Plug-and-Play Denoising Priors in Magnetic Resonance Image Reconstruction

Ben Hartwick, Department of Computer Science

Abstract

This project evaluates deep denoising models as plug-and-play priors for accelerated MR image reconstruction using an unrolled HQS-based framework. DiffUNet is introduced as a new denoising prior in this framework and is compared with two established models within the same reconstruction pipeline. Experiments using the fastMRI knee dataset show that DiffUNet produces sharper and more accurate reconstructions than the other priors. These findings indicate that DiffUNet is a promising addition to plug-and-play MRI reconstruction methods.

Introduction

Magnetic Resonance Imaging (MRI) is an essential tool in neurology, oncology, and musculoskeletal medicine. The technique uses magnetic fields and radio-frequency pulses to acquire detailed structural information within the human body. However, the technology samples the Fourier transform domain (k-space), requiring long scan times. Long acquisition periods remain a major limitation, as patients must stay still to obtain clear images, increasing discomfort and the risk of motion artifacts [1]. To reduce scan time, proposed solutions include advances in pulse sequence technology [2] and parallel imaging [3]. While effective, these approaches often come with high costs and technical barriers. An alternative strategy, and the focus of this project, is to capture less information by undersampling k-space and reconstructing the fully-sampled image computationally [4,5].

Image reconstruction from limited measurements is an inherently ill-posed problem, and ensuring accurate recovery of anatomical structures is challenging. Researchers have therefore developed methods to preserve image fidelity, including low-rank matrix recovery for cardiac MRI [6]. However, *these methods may struggle with complex textures and fine details*. An alternative is the plug-and-play (PnP) framework for MR image reconstruction, which incorporates a denoiser model as a regularization prior [7]. Prior PnP work has used denoisers such as Swin-Conv U-Net (SCUNet), an architecture combining local and non-local modeling through convolutional layers and Swin Transformer blocks [8]. Recent work has implemented PnP frameworks with multiple deep network priors, leveraging modern deep learning models that demonstrate improvements over earlier approaches [9]. The goal of this project is to test whether a new denoising prior, DiffUNet, can further improve image reconstruction. Although the project was limited in time and computational resources, results show that DiffUNet is a promising prior for MR image reconstruction.

Related Work

Serving as an authoritative dataset in MR image reconstruction, the NYU fastMRI Initiative database has been widely used for training and testing reconstruction models [10]. The dataset provides fully sampled single-coil and multi-coil MR images across several anatomies, including brain, knee, prostate, and breast imaging, and has supported benchmarking through a public reconstruction challenge comparing methods across various undersampling factors. This resource established a standardized foundation for evaluating reconstruction methods and remains central in current research.

The plug-and-play framework (PnP) for MR image reconstruction was first introduced with convolutional denoisers as a regularization prior [7]. DnCNN continues to act as a baseline prior in current PnP work, including in this project. Recent methods implement PnP reconstruction frameworks with multiple deep network priors. Wang et al. employed both SCUNet and BM3D within an HQS-based PnP algorithm, demonstrating improved performance over approaches using a single prior [9]. These results motivate the exploration of more advanced denoisers, such as DiffUNet, within PnP reconstruction pipelines.

Methods

DeepInverse Library:

This project leveraged the DeepInverse (deepinv) library to streamline all stages of the MRI reconstruction pipeline, including dataset creation, model definition, training, and evaluation [11]. DeepInverse was used to generate undersampling masks, apply the forward MRI operator, and assemble paired training data consisting of fully sampled images and their masked k-space measurements. Its modular physics operators and model wrappers allowed different denoiser priors (DnCNN, SCUNet, and DiffUNet), to be inserted directly into the unrolled reconstruction network. The training loop, optimization routines, and checkpointing were implemented using DeepInverse utilities, enabling consistent training across all models. For evaluation, DeepInverse provided convenient tools for running reconstructions, computing quantitative metrics such as PSNR and SSIM, and visualizing reconstructions for qualitative comparison. Together, these components allowed the entire reconstruction pipeline to be implemented cleanly and reproducibly within a unified framework, from data generation to results visualization.

