Diffusion Posterior Sampling for Fourier Compressed Sensing

Jaehyeok Bae

Abstract—This study introduces Diffusion Posterior Sampling (DPS) as an effective approach to Fourier compressed sensing, specifically targeting accelerated MRI reconstruction. Utilizing a pretrained latent diffusion model (LDM), DPS significantly enhances reconstruction quality compared to conventional methods such as GRAPPA, achieving superior structural similarity (SSIM) and peak signal-to-noise ratio (PSNR) across various acceleration factors. Performance analyses highlight the critical role of undersampling patterns, with the Variable Density (VD) Poisson Disc pattern consistently yielding optimal results. Ablation studies emphasize the importance of hyperparameter tuning, notably inference step size and gradient descent step size (ζ), in optimizing reconstruction performance. The results confirm DPS as a promising method for high-quality MRI reconstruction, paving the way for future work on the joint optimization of undersampling patterns and diffusion models.

Index Terms—Fourier Compressed Sensing, Accelerated MRI, Diffusion Posterior Sampling, Latent Diffusion Model

1 Introduction

Compressed sensing (CS) has transformed the field of image acquisition by enabling high-quality image reconstruction from significantly fewer measurements. This is achieved by exploiting the inherent sparsity of natural images in specific transform domains. The foundational CS theory [1] ensures that an image can be accurately reconstructed from a randomly sampled subset of its frequency components, provided it has a sparse representation.

Fourier compressed sensing (Fourier CS) is a specialized form of CS where measurements are taken in the discrete Fourier transform (DFT) domain. Given that most natural signals and medical images exhibit concentrated energy in low-frequency (LF) components, Fourier CS achieves high-fidelity reconstructions by prioritizing LF sampling. This property has led to its extensive application in various electromagnetic imaging modalities, including magnetic resonance imaging (MRI) [2].

The effectiveness of Fourier CS depends on two key factors: the reconstruction algorithm and the undersampling pattern used in k-space. Conventional parallel imaging methods like GRAPPA [3] leverage coil sensitivity profiles to interpolate missing k-space data but struggle under high acceleration rates. Recent advancements have focused on optimizing both the reconstruction model and the undersampling pattern to improve performance [4].

In this work, we introduce Diffusion Posterior Sampling (DPS) as an alternative approach for solving the Fourier CS problem. DPS employs a pretrained latent diffusion model (LDM) to iteratively refine the reconstructed image by leveraging the prior knowledge embedded in the model. Unlike conventional methods, DPS does not require explicit training on MRI-specific data, yet it effectively mitigates aliasing artifacts and enhances reconstruction quality.

We conduct extensive evaluations of DPS on accelerated MRI reconstruction, comparing it with the GRAPPA method across different acceleration factors. Additionally, we analyze the impact of undersampling masks, demonstrating that the Variable Density (VD) Poisson Disc pattern consistently outperforms uniform and Gaussian random sampling patterns. Furthermore, ablation studies reveal the importance of inference step size and gradient descent step size (ζ) in optimizing reconstruction quality.

In conclusion, DPS provides a promising alternative to conventional CS reconstruction methods, leveraging diffusion models for high-quality image restoration. Future work will focus on refining the undersampling pattern and integrating learnable sampling strategies to further enhance reconstruction performance.

2 Related Work

2.1 Latent Diffusion Models for Image Reconstruction

Diffusion models have emerged as powerful generative priors for solving inverse problems in imaging. Among these, Latent Diffusion Models (LDMs) [5] operate in a compressed latent space, significantly reducing computational complexity while preserving high-quality outputs. This characteristic makes LDMs particularly effective for MRI reconstruction, as MR images naturally exhibit redundancy that can be efficiently represented in lower-dimensional manifolds.

