ROI-specific k-space Reconstruction

Chiadika Obinwa

Abstract—Magnetic resonance imaging (MRI) acquisition time is fundamentally constrained by the need to collect sufficient k-space samples, creating a tradeoff between scan duration and image quality. For applications focusing on a specific region of interest (ROI), conventional uniform k-space sampling is inherently inefficient, as it allocates equal resources to all image regions regardless of their clinical relevance. We present a novel framework that accelerates MR acquisition by combining a ROI-specific k-space sampling scheme with compressed sensing to doubly improve acquisition times. Experiments on phantom data with different specified ROIs demonstrate that our approach outperforms standard reconstruction methods in the undersampled case both quantitatively and visually. This approach enables substantially accelerated acquisition without significantly compromising quality in the target region, offering significant advantages for clinical applications where speed and ROI-specific image quality are both needed.

Index Terms—MRI, region of interest, sampling, k-space, compressed sensing, ADMM, total variation regularization

1 Introduction

Image acquisition in magnetic resonance imaging (MRI) is typically an expensive, time-consuming process - as such, current MRI research focuses on accelerating acquisitions to provide equivalent application-specific quality with reduced scan time. While there exist many methods to reduce acquisition times for an entire image, the realm of optimizing for specific regions of interest (ROIs) or images with limited regions of support (ROSs) remains relatively unexplored compared to the global case. Intuitively, one would not need the entirety of densely sampled k-space to reconstruct an image in only a specific region, however the question of which k-space points to select isn't as clear compared to cases where we only care about certain frequencies in the image.

In this work, we propose combining ROI-specific k-space sampling with compressed sensing (CS) principles for efficient MRI reconstruction. Our approach consists of two main components: (1) a sequential forward selection (SFS) algorithm that greedily selects k-space samples to minimize mean-square error for a given ROI, and (2) reconstruction using ADMM with TV regularization that effectively preserves image structure within the region of interest, and additionally allows us to randomly undersample our chosen k-space points for further acceleration. This approach allows for potentially massive acceleration in specific clinical scenarios.

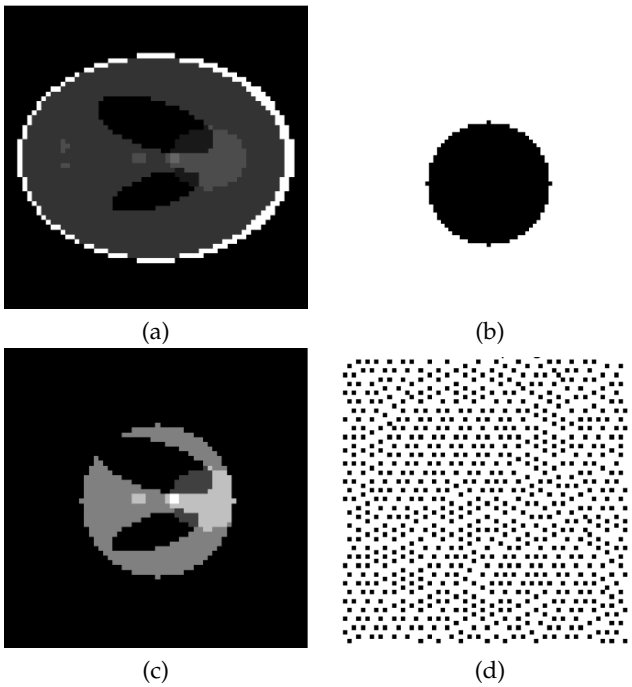

Fig. 1. Experimental setup: (a) Simulated Shepp-Logan phantom used in our experiments. (b) Example of a target region of interest (circular ROI). (c) The Shepp-Logan phantom masked by the target ROI, with brightness enhanced to show contrast. (d) K-space points optimized to minimize MSE in the least-squares reconstruction of the masked phantom, shown after 4x undersampling.

2 Related Work

2.1 Region of Support K-Space Sampling

When MR images have a limited region of support (ROS), it is possible to reconstruct the image from a subset of k-space samples. The authors of (1) introduced an approach for optimal k-space sampling in MRI for images with a limited ROS, using SFS to identify the optimal combination of k-space samples that minimizes noise amplification during reconstruction.

The SFS algorithm starts with an empty set and sequentially adds the k-space sample that provides the greatest improvement in reconstruction quality. However, the standard SFS algorithm cannot be directly applied when the number of selected samples is less than the number of unknowns - Gao and Reeves introduced a modified form of the criterion that overcomes this limitation, enabling the selection of a minimum number of k-space samples (equal to the number of pixels in the ROS) while maintaining exact reconstruction capability in the absence of noise.

