Abstract #00081 WORKSHOP ABSTRACT

# Patch-Based Diffusion Inverse Solver for T2-Weighted Prostate Imaging Reconstruction

**Traditional Poster** · 150 min | Session 7: Outside Poster Session · Tuesday, January 13 at 03:30 PM

**Hongze Yu** [1], **Jason Hu**[1], **Michael J Jaroszewicz**[2], **Hero K Hussain**[2], **Vikas Gulani**[2], **Jeffrey A Fessler**[1,2,3], **Yun Jiang**[2,3]

[1]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, United States of America
[2]Department of Radiology, University of Michigan, Ann Arbor, United States of America
[3]Department of Biomedical Engineering, University of Michigan, Ann Arbor, United States of America

***Presenting Author:*** *Hongze Yu (hongze@umich.edu)*

## Introduction

Prostate MRI benefits from fast, high-quality T2-weighted imaging, but acceleration is limited by SNR loss and undersampling artifacts, especially at 0.55T. Long scans and noisy images can introduce motion artifacts, reducing diagnostic confidence.

Diffusion models learn powerful image priors and have been adapted for accelerated MRI reconstruction[1-5,7], but whole-image diffusion requires large curated datasets, long training, and high GPU memory. The Patch-based Diffusion Inverse Solver (PaDIS[6]) addresses these limits by training on small spatially encoded patches and assembling local patch scores into a global prior at inference. This enables strong priors from limited data and lower memory, as has been shown for brain MRI[7].

In this work, we extend PaDIS to prostate T2-weighted imaging and evaluate accelerated reconstruction at 3T and 0.55T. We compare PaDIS with GRAPPA[8], $\ell 1$-wavelet[9], and whole-image diffusion[4], including blinded 0.55T reader scoring on artifact level, SNR, sharpness, and overall quality.

## Methods

PaDIS has two stages (**Figure 1**): (1) learning an image prior from prostate data using patches only, and (2) performing data-consistent reconstruction with the learned prior.

**Stage 1 – Patch prior.** From complex images (real/imaginary channels), we extract random patches and concatenate normalized positional coordinates. A U-Net denoiser ( $D_\theta$ ) is trained by denoising score matching: given a noisy patch and its noise level, the network predicts the clean patch, yielding a patch score

$$\mathbf{s}_{i,j,r}(\mathbf{x}_{i,j,r}) \approx \nabla_{\mathbf{x}_{i,j,r}} \log p(\mathbf{x}_{i,j,r}).$$

Whole-image scores are then formed by summing scores from tiled patches

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}) = \sum_{i,j=1}^{M} \left( \mathbf{s}_{i,j,B}(\mathbf{x}_{i,j,B}) + \sum_{r=1}^{(k+1)^2} \mathbf{s}_{i,j,r}(\mathbf{x}_{i,j,r}) \right),$$

where $i, j$ denote the patch indices along the X and Y dimensions and $S_{i,j,B}$ denotes scores for patches at zero-padded borders.

**Stage 2 – Reconstruction.** Given undersampled data y, PADIS alternates between enforcing data consistency $\|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_2^2$ with multi-coil forward model $\boldsymbol{A}$ and applying the patch score. Starting from complex noise, the current image is zero-padded and divided into randomly shifted patch tilings. The U-Net denoiser is applied to each position-encoded patch to produce denoised predictions and patch scores, which are then combined to form a whole-image log-prior gradient. We then update the current image estimate by stepping along this prior gradient while also enforcing multi-coil k-space data consistency. This update is repeated across iterations until convergence, with random shifts across iterations preventing persistent patch seams.

**Data and comparisons.** PaDIS and the whole-image diffusion baseline were trained on 6,000 fastMRI 3T prostate[10] slices, and tested at 3T and 0.55T. We compared against GRAPPA[8], $\ell 1$-wavelet[9], and whole-image diffusion[4], with blinded radiologist assessments on 0.55T images.