Given a medical image $x \in \mathbb{R}^{H\times W\times C}$, where H, W, C denote the height, width, and coil dimensions, respectively, an encoder $\mathcal{E}$ maps x to a latent representation $z = \mathcal{E}(x) \in \mathbb{R}^{h\times w\times C}$. The decoder $\mathcal{D}$ then reconstructs the image as $\tilde{x} = \mathcal{D}(z)$. Typically, a Variational Autoencoder (VAE)-based structure [6] is employed to learn this encoder-decoder pair.

In LDMs, the model $\epsilon_\theta(z_t, t)$ is trained to estimate a noise-free reconstruction of the input latent features z_t, where z_t represents a noisy transformation of the original input at time step $t \in [0, T]$. The diffusion model is trained using the following objective:

$$\mathbb{E}_{\mathcal{E}(x),\epsilon\sim\mathcal{N}(0,1),t}\left[\|\epsilon - \epsilon_\theta(z_t,t)\|_2^2\right]. \quad (1)$$

DPS [7] extends the application of diffusion models to inverse problems, including Fourier compressed sensing. By

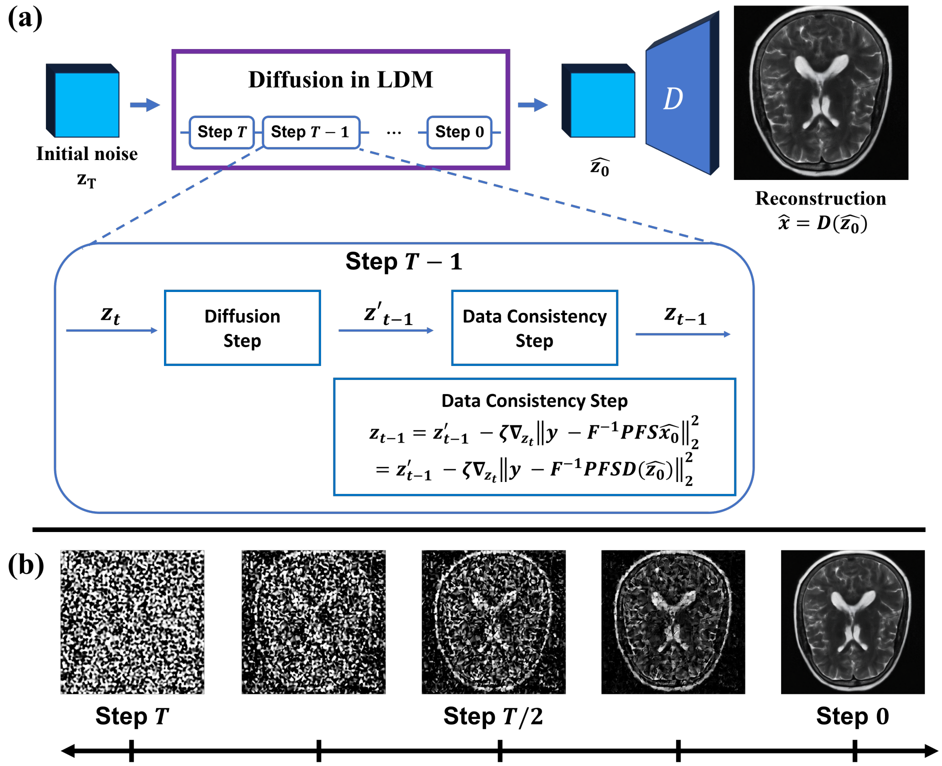

Fig. 1. Overview of the proposed method. (a) Diffusion step for solving the Fourier compressed sensing problem. (b) Visualization of the estimated reconstructed image through the diffusion process.

iteratively refining the reconstruction through a diffusion-based prior and a data consistency term, DPS enables high-fidelity recovery of MR images without requiring explicit retraining on MRI datasets.

2.2 GRAPPA and Parallel Imaging Techniques

Traditional MRI acceleration techniques rely on parallel imaging, where multiple receiver coils acquire k-space data simultaneously. A widely used method, Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [3], estimates missing k-space lines by exploiting the spatial redundancy across multiple coils.