3 Preliminaries

3.1 SFS Signal Model and Selection Criteria

We model the observed k-space samples $\mathbf{y}$ as a vector resulting from a linear transformation of the original spatial-

domain image $\mathbf{x}$ in the presence of additive noise $\mathbf{w}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \tag{1}$$

where $\mathbf{w}$ is zero-mean IID noise and $\mathbf{A} \in \mathbb{C}^{m \times n}$. Given the observed signal $\mathbf{y}$, our goal is to reconstruct a good estimate of $\mathbf{x}$ by selecting a limited set of observations that yield the best possible reconstruction. In our ROI-specific case, $\mathbf{A}$ represents a Fourier transform matrix with columns removed corresponding to the voxels outside the ROI.

From (2) and and (3), if the noise is zero-mean, i.i.d., and the reconstruction of $\mathbf{x}$ is performed via least squares, the mean square error (MSE) in the reconstruction is proportional to:

$$E(\mathbf{A}) = \text{tr}\left(\mathbf{A}^H\mathbf{A}\right)^{-1} \tag{2}$$

Our goal, then, is to select rows of A (corresponding to k-space samples) that minimize the quantity in (2). This standard criterion, however, is not defined when $\mathbf{A}$ has fewer rows than columns (i.e., when there are fewer k-space samples than unknowns in the ROS). To address this limitation, the authors of (1) use a modified criterion based on the pseudoinverse:

$$E(\mathbf{A}) = \text{tr}\left(\mathbf{A}^H\mathbf{A}\right)^{+} \tag{3}$$

A key computational insight is that this expression can be reformulated as:

$$E(\mathbf{A}) = \text{tr}\left(\mathbf{A}^H\mathbf{A}\right)^{+} = \sum_{i=1}^{r} \frac{1}{\sigma_i^2} = \text{tr}\left(\mathbf{A}\mathbf{A}^H\right)^{-1} \tag{4}$$

where σ_i are the singular values of $\mathbf{A}$ and $r = \text{rank}(\mathbf{A})$. This transformation significantly improves computational efficiency since $m \ll n$ during the early stages of the selection process.

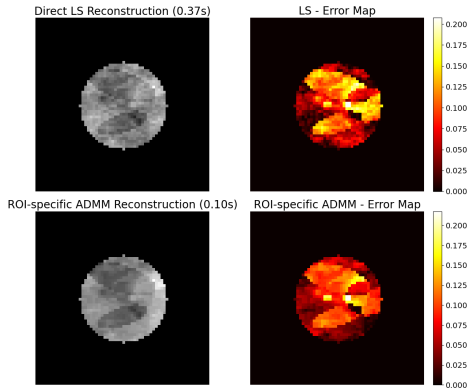

(a) 4x acceleration, optimized sampling

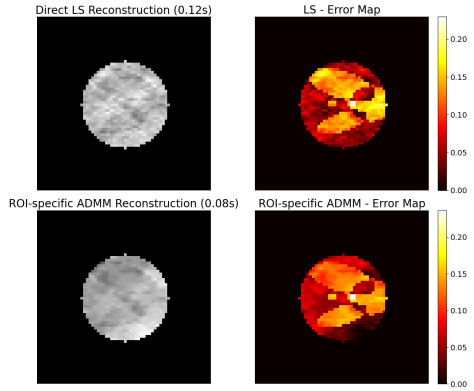

(b) 8x acceleration, optimized sampling

Fig. 2. Reconstruction results for centered circular ROI at different acceleration factors. Each panel shows direct least-squares reconstruction (top left), error map for least-squares (top right), ADMM reconstruction with TV regularization (bottom left), and error map for ADMM reconstruction (bottom right).

3.2 Sequential Forward Selection Algorithm

The SFS algorithm begins with the modified MSE criterion in Equation 4 when there are fewer selected samples than unknowns and then switches to the standard criterion in Equation 2 once the number of selected samples equals the number of unknowns. When matrix $\mathbf{A}$ is underdetermined, there are two parts in the error: one due to the loss of signal components and the other due to noise. The modified criterion reflects the noise error, providing a useful indicator of the degree to which the available components suppress or amplify noise.

For computational efficiency, we implement the recursive formulation proposed in (1) that allows us to efficiently update the selection metrics as new k-space samples are chosen. This approach significantly reduces the computational burden, making it feasible to select optimal k-space sampling patterns in practical scenarios.

