Classical Compressed Sensing MRI Reconstruction with Sparse Priors and Undersampling Mask Analysis

Haichuan Liu, Anisha Tehim

Abstract—We study compressed sensing for MRI reconstruction from undersampled k-space data. A baseline wavelet-based method is implemented and extended using alternative sparsity priors, including DCT, total variation (TV), and hybrid approaches. Results show that Wavelet provides an efficient and strong baseline, while TV-POCS and Wavelet+TV-POCS further improve reconstruction quality at higher computational cost.

Index Terms—Compressed sensing, MRI reconstruction, sparsity priors, computational imaging

1 Introduction

Magnetic Resonance Imaging (MRI) is a medical imaging technique that produces images of soft tissue by measuring radiofrequency signals emitted by hydrogen protons in a magnetic field. These measurements are acquired in the Fourier (k-space) domain. Conventional MRI requires dense sampling of k-space, resulting in long acquisition times that increase patient discomfort and potentially limits clinical throughput. Undersampling can accelerate imaging but introduces severe aliasing artifacts under direct inverse Fourier reconstruction. Therefore, compressed sensing aims to reconstruct MRI images from undersampled k-space by enforcing data consistency in the Fourier domain and sparsity in a transform domain, reducing aliasing artifacts relative to zero-filled reconstruction. While wavelet-based reconstruction improves results, performance depends on the choice of sparsity prior. Evaluating multiple priors reveals trade-offs between artifact suppression and detail preservation, showing that prior selection strongly impacts reconstruction quality.

2 Forward Model

Magnetic Resonance Imaging (MRI) acquires data in the spatial frequency domain, commonly referred to as k-space. Unlike conventional imaging systems that directly measure pixel intensities, MRI scanners measure Fourier coefficients of the underlying image. Specifically, the measured signal corresponds to samples of the Fourier transform of the image, which encodes global and local spatial structure in terms of frequency components.

This acquisition process can be modeled as

$$y = M\mathcal{F}x + n \tag{1}$$

where $x \in \mathbb{C}^N$ denotes the unknown image in the spatial domain, $\mathcal{F}$ is the discrete Fourier transform operator, M is a binary sampling mask that specifies which k-space locations are measured, y represents the acquired k-space data, and n models measurement noise.

Undersampling is commonly used to accelerate MRI acquisition, as collecting fewer k-space measurements reduces scan time. However, this results in incomplete data, since the number of acquired measurements is smaller than the number of unknown image pixels. Consequently, the forward model defines an underdetermined inverse problem, where infinitely many images can produce the same observed measurements.

It is important to note, that undersampling in MRI is not the same as computing the full Fourier transform and discarding measurements. In practice, only a subset of Fourier coefficients is directly acquired according to a sampling trajectory in k-space, and the remaining coefficients are never measured. The sampling mask M provides a mathematical representation of this process by indicating which frequency components are observed.

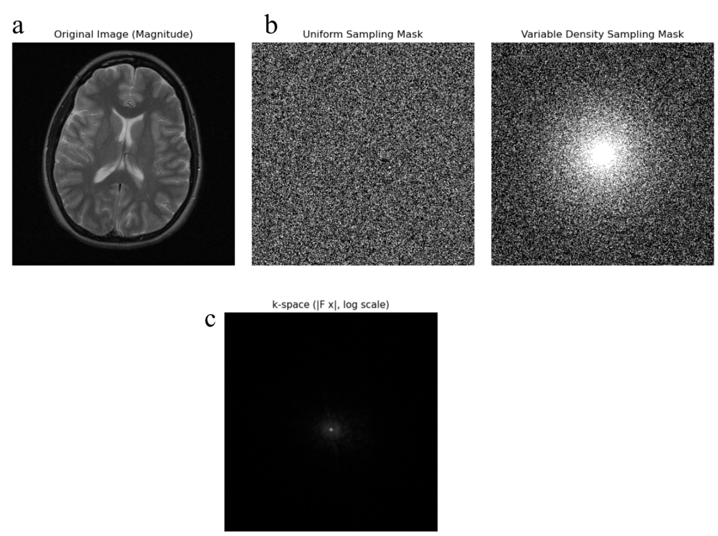

Fig. 1. MRI forward model components. (a) Original image and sampling masks (uniform and variable-density). (b) Fourier transform of the image (k-space).