Dataset Creation:

Datasets were created by simulating undersampling on fully sampled MR images. A Gaussian mask was applied in the Fourier domain at an acceleration factor of 4, meaning that only one quarter of the k-space frequencies were retained for reconstruction. This process generated paired fully sampled and undersampled images for training and evaluation. The knee subset of the fastMRI training and validation sets was used for this project, with 973 images allocated for training and 199 images for testing. To ensure consistent anatomical coverage, the middle slice from each volumetric scan was selected.

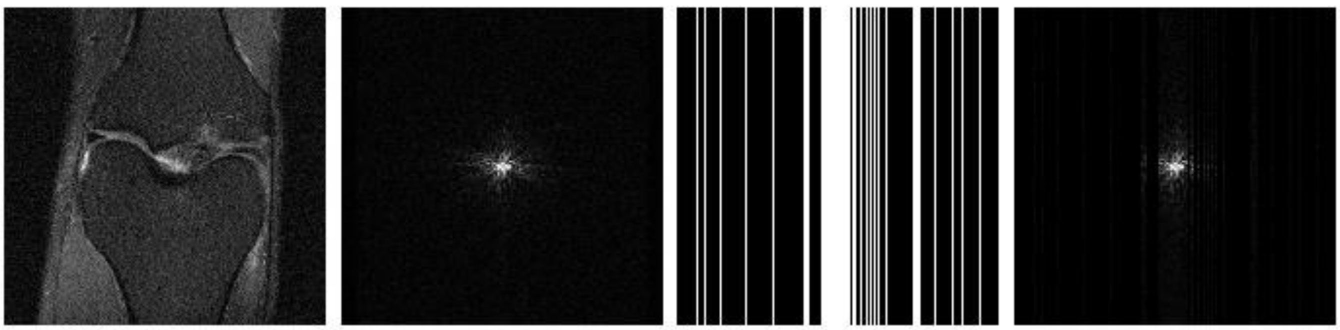

Fully Sampled Image	Fully Sampled K-Space Image	Gaussian Mask	Masked K-space Image

Figure 1. Illustration of the dataset creation process. Fully sampled images were Fourier-transformed to obtain fully sampled k-space data, which was then multiplied by a Gaussian undersampling mask to produce masked k-space measurements used for reconstruction.

Half Quadratic Splitting Algorithm:

$$\underset{\{x\}}{minimize}\ \tfrac{1}{2}\|Ax-b\|_2^2+\lambda\Psi(z),\ \ subject\ to\ x-z=0$$

Half-Quadratic Splitting (HQS) is widely used in inverse imaging because it cleanly separates the data-fidelity and prior terms, allowing each to be optimized with efficient sub-problems. In each iteration, HQS alternates between enforcing data consistency with the measured k-space and applying a learned prior through a denoising step, which makes it particularly suitable for plug-and-play reconstruction methods [12].

Unfolded Reconstruction Framework (MoDL):

MoDL is implemented here as a simple unrolled architecture based on half-quadratic splitting, where each iteration alternates between a physics-based data-consistency update and a proximal step. In this formulation, the proximal operator is replaced by a trainable denoising prior, yielding an efficient model-based reconstruction framework [13]. The original MoDL design uses a CNN denoiser with shared weights across cascades; in this study, this denoising block is replaced with the three priors under evaluation (DnCNN, SCUNet, and DiffUNet). Our network uses two cascades, each consisting of a data-consistency module and a denoising module. The model was trained for each evaluation with DeepInverse's Trainer module with the Pytorch Adam optimizer [11, 14].

Denoiser Priors:

DnCNN is a classic CNN-based denoiser composed of a deep stack of convolutional layers equipped with ReLU activations [15]. Instead of predicting the clean image directly, DnCNN learns to estimate the noise present in the input, allowing the clean output to be recovered by subtraction. Its advantages come from using many relatively simple convolutional layers to capture increasingly abstract features of the noise distribution, enabling it to generalize well across different noise levels. This architecture became a foundational baseline for image denoising because it is efficient, easy to train, and highly effective for Gaussian noise.

SCUNet extends the traditional U-Net design by combining local feature extraction from convolutional layers with non-local modeling provided by Swin Transformer blocks [8]. The U-Net structure supplies strong multi-scale encoding and decoding pathways—important for capturing both fine details and global structure—while the Swin Transformers introduce windowed self-attention that can model long-range dependencies that pure CNNs struggle with. This hybrid architecture enables SCUNet to handle complex, spatially varying noise patterns and produce more perceptually consistent reconstructions, making it a strong modern denoiser for diverse real-world degradations.