GRAPPA reconstructs the missing k-space data through a linear combination of acquired k-space signals, with coil-specific weights derived from autocalibration signal (ACS) lines. While GRAPPA effectively reduces scan time and mitigates aliasing, its performance deteriorates under high acceleration factors due to the increased interpolation error. Additionally, GRAPPA is inherently limited by the availability and quality of ACS lines, making it less robust in cases where autocalibration data is sparse or undersampled.

Recent advancements in deep learning have explored alternative reconstruction approaches that incorporate learned priors from generative models. Unlike GRAPPA, DPS does not rely on explicit coil sensitivity calibration but instead leverages the powerful prior of a pretrained diffusion model to iteratively refine the reconstructed image. This makes DPS particularly advantageous for high-acceleration MRI reconstruction, where conventional methods struggle to maintain image quality.

3 Theory and Methods

3.1 Background Theory

Magnetic Resonance Imaging (MRI) acquisition relies on sampling k-space data, where undersampling accelerates imaging at the cost of introducing aliasing artifacts. The MRI forward model for an input image x can be expressed as:

$$y = F^{-1}PFS(x) + n, \tag{2}$$

where $y \in \mathbb{R}^{H\times W}$ represents the reconstruction from acquired k-space measurements, P is the undersampling mask, F is the Fourier transform operator, and S is the sensitivity map accounting for multiple coil acquisitions. The noise term $n \sim \mathcal{N}(0, \sigma^2)$ models general measurement noise.

The goal of reconstruction is to recover the original image x by solving:

$$\hat{x} = \arg\min_x \|F^{-1}PFS(x) - y\|_2^2 + \mathcal{R}(x), \tag{3}$$

where $\mathcal{R}(x)$ is a regularization term acting as a prior $p(x)$ to enforce realistic image constraints. Traditional approaches, such as GRAPPA, estimate missing k-space lines through coil sensitivity interpolation. In contrast, recent methods employ learned priors, such as diffusion models, to guide reconstruction.

3.2 Diffusion Posterior Sampling for Fourier Compressed Sensing

Diffusion models define a stochastic process that transforms data into noise and learns to reverse this process. In Dif-

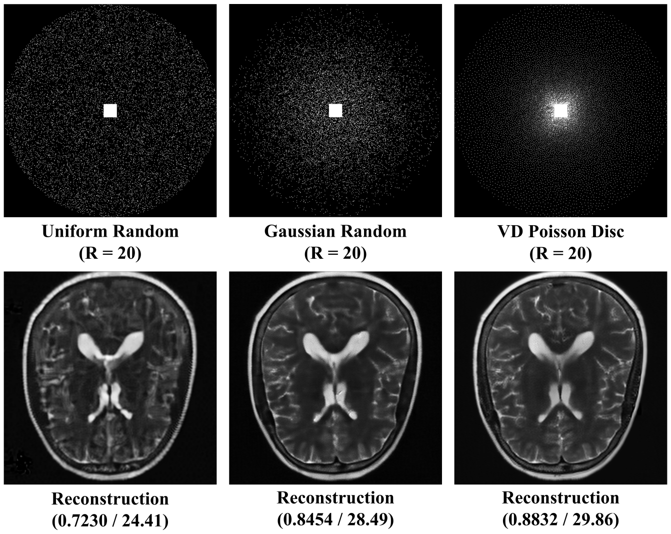

Fig. 2. (Top row) Undersampling patterns corresponding to an acceleration factor of $R = 20$. (Bottom row) Visualization of the corresponding reconstructed image using the proposed method.

fusion Posterior Sampling (DPS) [7], a pretrained latent diffusion model (LDM) is utilized to iteratively refine the reconstruction. The prior $p(x)$ is modeled in the latent space, $p(z)$, where $z = \mathcal{E}(x)$ is a lower-dimensional representation of x.

