End-to-End Sequential Sampling and Reconstruction for MR Imaging

Tianwei Yin*,[2] Zihui Wu*,[1] He Sun[1] Adrian V. Dalca[3,4]
Yisong Yue[1] Katherine L. Bouman[1]
[1]California Institute of Technology [2]University of Texas at Austin
[3]Harvard Medical School [4]Massachusetts Institute of Technology *Equal contribution

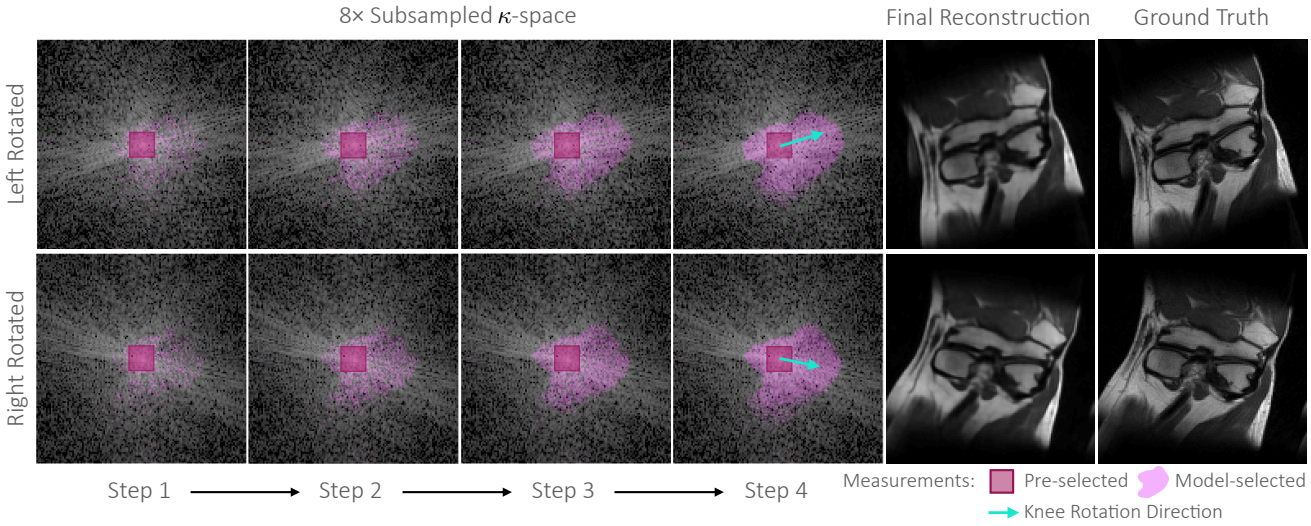

Figure 1: We propose a sequential sampling and reconstruction co-design framework for accelerated MRI that adapts to a target during acquisition. Here, we visualize the sampling policy and final reconstruction of rotated knees in a single-coil imaging setting with 8× acceleration (8× subsampling). The first four columns show the cumulative κ-space measurements selected by the proposed learned sampler (pink) in acquisition steps 1 through 4 (during a 4-step acquisition). The fifth column shows the final image recovered by the proposed learned reconstructor, and the last column is the ground truth. This example illustrates how our model has learned to adapt to different κ-space distributions: the final sampling patterns in the fourth column contain visible directional structure that aligns with the κ-space power spectrum. Rotated anatomical images, such as these rotated knee images, were *not* included in the training set (or quantitatively evaluated test set).

Abstract

Accelerated MRI shortens acquisition time by subsampling in the measurement κ-space. Recovering a high-fidelity anatomical image from subsampled measurements requires close cooperation between two components: (1) a sampler that chooses the subsampling pattern and (2) a reconstructor that recovers images from incomplete measurements. In this paper, we leverage the sequential nature of MRI measurements, and propose a fully differentiable framework that jointly learns a sequential sampling policy simultaneously with a reconstruction strategy. This co-designed framework is able to adapt during acquisition in order to capture the most informative measurements for a particular target (Figure 1). Experimental results on the fastMRI knee dataset demonstrate that the proposed approach successfully utilizes intermediate information during the sampling process to boost reconstruction performance. In particular, our proposed method outperforms the current state-of-the-art learned κ-space sampling baseline on up to 96.96% of test samples. We also investigate the individual and collective benefits of the sequential sampling and co-design strategies. Code and more visualizations are available at http://imaging.cms.caltech.edu/seq-mri.

1. Introduction

Magnetic Resonance Imaging (MRI) is a widely used imaging technology for clinical diagnosis and biomedical research. MRI is non-invasive, requires zero radiation, and can result in images with strong tissue contrast and excellent quality. However, a central challenge of MRI is its slow acquisition process. Standard MRI scans can take up to half an hour as measurements in *κ-space* are being collected, es-

pecially during research studies [46]. This long acquisition time leads to high cost, patient discomfort, and significant reconstruction artifacts when patients move. Thus, there is strong motivation to accelerate the MRI acquisition process.

One way to accelerate MRI is to collect fewer measurements and reconstruct anatomical images from only partial κ-space data. This acceleration requires: (a) a carefully designed κ-space subsampling pattern to collect informative measurements, and (b) a reconstruction method that accurately recovers high-quality images from undersampled data. Current MRI protocols collect measurements over time using *static* subsampling patterns that were designed *a priori*. To further accelerate a scan, we are interested in *sequential* sampling patterns that adapt to a target based on intermediate information collected during acquisition.

A high-fidelity MRI reconstruction stems from cooperation between the κ-space sampling strategy and the reconstruction method. Traditionally, MRI subsampling patterns and reconstruction methods have been largely *independently designed*. We are instead interested in *co-design*, where jointly designing the two components can synergistically boost reconstruction quality. Our approach builds on neural network based co-design frameworks that have shown strong empirical performance and take advantage of efficient differentiable training [1, 15, 16, 34].

In this paper, we propose an end-to-end differentiable framework that successfully combines co-design and sequential sampling. Specifically, we design an explicit sequential structure of T steps, with each step consisting of a jointly learned κ-space sampler and reconstructor. Comparing our model with prior work in accelerated MRI, we investigate the individual and collective benefits of sequential sampling and co-design. We evaluate the proposed model on the NYU fastMRI datasets and find that: (1) even a single sequential step consistently improves performance compared to using a pre-designed sampling pattern; (2) more sequential steps can improve reconstruction quality, but with diminishing returns; and (3) a fully differentiable approach enables more efficient and effective co-design than non-differentiable methods (e.g., prior approaches that optimize via reinforcement learning).