A direct reconstruction via the inverse Fourier transform implicitly assumes that missing k-space values are zero, leading to aliasing artifacts in the image domain. These artifacts are a result of the loss of high-frequency information and appear as overlapping or distorted structures in the reconstructed image. Therefore, recovering the original image from undersampled measurements requires incorporating additional constraints that reflect prior knowledge about the structure of the images.

3 Inverse Problem

Let x denote the unknown image in the spatial (image) domain that we aim to reconstruct. The measured data y consists of under-sampled Fourier (k-space) measurements of x, obtained according to the forward model, where $\mathcal{F}$ is the discrete Fourier transform, M is a binary sampling mask that selects which k-space locations are measured, and n represents measurement noise.

Because the sampling mask M removes a subset of Fourier coefficients, the number of observed measurements in y is smaller than the number of unknown pixel values in x. As a result, the inverse problem of recovering x from y is underdetermined.

A direct reconstruction is obtained by applying the inverse Fourier transform to y, which corresponds to a zero-filled reconstruction where missing k-space values are set to zero. This produces aliasing artifacts in the image domain due to the loss of frequency information. To address this, compressed sensing assumes that the image x is sparse or compressible in a transform domain. Let W denote a transform operator (e.g., a wavelet transform), such that Wx represents the transform coefficients of the image.

The reconstruction is then formulated as the optimization problem

$$\hat{x} = \arg\min_{x} \frac{1}{2}\|M\mathcal{F}x - y\|_2^2 + \lambda\|Wx\|_1, \tag{2}$$

where $\hat{x}$ is the reconstructed image. The first term enforces data consistency, requiring that the Fourier transform of the reconstructed image matches the measured k-space data at sampled locations defined by M. The second term is an L_1 regularization term that promotes sparsity of the transform coefficients Wx. The parameter $\lambda > 0$ controls the trade-off between enforcing agreement with the measured data and promoting sparsity. By restricting the solution to images that are both consistent with the measured k-space data and sparse in the transform domain, this reduces the number of possible solutions and enables accurate reconstruction from undersampled measurements.

4 Methods

We implemented a baseline compressed sensing MRI reconstruction using a POCS-based algorithm with wavelet sparsity. The reconstruction alternates between enforcing sparsity via soft-thresholding in the wavelet domain and enforcing data consistency in k-space.

We then evaluated the effect of different design choices by varying:

- Sampling masks: uniform, variable-density, Cartesian, and radial undersampling patterns
- Sparsity priors: wavelet, discrete cosine transform (DCT), total variation (TV), and hybrid combinations
- Regularization parameter: λ values controlling the sparsity strength

All reconstructions were performed using the same iterative framework, enabling direct comparison across masks, priors, and parameter settings.

5 Results and Findings

5.1 Baseline Reconstruction

A baseline reconstruction is obtained by applying the inverse Fourier transform to the undersampled k-space data:

$$x_{\text{zf}} = \mathcal{F}^{-1}(y) \tag{3}$$

This zero-filled reconstruction exhibits strong aliasing artifacts due to missing frequency components.

We define the baseline reconstruction as using a variable undersampling mask and a wavelet-based sparsity prior, with regularization parameter $\lambda = 0.025$. To improve reconstruction quality, we solve the compressed sensing formulation introduced in the previous section using a POCS-based iterative algorithm.

Reconstruction Algorithm

1) Transform to wavelet domain: $x_w = Wx$
2) Apply soft-thresholding: $x_w = S(x_w, \lambda)$
3) Inverse transform: $x = W^{-1}x_w$
4) Enforce data consistency in k-space
5) Repeat for multiple iterations

Soft-thresholding promotes sparsity by removing small wavelet coefficients associated with artifacts, while the data consistency step ensures agreement with the measured k-space data at sampled locations. The parameter λ controls the strength of the sparsity constraint.

Figure 2 shows the resulting reconstruction. Compared to the zero-filled image, the compressed sensing reconstruction significantly reduces aliasing artifacts. The relative error curve demonstrates convergence over successive iterations.