At each iteration t, the log-likelihood gradient is computed as:

$$\nabla_{x_t} \log p(x_t|y) = \nabla_{x_t} \log p(x_t) + \nabla_{x_t} \log p(y|x_t), \quad (4)$$

where $\nabla_{x_t} \log p(x_t)$ is the diffusion step guided by the pretrained model and $\nabla_{x_t} \log p(y|x_t)$ enforces data consistency using the MRI forward model. The clean image estimate $\hat{x}_0$ is obtained as:

$$\hat{x}_0 = \mathbb{E}[x_0|x_t], \quad (5)$$

which is further refined in the latent space through gradient updates:

$$\nabla_{z_t} \log p(y|z_t) = \nabla_{z_t} \log p(y|\mathcal{D}(\mathbb{E}[z_0|z_t])). \quad (6)$$

3.3 Proposed Method

Figure 1 illustrates the proposed DPS-based MRI reconstruction framework. The process starts with an initial noise representation z_T in latent space, which undergoes iterative refinement via diffusion and data consistency steps.

In the Data Consistency step, the ℓ_2-norm between the acquired measurement y and the MRI forward operation of the estimated reconstruction, $F^{-1}PFS\hat{x}_0$, is minimized. This ensures alignment with the acquired k-space data while preventing overfitting to noise. The latent variable z_t is updated via gradient descent to improve fidelity. The final reconstruction is obtained by decoding the refined latent representation:

$$\hat{x}_0 = \mathcal{D}(\hat{z}_0), \quad (7)$$

where $\mathcal{D}$ maps the latent representation back to the image space. Figure 1(b) visualizes the evolution of the reconstruction, demonstrating how the proposed method progressively refines the MRI image, effectively mitigating aliasing artifacts and enhancing structural details.

4 Analysis and Evaluation

In this work, we evaluate the proposed Diffusion Posterior Sampling (DPS) method against the conventional GRAPPA method for accelerated MRI reconstruction. Quantitative and qualitative analyses are conducted across various acceleration factors to access reconstruction quality.

4.1 Undersampling Masks

To analyze the impact of different undersampling patterns, we generated three distinct masks for each acceleration rate. Figure 2 illustrates these masks for an acceleration factor of $R = 20$:

- **Uniform Random Mask**: Samples points randomly from a uniform distribution without considering spatial constraints. - **Gaussian Random Mask**: Samples points randomly from a Gaussian distribution, concentrating more samples in the central k-space region. - **Variable Density (VD) Poisson Disc Pattern**: Ensures an adaptive spatial distribution, maintaining a balance between dense sampling in low-frequency regions and sparse sampling in high-frequency regions.

All undersampling patterns operate in Cartesian space. The mask size matches the image dimensions, 512×512, with an autocalibration signal (ACS) region of 32×32. The VD Poisson Disc pattern has been observed to yield superior reconstruction performance due to its optimal coverage of k-space information.

4.2 Evaluation Datasets

To comprehensively evaluate the proposed method, we conducted experiments on three datasets:

1. **Brain MRI images** from the FastMRI dataset [8]. 2. **Knee MRI images** from the FastMRI dataset. 3. **Sample images** from the CelebA dataset [9].

For the FastMRI dataset, we converted multi-coil acquisitions into a single-coil format to align with the reconstruction framework. We evaluate reconstruction performance using two key metrics; Strurctural Similarity Index (SSIM) and Peak Signal-to-Noise Ratio (PSNR).

The reported SSIM and PSNR values are averaged across all test samples to provide a robust assessment of the general reconstruction performance.

4.3 Diffusion Model Configuration

We employ a pretrained Stable Diffusion v2.1-base model [5], originally trained on a general large image dataset. The model operates in a compressed latent space, significantly reducing computational complexity while maintaining high-fidelity reconstructions.