The paper is organized as follows. In Section 2, we review past literature in accelerated MRI from the perspectives of co-design and sequential sampling. In Section 3, we mathematically formulate the accelerated MRI problem. We then introduce our proposed framework and its training procedure in Section 4. Section 5 presents our experimental settings, comparisons between our model and other baselines, and ablation studies. Finally, we conclude with a discussion on future directions of our framework in Section 6.

2. Related Work in Accelerated MRI

Prior work in accelerated MRI can be organized into four quadrants, split across two dimensions: methods that (1) independently (and/or manually) design the sampler and reconstructor versus data-driven co-design, and (2) specify the sampling pattern prior to a scan (pre-designed) versus adapt samples to the target during acquisition. In Section 2.1, we cover traditional methods that independently (and/or manually) design the sampler and reconstructor. In Section 2.2, we discuss previous methods that perform pre-designed acquisition in a co-design framework. In Section 2.3, we introduce recent work on sequential sampling for accelerated MRI. We conclude in Section 2.4 with an overview of methods that attempt to combine co-design and sequential sampling, but without end-to-end learning. In this paper, we propose an end-to-end framework that efficiently combines co-design and sequential sampling, successfully inheriting the advantages of both approaches.

2.1. Traditional Methods

Accelerated MRI sampling patterns implemented on commercial scanners are motivated by ideas in compressed sensing (CS) [7]. Since anatomical images are sparse in a linearly transformed space, it is possible to reconstruct a high-fidelity image with incoherent κ-space data sampled below the Nyquist-Shannon rate [23]. In the context of 2D CS-MRI, prior work has investigated uniform density random sampling, variable density sampling [22], Poisson-disc sampling [37], continuous-trajectory variable density sampling [8], and equi-spaced sampling [9]. These sampling patterns are easy to implement, but not adaptive to specific datasets or target images.

Once sparse κ-space measurements have been acquired, an image is typically reconstructed via an optimization problem that involves two objectives: the first encourages a reconstruction that matches the observed data, while the second addresses the ill-posed nature of the underdetermined system through image regularization. Common regularization terms include total variation (TV) [6] and the ℓ_1-norm after a sparsifying transformation (obtained using wavelets [22, 30] or dictionary decompositions [12, 27, 47]).

Recently, convolutional neural networks (CNNs) have demonstrated impressive performance in MRI reconstruction. Strategies include unrolled networks [10, 21, 29, 45], UNet-based networks [13, 19], GAN-based networks [25, 44], among others [20, 40, 50]. These learning methods have achieved state-of-the-art performance on public MRI challenge datasets [46]. In our proposed co-design model, we employ a convolutional UNet for image reconstruction.

2.2. Co-design

The goal of co-design is to jointly identify the optimal sampling and reconstruction strategies. This is an NP-hard

combinatorial optimization problem due to the discrete nature of the sampling pattern. Theoretically, one could identify an optimized reconstructor for every possible sampling strategy, and then pick the overall strategy that performs best. However, this brute-force optimization approach is not practical, as it requires enumerating an exponential number of possible sampling combinations. Early work formulated the co-design as a nested (or bi-level) optimization problem and alternated between optimizing a sampler and a reconstructor [26].

More recently, deep learning has enabled a data-driven solution to the co-design problem, where the sampler and reconstructor can be jointly learned through end-to-end training. For example, [1,43,48] proposed co-design frameworks for 2D Cartesian κ-space sampling and [39, 42] applied co-design to 2D radial κ-space sampling.[1] These methods have shown superior performance over previous baselines that combine an individually-optimized sampler and reconstructor pair [1, 34, 39, 43, 48]. However, these methods do not take advantage of the sequential nature of data collection during an MRI scan, and only solve for a generic sampling pattern for an entire dataset.

2.3. Sequential Sampling

Since MRI scanners acquire measurements over time, recent work has modeled the sampling process in the context of sequential decision making. Sequential decisions enable the sampling pattern to adapt to different input images by choosing the next κ-space sample based on prior measurements. Reinforcement learning (RL) methods have primarily been employed for this purpose. For example, [2, 24] formulate the sampling problem as a Partially Observable Markov Decision Process (POMDP) and use Policy Gradient [4] and DDQN [36] methods, respectively. These RL methods heavily rely on a pre-trained reconstructor, which leads to a training mismatch (and thus potentially suboptimal performance), since the reconstructor was trained with a sampling strategy that does not match the strategy eventually employed by the RL-learned sampler. Furthermore, these RL methods are difficult and costly to train, as they are non-differentiable. As a consequence, in the context of accelerated MRI, these methods either fail to be adaptive to different input images or have only limited improvement over simple baselines [2, 24].

2.4. Co-design & Sequential Sampling

Approaches that seek to combine co-design and sequential sampling strategies have been proposed, however with only limited success thus far. The work of [14] draws inspiration from AlphaGo [31] and trains a sampler to emulate the policy distribution obtained through a Monte Carlo Tree Search (MCTS); the reconstructor is trained during alternating optimization steps. However, according to the results in [2], the MCTS method in [14] has limited improvement over simple baselines, and is outperformed by the sequential sampling method in [2] without co-design. This poor performance may be due to the overall MCTS framework not being end-to-end differentiable. Alternatively, [49] proposes a framework that trains a ResNet to reconstruct the anatomical image simultaneously with an evaluator network that is trained to select the most uncertain measurement in κ-space. Although the authors demonstrate how this framework can be used to sequentially choose the next sample, it is not explicitly trained end-to-end and is outperformed by [24], which does not use co-design. This training-testing mismatch limits the potential improvement of sequential sampling. In contrast, we design a fully differentiable end-to-end framework that leverages the sequential nature of κ-space MRI acquisition during both training and testing.

3. MRI Fundamentals

MRI acquires measurements in the Fourier space (i.e. κ-space). Let $\mathbf{y} \in \mathbb{C}^{M \times N}$ be the complex-valued matrix representing the full κ-space data of an $M \times N$ target image $\mathbf{x} \in \mathbb{R}^{M \times N}$. In the case of no noise, the true image can be simply recovered through an inverse Fourier transform: $\mathbf{x} = \mathcal{F}^{-1}(\mathbf{y})$. However, in accelerated MRI scanning, only a subset of κ-space samples, $\hat{\mathbf{y}}$, are measured:

$$\hat{\mathbf{y}} = \mathbf{M} \odot \mathbf{y} = \mathbf{M} \odot \mathcal{F}(\mathbf{x}), \tag{1}$$

where $\odot$ indicates element-wise multiplication and $\mathbf{M} \in \{0, 1\}^{M \times N}$ is a binary sampling mask.