3.3 Compressed Sensing Framework

We formulate the MRI reconstruction problem using compressed sensing principles, where we aim to recover an image $\mathbf{x}$ from undersampled k-space measurements $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{A}$ represents the undersampled Fourier encoding matrix. The reconstruction problem can be posed as:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda R(\mathbf{x}) \tag{5}$$

where the first term enforces data consistency with the acquired measurements, and $R(\mathbf{x})$ is a regularization term that promotes sparsity in a transform domain, with λ controlling the regularization strength.

3.4 Total Variation Regularization

For our ROI-specific reconstruction, we employ anisotropic total variation (TV) regularization, which promotes sparsity in the gradient domain while preserving important structural features. While oftentimes wavelet sparsity might be enforced in a compressed sensing scenario, TV regularization can be implemented more easily with arbitrarily-shaped ROIs. The TV regularization term is defined as:

$$\text{TV}(\mathbf{x}) = \sum_i |(\nabla_x \mathbf{x})_i| + |(\nabla_y \mathbf{x})_i| \tag{6}$$

where ∇_x and ∇_y represent finite difference operators in the horizontal and vertical directions, respectively.

3.5 ADMM-Based Reconstruction

To solve the TV-regularized reconstruction problem efficiently, we employ the alternating direction method of

TABLE 1
Mean Squared Error (MSE) Comparison of Phantom Reconstructions

Mask Type	K-space Distribution	Reconstruction Method	Acceleration Factor		
			8x	4x	2x
Centered circular	Optimized	Direct least-squares	0.090551	0.079092	0.051512
		ADMM	0.086187	0.073979	0.042529
	Gaussian	Direct least-squares	0.064172	0.041066	0.016410
		ADMM	0.055514	0.035668	0.016423
Off-center circular	Optimized	Direct least-squares	0.106115	0.099383	0.081342
		ADMM	0.109282	0.093218	0.072694
	Gaussian	Direct least-squares	0.094492	0.080551	0.057094
		ADMM	0.090292	0.073008	0.056342
Centered rectangular	Optimized	Direct least-squares	0.089958	0.067144	0.056687
		ADMM	0.091675	0.058826	0.048437
	Gaussian	Direct least-squares	0.082718	0.057736	0.019545
		ADMM	0.080594	0.048930	0.021190

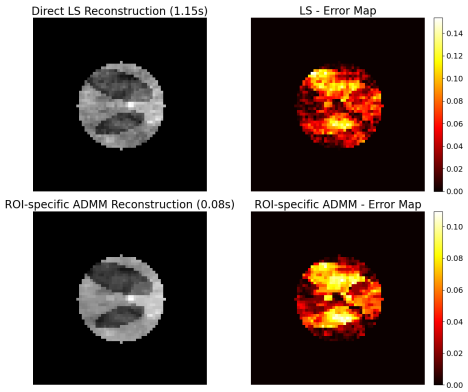

(c) 2x acceleration, optimized sampling

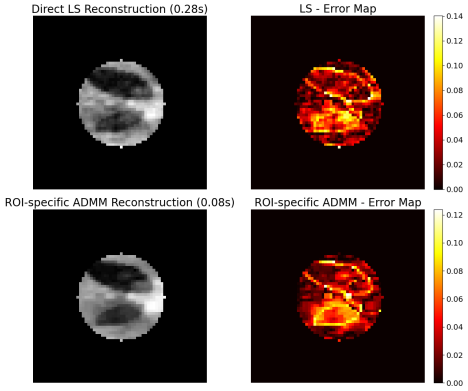

(d) 4x acceleration, Gaussian sampling

Fig. 3. Additional reconstruction results for centered circular ROI. (c) Results at 2x acceleration with optimized sampling, showing improved quality; (d) Results at 4x acceleration with Gaussian sampling for comparison with optimized sampling at the same acceleration factor.

multipliers (ADMM) (4). We introduce auxiliary variables to decouple the data consistency and regularization terms, resulting in the following reformulation:

$$\min_{\mathbf{x},\mathbf{v},\mathbf{z}} \frac{1}{2}\|\mathbf{A}\mathbf{x}-\mathbf{y}\|_2^2 + \lambda\|\mathbf{z}\|_1 \quad \text{subject to } \mathbf{v} = \nabla\mathbf{x}, \mathbf{z} = \mathbf{v} \tag{7}$$