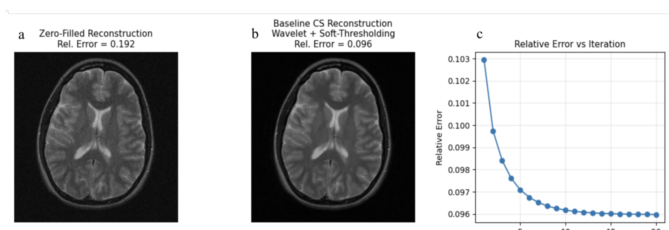

Fig. 2. Baseline reconstruction. (a) Zero-filled reconstruction showing aliasing artifacts. (b) Compressed sensing reconstruction using wavelet sparsity and soft-thresholding. (c) Relative error versus iteration, demonstrating convergence.

5.2 Extension Reconstruction

In the extension reconstruction part, three factors that can affect the reconstruction quality are further investigated: the undersampling mask, the sparse prior, and the regularization parameter λ. The goal is to test whether the baseline reconstruction can be improved by changing the sampling strategy and sparse priors.

Undersampling Mask

The original dataset provides two undersampling masks. However, since the sampling pattern in k-space directly affects the information preserved during acquisition, it is worth exploring other possible masks.

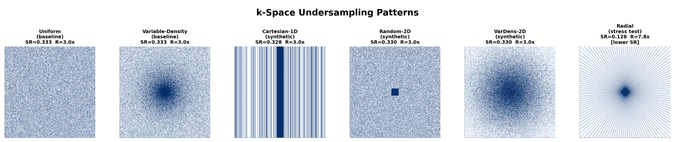

Fig. 3. k-space undersampling patterns. Uniform, Variable-Density, Cartesian-1D, Random-2D and Variable-Density-2D patterns all have sampling rates close to 33% ($R \approx 3\times$). The radial pattern has a lower sampling rate of 12.8% ($R \approx 7.8\times$) and is included as a stress test.

The performance of different undersampling masks is shown in Figure 4. For each mask, several reconstruction methods are compared, including zero-filled IFFT, Wavelet-POCS, DCT-POCS, TV-POCS, and Wavelet+TV-POCS.

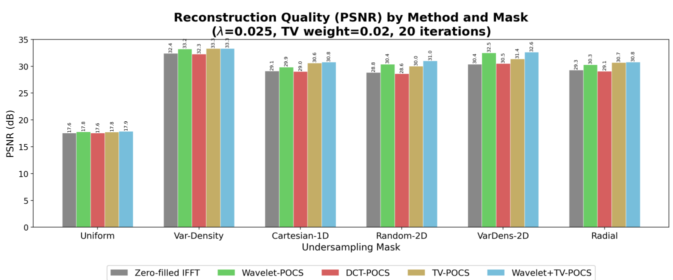

Fig. 4. PSNR comparison of different reconstruction methods under different undersampling masks.

Among the masks with similar sampling rates, the Variable-Density mask gives the best reconstruction quality on this dataset. This is consistent with the structure of MRI k-space, where low-frequency coefficients near the center contain most of the image contrast and global anatomical structure. Therefore, preserving more central k-space samples helps produce a stronger initial reconstruction.

Sparse Priors

In the baseline reconstruction, Wavelet sparsity is used as the main prior. In this extension, DCT sparsity and TV regularization are also tested. In addition, two combined priors are examined: Hybrid-POCS, which combines Wavelet and DCT sparsity, and Wavelet+TV-POCS, which combines wavelet thresholding with TV denoising.

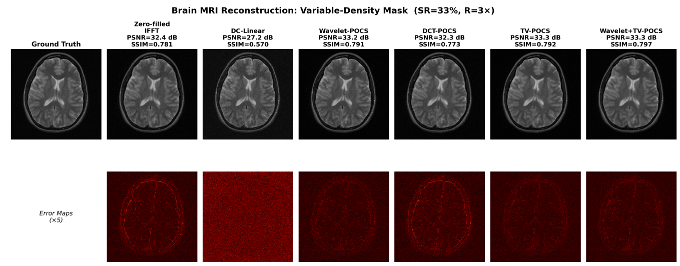

Fig. 5. Representative reconstruction results under the Variable-Density mask with sampling rate 33.3% ($R \approx 3\times$). The first row shows reconstructed images, and the second row shows error maps amplified by $5\times$.

Figure 5 shows that POCS-based methods improve the zero-filled IFFT reconstruction by reducing aliasing artifacts and improving structural similarity. Among the tested methods, Wavelet+TV-POCS gives the best visual quality and the highest PSNR/SSIM in this example.

