Learning to Sample MRI via Variational Information Maximization

Cagan Alkan
Stanford University
calkan@stanford.edu

Morteza Mardani
Stanford University
morteza@stanford.edu

Shreyas S. Vasanawala
Stanford University
vasanawala@stanford.edu

John M. Pauly
Stanford University
pauly@stanford.edu

Abstract

Accelerating MRI scans requires optimal sampling of k-space data. This is however a daunting task due to the discrete and non-convex nature of sampling optimization. To cope with this challenge, we put forth a novel deep learning framework that leverages uncertainty autoencoders to enable joint optimization of sampling pattern and reconstruction of MRI scans. We represent the encoder as a non-uniform Fast Fourier Transform that allows *continuous* optimization of k-space samples on a non-Cartesian plane, while the decoder is a deep reconstruction network. Our approach is universal in a sense that it can be used with any reconstruction network. Experiments with knee MRI shows improved reconstruction quality of our data-driven sampling over the prevailing variable-density sampling.

1 Introduction

Deep learning methods has shown great promise at solving inverse problems. Of particular interest is MR reconstruction problem, where the goal is to reconstruct high fidelity images from a subset of spatial Fourier domain (k-space) representations. Even though deep learning methods have shown improved image quality when solving MR reconstruction problems [1–9], undersampling patterns that specify the selected subset of Fourier domain representation are typically chosen heuristically. The reconstruction models are optimized for a pre-determined acquisition (encoding) model without taking advantage of possible gains that can be obtained from learning undersampling patterns.

Recently, end-to-end deep learning methods have been proposed for learning undersampling patterns for MRI reconstruction problem. Active acquisition strategies [10–12] attempt to predict the next k-space samples to be acquired using information from existing samples. The non-active strategies can be grouped into two categories. [13], [14] focus on the Cartesian sampling case and model binary sampling masks probabilistically. These methods employ certain relaxations to maintain differentiability with respect to sampling probabilities. The second type of methods [15, 16] directly optimize for the k-space coordinates instead of estimating sampling probabilities. Authors of [15] restrict the optimization to a set of separable variables such as horizontal and vertical directions in 2-D plane to reduce the search space. As a result, their optimized sampling masks are restricted to a uniform *grid* pattern and do not cover the more general sampling patterns.

In this work, we present a variational information maximization method that enables joint optimization of acquisition and reconstruction of MRI scans in a data-driven manner. Our method enables learning an undersampling pattern tuned spefically to the reconstruction network, and vice versa, to obtain improved reconstruction performance. We represent the acquisition (encoder) model with the non-uniform Fast Fourier Transform (nuFFT) operator [17] that is parameterized by the sampling locations

NeurIPS 2020 Workshop on Deep Learning and Inverse Problems, virtual.

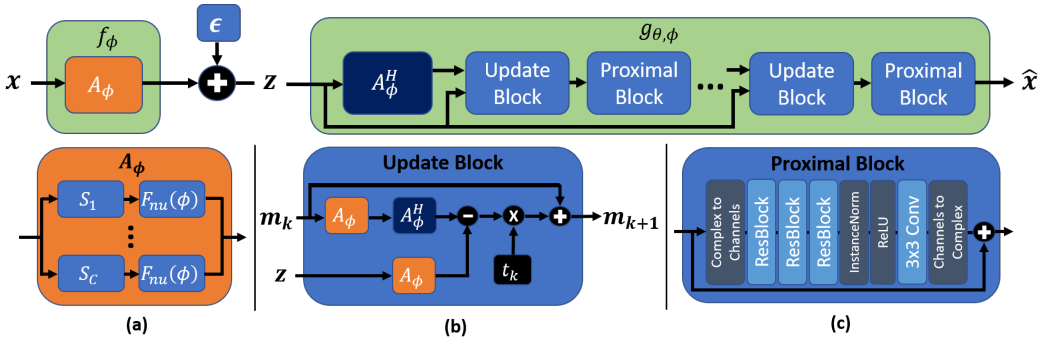

Figure 1: Network architecture consisting of nuFFT based encoder (a) and unrolled reconstruction network (b, c). Sampling locations ϕ are shared between encoder and decoder.

in k-space. This allows interpolation of non-cartesian coordinates and enables continuous optimization of the sampling pattern. On the reconstruction (decoder) side, we use an unrolled reconstruction network which mimics the proximal gradient based solutions to compressed sensing problems.