We can compute a *zero-filled image* reconstruction by applying an inverse Fourier transform to the under-sampled κ-space, where zeros occupy the unobserved κ-space samples: $\hat{\mathbf{x}} = \mathcal{F}^{-1}(\hat{\mathbf{y}})$. This zero-filled reconstruction contains aliasing artifacts, and a reconstruction algorithm is often used to recover a clean target image [1, 2, 10, 40, 49]. We define the *acceleration factor* α as the ratio between the total number of possible κ-space samples K and the number of acquired measurements (i.e., $\alpha = K / \sum \mathbf{M}$).

4. Method

Figure 2 summarizes the co-design framework for our sequential sampling and reconstruction model. We partition the κ-space sampling budget into T steps. At each step, t, the pipeline applies a reconstructor, $A_w(\cdot)$, and a sampler, $\pi_\theta(\cdot)$. The goal of the reconstructor is to remove aliasing artifacts that appear in the zero-filled reconstruction, $\hat{\mathbf{x}}_t$:

$$\tilde{\mathbf{x}}_t = A_w(\hat{\mathbf{x}}_t). \tag{2}$$

The goal of the sampler is to intelligently select which κ-space samples to observe next, based on previously ob-

[1] Differentiable co-design of sensing and reconstruction methods has also been successfully applied to other imaging domains as well [34].

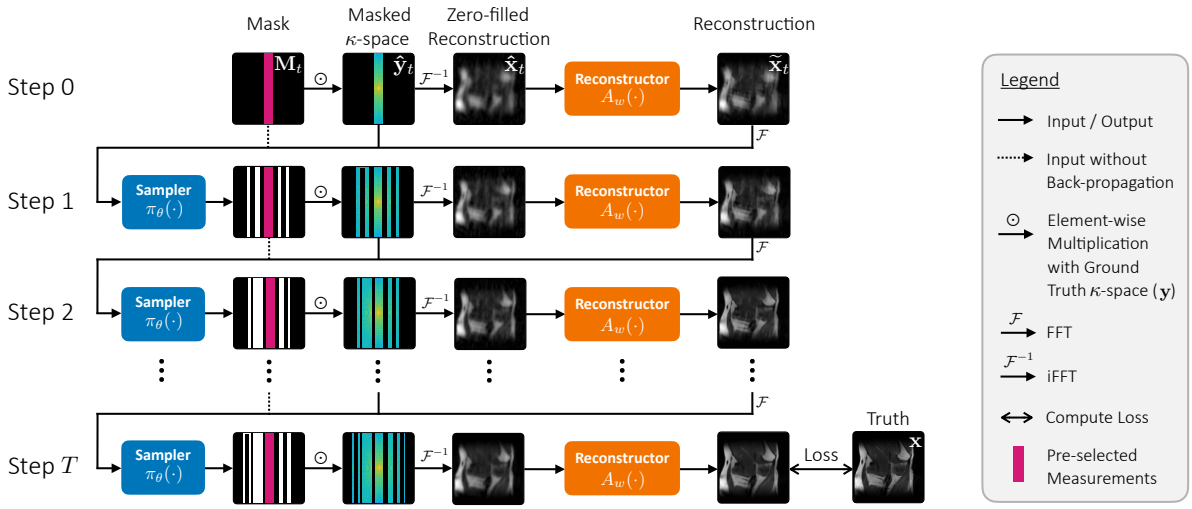

Figure 2: Overview of the proposed sequential sampling framework. Low-frequency samples are pre-selected and measured in κ-space. The subsampled κ-space is transformed into a zero-filled image, which is fed into a reconstructor $A_w(\cdot)$ to produce an intermediate image reconstruction (Equation (2)). The intermediate reconstruction and measurements are passed into a sampler network $\pi_\theta(\cdot)$, which outputs a discrete probability distribution representing suggested samples for the next iteration. An action is sampled from this distribution (Equation (3)), and the corresponding κ-space measurements are acquired. The sampling and reconstruction process is repeated for T steps. The sampler and reconstructor are neural networks learned via end-to-end training with a loss on the final reconstructed image. Weights are shared across all T acquisition steps.

served measurements and a preliminary reconstruction:

$$\begin{aligned} \mathbf{M}_{t+1} \sim & \pi_\theta(\hat{\mathbf{y}}_t, \widetilde{\mathbf{y}}_t, \mathbf{M}_t) \\ & s.t. \sum (\mathbf{M}_{t+1} - \mathbf{M}_t) = S \end{aligned} \tag{3}$$

where $\hat{\mathbf{y}}_t$ and $\widetilde{\mathbf{y}}_t = \mathcal{F}(\tilde{\mathbf{x}}_t)$ denote the κ-space representation of the zero-filled image ($\hat{\mathbf{x}}_t$) and the reconstructed image ($\tilde{\mathbf{x}}_t$), respectively, $\mathbf{M}_t$ is a binary mask representing the sampling pattern collected up until step t, and S is the sampling budget at each step.

We model the sampler, $\pi_\theta(\cdot)$, and reconstructor, $A_w(\cdot)$, as neural networks, and co-optimize the network weights, θ and w, by minimizing the image reconstruction error between the final step reconstruction $\tilde{\mathbf{x}}_T$ and the ground truth target image $\mathbf{x}$:

$$\theta^*, w^* = \arg\min_{\theta,w} \mathcal{D}(\tilde{\mathbf{x}}_T, \mathbf{x}), \tag{4}$$

where $\mathcal{D}$ is an image distance metric, such as the structural similarity index measure (SSIM) [41] or peak signal-to-noise ratio (PSNR). We choose to share sampler and reconstructor weights across all T steps. The sampler and reconstructor are described in more detail in Sections 4.1 and 4.2, respectively.

4.1. Sampler

In the design of the sampler, $\pi_\theta(\cdot)$, we consider two types of κ-space sampling: 1D line sampling and unconstrained 2D point sampling. Figure 3 illustrates these two sampling scenarios, which enable different levels of sampling flexibility. 1D line sampling represents one of the most widely used κ-space sampling strategies on commercial scanners due to its fast acquisition time [22]. In 2D point sampling any measurement on the $M \times N$ frequency grid in κ-space can be acquired. Unconstrained 2D point sampling represents an upper bound on sampling flexibility, can be physically feasible in specific scenarios [1], and is often explored as a methodological building block.

