times$. For experimental fairness, the MoDL model was trained on the FastMRI knee dataset using different sampling operators at these acceleration factors ($4\times$ and $8\times$). Table 2 presents the quantitative results. Our method outperforms others in 10 out of 12 studies and achieves the second best performance in the rest 2 studies, demonstrating greater robustness across various sampling patterns and acceleration factors.

Table 2. Robustness against different sampling patterns and acceleration factors on open-world FastMRI-brain dataset. Comparison includes both Gaussian 1D and Uniform 1D sampling at $4\times$ and $8\times$ acceleration rates.

Method	Metric	Gaussian 1D		Uniform 1D	
		$4\times$	$8\times$	$4\times$	$8\times$
TV	PSNR	33.65 ± 1.29	28.14 ± 1.86	28.59 ± 1.69	25.71 ± 1.61
	SSIM	0.868 ± 0.017	0.817 ± 0.027	0.801 ± 0.031	0.725 ± 0.043
	LPIPS	0.136 ± 0.017	0.187 ± 0.024	0.174 ± 0.028	0.254 ± 0.043
ConvDecoder	PSNR	33.19 ± 2.41	31.27 ± 2.73	31.94 ± 2.35	26.27 ± 2.86
	SSIM	0.892 ± 0.020	0.876 ± 0.022	0.860 ± 0.034	0.728 ± 0.070
	LPIPS	0.116 ± 0.026	0.141 ± 0.022	0.135 ± 0.031	0.247 ± 0.052
MoDL	PSNR	33.59 ± 1.637	32.53 ± 0.99	31.96 ± 2.056	27.08 ± 1.919
	SSIM	0.871 ± 0.022	0.807 ± 0.064	0.797 ± 0.064	0.659 ± 0.077
	LPIPS	0.120 ± 0.032	0.137 ± 0.019	0.125 ± 0.021	0.200 ± 0.040
DDS	PSNR	34.10 ± 1.97	32.56 ± 1.63	30.99 ± 1.79	27.23 ± 2.03
	SSIM	0.883 ± 0.028	0.864 ± 0.029	0.851 ± 0.031	0.736 ± 0.049
	LPIPS	0.114 ± 0.015	0.140 ± 0.016	0.135 ± 0.020	0.223 ± 0.040
Ours	PSNR	35.03 ± 2.01	33.65 ± 1.42	32.72 ± 1.92	28.40 ± 2.22
	SSIM	0.889 ± 0.032	0.885 ± 0.019	0.869 ± 0.033	0.743 ± 0.041
	LPIPS	0.110 ± 0.018	0.136 ± 0.025	0.134 ± 0.023	0.211 ± 0.046

Table 3. Ablation study on Cross-Center shift under $4\times$ Gaussian 1D acceleration. FPC: Frequency-guided Prior Calibration; RPA: Regularization Parameter Adaptation.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Baseline	30.23 ± 3.57	0.844 ± 0.035	0.187 ± 0.028
w/o RPA	$31.28 \pm 3.22(\uparrow 1.05)$	$0.855 \pm 0.035(\uparrow .011)$	$0.183 \pm 0.029(\downarrow .004)$
w/o FPC	$32.51 \pm 4.04(\uparrow 2.28)$	$0.853 \pm 0.058(\uparrow .009)$	$0.157 \pm 0.065(\downarrow .030)$
Ours (FPC + RPA)	$33.17 \pm 3.11(\uparrow 2.94)$	$0.866 \pm 0.041(\uparrow .022)$	$0.138 \pm 0.027(\downarrow .049)$

3.3 Parameter Analysis

Fig. 4 shows that while α and β in top layers remain stable, bottom layers exhibit pronounced changes: α decreases while β increases, indicating suppression of domain-specific low-frequency components while preserving high-frequency details. The gradual increase in γ reflects the transition from prior-guided denoising to measurement-consistent refinement as reconstruction progresses.

3.4 Ablation Study

Table 3 validates the effectiveness of our two adaptation components under cross-center distribution shift. Both components contribute meaningfully to the reconstruction quality, demonstrating their complementary nature. *Frequency-guided prior calibration* (+1.05 dB) adapts the learned prior by calibrating low-frequency components to accommodate variations in contrast, intensity distributions, and scanner-specific artifacts across different imaging centers. *Regularization parameter adaptation* (+2.28 dB) dynamically adjusts the prior-measurement trade-off to match test-time noise characteristics and k-space sampling quality. The full BiasRecon framework achieves +2.94 dB improvement, demonstrating that both components work synergistically.

