

---

# [Supplementary Materials] Coupled Data and Measurement Space Dynamics for Enhanced Diffusion Posterior Sampling

---

Anonymous Author(s)

Affiliation

Address

email

## 1 Evaluation on Medical Imaging Inverse Problems

In this section, we extend our study by applying the proposed diffusion-based approach to inverse problems in the *medical imaging* domain. Medical images often contain complex structures and are typically acquired under constraints that make problems such as reconstruction, denoising, or inpainting particularly challenging. By evaluating our method on representative medical imaging tasks, we aim to demonstrate its effectiveness and generalization ability in real-world scenarios where accurate recovery of fine details is critical.

### 1.1 Datasets and Preprocessing.

We evaluate our proposed method, C-DPS, on three representative medical imaging tasks using public datasets: undersampled MRI, sparse-view CT, and super-resolution. For MRI reconstruction, we use the fastMRI knee dataset [1]. Following standard preprocessing [2], we crop the raw k-space data to  $320 \times 320$  resolution and reconstruct single-coil images using a minimum variance unbiased estimator. Undersampling is simulated using 1D Gaussian and Uniform sampling masks with acceleration rates (ACR) of 4 and 8.

For sparse-view CT, we utilize the LIDC dataset [3]. Two-dimensional CT slices of size  $320 \times 320$  are extracted from volumetric scans. Sinograms are generated using a parallel-beam geometry with either 23 or 10 projections evenly spaced over 180 degrees to simulate sparse-view acquisition.

The super-resolution task is evaluated on the fastMRI brain dataset [1]. We downsample 2D full-resolution images to generate low-resolution inputs for  $2 \times 2$  and  $4 \times 4$  upscaling. We select approximately 63% of the slices labeled as `reconstruction_rss`, yielding a total of 34,698 training samples.

### 1.2 Baselines and Evaluation Metrics

We compare our method, C-DPS, with strong training-free diffusion-based baselines including DPS [4], Score-MRI [5], DDS [6], and ScoreMed [7], depending on the task. All methods are implemented with their publicly available code and evaluated under consistent experimental conditions.

We report results using standard image quality metrics: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). Each task is evaluated on a test set of 1,000 images. For CT experiments, we follow prior work and set the likelihood step size  $\zeta = 0$  where applicable. In all experiments, our method outperforms the baselines across all configurations in terms of both PSNR and SSIM, as shown in Tables 1, 3, and 2.

### 1.3 Architecture and Sampling Settings

All models are based on the ADM architecture [8], without classifier-free guidance or dropout. Separate diffusion priors are trained for each task. For sampling, we use 100 timesteps for MRI and super-resolution, and 350 timesteps for sparse-view CT.

Our method introduces a measurement-consistent update through a bi-level guidance scheme. We tune the likelihood step size  $\zeta$  and refinement weight  $\gamma$  within a restricted search space,  $\zeta \in [0, 2]$ ,  $\gamma \in [0, 4.5]$ , and fix  $\lambda = 10^{-3}$  for the outer optimization. No extensive hyperparameter tuning is required to achieve strong performance.

All experiments are conducted on a NVIDIA P100 GPU with 12 GB of memory.

### 1.4 Results

We assess the effectiveness of our method, C-DPS, across three key inverse problems: MRI reconstruction, CT reconstruction, and super-resolution. In all cases, C-DPS consistently outperforms prior training-free diffusion-based baselines, demonstrating both robustness and accuracy across varying acquisition settings.

**MRI Reconstruction.** Table 1 reports results on the fastMRI knee dataset under Uniform1D and Gaussian1D masks at acceleration rates of  $4\times$  and  $8\times$ . Our method achieves the highest PSNR and SSIM across all configurations. Notably, under more aggressive undersampling ( $8\times$ ), C-DPS shows significant improvements over both DPS and DDS, highlighting its resilience in challenging reconstruction scenarios. The gains in SSIM, particularly under the Gaussian1D mask, indicate improved structural fidelity and reduced aliasing artifacts.

