

# LEARNING-BASED OPTIMIZATION OF THE UNDER-SAMPLING PATTERN IN MRI

A Thesis

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

M.Sc.

by

Cagla Deniz Bahadir

May 2019

© 2019 Cagla Deniz Bahadir

ALL RIGHTS RESERVED

## ABSTRACT

The long scan times of Magnetic Resonance Imaging (MRI) create a bottleneck in patient care and acquisitions can be accelerated by under-sampling in k-space (i.e., the Fourier domain). In this thesis, we focus on the optimization of the sub-sampling pattern with a data-driven framework. Since the reconstruction quality of the models are shown to be strongly dependent on the sub-sampling pattern, we combine the two problems. For a provided sparsity constraint, our method optimizes the sub-sampling pattern and reconstruction model, using an end-to-end unsupervised learning strategy. Our algorithm is trained on full-resolution data that are under-sampled retrospectively, yielding a sub-sampling pattern and reconstruction model that are customized to the type of images represented in the data set. The proposed method, which we call LOUPE (Learning-based Optimization of the Under-sampling Pattern), was implemented by modifying a U-Net, a widely-used convolutional neural network architecture, that we append with the forward model that encodes the under-sampling process. Our experiments with T1-weighted structural brain MRI scans, PD and PDFS weighted knee MRI scans show that the optimized sub-sampling pattern can yield significantly more accurate reconstructions compared to standard random uniform, variable density or cartesian under-sampling schemes. The code is made available at: <https://github.com/caglabahadir/LOUPE>.

**Keywords:** k-space Under-sampling, Convolutional Neural Networks, Compressed Sensing.

## BIOGRAPHICAL SKETCH

Cagla Deniz Bahadir is a Fulbright Scholar and a second year graduate student in the Department of Biomedical Engineering at Cornell University. She graduated with a Bachelor's Degree in Electrical and Electronics Engineering from Bilkent University, Turkey. During her undergraduate studies she worked on Magnetic Fluid Hyperthermia and a new imaging modality: Magnetic Particle Imaging. She further conducted research in National Magnetic Resonance Research Center (UMRAM) in Turkey, working on 2D RF Pulse Design for Magnetic Resonance Imaging. She is working under the supervision of Dr. Mert Rory Sabuncu. Her current research interests are application of Machine Learning to Magnetic Resonance Imaging, particularly in the area of Compressed Sensing. She plays the piano and enjoys writing short stories.

This thesis is dedicated to my family and friends, for their endless support and to my late grandfather who would have been extremely proud to see me become a research scientist.

## ACKNOWLEDGEMENTS

This thesis has been supervised by the chairperson Dr. Mert Rory Sabuncu and two committee members: Dr. Yi Wang and Dr. Peter Doerschuk.

This work was accepted as a conference paper in the 26<sup>th</sup> international conference on Information Processing in Medical Imaging (IPMI), June 2019, Hong Kong and co-authored by Dr. Adrian V. Dalca and Dr. Mert Rory Sabuncu. The text is available on arxiv as pre-print [2].

This work was supported by NIH R01 grants (R01LM012719 and R01AG053949), the NSF NeuroNex grant 1707312, and NSF CAREER grant (1748377). I would also like to thank the Fulbright Commission for making this research possible.

## TABLE OF CONTENTS

Biographical Sketch . . . . .	iii
Dedication . . . . .	iv
Acknowledgements . . . . .	v
Table of Contents . . . . .	vi
List of Tables . . . . .	vii
List of Figures . . . . .	viii
<b>1 INTRODUCTION</b>	<b>1</b>
1.1 Machine Learning for Under-sampled Image Reconstruction . . .	2
1.2 Optimization of the Sub-sampling Pattern . . . . .	3
<b>2 METHOD</b>	<b>5</b>
2.1 Learning-based Optimization of the Under-sampling Pattern . . .	5
2.2 Implementation . . . . .	7
<b>3 Empirical Analysis</b>	<b>10</b>
3.1 Data . . . . .	10
3.1.1 ABIDE-I . . . . .	10
3.1.2 NYU fastMRI . . . . .	11
3.2 Evaluation . . . . .	12
3.3 Benchmark Reconstruction Methods . . . . .	12
3.4 Sub-sampling Masks . . . . .	14
<b>4 Results</b>	<b>17</b>
<b>5 Discussion</b>	<b>26</b>
<b>Bibliography</b>	<b>27</b>

## List of Tables

4.1 Average per volume run times (in sec) for different reconstruction methods. All except U-Net (GPU) were evaluated on a CPU - a dual Intel Xeon (E5-2640, 2.4GHz). . . . .	19
---	----