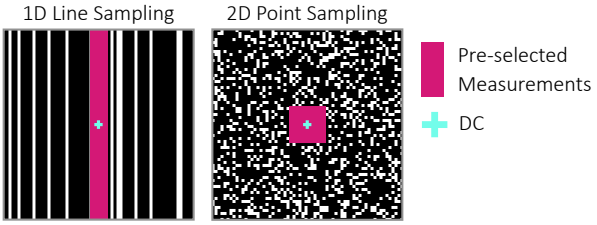

Figure 3: Visualizations of two types of κ-space sampling patterns: 1D line sampling and 2D point sampling. White regions are sampled from a uniform distribution over the space of possible actions. The center low-frequency samples are pre-selected in all experiments before any further sampling. DC corresponds to the $(0, 0)$ frequency.

As low-frequency κ-space measurements contain the most information about large-scale anatomical structure, it

is common practice in accelerated MRI to fix a small number of low-frequency κ-space samples to always be collected [2, 24, 49]. We follow this strategy by allocating $\frac{1}{8}$ of the total sampling budget to the central low-frequency region in all experiments.

4.1.1 Neural Sampler Architecture

The action space at each step t is the set of possible sampling indices (i.e., $K = N$ in the line sampler and $K = N \times M$ in the point sampler). As shown in Equation (3), the input to the sampler is the past κ-space measurements, $\hat{\mathbf{y}}_{t-1}$, κ-space reconstruction, $\widetilde{\mathbf{y}}_{t-1}$, and sampling mask, $\mathbf{M}_{t-1}$. The output is a binary sampling mask $\mathbf{M}_t \in \{0,1\}^K$. New samples acquired at time step t are indicated by $\mathbf{M}_t - \mathbf{M}_{t-1}$.

An intermediate output of the network-based sampler is a heatmap $\mathbf{P}_t \in [0,1]^K$, which defines the probablity that a sample will be selected at acquisition time t. In order to ensure that $\mathbf{P}_t$ is between zero and one, a softplus and normalization are applied.[2] Additionally, to avoid reacquiring previous measurements, the sampling probability of previously acquired lines is set to zero:

$$\mathbf{P}'_t = \mathbf{P}_t \odot (1 - \mathbf{M}_{t-1}) \tag{5}$$

Inspired by the stochastic strategy in [1], we sample from the distribution $\mathbf{P}'_t$ to obtain the κ-space sampling mask $\mathbf{M}_t$ for acquisition step t:

$$\mathbf{M}_t = \mathbb{1}_{\mathbf{U} \leq \mathbf{P}'_t} + \mathbf{M}_{t-1}, \tag{6}$$

where $\mathbf{U}$ is a vector of N independent realizations of the uniform distribution on the interval $[0,1]$. We use rejection sampling to guarantee the exact number of specified κ-space samples is obtained at each step t.

The indicator function $\mathbb{1}_{\mathbf{U} \leq \mathbf{P}}$ is not differentiable, which hinders the training of the model through back-propagation. In this paper, we follow [5, 48] and use a straight-through estimator that applies the indicator function in the forward pass to generate the binary sampling mask $\mathbf{M}_{t+1}$, while approximating its gradients by treating the binary indicator function as a sigmoid during back-propagation. In this way, we are able to capture binary sampling in real MR scanning, while retaining useful gradients for end-to-end training.

We instantiate the 1D line sampler as a Multilayer Perceptron (MLP) with five layers separated by ReLU activation functions. In the 2D point sampler we replace the MLP with a 8-block convolutional UNet network design with ReLU activation functions. We find the convolutional architecture more efficient on the higher dimensional action space. Further details of the network architectures for both sampler networks are included in the supplemental material.

[2] To help enforce the sampling budget constraint in Equation (3), $\mathbf{P}_t$ is also rescaled to obtain the desired average value following [1].

4.2. Reconstructor

Our proposed co-design sequential framework learns the parameters of a reconstructor, $A_w(\cdot)$, jointly with the sampler. The only requirement for the reconstructor is that it is differentiable with respect to parameters w. We model the reconstructor as a neural network. Although many networks have been proposed for MR image reconstruction [10, 29, 32, 45], in this paper, we adopt a standard 8-block U-Net architecture [28] following [1, 2, 46]. The input to the reconstructor at each time t is the complex-valued zero-filled image, $\hat{\mathbf{x}}_t$, and the output is a single channel real-valued image, $\tilde{\mathbf{x}}_t$. The UNet reconstructor contains four downsampling blocks and four upsampling blocks, each consisting of two 3×3 convolutions separated by ReLU and instance normalization [35]. Our framework is agnostic to the specific reconstructor architecture.

5. Experimental Results

5.1. Setup and Implementation Details

We evaluate our sequential sampling and reconstruction method on the NYU fastMRI open dataset [46][3]. The dataset provides RAW single-coil κ-space measurements for knee images, with 973 training set volumes and 97 validation set volumes [46]. We follow the setup of [24] and split the original validation set into a new validation set with 48 volumes and a test set with 49 volumes, which results in 34,742 2D slices for training, 1,785 slices for validation, and 1,851 slices for testing. To save on computation, we crop the k-space to the center 128×128 region, as is done in [2, 14, 49].

We use the structural similarity index measure (SSIM) for our model's training loss, following [2, 24, 32]. SSIM is computed using a window size of 7×7 and hyperparmeters $k_1 = 0.01, k_2 = 0.03$ following the fastMRI challenge's official implementation. We use the Adam optimizer [17] and train our model for 50 epochs with a learning rate of $1e-3$ for 2D point sampling experiments and $5e-5$ for 1D line sampling experiments. The learning rate is decreased by half every ten epochs. Training each model takes at most one day on a single RTX 2080Ti GPU. We use SSIM as the primary evaluation metric, which has been found to correlate well with expert evaluations [18].

