EE367 Project Proposal: Optimized ROI-specific k-space sampling

Chiadika O.

February 2025

1 Motivation

Magnetic resonance imaging (MRI) reconstruction from undersampled k-space measurements remains an active area of research, with particular interest in optimizing sampling patterns to enable faster acquisition. While traditional approaches aim to maintain uniform reconstruction quality across the entire field of view, oftentimes in clinical applications only a particular region of interest (ROI) is of relevance. This project aims to develop an end-to-end optimization approach for directly optimizing k-space sampling patterns to prioritize quality in user-specified ROIs, using an unrolled reconstruction method to maintain differentiability throughout the entire pipeline. By incorporating region-specific weights in our objective function and leveraging the differentiability of an unrolled reconstruction method, we can directly optimize for sampling patterns that preserve critical image features in arbitrary ROIs while maintaining acceptable global reconstruction quality.

2 Related Work

Greedy selection of k-space samples for ROI imaging. The authors of [3] propose a method to choose k-space samples to optimize quality in a specified ROI using an extension of sequential forward selection (SFS). SFS essentially iteratively selects rows of the forward operator A to minimize an error criterion - our approach is similar in spirit, but here we jointly optimize our k-space selection rather than going for a greedy approach, which might find a suboptimal solution.

Jointly learning subsampling and reconstruction. The authors of [1] present a method for jointly optimizing k-space sampling patterns and reconstruction networks for MRI using a differentiable relaxation of the binary sampling mask and train an end-to-end pipeline that learns both the sampling pattern and reconstruction. The end-to-end approach is similar to this project, however the authors gear their work towards optimizing quality globally, rather than in a specific ROI. In addition, our approach of using an unrolled network can allows us to potentially constrain the reconstruction step, making it more interpretable compared to a deep U-Net.

3 Proposed Method

We formulate an end-to-end optimization framework where the binary sampling pattern is relaxed to be able to take values in $[0, 1]$ that gets thresholded to produce binary sampling decisions at the end of the optimization process. This mask is optimized jointly with an unrolled reconstruction network (such as ADMM [2], but could also include other iterative methods like conjugate gradient) that performs a fixed number of iterations. More formally, we seek to solve the following optimization problem:

$$\min_{M} ||W(A(MFx) - x)||_2^2 + \lambda||M||_1$$

Where M is our sampling mask (relaxed to take on values between 0 and 1 for differentiability), F is the Fourier transform operator, A represents the unrolled iterative reconstruction routine, and W is a diagonal matrix that ideally weights our desired ROI more heavily.

4 Evaluation

To evalutate the performance of this method, we'll look at a) ROI quality via PSNR in the relevant region compared to our baseline, b) global quality via PSNR of the reconstructed image compared to our baseline, and c) efficiency measured as the ratio of k-space samples used in our method compared to our baseline for a given PSNR. Our baseline will just be the same reconstruction method using fully sampled k-space. Since our optimized sampling mask is continuous, we can additionally investigate the quality as a function of our threshold value to find how many samples can sufficiently represent a desired ROI in practice.

5 Timeline

- Week 1: Implement differentiable sampling mask optimization and unrolled ADMM/CG reconstruction, collect images and label ROIs to be used for evaluation
- Week 2: Add ROI weighting and optimize, evaluate and compare with baselines
- Week 3: Summarize results and compile report

Stretch goals: Jointly optimize the unrolled optimization routine A in conjunction with our sampling mask M.

References

[1] Cagla Deniz Bahadir, Adrian V. Dalca, and Mert R. Sabuncu. Learning-based Optimization of the Under-sampling Pattern in MRI, April 2019. URL `http://arxiv.org/abs/1901.01960`. arXiv:1901.01960 [eess].

[2] Stephen Boyd. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2010. ISSN 1935-8237, 1935-8245. doi: 10.1561/2200000016. URL `http://www.nowpublishers.com/article/Details/MAL-016`.

[3] 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*, 19(12): 1168–1178, December 2000. ISSN 1558-254X. doi: 10.1109/42.897808. URL `https://ieeexplore.ieee.org/document/897808/?arnumber=897808tag=1`. Conference Name: IEEE Transactions on Medical Imaging.