2 Methods

We consider the MR signal model under the additive white complex Gaussian noise as

$$z = f_\phi(x) + \epsilon \tag{1}$$

where $x \in \mathbb{C}^N$ is the image, $z \in \mathbb{C}^M$ is the measured data in k-space domain, $f_\phi(\cdot)$ is the forward model that describes the imaging system parameterized by k-space sample coordinates $\phi \in [-0.5, 0.5]^M$, and $\epsilon \sim \mathcal{N}_c(0, \sigma^2 I)$ is the measurement noise. The imaging model f_ϕ includes non-uniform Fast Fourier Transform (nuFFT) operation F_{nu}, and for the multi-coil scenario, contains signal modulation by coil sensitivity maps S. Given an acceleration factor $R = N/M$, the goal is to find the optimal set of samples ϕ along with a reconstruction function $g(z)$ that maintains the full k-space data image quality.

2.1 Variational Information Maximization for Acquisition and Reconstruction

Under the setting in 1, the distribution of measurements Z given image X has a complex valued Gaussian distribution $q_\phi(Z|X) = \mathcal{N}_c(f_\phi(X), \sigma^2 I)$. We adapt the uncertainty autoencoder framework defined in [18] and we make use of the InfoMax principle [19, 20] that maximizes the mutual information between the measurements and noisy latent representations

$$\max_\phi I_\phi(X, Z) = \max_\phi H(X) - H_\phi(X|Z) \tag{2}$$

Since $H(X)$ does not depend on k-space coordinates ϕ, this is equivalent to

$$\max_\phi -H_\phi(X|Z) = \max_\phi \mathbb{E}_{q_\phi(X,Z)}[\log q_\phi(X|Z)] \tag{3}$$

$$\geq \max_{\phi,\theta} \mathbb{E}_{q_\phi(X,Z)}[\log p_\theta(X|Z)] \tag{4}$$

where in the last line we introduce a variational approximation to the model posterior $q_\phi(X|Z)$ to obtain a lower bound. The variational parameters θ correspond to the weights of the reconstruction network. Using the bound on 4 we express our objective as

$$\mathcal{L}(\phi, \theta; \mathcal{D}) = \max_{\phi,\theta} \sum_{x \in \mathcal{D}} \mathbb{E}_{q_\phi(Z|x)}[\log p_\theta(x|z)] \tag{5}$$

where we estimate the expectation with respect to $q_\phi(X)$ using Monte-Carlo sampling from the dataset $\mathcal{D}$. Thus, we propose optimizing shared ϕ and θ for the entire dataset rather than specific acquisition and reconstruction parameters for each example in the dataset. In addition, for a different anatomy (hence for a different dataset), the optimization must be repeated. Similarly, we estimate the

expectation with respect to $q_\phi(Y|x)$ using Monte-Carlo methods, and estimate the gradients with respect to ϕ using reparameterization trick.

Depending on the observation model used for $p_\theta(\cdot)$, we end up with a different loss function. For example, we can have ℓ_2 or ℓ_1 distances between reconstruction and ground truth as loss functions for Gaussian and Laplacian observation models, respectively. The main difference between [18] and our formulation is that we enforce the latent space to correspond to the Fourier domain coefficients by using a nuFFT operator rather than a fully parameterized neural network.

2.2 Multi-Channel Acquisition and Reconstruction Models

Our overall network architecture is illustrated in Figure 1. We represent the acquisition (encoding) model with the nuFFT operator that is parameterized by the sampling locations ϕ in k-space. This enables us to represent k-space sampling locations as continuous variables and directly optimize for the sampling pattern by backpropagating through ϕ. In the multi-channel MRI setting, the acquisition model admits

$$f_\phi(x) = A_\phi x = \left[(F_{nu}(\phi)S_1)^H \cdots (F_{nu}(\phi)S_C)^H\right]^H x \tag{6}$$

where C is the number of channels, $S_i \in \mathbb{C}^{N\times N}$ is a diagonal matrix containing coil sensitivity profiles for coil i, and $F_{nu}(\phi) : \mathbb{C}^N \to \mathbb{C}^{M/C}$ is the nuFFT operator at sampling locations ϕ.