5.2. Results

In Figure 1, we visualize our framework's sequential sampling trajectory and final reconstruction for rotated knees in the $8\times$ acceleration setting. Starting from preselected measurements, our model sequentially samples 2D κ-space measurements based on previous observations. Here, we demonstrate that our model is able to accurately

[3] https://fastmri.org/

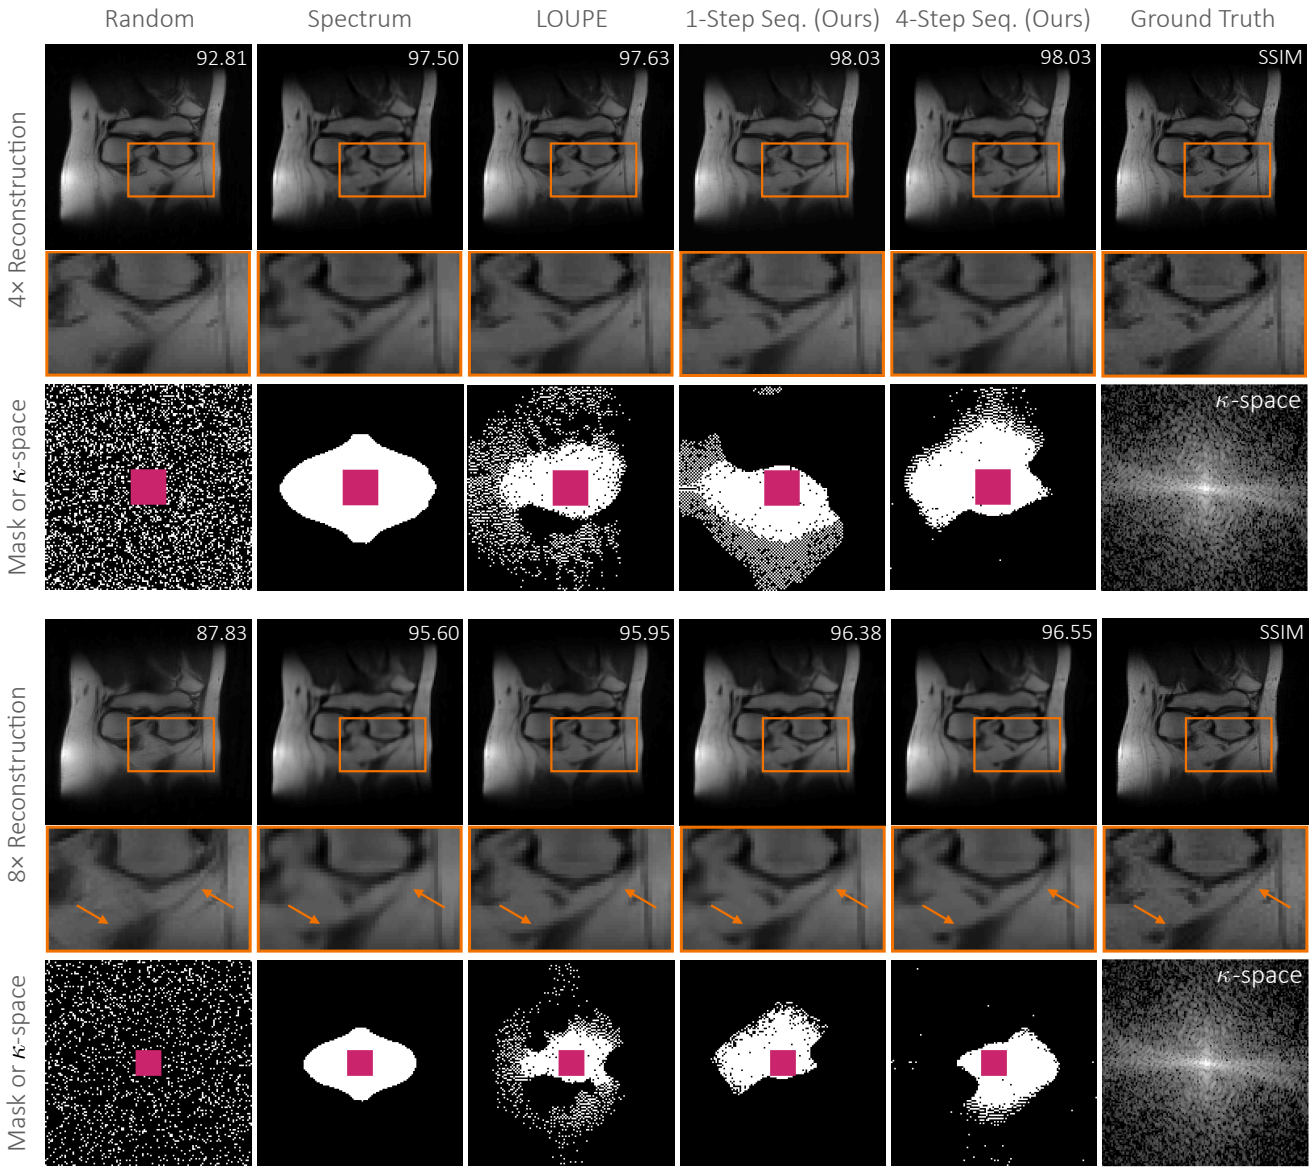

Figure 4: Visualizations of example reconstructions with an acceleration factor of 4× (top) and 8× (bottom) for 2D point sampling. A zoomed-in image patch is shown along with the cumulative κ-space measurements selected by each policy. Our sequential approach often provides more accurate reconstructions with detailed local structures. More visualizations are included in the supplemental material.

Methods	Random	Equispaced [46]	Evaluator [49]	PG-MRI [2]	LOUPE [1]	4-Step Seq. (Ours)
SSIM	85.95 ± 0.05	86.86 ± 0.06	85.99 ± 0.04	87.97 ± 0.09	89.52 ± 0.02	**91.08 ± 0.09**

Table 1: The SSIM comparison of 1D line sampling with a 4× acceleration factor. Our 4-step sequential model outperforms the previous approaches when tested on the fastMRI knee test set. A paired t-test shows a statistically significant difference between our 4-step sequential model and LOUPE [1], with a p-value smaller than 10^{-300}. For each model, we compute the test average and standard deviation obtained across three trained models with independent initialization.

estimate and leverage the κ-space structure during the sequential sampling steps. In particular, the final sampling patterns contain visible directional structures that align with the true κ-space power spectrum induced by knee rotation.