**In vivo experiments.** The fastMRI prostate data were acquired at 3T. Additional in-house scans included one healthy subject at 3T (Siemens Vida), and two healthy subjects and five patients with suspected prostate cancer at 0.55T (Siemens FreeMax). Axial T2-weighted turbo-spin-echo (TSE) images were acquired with following parameters: FOV=180×180mm$^2$, resolution=0.56×0.56mm$^2$, slice thickness 3mm, 6x-retrospectively undersampled (fastMRI); FOV=200×200mm$^2$, resolution=0.52×0.52mm$^2$, slice thickness 3mm, 3×-prospective undersampled (3T healthy subject); FOV=220x220mm$^2$, resolution=0.85×0.85mm$^2$, slice thickness 3mm, 4×-prospectively undersampled (0.55T health subjects and patients).

## Results

**Figure 2a** shows 6×-undersampled fastMRI prostate T2-weighted reconstructions; **Table I** summarizes metrics. PaDIS matches whole-image diffusion performance while yielding sharper prostate detail in the zoom, trains ~3× faster and uses ~2× less GPU memory. Quantitatively, PaDIS averaged 32.66 dB PSNR and 0.915 SSIM, outperforming $\ell 1$-wavelet (31.41 dB, 0.879) and GRAPPA (30.05 dB, 0.869).

**Figure 3** shows images from in-house scans at 3T. GRAPPA and $\ell 1$-wavelet exhibit residual aliasing anteriorly near the pubic symphysis, and GRAPPA appears noisier overall; PaDIS suppresses both artifacts and noise.

**Figure 4** shows images from two healthy subjects at 0.55T. PaDIS improves SNR over other methods, whereas GRAPPA shows fold-over in the femoral

heads.

Figure 5 shows five 0.55T patient cases. Across all patients, PaDIS improved SNR while maintaining low aliasing. Two board-certified radiologists scored artifact level, SNR, overall diagnostic quality, and blurriness in a blinded read. PaDIS had similar artifact and overall quality scores to $\ell 1$-wavelet (4.4 and 3.6 vs. 4.4 and 3.6) and outperformed GRAPPA (3.3 and 2.4). PaDIS achieved the highest SNR score (3.9 vs. 3.5 for $\ell 1$-wavelet and 3.2 for GRAPPA), but was rated slightly blurrier (3.2) than $\ell 1$-wavelet (3.6) and similar to GRAPPA (3.2).

## Discussion and Conclusion

PaDIS introduces a data-driven diffusion prior for accelerated prostate MRI reconstruction that requires only patch-based training, reducing memory requirements and data demands compared to whole-image diffusion models. Across both 3T and 0.55T acquisitions, PaDIS demonstrated improved SNR, reduced aliasing artifacts compared to images reconstructed using GRAPPA and $\ell 1$-wavelet methods, and received similar or higher preference ratings from radiologists. This approach can potentially enable higher-quality, faster prostate MRI in routine clinical workflows, especially for low-field imaging, and is readily extensible to imaging of other anatomical regions.

## References


1. Y. Song and S. Ermon, "Generative Modeling by Estimating Gradients of the Data Distribution," in Advances in Neural Information Processing Systems, vol. 32, 2019.

2. J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," in Advances in Neural Information Processing Systems, vol. 33, pp. 6840-6851, 2020.

3. H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, "Diffusion Posterior Sampling for General Noisy Inverse Problems," in Proceedings of the International Conference on Learning Representations, 2023.

4. T. Karras, M. Aittala, T. Aila, and S. Laine, "Elucidating the Design Space of Diffusion-Based Generative Models," in Proceedings of the Conference on Neural Information Processing Systems, 2022.

5. C. Cao et al., "High-Frequency Space Diffusion Model for Accelerated MRI," IEEE Transactions on Medical Imaging, vol. 43, no. 5, pp. 1853-1865, May 2024. https://doi.org/10.1109/TMI.2024.3351702

6. J. Hu, B. Song, X. Xu, L. Shen, and J. A. Fessler, "Learning Image Priors Through Patch-Based Diffusion Models for Solving Inverse Problems," in Advances in Neural Information Processing Systems, vol. 37, 2024.