Notice that 5 allows any deep neural network for reconstruction as the decoder. In this work, we used an unrolled reconstruction network similar to [7, 8] that mimic the proximal gradient based solutions for non-smooth compressed sensing. The unrolled architecture also incorporates MR physics, where the reconstruction is performed by alternating between data consistency (Fig. 1b) and proximal steps (Fig. 1c) for a fixed number of iterations. The data-consistency renders the reconstruction a function of both ϕ and θ, hence we share the parameters of the nuFFT encoder with the decoder network. The unrolled nature of the network also allows for multiple gradient paths to ϕ from the decoder.

3 Experiments

3.1 Data

We used the "Stanford Fully Sampled 3D FSE Knees" dataset available in mridata.org [21] which contains 3D knee scans from 20 subjects. Each 3D volume consists of 320 slices with a matrix size of 256×320 and was acquired with 8 channels. The sensitivity maps of each coil were estimated using ESPIRiT [22]. Each of the slices were treated as separate examples during training and validation. The datasets were divided according to subjects where 15 subjects (4800 slices) were used for training, 2 subjects (640 slices) were used for validation, and 3 subjects (960 slices) were used for testing.

3.2 Network Implementation and Training Details

All of our models were implemented on TensorFlow and trained on a NVIDIA Titan Xp graphics card with 12GB of memory. We used the TensorFlow implementation of nuFFT available in [23]. As the sampling locations change during the training, we did not use density compensation while calculating the adjoint nuFFT. Instead we relied on the proximal block to compensate for the density. The proximal block consists of 3 Residual Blocks each having 64 channels and we unrolled the network for 4 iterations. We also used instance normalization at the output of convolutional layers. We considered the Gaussian observation model for $p_\theta(\cdot)$ in Equation 5 which corresponds to the ℓ_2 reconstruction loss. Due to GPU memory constraints, we used a batch size of 2. We also used the Adam optimizer with a learning rate of 0.01, $\beta_1 = 0.9$, $\beta_2 = 0.999$ for network optimization. To assess the image quality, we adopt peak signal-to-noise ratio (pSNR) and structural similarity index (SSIM) between the reconstruction and fully-sampled ground truth. pSNR was evaluated directly on complex-valued images, while for SSIM complex-valued images are treated as two channels.

3.3 Experiment Details

We considered four acceleration factors in our experiments: $R = 5, 10, 15, 20$. For each acceleration factor, we initialized the k-space sampling pattern (ϕ) by sampling realizations from a 2D Gaussian

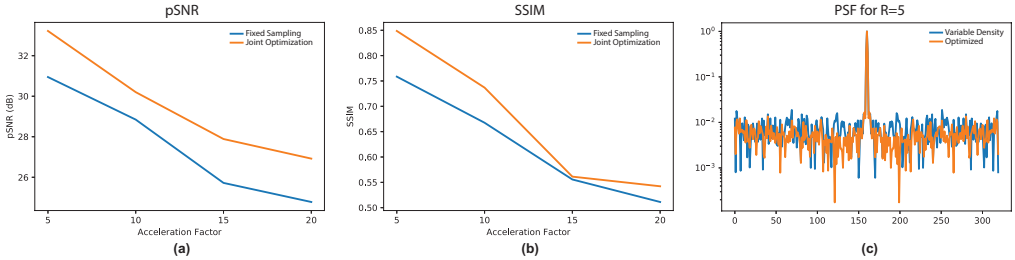

Figure 2: pSNR (a) and SSIM (b) evaluated on test set for different acceleration factors. Point spread function (c) of sampling trajectories before and after optimization for $R = 5$.