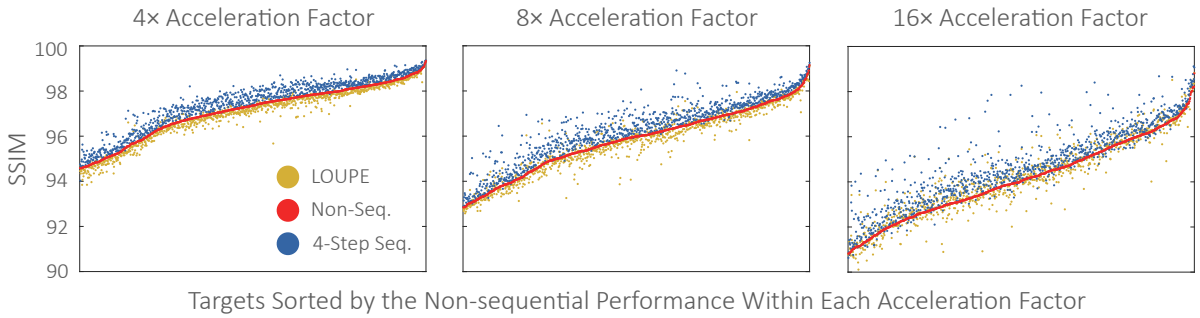

Figure 5: Pair-wise SSIM comparison of fastMRI knee targets using 2D point sampling for different acceleration factors. Results are sorted for the top 1,000 targets according to the performance of the non-sequential baseline on the test set. Our sequential model outperforms the LOUPE [1] and non-sequential baselines for most subjects. This performance pattern holds on all target reconstructions and is shown in the supplemental material. More quantitative results are shown in Table 3.

Acceleration	4×	8×	16×
Random	90.40 ± 0.02	87.43 ± 0.05	84.25 ± 0.00
Spectrum [38]	92.39 ± 0.01	90.38 ± 0.01	88.37 ± 0.01
LOUPE [1]	92.44 ± 0.01	90.60 ± 0.03	88.73 ± 0.04
4-Step Seq. (Ours)	**92.91 ± 0.01**	**91.07 ± 0.02**	**89.10 ± 0.03**

Table 2: SSIM comparison of 2D point sampling for 4×, 8×, and 16× accelerations. Our 4-step sequential model outperforms the previous approaches when tested on the fastMRI knee test set. For each model, we compute the test average and standard deviation obtained across three trained models with independent initialization.

This highlights the adaptivity of our sequential model, as no rotated anatomical images were included in the training set.

2D Point Sampling Figure 4 shows sample reconstructions obtained with an 8× and 16× acceleration ratio for multiple approaches. In each approach, the reconstruction network has been trained jointly with the specified sampling policy. See the supplementary material for the implementation details of these baseline methods. Using the same number of κ-space samples, our 1-step and 4-step sequential methods are able to recover structure that was lost in the baseline Random, Spectrum [38], and LOUPE [1] methods. Table 2 summarizes the quantitative comparison between methods for 4×, 8×, and 16× acceleration. Our 4-step sequential framework achieves the best reconstruction performance across different acceleration ratios for 2D point sampling. Additionally, a paired t-test between our method and the previous state-of-the-art pre-designed sampling approach, LOUPE [1], indicates a statistically significant difference in performance, with a p-value less than 10^{-160} for all acceleration ratios.

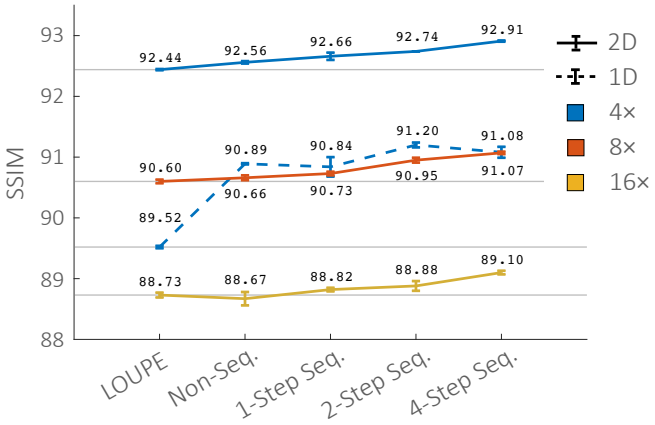

Figure 6: Comparison between our sequential model and the LOUPE model on the fastMRI knee test set. Our sequential model outperforms LOUPE for all acceleration ratios with an improvement comparable to 25% of the benefit of doubling the number of κ-space measurements. The performance of our sequential model in the 1D line sampling case significantly outerforms LOUPE but plateaus after 2 sequential sampling steps, possibly due to the restricted action space of 1D line sampling.

1D Line Sampling Table 1 compares our model to previous methods for 1D line sampling with a 4× acceleration factor. The baselines we consider include: (1) Random: randomly select κ-space lines from a uniform distribution, (2) Equispaced: select lines that are equidistant from each other [46], (3) Evaluator: sequentially select lines following the evaluation scoring function introduced in [49], (4) PG-MRI: sequentially select lines using conditional distribution trained using a policy gradient algorithm [2], (5) LOUPE: select lines prior to acquisition using a distribution learned via co-design [1]. The implementation details of these baselines are included in our supplemental material. Our 4-step sequential framework significantly outperforms prior meth-

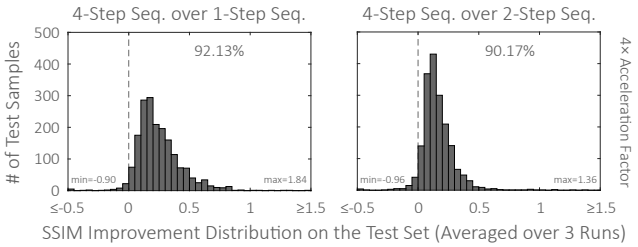

Figure 7: Histograms of pair-wise SSIM comparison on all 1,851 test images with a different number of sequential steps (T), using 2D point sampling with a $4\times$ acceleration factor. The relative error between the 4-step and 1-step (left) or 2-step(right) demonstrates that additional sequential steps help to boost performance, but with diminishing returns as T increases.

ods, with an SSIM improvement of 1.56 over the previous state-of-the-art learning-based method, LOUPE [1]. A paired t-test also indicates a highly statistically significant boost in performance compared to LOUPE with a t-score of 64.01 and a p-value smaller than 10^{-300}. Note that our differentiable end-to-end framework also significantly outperforms a sequential reinforcement learning optimization approach, PG-MRI [2].

Adaptive versus Pre-designed Sampling Figure 5 shows a pair-wise comparison between our 4-step sequential framework and the pre-designed sampling baselines. The baselines include previous state-of-the-art LOUPE method and a non-sequential baseline that uses the same network architecture as our sequential model, but replaces the prior κ-space measurements used as input with a random tensor (referred to as "Non-seq."). The non-sequential baseline is comparable to or outperforms the LOUPE baseline [1] (refer to Table 3). Across all three different acceleration ratios studied, our 4-step sequential model outperforms the pre-designed sampling baselines for most subjects in the fastMRI knee test set.