The augmented Lagrangian for this problem is:

$$L_\rho(\mathbf{x},\mathbf{v},\mathbf{z},\mathbf{u},\mathbf{w}) = \frac{1}{2}\|\mathbf{A}\mathbf{x}-\mathbf{y}\|_2^2 + \lambda\|\mathbf{z}\|_1 + \frac{\rho}{2}\|\mathbf{v}-\nabla\mathbf{x}+\mathbf{u}\|_2^2 + \frac{\rho}{2}\|\mathbf{z}-\mathbf{v}+\mathbf{w}\|_2^2 \tag{8}$$

where $\mathbf{u}$ and $\mathbf{w}$ are scaled dual variables, and $\rho > 0$ is a penalty parameter. The ADMM algorithm alternately minimizes L_ρ with respect to each primal variable and then updates the dual variables. Each subproblem has an efficient solution: the $\mathbf{x}$-update involves solving a linear system, the $\mathbf{v}$-update has a closed-form solution, and the $\mathbf{z}$-update is a simple soft-thresholding operation.

Typically the algorithm iterates until a convergence criterion is met, typically based on the primal and dual residuals falling below specified tolerances, however for the results of this paper we chose a fixed number of iterates for computational speed.

3.5.1 *Proposed Method*

We propose to combine the strengths of both ROI-specific sampling and compressed sensing to doubly reduce the number of k-space samples needed to reconstruct an image - firstly, we use SFS to select ROI-optimal k-space samples, then randomly select a subset of samples to use for reconstruction. To reconstruct from these undersampled measurements, we use 20 iterations of ADMM ($\rho = 0.5, \lambda = 0.001$) to find our estimate of the iamge, using anisotropic TV as our regularizer. As a baseline, we use a naive matrix approach to construct the least-squares solution, giving us a guaranteed lower bound on the MSE compared to iterative methods like our ADMM algorithm.

4 Analysis and Evaluation

All experiments were done on a simulated Shepp-Logan phantom, meant to loosely emulate structures one might see in an axial scan of the human brain. Results were compared across different ROI configurations, k-space sampling strategies (ROI-optimized sampling vs. a Gaussian baseline), and accelerations factors at 2x, 4x, and 8x acceleration. Table 1 provides a comprehensive comparison for MSE values

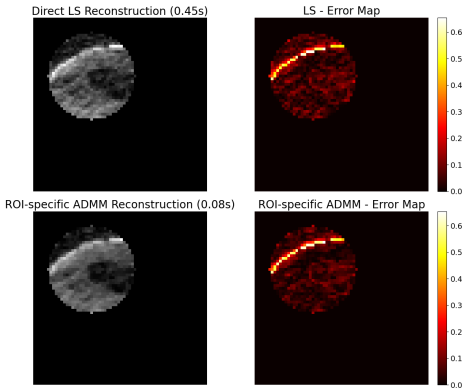

(a) Off-center circular ROI, 4x acceleration, Gaussian sampling

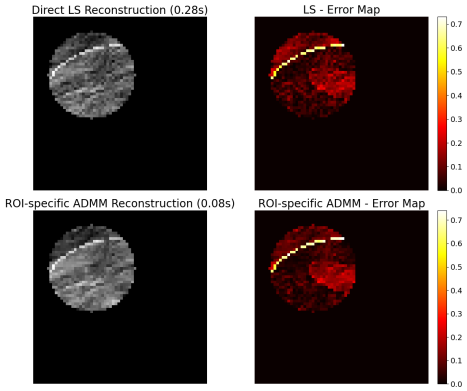

(b) Off-center circular ROI, 4x acceleration, optimized sampling

Fig. 4. Results for off-center circular ROI showing reconstruction performance for different sampling patterns. Each panel shows direct least-squares reconstruction (top left), error map for least-squares (top right), ADMM reconstruction with TV regularization (bottom left), and error map for ADMM reconstruction (bottom right).

across all of these configurations, as well as between our ADMM method and a direct least-squares reconstruction.

Looking at Table 1, we can see that MSE increases with increasing acceleration factor across all configurations (as expected), though our ADMM-based reconstruction performs slightly better across most scenarios. Interestingly, the Gaussian sampling baseline often does better at equal acceleration factors compared to our optimized sampling scheme - this is most likely due to the low-frequency bias of our Gaussian distribution, allowing models to reconstruct coarse, large-scale image features better compared to the optimized samples, which do not have the same spectral bias.

From a qualitative standpoint, more interesting differences can be seen. For most cases with higher acceleration ($>$ 4x), the reconstructed image seems totally unrecognizable compared to the ground truth (see Figure 2b), despite the moderate-seeming MSE numbers in Table 1. Across all scenarios (Figures 2-5), we see that the ADMM-reconstructed image maintains a similar level of large-scale structure preservation to the least-squares baseline, with notably less noise present in the reconstructed image. Additionally, computation times for ADMM are much lower than direct least-squares (as expected), though both are in the range of seconds for this image scale and as such, may be clinically feasible.