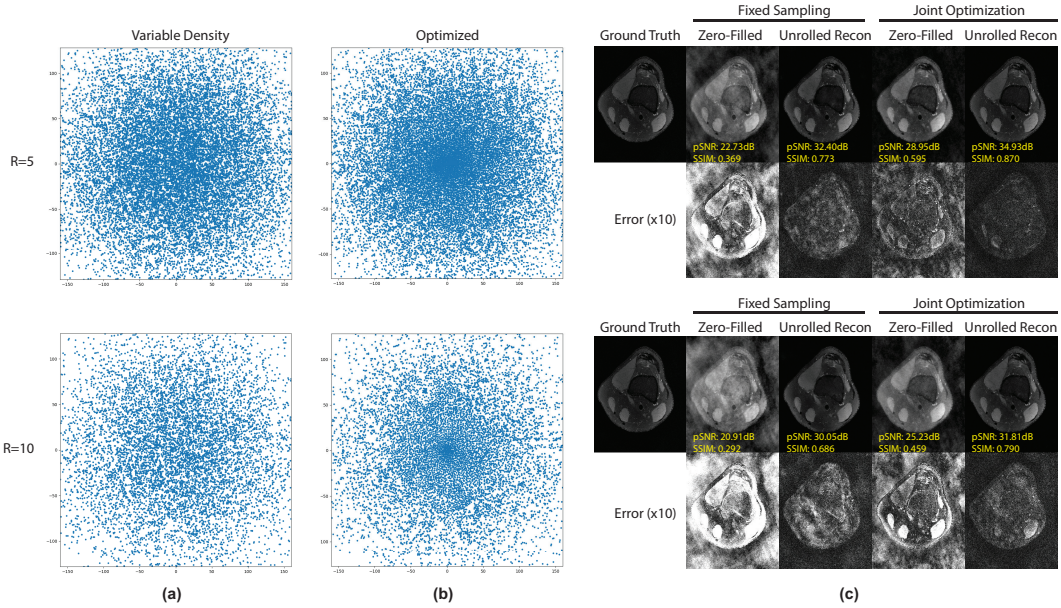

Figure 3: Variable density (a) and optimized (b) trajectories along with the reconstruction result on a slice in the test set (c). Zero-filled reconstruction corresponds to applying A_ϕ^H on z directly.

distribution centered at the origin (Variable Density Sampling [24]) and picked a variance $\sigma = 0.001$ for the measurement noise ϵ. We then ran our model and compared the result with the case where we only optimized the reconstruction network. Figure 2a,b shows pSNR and SSIM metrics, respectively. We observe that joint optimization using our framework improves reconstruction quality.

Figure 3 shows the initial and optimized sampling patterns along with the reconstruction results on a sample test slice for $R = 5$ and $R = 10$. Sampling optimization accumulates more samples near k-space center and reduces the gaps between them which enhances reconstruction quality. In addition, zero-filling reconstruction using the optimized samples have less severe aliasing artifacts, making it easier for the decoder network to remove these artifacts.

Figure 2c shows the point spread function (PSF) of the sampling pattern before and after optimization for $R = 5$. The sidelobes of the PSF are suppressed which reduces the amount of aliasing due to undersampling.

4 Conclusion and Discussion

This work explores the use of variational information maximization for learning the sampling pattern in MRI jointly with image reconstruction. Using the nuFFT operator in the encoder enables

continuous optimization of sampling variables. On a knee MR dataset, we have shown that the optimized sampling pattern improves reconstruction quality highlighting possible benefits that can be obtained by learning undersampling patterns.

One observation is that the peripheral samples do not change as much as the samples near origin during training. Incorporating additional regularizers that emphasize high frequency information can improve the gradient flow to these samples.

Our framework allows incorporating other MR physics related artifacts beyond the Gaussian measurement noise such as motion and off-resonance effects. These artifacts can be simulated as a part of the forward model to find the optimal sampling locations in the presence of such effects.

Acknowledgments and Disclosure of Funding

This work is supported by NIH R01EB009690 and R01EB026136, as well as GE Healthcare.

References

[1] Bo Zhu, Jeremiah Z Liu, Stephen F Cauley, Bruce R Rosen, and Matthew S Rosen. Image reconstruction by domain-transform manifold learning. *Nature*, 555(7697):487–492, 2018.

[2] Dongwook Lee, Jaejun Yoo, and Jong Chul Ye. Deep residual learning for compressed sensing mri. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 15–18. IEEE, 2017.

[3] Guang Yang, Simiao Yu, Hao Dong, Greg Slabaugh, Pier Luigi Dragotti, Xujiong Ye, Fangde Liu, Simon Arridge, Jennifer Keegan, Yike Guo, et al. Dagan: Deep de-aliasing generative adversarial networks for fast compressed sensing mri reconstruction. *IEEE transactions on medical imaging*, 37(6):1310–1321, 2017.