For inference, we set the default diffusion step size to $T = 1000$ and use the Euler Discrete Scheduler to control the iterative denoising process. Hyperparameter tuning, including gradient step size (ζ), was conducted to optimize reconstruction stability and accuracy.

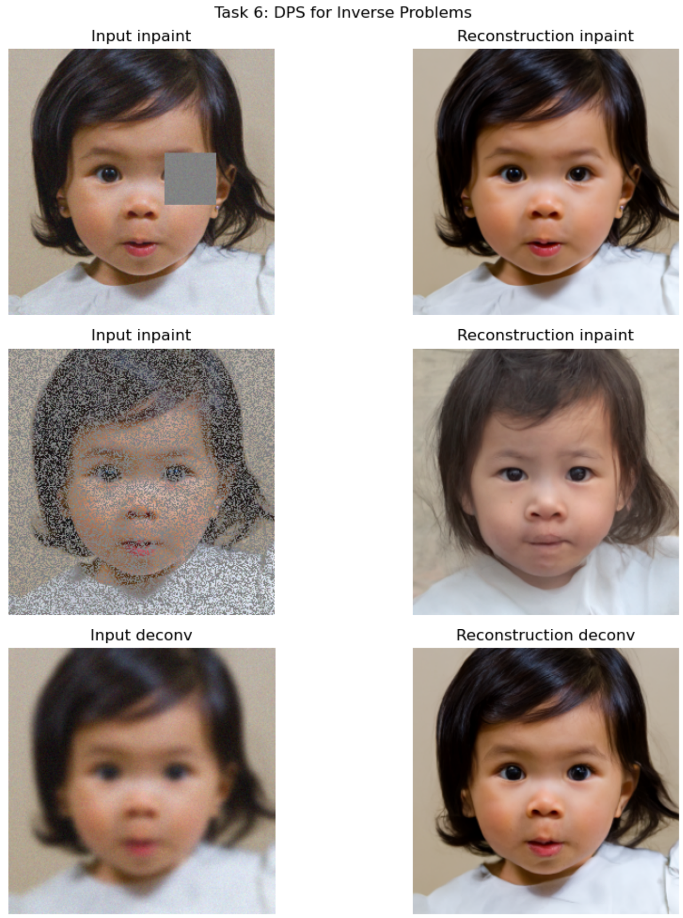

Fig. 3. Results of DPS applied to several inverse problems.

5 Experimental Results

5.1 Solving Inverse Problems with DPS

Before evaluating DPS on Fourier Compressed Sensing, we first assessed its performance on simpler inverse problems, including image inpainting and deconvolution (Task 6 of the default Diffusion Project). These preliminary experiments provided insights into the model's capability to restore missing or degraded image structures.

The results of DPS applied to inverse problems are shown in Figure 3:

- **First row**: Image inpainting with a rectangular mask. - **Second row**: Image inpainting with 50% of pixels randomly masked. - **Third row**: Image deconvolution with Gaussian blur (kernel size = 61, standard deviation = 3.0).

For quantitative evaluation, we measured Peak Signal-to-Noise Ratio (PSNR) and Learned Perceptual Image Patch Similarity (LPIPS):

- **Image inpainting with a rectangular mask**: PSNR = 35.85, LPIPS = 0.0097. - **Image inpainting with a random mask**: PSNR = 15.71, LPIPS = 0.2757. - **Image deconvolution**: PSNR = 28.17, LPIPS = 0.0594.

The results indicate that DPS performs well for structured inpainting and deconvolution tasks. However, performance degrades for randomly masked inpainting due to the lack of structured context, demonstrating the limitations of diffusion-based priors when large portions of image content are missing.