Number of Sequential Steps Figure 6 shows a comparison between our sequential models and the previous state-of-the-art LOUPE model [1]. For the case of 2D point sampling, the accuracy consistently increases as the number of sequential sampling steps increases. In fact, our 4-step sequential model achieves a 0.37-0.48 SSIM improvement over the LOUPE baseline for all acceleration factors, which is approximately 25% of the benefit of doubling the number of κ-space measurements (when switching from $8\times$ to $4\times$ acceleration for the LOUPE model SSIM increases by 1.84). To further understand the improvements seen with additional sequential steps, we perform a pair-wise SSIM comparison between our sequential models; Figure 7 shows the result of 2D point sampling with a $4\times$ acceleration ratio. Additional sequential steps boost the reconstruction performance for almost all subjects, with diminishing returns as T increases. Table 3 shows quantitative results that compare the percentage of test samples that outperform the LOUPE baseline. On 2D point sampling, our 4-step sequential model outperforms LOUPE roughly 97%, 89%, and 77% of the time for the $4\times$, $8\times$ and $16\times$ acceleration factors, respectively.

Acceleration	$4\times$	$8\times$	$16\times$
Non-Seq.	74.05 ± 2.56	60.18 ± 3.03	46.98 ± 8.58
1-Step Seq.	77.42 ± 7.89	57.05 ± 4.36	51.09 ± 4.16
2-Step Seq.	88.74 ± 0.45	83.04 ± 3.78	56.42 ± 4.62
4-Step Seq.	$\mathbf{96.96 \pm 0.73}$	$\mathbf{92.62 \pm 0.46}$	$\mathbf{76.91 \pm 2.29}$

Table 3: The percentage of test samples that outperform the LOUPE [1] baseline, demonstrating the performance of our framework as a function of the number of sequential sampling steps (T) for 2D point sampling. The percentage average and standard deviation are obtained using results from three trained models with independent initialization.

Co-design	1-Step Seq.	4-Step Seq.
Yes	$\mathbf{92.66 \pm 0.06}$	$\mathbf{92.91 \pm 0.01}$
No	90.33 ± 0.01	90.40 ± 0.02

Table 4: Ablation results showing the advantage of co-design with a $4\times$ acceleration ratio and 2D point sampling. When co-design is specified as "Yes" the reconstruction network has been jointly optimized with the sampler. Otherwise, the sampler was optimized with a fixed reconstructor that was pre-trained with a random sampling policy.

Co-design Ablation We demonstrate the advantage of co-designing the sampler and reconstructor in Table 4. Specifically, we pre-train a reconstructor using a uniform sampling policy and demonstrate the improvement in performance that occurs when jointly learning the reconstructor weights with the sampler. Co-designing the reconstructor with the sampler significantly improves performance, with an increase of 2.33-2.51 SSIM for 2D point sampling with a $4\times$ acceleration factor.

6. Conclusion

Accelerating the MRI acquisition process has the potential to reduce patient discomfort, increase throughput, and expand the use of MRI worldwide. In this paper, we have proposed an end-to-end sequential sampling and reconstruction framework for accelerated MR imaging. We leverage the sequential nature of MRI acquisition and design a model with explicit sequential structure that jointly

optimizes a neural network-based sampler simultaneously with a network-based reconstructor. In our experiments, this simple framework outperforms previous state-of-the-art MR sampling approaches for up to nearly 97% of the test samples on the fastMRI single-coil knee dataset. Our framework is trained through end-to-end differentiable learning, making our method easy to implement in standard machine learning libraries and train on new datasets.

In the future we plan to expand our general framework to handle more realistic experimental settings. In particular, by replacing our discrete 2D sampler with one that samples from a continuous 2D trajectory space, we can model more complex but feasible trajectories. Additionally, realistic subject motion can be simply incorporated into our model by perturbing the 3D target volume according to a motion model during the T acquisition steps. Other future directions include incorporating uncertainty quantification [33, 49] and integrating with tasks such as anatomical registration [3] or image segmentation [11] to arrive at more unified end-to-end frameworks. Overall, our results suggest that future methods for MRI sampling can benefit from the collaboration of sequential sampling and co-design via end-to-end learning.

References

[1] 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. 2, 3, 4, 5, 6, 7, 8

[2] Tim Bakker, Herke van Hoof, and Max Welling. Experimental design for mri by greedy policy search. *arXiv preprint arXiv:2010.16262*, 2020. 3, 5, 6, 7, 8

[3] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019. 9

[4] Jonathan Baxter and Peter L. Bartlett. Infinite-horizon policy-gradient estimation. *J. Artif. Int. Res.*, 15(1):319–350, Nov. 2001. 3

[5] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 5

[6] C. Bouman and K. Sauer. A generalized gaussian image model for edge-preserving map estimation. *IEEE Transactions on Image Processing*, 2(3):296–310, 1993. 2

[7] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*, 52(2):489–509, 2006. 2

[8] Nicolas Chauffert, Philippe Ciuciu, Jonas Kahn, and Pierre Weiss. Variable density sampling with continuous trajectories. *SIAM Journal on Imaging Sciences*, 7(4):1962–1992, 2014. 2

[9] J. P. Haldar, D. Hernando, and Z. Liang. Compressed-sensing mri with random encoding. *IEEE Transactions on Medical Imaging*, 30(4):893–903, 2011. 2

[10] 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. *CoRR*, abs/1704.00447, 2017. 2, 3, 5

[11] Mohammad Hesam Hesamian, Wenjing Jia, Xiangjian He, and Paul Kennedy. Deep learning techniques for medical image segmentation: achievements and challenges. *Journal of digital imaging*, 32(4):582–596, 2019. 9

[12] Yue Huang, John W. Paisley, Qin Lin, Xinghao Ding, Xueyang Fu, and Xiao-Ping (Steven) Zhang. Bayesian nonparametric dictionary learning for compressed sensing MRI. *IEEE Trans. Image Process.*, 23(12):5007–5019, 2014. 2

[13] Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, Sungchul Lee, and Jin Keun Seo. Deep learning for undersampled MRI reconstruction. *CoRR*, abs/1709.02576, 2017. 2

[14] Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547*, 2019. 3, 5

[15] Michael Kellman, Emrah Bostan, Michael Chen, and Laura Waller. Data-driven design for fourier ptychographic microscopy. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–8. IEEE, 2019. 2

[16] Michael R Kellman, Emrah Bostan, Nicole A Repina, and Laura Waller. Physics-based learned design: optimized coded-illumination for quantitative phase imaging. *IEEE Transactions on Computational Imaging*, 5(3):344–353, 2019. 2