[4] Morteza Mardani, Enhao Gong, Joseph Y Cheng, Shreyas S Vasanawala, Greg Zaharchuk, Lei Xing, and John M Pauly. Deep generative adversarial neural networks for compressive sensing mri. *IEEE transactions on medical imaging*, 38(1):167–179, 2018.

[5] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P Recht, Daniel K Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine*, 79(6):3055–3071, 2018.

[6] Morteza Mardani, Qingyun Sun, Shreyas Vasawanala, Vardan Papyan, Hatef Monajemi, John Pauly, and David Donoho. Neural proximal gradient descent for compressive imaging. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, page 9596–9606, 2018.

[7] Joseph Y Cheng, Feiyu Chen, Marcus T Alley, John M Pauly, and Shreyas S Vasanawala. Highly scalable image reconstruction using deep neural networks with bandpass filtering. *arXiv preprint arXiv:1805.03300*, 2018.

[8] Christopher M Sandino, Joseph Y Cheng, Feiyu Chen, Morteza Mardani, John M Pauly, and Shreyas S Vasanawala. Compressed sensing: From research to clinical practice with deep neural networks: Shortening scan times for magnetic resonance imaging. *IEEE Signal Processing Magazine*, 37(1):117–127, 2020.

[9] Hemant K Aggarwal, Merry P Mani, and Mathews Jacob. Modl: Model-based deep learning architecture for inverse problems. *IEEE transactions on medical imaging*, 38(2):394–405, 2018.

[10] Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547*, 2019.

[11] Zizhao Zhang, Adriana Romero, Matthew J Muckley, Pascal Vincent, Lin Yang, and Michal Drozdzal. Reducing uncertainty in undersampled mri reconstruction with active acquisition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2049–2058, 2019.

[12] Luis Pineda, Sumana Basu, Adriana Romero, Roberto Calandra, and Michal Drozdzal. Active mr k-space sampling with reinforcement learning. *arXiv preprint arXiv:2007.10469*, 2020.

[13] Cagla Deniz Bahadir, Adrian V Dalca, and Mert R Sabuncu. Learning-based optimization of the under-sampling pattern in mri. In *International Conference on Information Processing in Medical Imaging*, pages 780–792. Springer, 2019.

[14] Iris AM Huijben, Bastiaan S Veeling, and Ruud JG van Sloun. Learning sampling and model-based signal recovery for compressed sensing mri. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8906–8910. IEEE, 2020.

[15] Hemant Kumar Aggarwal and Mathews Jacob. J-modl: joint model-based deep learning for optimized sampling and reconstruction. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1151–1162, 2020.

[16] Tomer Weiss, Ortal Senouf, Sanketh Vedula, Oleg Michailovich, Michael Zibulevsky, and Alex Bronstein. Pilot: Physics-informed learned optimal trajectories for accelerated mri. *arXiv preprint arXiv:1909.05773*, 2019.

[17] Leslie Greengard and June-Yub Lee. Accelerating the nonuniform fast fourier transform. *SIAM review*, 46(3):443–454, 2004.

[18] Aditya Grover and Stefano Ermon. Uncertainty autoencoders: Learning compressed representations via variational information maximization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2514–2524, 2019.

[19] Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988.

[20] Ralph Linsker. How to generate ordered maps by maximizing the mutual information between input and output signals. *Neural computation*, 1(3):402–411, 1989.

[21] F Ong, S Amin, SS Vasanawala, and M Lustig. An open archive for sharing mri raw data. In *ISMRM & ESMRMB Joint Annu. Meeting*, page 3425, 2018.

[22] Martin Uecker, Peng Lai, Mark J Murphy, Patrick Virtue, Michael Elad, John M Pauly, Shreyas S Vasanawala, and Michael Lustig. Espirit—an eigenvalue approach to autocalibrating parallel mri: where sense meets grappa. *Magnetic resonance in medicine*, 71(3):990–1001, 2014.

[23] Frank Ong. tf-nufft. `https://github.com/mikgroup/tf-nufft`, 2018.

[24] Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.