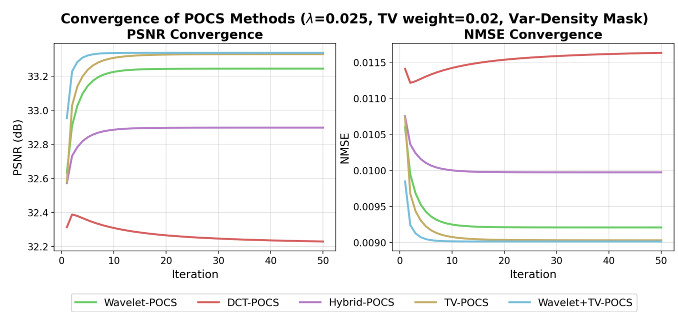

Fig. 6. Convergence of POCS methods. $\lambda = 0.025$, $TV_{weight} = 0.02$, undersampled by the Variable-Density mask.

The convergence curves in Figure 6 show that Wavelet+TV-POCS increases PSNR rapidly in the first few iterations and reaches the best final performance. TV-POCS also performs well, while DCT-POCS gives lower quality and is less stable in terms of NMSE. This suggests that wavelet sparsity and TV regularization are complementary: wavelet thresholding removes sparse-domain artifacts, while TV regularization helps preserve smooth regions and suppress local noise.

PSNR (dB) Summary by Mask and Method [green = best per row]

	Zero-filled IFFT	Wavelet-POCS	DCT-POCS	TV-POCS	Wavelet+TV POCS
Uniform	17.58	17.81	17.56	17.76	**17.90**
Var-Density	32.42	33.24	32.27	33.33	**33.34**
Cartesian-1D	29.11	29.86	29.02	30.63	**30.79**
Random-2D	28.84	30.40	28.61	30.03	**31.02**
VarDens-2D	30.41	32.53	30.51	31.38	**32.64**
Radial	29.29	30.31	29.09	30.69	**30.79**

Fig. 7. PSNR summary by mask and reconstruction method. Green cells indicate the best method for each mask.

The full PSNR summary is shown in Figure 7. Wavelet+TV-POCS achieves the best PSNR for all tested masks. However, under the Variable-Density mask, the improvement over TV-POCS is relatively small. This means that TV alone is already a strong prior for this dataset, while Wavelet+TV gives slightly better quality at the cost of longer runtime.

Regularization Parameter λ

Another important parameter is the regularization parameter λ. This parameter controls the strength of sparsity enforcement. If λ is too small, the reconstruction may keep more artifacts and noise. If λ is too large, useful image details may also be removed. Therefore, we test different λ values and compare their influence on reconstruction quality.

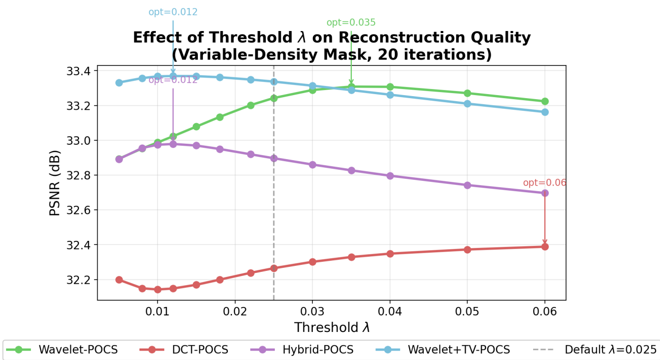

Fig. 8. Effect of the threshold parameter λ on reconstruction quality under the Variable-Density mask.

Figure 8 shows that the best λ value depends on the reconstruction prior. The default value $\lambda = 0.025$ gives reasonable results, but it is not optimal for all methods. For Wavelet+TV-POCS, the best performance is achieved around $\lambda = 0.012$, which is smaller than the grid $\lambda = 0.035$ of simply Wavelet. This is reasonable because TV regularization already provides additional smoothing, so the wavelet threshold should be weaker to avoid removing useful details.

Based on this observation, we further test a simple mask-aware heuristic for choosing λ. Instead of using a fixed threshold for all settings, the heuristic adjusts λ according to the mask and reconstruction prior.