4 Conclusion

We present BiasRecon, a bias-calibrated adaptation framework for open-world MRI reconstruction grounded in the *minimal intervention principle*: preserve transferable knowledge while calibrating only bias-prone components. Through transferability analysis, we observe that deeper layers and low-frequency components exhibit domain-specific characteristics while shallow layers and high-frequency components remain domain-agnostic, motivating frequency-guided prior calibration and regularization parameter adaptation. By intervening minimally and precisely, BiasRecon achieves state-of-the-art performance across cross-center, cross-anatomy, and cross-modality scenarios with fewer than 100 parameters.

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

References

1. Aggarwal, H.K., Mani, M.P., Jacob, M.: Modl: Model-based deep learning architecture for inverse problems. *IEEE transactions on medical imaging* **38**(2), 394–405 (2018)
2. 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: An Official Journal of the International Society for Magnetic Resonance in Medicine* **57**(6), 1086–1098 (2007)
3. Chung, H., Lee, S., Ye, J.C.: Decomposed diffusion sampler for accelerating large-scale inverse problems. In: *The Thirteenth International Conference on Learning Representations* (2024)
4. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated mri. *Medical image analysis* **80**, 102479 (2022)
5. Darestani, M.Z., Heckel, R.: Accelerated mri with un-trained neural networks. *IEEE Transactions on Computational Imaging* **7**, 724–733 (2021)
6. Geethanath, S., Reddy, R., Konar, A.S., Imam, S., Sundaresan, R., DR, R.B., Venkatesan, R.: Compressed sensing mri: a review. *Critical Reviews in Biomedical Engineering* **41**(3) (2013)
7. 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)
8. Knoll, F., Hammernik, K., Kobler, E., Pock, T., Recht, M.P., Sodickson, D.K.: Assessment of the generalization of learned image reconstruction and the potential for transfer learning. *Magnetic resonance in medicine* **81**(1), 116–128 (2019)
9. Liu, Q., Yang, Q., Cheng, H., Wang, S., Zhang, M., Liang, D.: Highly undersampled magnetic resonance imaging reconstruction using autoencoding priors. *Magnetic resonance in medicine* **83**(1), 322–336 (2020)
10. Ong, F., Amin, S., Vasanawala, S., Lustig, M.: Mridata. org: An open archive for sharing mri raw data. In: *Proc. Intl. Soc. Mag. Reson. Med.* vol. 26 (2018)
11. Qu, X., Guo, D., Ning, B., Hou, Y., Lin, Y., Cai, S., Chen, Z.: Undersampled mri reconstruction with patch-based directional wavelets. *Magnetic resonance imaging* **30**(7), 964–977 (2012)

12. Ramani, S., Blu, T., Unser, M.: Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms. *IEEE Transactions on image processing* **17**(9), 1540–1554 (2008)
13. 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)
14. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *International Conference on Learning Representations* (2021)
15. Uecker, M., Lai, P., Murphy, M.J., Virtue, P., Elad, M., Pauly, J.M., Vasanawala, S.S., Lustig, M.: Espirit—an eigenvalue approach to autocalibrating parallel mri: where sense meets grappa. *Magnetic resonance in medicine* **71**(3), 990–1001 (2014)
16. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9446–9454 (2018)
17. Wang, C., Lyu, J., Wang, S., Qin, C., Guo, K., Zhang, X., Yu, X., Li, Y., Wang, F., Jin, J., et al.: Cmrrecon: A publicly available k-space dataset and benchmark to advance deep learning for cardiac mri. *Scientific Data* **11**(1), 687 (2024)
18. Wang, S., Lv, J., He, Z., Liang, D., Chen, Y., Zhang, M., Liu, Q.: Denoising auto-encoding priors in undecimated wavelet domain for mr image reconstruction. *Neurocomputing* **437**, 325–338 (2021)
19. Yaman, B., Hosseini, S.A.H., Akcakaya, M.: Zero-shot self-supervised learning for mri reconstruction. *International Conference on Learning Representations* (2022)
20. Yang, A.C., Kretzler, M., Sudarski, S., Gulani, V., Seiberlich, N.: Sparse reconstruction techniques in magnetic resonance imaging: methods, applications, and challenges to clinical adoption. *Investigative radiology* **51**(6), 349–364 (2016)
21. Yang, Y., Sun, J., Li, H., Xu, Z.: Admm-csnet: A deep learning approach for image compressive sensing. *IEEE transactions on pattern analysis and machine intelligence* **42**(3), 521–538 (2018)
22. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., Defazio, A., Stern, R., Johnson, P., Bruno, M., et al.: fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839* (2018)
23. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *CVPR* (2018)
24. Zhussip, M., Soltanayev, S., Chun, S.Y.: Extending stein’s unbiased risk estimator to train deep denoisers with correlated pairs of noisy images. *Advances in neural information processing systems* **32** (2019)