DiffUNet is a U-Net-based denoising model originally developed for use in diffusion frameworks, where it typically predicts noise or clean image structure at different stages of a diffusion process [16]. Although diffusion models usually rely on many iterative timesteps, DiffUNet can also be used as a single-step denoiser when the timestep input is fixed or omitted. In this project, DiffUNet is used in this simplified way: it serves as the denoising prior within the Half-Quadratic Splitting (HQS) reconstruction algorithm. Instead of running a full diffusion pipeline, the model simply takes intermediate image estimates and returns a cleaned version that helps guide the reconstruction toward more accurate and natural-looking MR images. This repurposing allows the project to benefit from DiffUNet's strong denoising capability.

Evaluation:

Reconstruction quality was evaluated using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). These metrics are standard in MR image reconstruction, and the fastMRI challenge ranked submissions using SSIM in conjunction with radiologist evaluations [17]. Although radiologist scores were not available for this project, PSNR and SSIM provided consistent quantitative measures for comparing denoiser priors.

Results

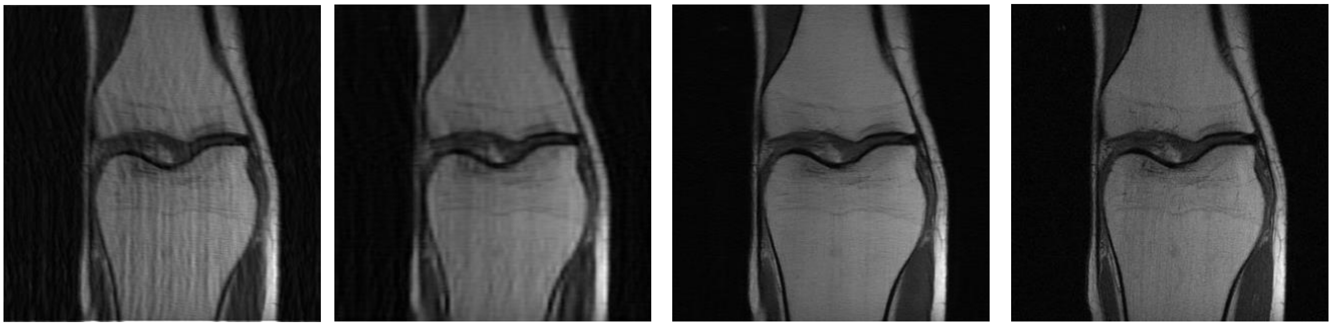

DnCNN Reconstruction | SCUNet Reconstruction | DiffUNet Reconstruction | Fully Sampled Image

Figure 2. Qualitative comparison of reconstructions using different denoiser priors.

Denoiser Prior	PSNR (dB)	SSIM
DnCNN	24.6	0.518
SCUNet	26.8	0.533
DiffUNet	**29.1**	**0.711**

Table 1. Quantitative reconstruction metrics from each denoiser prior.

In the qualitative comparison shown in Figure 2, DnCNN and SCUNet exhibit a loss of fine structural detail and artifacts, including patterned vertical streaking. In contrast, DiffUNet produces sharper and more anatomically faithful reconstructions that more closely resemble the fully sampled ground-truth image.

Table 1 summarizes the quantitative metrics for each denoiser prior. DiffUNet achieves the highest PSNR and SSIM values, outperforming both DnCNN and SCUNet under the same reconstruction settings. Numerical results indicate a consistent improvement in reconstruction fidelity when DiffUNet is used as the plug-and-play prior.

Discussion

DiffUNet produced the most accurate reconstructions, with fewer visible artifacts and more consistent recovery of anatomical details compared to DnCNN and SCUNet. However, all methods struggled to reproduce high-frequency structures present in the ground-truth images, suggesting that additional training data, stronger priors, or improved masking strategies may be required. Performance was also limited by the relatively small dataset size (~1000 images of a single anatomy type) and the use of a fixed acceleration factor of 4, both of which constrained model generalization. Compared to previously reported results on the fastMRI single-coil benchmark, where the highest-performing methods achieved SSIM scores around 0.754 and PSNR values above 32 dB, this project's DiffUNet reconstructions fall slightly short, likely due to limited training time, restricted model capacity, and the use of single-slice datasets rather than full volumetric training. More advanced variants of the evaluated priors, as well as training on multi-coil data or larger cohorts, could further improve reconstruction fidelity.