Acceleration (R)	Undersampling Pattern	SSIM	PSNR
5x	Uniform Random	0.8728	29.63
	Gaussian Random	0.9015	29.72
	VD Poisson Disc	**0.9119**	**30.88**
10x	Uniform Random	0.8175	27.68
	Gaussian Random	0.8783	**29.57**
	VD Poisson Disc	**0.8981**	28.88
15x	Uniform Random	0.7749	26.15
	Gaussian Random	0.8648	29.38
	VD Poisson Disc	**0.8852**	**30.31**
20x	Uniform Random	0.7230	24.41
	Gaussian Random	0.8454	28.49
	VD Poisson Disc	**0.8832**	**29.86**

TABLE 1
Quantitative results for MRI reconstruction using DPS with various acceleration factors and undersampling patterns.

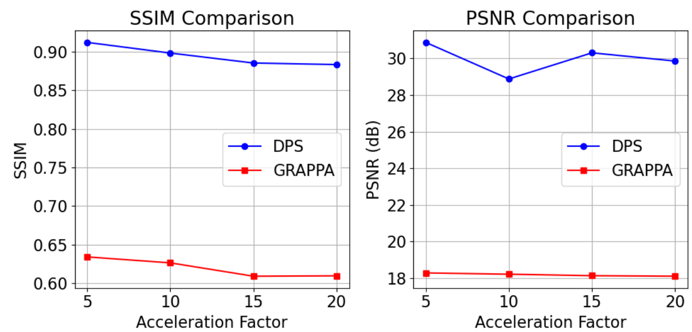

Fig. 4. Comparison of DPS with conventional GRAPPA.

5.2 Fourier Compressed Sensing Reconstruction

DPS was evaluated for MRI reconstruction using different undersampling patterns and acceleration factors. Figure 2 presents samples of visualization of the reconstructed images, while Table 1 summarizes the quantitative results.

VD Poisson Disc consistently achieved the best reconstruction performance, suggesting that it provides optimal k-space coverage for DPS-based reconstruction.

5.3 Comparison with GRAPPA

We compared DPS with GRAPPA across different acceleration factors. Figure 4 highlights that DPS consistently outperforms GRAPPA, particularly at higher acceleration rates. For instance, at $R = 20$, DPS achieves an SSIM of 0.8832 and a PSNR of 29.86 dB, while GRAPPA yields 0.6097 and 18.09 dB, respectively. The results confirm that DPS better preserves fine structures and reduces aliasing artifacts, reinforcing its suitability for high-quality MRI reconstruction.

5.4 Ablation Studies

We conducted ablation studies on inference step size and gradient descent step size (ζ) to assess their impact on reconstruction quality.

Inference Step Size: As shown in Figure 5, increasing the number of steps improves reconstruction quality. Performance peaks at 1000 steps (SSIM = 0.8728, PSNR = 29.63

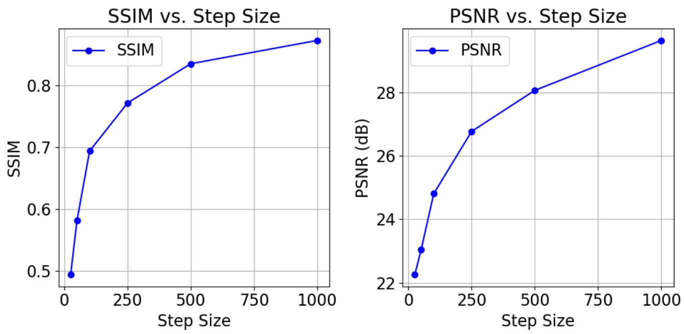

Fig. 5. Ablation study on inference step size.

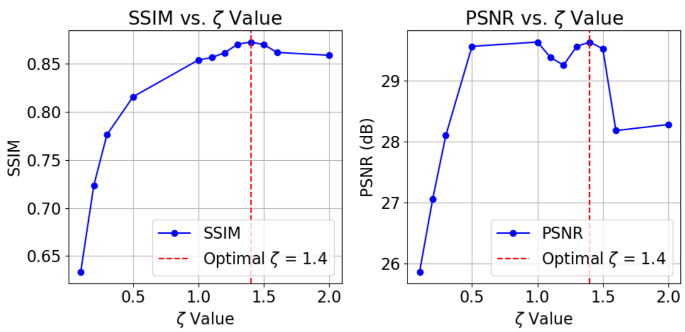

Fig. 6. Ablation study on gradient descent step size (ζ).

dB), which is the default hyperparameter, beyond which computational costs outweigh further gains.