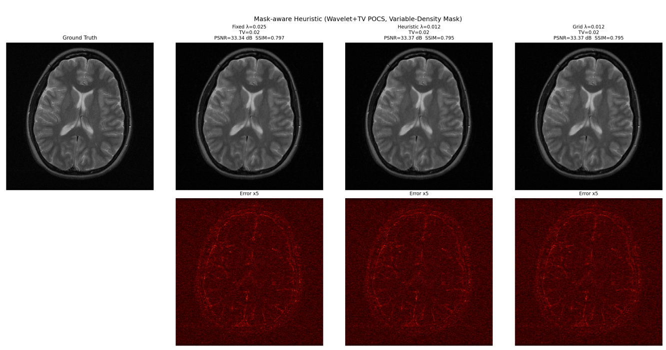

Fig. 9. Comparison between fixed λ, heuristic λ, and grid-search λ for Wavelet+TV-POCS under the Variable-Density mask.

In the future, we can explore the heuristic λ based on datasets, masks and priors, which is helpful to improve the performance.

Final Result

With the best setting found in the extension experiments, we use the Variable-Density mask, Wavelet+TV-POCS prior, and the tuned threshold parameter. The final quantitative comparison is shown in Figure 10.

Method	Best Mask	Sample Rate	PSNR	SSIM	NMSE	Runtime
Wavelet	Var-Den	33.3%	33.2425	0.7910	0.0092	1.9550 sec
DCT	Var-Den	33.3%	32.2652	0.7731	0.0115	1.7874 sec
TV	Var-Den	33.3%	33.3263	0.7924	0.0090	6.3667 sec
Wavelet+TV	Var-Den	33.3%	33.3367	0.7966	0.0090	8.6599 sec

Fig. 10. Final quantitative comparison under the best mask setting.

Wavelet+TV-POCS achieves the best overall reconstruction quality, with PSNR of 33.3367 dB and SSIM of 0.7966. It also obtains one of the lowest NMSE values. However, this method has the longest runtime because it applies both wavelet thresholding and TV denoising during the POCS iterations. TV-POCS gives a very similar PSNR with shorter runtime, so it can be considered a good trade-off between quality and efficiency. However, the most time-consuming part of MRI is scanning, which is taken in unit of minutes. So the longer runtime here is acceptable.

Overall, the extension experiments show that the undersampling mask is the most important factor. The Variable-Density mask already gives a strong zero-filled reconstruction because it preserves more low-frequency k-space information. Sparse priors further improve the result by reducing artifacts and enforcing image structure. Among the tested priors, Wavelet+TV gives the best reconstruction quality, while Wavelet alone provides a strong and faster alternative.

6 Conclusion

Compressed sensing MRI (CS-MRI) reconstructs images from undersampled k-space by enforcing Fourier data consistency and sparsity priors, reducing aliasing artifacts compared to zero-filled reconstruction. Variable-density undersampling achieved the best reconstruction quality across all priors. Wavelet-based reconstruction provides an efficient and strong baseline, while TV-POCS and Wavelet+TV-POCS further improve reconstruction quality at the cost of increased computational complexity.

7 Future Work

Future work includes evaluating the proposed methods on more diverse datasets, including images corresponding to different anatomical structures and disease conditions, to assess robustness and generalization. Additionally, developing adaptive or heuristic strategies for selecting the regularization parameter λ could improve reconstruction quality and reduce the need for manual tuning.

References

[1] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete Cosine Transform," *IEEE Transactions on Computers*, vol. C-23, no. 1, pp. 90–93, 1974.

[2] A. Chambolle, "An Algorithm for Total Variation Minimization and Applications," *Journal of Mathematical Imaging and Vision*, vol. 20, pp. 89–97, 2004.

[3] M. Lustig, D. L. 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.

[4] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image Quality Assessment: From Error Visibility to Structural Similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[5] M. Lustig, D. Donoho, J. Santos, and J. Pauly, "Compressed Sensing MRI," Compressive Sampling, 2008. [Online]. Available: https://people.eecs.berkeley.edu/ mlustig/CS/CSMRI.pdf

[6] M. Lustig, "Assignment 6, Medical Image Reconstruction, EE369C," 2007. [Online]. Available: https://people.eecs.berkeley.edu/~mlustig/CS/CS_ex.pdf