**Sparse-View CT Reconstruction.** As shown in Table 3, C-DPS achieves state-of-the-art results on the LIDC dataset with both 23 and 10 projections. It outperforms ScoreMed and BGDM by a notable margin in PSNR and SSIM, despite using the same number of diffusion steps (350). The improvement is especially pronounced under the extremely sparse 10-view setting, suggesting that C-DPS maintains strong data consistency even under severe measurement limitations.

**Super-Resolution.** On the fastMRI brain dataset, Table 2 demonstrates that C-DPS achieves the highest reconstruction quality for both  $2\times 2$  and  $4\times 4$  super-resolution tasks. Compared to DPS and DDS, our method produces sharper images with higher structural similarity, particularly in the more difficult  $4\times 4$  setting. These results affirm the generalization capability of C-DPS across different types of inverse problems.

**Overall Observations.** Across all tasks, C-DPS improves upon existing baselines without requiring task-specific training or extensive parameter tuning. The results highlight the advantage of incorporating measurement-aware updates through our bi-level guidance framework, which enhances both reconstruction accuracy and robustness under limited data.

Table 1: Quantitative results for the fastMRI knee dataset across different sampling masks and acceleration rates.

Method	Uniform1D $4\times$ ACR		Uniform1D $8\times$ ACR		Gaussian1D $4\times$ ACR		Gaussian1D $8\times$ ACR	
	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$
DPS [4]	32.40 $\pm$ 2.19	0.843 $\pm$ 0.063	31.07 $\pm$ 2.32	0.804 $\pm$ 0.073	34.93 $\pm$ 1.90	0.882 $\pm$ 0.063	33.72 $\pm$ 1.97	0.853 $\pm$ 0.071
Score-MRI [5]	31.95 $\pm$ 1.45	0.812 $\pm$ 0.036	27.97 $\pm$ 2.03	0.738 $\pm$ 0.053	33.96 $\pm$ 1.27	0.858 $\pm$ 0.028	30.82 $\pm$ 1.37	0.762 $\pm$ 0.034
DDS [6]	33.83 $\pm$ 2.54	0.859 $\pm$ 0.045	32.09 $\pm$ 2.84	0.822 $\pm$ 0.060	37.61 $\pm$ 2.29	0.900 $\pm$ 0.045	35.82 $\pm$ 2.42	0.874 $\pm$ 0.052
C-DPS (ours)	<b>35.63<math>\pm</math>2.47</b>	<b>0.877<math>\pm</math>0.057</b>	<b>33.33<math>\pm</math>2.66</b>	<b>0.842<math>\pm</math>0.077</b>	<b>38.05<math>\pm</math>2.43</b>	<b>0.918<math>\pm</math>0.0428</b>	<b>36.52<math>\pm</math>1.88</b>	<b>0.892<math>\pm</math>0.051</b>

### References

- [1] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.

Table 2: Super-resolution results on the fastMRI brain dataset.

Method	2×2 SR		4×4 SR	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
DPS [4]	35.44±3.71	0.931±0.027	30.29±2.84	0.854±0.034
DDS [6]	36.15±3.75	0.943±0.023	32.04±2.92	0.869±0.031
C-DPS (ours)	<b>36.42±3.02</b>	<b>0.951±0.044</b>	<b>32.48±2.32</b>	<b>0.889±0.051</b>

Table 3: Quantitative results of sparse-view CT reconstruction on the LIDC dataset using 350 NFEs.

Method	23 Projections		10 Projections	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
ScoreMed [7]	35.24±2.71	0.905±0.046	29.52±2.63	0.823±0.061
C-DPS (ours)	<b>35.92±2.33</b>	<b>0.925±0.046</b>	<b>30.59±2.21</b>	<b>0.842±0.058</b>

- [2] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *NeurIPS 2021 Workshop on Deep Learning and Inverse Problems*.
- [3] Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer, Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman, et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011.
- [4] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations (ICLR)*, 2023.
- [5] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated MRI. *Medical Image Analysis*, 76:102479, 2022.
- [6] Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating large-scale inverse problems. In *International Conference on Learning Representations (ICLR)*, 2024.
- [7] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations (ICLR)*, 2022.
- [8] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.