Gradient Descent Step Size (ζ**)**: Figure 6 shows that an optimal ζ value of 1.4 achieves the best reconstruction quality. Smaller values lead to insufficient updates, while excessively large values cause instability.

6 Discussion and Conclusion

This study introduced Diffusion Posterior Sampling (DPS) as an effective approach for solving Fourier compressed sensing problems, with a particular focus on MRI reconstruction. Compared to conventional parallel imaging methods such as GRAPPA, DPS demonstrated superior performance across multiple acceleration factors, achieving consistently higher Structural Similarity Index (SSIM) and Peak Signal-to-Noise Ratio (PSNR). The ability of DPS to leverage a pretrained latent diffusion model (LDM) allowed it to effectively mitigate aliasing artifacts and restore fine structural details without requiring additional training on MRI-specific datasets.

The choice of undersampling patterns played a significant role in reconstruction performance. Among the tested sampling strategies, the Variable Density (VD) Poisson Disc pattern consistently outperformed uniform and Gaussian random masks, highlighting the importance of structured k-space sampling. The effectiveness of VD Poisson Disc sampling suggests that strategic undersampling can enhance the compatibility between acquired data and the generative priors used in DPS reconstruction.

Ablation studies provided deeper insights into the impact of key hyperparameters on DPS performance. Increasing the inference step size improved reconstruction quality by allowing more refined denoising steps, but computational cost increased significantly beyond 1000 steps with diminishing returns. Additionally, tuning the gradient descent step size (ζ) in the data consistency step proved essential for stability, with an optimal balance found at $\zeta = 1.4$. These findings reinforce the necessity of careful hyperparameter selection to maximize reconstruction fidelity while maintaining efficiency.

In conclusion, DPS emerges as a promising alternative to conventional compressed sensing reconstruction methods by integrating diffusion-based priors for enhanced image quality. Future research directions should explore joint optimization of undersampling patterns and diffusion models, as well as the incorporation of learnable sampling strategies. Further advancements may include adapting DPS to multi-coil MRI reconstruction and extending its application to other medical imaging modalities.

References

[1] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[2] M. Lustig, D. Donoho, and J. M. Pauly, “Sparse mri: The application of compressed sensing for rapid mr imaging,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[3] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, “Generalized autocalibrating partially parallel acquisitions (grappa),” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.

[4] C. Alkan, M. Mardani, C. Liao, Z. Li, S. S. Vasanawala, and J. M. Pauly, “Autosamp: autoencoding k-space sampling via variational information maximization for 3d mri,” *IEEE Transactions on Medical Imaging*, 2024.

[5] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10 684–10 695.

[6] D. P. Kingma, M. Welling *et al.*, “Auto-encoding variational bayes,” 2013.

[7] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, “Diffusion posterior sampling for general noisy inverse problems,” *arXiv preprint arXiv:2209.14687*, 2022.

[8] J. Zbontar, F. Knoll, A. Sriram, T. Murrell, Z. Huang, M. J. Muckley, A. Defazio, R. Stern, P. Johnson, M. Bruno *et al.*, “fastmri: An open dataset and benchmarks for accelerated mri,” *arXiv preprint arXiv:1811.08839*, 2018.

[9] Y. Zhang, Z. Yin, Y. Li, G. Yin, J. Yan, J. Shao, and Z. Liu, “Celeba-spoof: Large-scale face anti-spoofing dataset with rich annotations,” in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*. Springer, 2020, pp. 70–85.