7. R. Sanda, A. Aali, A. Johnston, E. Reis, J. Singh, G. Wetzstein, and S. Fridovich-Keil, "Patch-Based Diffusion for Data-Efficient, Radiologist-Preferred MRI Reconstruction," arXiv preprint arXiv:2509.21531, 2025.

8. 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, vol. 47, no. 6, pp. 1202-1210, 2002. https://doi.org/10.1002/mrm.10171

9. M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," Magnetic Resonance in Medicine, vol. 58, no. 6, pp. 1182-1195, 2007. https://doi.org/10.1002/mrm.21391

10. R. Tibrewala, T. Dutt, A. Tong, et al., "FastMRI Prostate: A public, biparametric MRI dataset to advance machine learning for prostate cancer imaging," Scientific Data, vol. 11, p. 404, 2024. https://doi.org/10.1038/s41597-024-03252-w


## Figures and Tables

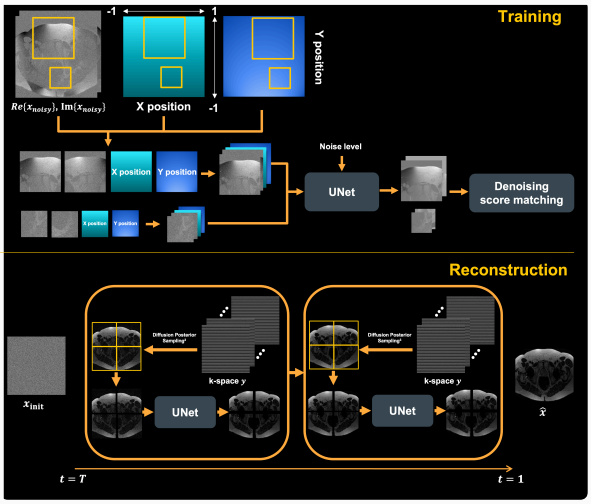


Scan for high-resolution version

***Figure 1:*** *(a) PaDIS training for prostate MRI. Random patches of the complex image are extracted and stacked with normalized X/Y position arrays as two extra channels, then fed to a UNet with denoising score matching. (b) Reconstruction. Starting from noise at t = T, Diffusion Posterior Sampling (DPS)[3] data consistency and the patch UNet are applied to random tilings of the zero-padded image to estimate the full-image score.*

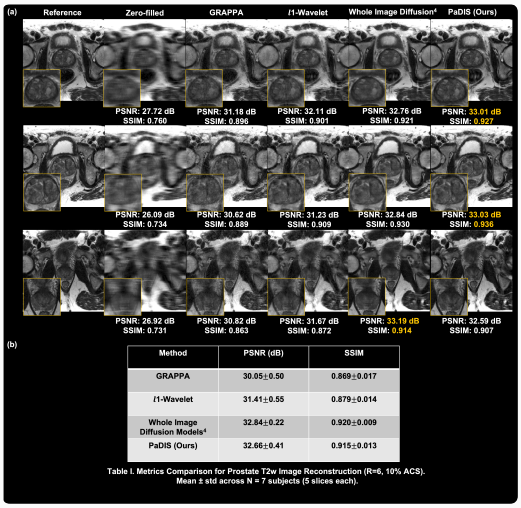


Scan for high-resolution version

***Figure 2:*** *(a) Prostate T2-weighted images from the fastMRI 3T dataset. k-space was retrospectively undersampled 6× with 10% fully-sampled center. (b) Quantitative comparison (PSNR / SSIM) of GRAPPA, ℓ1-wavelet, whole-image diffusion, and PaDIS, averaged over 7 subjects (5 slices each).*