Lastly, the performance differences between circular and rectangular ROIs (Figure 5) reveal that ROI geometry may influence reconstruction quality. Circular ROIs benefit more from our approach, while rectangular ROIs show a smaller visual gap between our ADMM-based reconstruction and our least-squares baseline.

5 Conclusion

This paper presents a novel approach to MRI reconstruction that combines ROI-specific k-space sampling with compressed sensing to significantly accelerate acquisition times. Our experimental results on a simulated phantom demonstrate that by strategically selecting k-space samples optimized for a specific region of interest and applying ADMM with TV regularization, we can achieve substantial acceleration factors (up to 4x while maintaining decent image quality) within the target ROI. Our ADMM-based method generally outperforms direct least-squares reconstruction across acceleration factors, particularly in terms of visual smoothness. Additionally, the geometry of the ROI seems to have an impact on reconstruction quality, with circular ROIs benefiting more from our approach than rectangular ones.

Future work could explore the application of this framework to in vivo data, as well as the affects of different ROI types - in particular, investigating why circular ROIs see

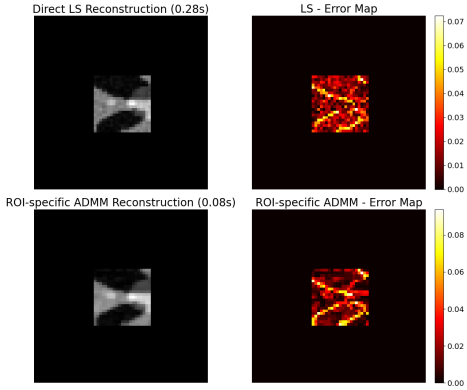

(c) Rectangular ROI, 2x acceleration, Gaussian sampling

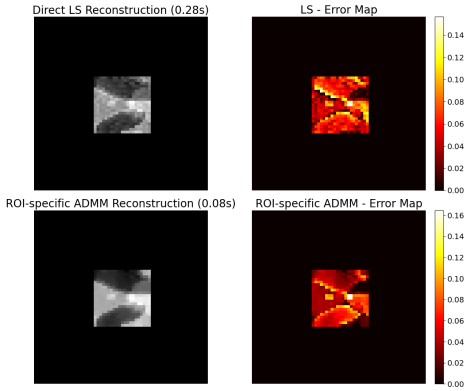

(d) Rectangular ROI, 2x acceleration, optimized sampling

Fig. 5. Results for rectangular ROI shape. The rectangular ROI shows different reconstruction characteristics compared to circular ROIs, with a smaller visual gap between ADMM-based reconstruction and the least-squares baseline.

more of an improvement from the ADMM reconstruction compared to the square/rectangular ROIs, and extending this investigation to differently shaped and potentially non-contiguous ROIs. As well, another avenue to be explored is the spectral bias of the Gaussian sampling scheme that seemed to lead to better reconstructions, and biasing the sampling of our ROI-optimized points to reflect the same spectral bias to potentially lead to better quality reconstructions. Optimization of the reconstruction algorithm itself could be investigated, with the ADMM parameters ρ and λ being able to be trained on data, given a fixed number of iterations. Lastly, extending this approach to multi-coil acquisitions and 3D imaging presents promising directions for further acceleration in clinical MRI applications.

REFERENCES

[1] Y. Gao and S. J. Reeves, "Optimal k-space sampling in MRSI for images with a limited region of support," *IEEE transactions on medical imaging*, vol. 19, no. 12, pp. 1168–1178, Dec. 2000.

[2] S. Reeves and L. Heck, "Selection of observations in signal reconstruction," *IEEE Transactions on Signal Processing*, vol. 43, no. 3, pp. 788–791, Mar. 1995, conference Name: IEEE Transactions on Signal Processing. [Online]. Available: https://ieeexplore.ieee.org/document/370637/?arnumber=370637

[3] S. Reeves and Z. Zhe, "Sequential algorithms for observation selection," *IEEE Transactions on Signal Processing*, vol. 47, no. 1, pp. 123–132, Jan. 1999, conference Name: IEEE Transactions on Signal Processing. [Online]. Available: https://ieeexplore.ieee.org/document/738245/?arnumber=738245

[4] S. Boyd, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010. [Online]. Available: http://www.nowpublishers.com/article/Details/MAL-016