[17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5

[18] Florian Knoll, Tullie Murrell, Anuroop Sriram, Nafissa Yakubova, Jure Zbontar, Michael Rabbat, Aaron Defazio, Matthew J Muckley, Daniel K Sodickson, C Lawrence Zitnick, et al. Advancing machine learning for mr image reconstruction with an open competition: Overview of the 2019 fastmri challenge. *Magnetic resonance in medicine*, 84(6):3054–3070, 2020. 5

[19] Dongwook Lee, Jae Jun Yoo, and Jong Chul Ye. Deep residual learning for compressed sensing MRI. In *14th IEEE International Symposium on Biomedical Imaging, ISBI 2017, Melbourne, Australia, April 18-21, 2017*, pages 15–18. IEEE, 2017. 2

[20] Jiaming Liu, Yu Sun, Cihat Eldeniz, Weijie Gan, Hongyu An, and Ulugbek S. Kamilov. RARE: image reconstruction using deep priors learned without groundtruth. *IEEE J. Sel. Top. Signal Process.*, 14(6):1088–1099, 2020. 2

[21] Jiaming Liu, Yu Sun, Weijie Gan, Xiaojian Xu, Brendt Wohlberg, and Ulugbek S. Kamilov. Sgd-net: Efficient model-based deep learning with theoretical guarantees. *CoRR*, abs/2101.09379, 2021. 2

[22] Michael Lustig, David Donoho, and John Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic resonance in medicine : official journal of*

the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine, 58:1182–95, 12 2007. 2, 4

[23] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing mri. *IEEE Signal Processing Magazine*, 25(2):72–82, 2008. 2

[24] 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. 3, 5

[25] Tran Minh Quan, Thanh Nguyen-Duc, and Won-Ki Jeong. Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss. *IEEE Trans. Medical Imaging*, 37(6):1488–1497, 2018. 2

[26] Saiprasad Ravishankar and Yoram Bresler. Adaptive sampling design for compressed sensing MRI. In *33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC 2011, Boston, MA, USA, August 30 - Sept. 3, 2011*, pages 3751–3755. IEEE, 2011. 3

[27] Saiprasad Ravishankar and Yoram Bresler. MR image reconstruction from highly undersampled k-space data by dictionary learning. *IEEE Trans. Medical Imaging*, 30(5):1028–1041, 2011. 2

[28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5

[29] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony N. Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans. Medical Imaging*, 37(2):491–503, 2018. 2, 5

[30] Shiqian Ma, Wotao Yin, Yin Zhang, and A. Chakraborty. An efficient algorithm for compressed mr imaging using total variation and wavelets. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 2

[31] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, Jan. 2016. 3

[32] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated mri reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 64–73. Springer, 2020. 5

[33] He Sun and Katherine L Bouman. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. *arXiv preprint arXiv:2010.14462*, 2020. 9

[34] H. Sun, A. V. Dalca, and K. L. Bouman. Learning a probabilistic strategy for computational imaging sensor selection. In *2020 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2020. 2, 3

[35] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. 5

[36] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 2094–2100. AAAI Press, 2016. 3

[37] S. Vasanawala, M. Murphy, M. Alley, P. Lai, K. Keutzer, J. Pauly, and M. Lustig. Practical parallel imaging compressed sensing mri: Summary of two years of experience in accelerating body mri of pediatric patients. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1039–1043, 2011. 2

[38] Jaganathan Vellagoundar and Ramasubba Reddy Machireddy. A robust adaptive sampling method for faster acquisition of mr images. *Magnetic resonance imaging*, 33(5):635—643, June 2015. 7

[39] Guanhua Wang, Tianrui Luo, Jon-Fredrik Nielsen, Douglas C. Noll, and Jeffrey A. Fessler. B-spline parameterized joint optimization of reconstruction and k-space trajectories (bjork) for accelerated 2d mri, 2021. 3

[40] Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, and Dong Liang. Accelerating magnetic resonance imaging via deep learning. In *13th IEEE International Symposium on Biomedical Imaging, ISBI 2016, Prague, Czech Republic, April 13-16, 2016*, pages 514–517. IEEE, 2016. 2, 3

[41] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4

[42] Tomer Weiss, Ortal Senouf, Sanketh Vedula, Oleg Michailovich, Michael Zibulevsky, and Alex Bronstein. Pilot: Physics-informed learned optimized trajectories for accelerated mri, 2020. 3

[43] Tomer Weiss, Sanketh Vedula, Ortal Senouf, Oleg V. Michailovich, Michael Zibulevsky, and Alexander M. Bronstein. Joint learning of cartesian under sampling andre construction for accelerated MRI. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 8653–8657. IEEE, 2020. 3

[44] Guang Yang, Simiao Yu, Hao Dong, Gregory G. Slabaugh, Pier Luigi Dragotti, Xujiong Ye, Fangde Liu, Simon R. Arridge, Jennifer Keegan, Yike Guo, and David N. Firmin. DAGAN: deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction. *IEEE Trans. Medical Imaging*, 37(6):1310–1321, 2018. 2

[45] Yan Yang, Jian Sun, Huibin Li, and Zongben Xu. Deep admm-net for compressive sensing MRI. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 10–18, 2016. 2, 5

[46] Jure Zbontar, Florian Knoll, Anuroop Sriram, Matthew J. Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: An open dataset and benchmarks for accelerated MRI. 2018. 2, 5, 6, 7

[47] Zhifang Zhan, Jian-Feng Cai, Di Guo, Yunsong Liu, Zhong Chen, and Xiaobo Qu. Fast multi-class dictionaries learning with geometrical directions in MRI reconstruction. *CoRR*, abs/1503.02945, 2015. 2

[48] Jinwei Zhang, Hang Zhang, Alan Wang, Qihao Zhang, Mert R. Sabuncu, Pascal Spincemaille, Thanh D. Nguyen, and Yi Wang. Extending LOUPE for k-space under-sampling pattern optimization in multi-coil MRI. In *Machine Learning for Medical Image Reconstruction - Third International Workshop, MLMIR 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 8, 2020, Proceedings*, volume 12450 of *Lecture Notes in Computer Science*, pages 91–101. Springer, 2020. 3, 5

[49] 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. 3, 5, 6, 7, 9

[50] Bo Zhu, Jeremiah Z. Liu, Stephen F. Cauley, Bruce R. Rosen, and Matthew S. Rosen. Image reconstruction by domain-transform manifold learning. *Nat.*, 555(7697):487–492, 2018. 2