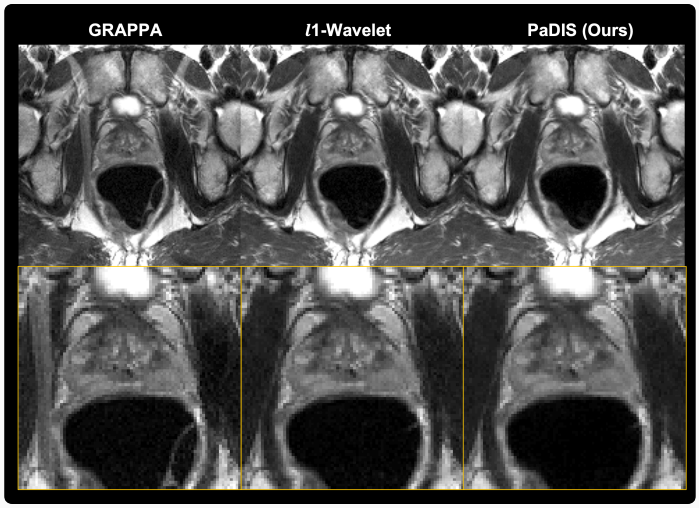


Scan for high-resolution version

***Figure 3:*** *3T healthy subject reconstruction comparisons (3× undersampled without fully sampled center k-space) on clinical axial T2 weighted acquisition with FOV=200×200mm$^2$, resolution=0.52×0.52mm$^2$ and 3mm slice thickness.*

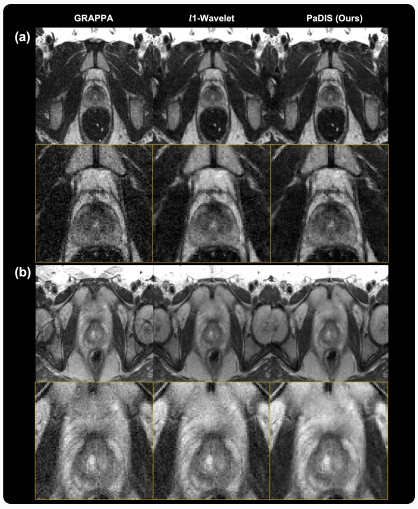


Scan for high-resolution version

***Figure 4:*** *0.55T healthy subjects reconstruction comparisons (3× undersampled without fully sampled center k-space) on clinical axial T2 weighted acquisition with FOV=220×220mm$^2$, resolution=0.85×0.85mm$^2$ and 3mm slice thickness.*

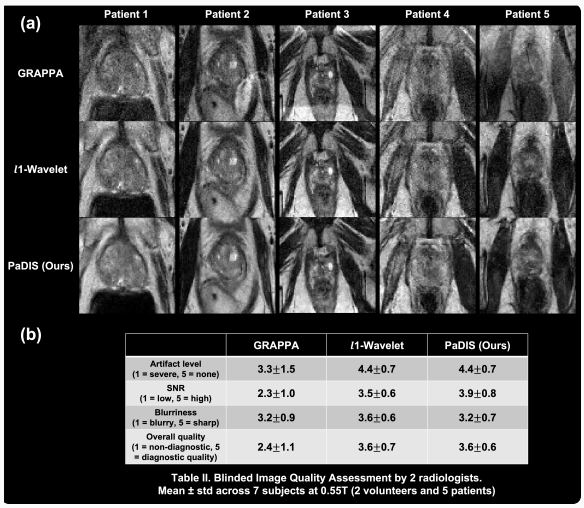


| | GRAPPA | *l*1-Wavelet | PaDIS (Ours) |
|---|---|---|---|
| Artifact level (1 = severe, 5 = none) | 3.3±1.5 | 4.4±0.7 | 4.4±0.7 |
| SNR (1 = low, 5 = high) | 2.3±1.0 | 3.5±0.6 | 3.9±0.8 |
| Blurriness (1 = blurry, 5 = sharp) | 3.2±0.9 | 3.6±0.6 | 3.2±0.7 |
| Overall quality (1 = non-diagnostic, 5 = diagnostic quality) | 2.4±1.1 | 3.6±0.7 | 3.6±0.6 |

Table II. Blinded Image Quality Assessment by 2 radiologists. Mean ± std across 7 subjects at 0.55T (2 volunteers and 5 patients)

Scan for high-resolution version

***Figure 5:*** *(a) 0.55T Patient Reconstruction comparisons (4x undersampled without fully-sampled center k-space) (f). Blinded radiologist Image quality assessment for volunteers and patients reconstruction in Figure 4 and 5 (a).*