Bibliography

[1] 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, doi: https://doi.org/10.1002/mrm.21391.

[2] S. B. Reeder and A. Z. Faranesh, "Ultrafast Pulse Sequence Techniques for Cardiac Magnetic Resonance Imaging," *Topics in Magnetic Resonance Imaging*, vol. 11, no. 6, pp. 312–330, Dec. 2000, doi: https://doi.org/10.1097/00002142-200012000-00002.

[3] K. P. Pruessmann, "Encoding and reconstruction in parallel MRI," *NMR in Biomedicine*, vol. 19, no. 3, pp. 288–299, May 2006, doi: https://doi.org/10.1002/nbm.1042.

[4] K. G. Hollingsworth, "Reducing acquisition time in clinical MRI by data undersampling and compressed sensing reconstruction," *Physics in Medicine and Biology*, vol. 60, no. 21, pp. R297–R322, Oct. 2015, doi: https://doi.org/10.1088/0031-9155/60/21/r297.

[5] S. Ravishankar and Y. Bresler, "MR Image Reconstruction From Highly Undersampled k-Space Data by Dictionary Learning," *IEEE Transactions on Medical Imaging*, vol. 30, no. 5, pp. 1028–1041, May 2011, doi: https://doi.org/10.1109/tmi.2010.2090538.

[6] J. Kent, *The Blah*. London: Abelard-Schuman Limited, 1974.

[7] R. Ahmad *et al.*, "Plug-and-Play Methods for Magnetic Resonance Imaging: Using Denoisers for Image Recovery," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 105–116, Jan. 2020, doi: https://doi.org/10.1109/msp.2019.2949470.

[8] K. Zhang *et al.*, "Practical Blind Image Denoising via Swin-Conv-UNet and Data Synthesis," *Machine Intelligence Research*, vol. 20, no. 6, pp. 822–836, Sep. 2023, doi: https://doi.org/10.1007/s11633-023-1466-0.

[9] J. Wang, C. Liu, Y. Zhong, X. Liu, and J. Wang, "Deep plug-and-play MRI reconstruction based on multiple complementary priors," *Magnetic Resonance Imaging*, vol. 115, p. 110244, Jan. 2025, doi: https://doi.org/10.1016/j.mri.2024.110244.

[10] F. Knoll *et al.*, "fastMRI: A Publicly Available Raw k-Space and DICOM Dataset of Knee Images for Accelerated MR Image Reconstruction Using Machine Learning," *Radiology: Artificial Intelligence*, vol. 2, no. 1, p. e190007, Jan. 2020, doi: https://doi.org/10.1148/ryai.2020190007.

[11] J. Tachella *et al.*, "DeepInverse: A Python package for solving imaging inverse problems with deep learning," *arXiv (Cornell University)*, May 2025, doi: https://doi.org/10.48550/arxiv.2505.20160.

[12] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," *IEEE transactions on image processing*, vol. 4, no. 7, pp. 932–946, Jul. 1995, doi: https://doi.org/10.1109/83.392335.

[13] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-Based Deep Learning Architecture for Inverse Problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, Feb. 2019, doi: https://doi.org/10.1109/tmi.2018.2865356.

[14] A. Paszke *et al.*, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," *arXiv (Cornell University)*, Dec. 2019, doi: https://doi.org/10.48550/arxiv.1912.01703.

[15] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017, doi: https://doi.org/10.1109/tip.2017.2662206.

[16] J.-S. Choi, S.-W. Kim, Y. Jeong, Y. Gwon, and S. Yoon, "ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models," *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2021, doi: https://doi.org/10.1109/iccv48922.2021.01410.

[17] F. Knoll *et al.*, "Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge," vol. 84, no. 6, pp. 3054–3070, Jan. 2020, doi: https://doi.org/10.1002/mrm.28338.