## List of Figures

2.1	The neural network architecture for LOUPE. Each vertical blue line represents a 2D image, with the number of channels indicated above and the size shown on the lower left side. The green line represents a 2D real-valued image of <i>weight</i> parameters, where one parameter is learned at each location, which is then passed through a sigmoid to yield the probability mask $\mathbf{p}$ . . . . .	7
3.1	Optimized and benchmark masks for two levels of sub-sampling rates for ABIDE-I data set: $R = 10$ and $R = 20$ . Figures are in 2D k-space and black dots indicate the points at which a sample is acquired. Representative instantiations are visualized for the random masks. . . . .	14
3.2	Optimized and benchmark masks for two levels of sub-sampling rates for NYU fastMRI data set: $R = 4$ and $R = 8$ . Figures are in 2D k-space and black dots indicate the points at which a sample is acquired. Representative instantiations are visualized for the random masks. . . . .	15
4.1	Optimized masks compared side to side for the knee anatomy and brain anatomy for sub-sampling rate of $R = 8$ and $R = 10$ respectively.	18
4.2	Quantitative evaluation of reconstruction quality for ABIDE-I data set. For each plot, we show four reconstruction methods using four acquisition masks, including the Optimized Mask obtained using LOUPE in green. Each dot is the PSNR value for a single test subject across slices. For each box, the blue line shows the median value, and the whiskers indicate the the most extreme (non-outlier) data points. . . . .	20
4.3	Quantitative evaluation of reconstruction quality for the NYU fastMRI data set. For each plot, we show four reconstruction methods using four acquisition masks, including the Optimized Mask obtained using LOUPE in green. Each dot is the PSNR value for a single test subject across slices. For each box, the blue straight line shows the median value for patients with Proton Density images while the blue dashed line shows the median value for patients with Proton Density Fat Supressed images The whiskers indicate the the most extreme (non-outlier) data points.. . . . .	21
4.4	Reconstructions for a representative slice from ABIDE-I experiments with $R = 10$ sub-sampling rate. Each row is a reconstruction method. Each column corresponds to a sub-sampling mask. We observe that our optimized mask yields reconstructions that capture more anatomical detail. Red arrows highlight some nuanced features that were often missed in reconstructions. . . . .	22
4.5	Reconstructions for a representative slice from ABIDE-I experiments with $R = 20$ sub-sampling rate. See caption of Figure 4.4 and text for more detail. . . . .	23

4.6 Reconstructions for a representative slice from NYU fastMRI experiments with $R = 4$ sub-sampling rate. Each row is a reconstruction method. Each column corresponds to a sub-sampling mask. We observe that our optimized mask yields reconstructions that capture more anatomical detail. . . . .	24
4.7 Reconstructions for a representative slice from NYU fastMRI experiments with $R = 8$ sub-sampling rate. See caption of Figure 4.6 and text for more detail. . . . .	25

# CHAPTER 1

## INTRODUCTION

MRI is a highly established, reliable and safe imaging technique that has been playing a central role in diagnosis. However, the long scan times have resulted in a decrease in availability of the MRI scanners and increase in costs. Hardware modifications have been proposed such as the parallel imaging, which depends on simultaneous multi-coil data acquisition for speeding purposes. Another widely used acceleration technique is Compressed Sensing (CS) [21], which does not demand changes in the MR hardware.

MRI acquires the data in the Fourier Domain, also known as k-space, and images are computed by solving the inverse Fourier transform that converts k-space data into the spatial domain. Medical images demonstrate continuity amongst the neighboring pixels and intensity values usually vary smoothly over space, except at a small number of boundary voxels. This regularity leads to redundancy in k-space and creates an opportunity for sampling below the Shannon-Nyquist rate [21]. Several Cartesian and non-Cartesian under-sampling patterns have been proposed in the literature and are widely used in practice, such as Random Uniform [8], Variable Density [33] and equispaced Cartesian [12] with skipped lines.

A linear reconstruction of under-sampled k-space data (i.e., a direct inverse Fourier) often falls short in terms of providing clinically acceptable results and yields aliasing artifacts, which are challenging to distinguish from real image features for regular sub-sampling patterns. Stochastic sub-sampling patterns, on the other hand, create incoherent, noise-like artifacts that are relatively easier to remove [21]. The classical reconstruction strategy in CS involves regu-

larized regression, where a non-convex objective function that includes a data fidelity term on the acquired k-space measurements and a regularization term is optimized for a given set of measurements. The regularization term reflects our *a priori* knowledge of regularity in natural images. Common examples include sparsity-encouraging penalties such as L1-norm on wavelet coefficients and total variation [22].

In regularized regression, optimization is achieved via iterative numerical strategies, such as gradient-based methods, which can be computationally demanding. Furthermore, the choice of the regularizer is often arbitrary and not optimized in a data-driven fashion. These drawbacks can be addressed using machine learning approaches, which enable the use of models that learn from data and facilitate very efficient and fast reconstructions.

## 1.1 Machine Learning for Under-sampled Image Reconstruction

A variety of penalty terms have been implemented with multiple dictionary learning techniques [15, 26, 28] in regularized regression-based reconstruction. A common strategy is to project the images (or patches) onto a “sparsifying” dictionary. Thus, a sparsity-inducing norm, such as L1, on the associated coefficients can be used as a regularizer. The drawback of such methods is that they still rely on iterative numerical optimization, which can be computationally expensive.

Recently, deep learning has been used to speed up and improve the quality of under-sampled MRI reconstructions [19, 25, 27, 31, 34]. These models are trained on data to learn to do anti-aliasing and denoising operations from under-sampled images to high quality reconstructions. The computation is ef-

efficient for a new data point due to non-iterative nature of the method and reconstruction is achieved via a single forward pass through the “anti-aliasing” neural network. However, these machine learning-based methods are typically optimized for a specific under-sampling pattern provided by the user. Furthermore, there are also techniques that are optimizing the sub-sampling patterns for given reconstruction methods [9, 3, 24, 23]. The reconstruction model’s performance will depend significantly on the sub-sampling pattern. In this paper, we are interested in optimizing the sub-sampling pattern in a data-driven fashion. Therefore, our method optimizes the sub-sampling pattern *and* reconstruction model *simultaneously*, using an end-to-end learning strategy. We are able to achieve this thanks to the two properties of deep learning based reconstruction models: their speed and differentiable nature. These properties enable us to rapidly evaluate the effect of small changes to the sub-sampling pattern on reconstruction quality.

## 1.2 Optimization of the Sub-sampling Pattern

Several studies suggested ways to optimize the sub-sampling pattern in compressed sensing MRI. The OEDIPUS framework [13] uses the information-theoretic Cramer-Rao bound to compute a deterministic sampling pattern that is tailored to the specific imaging context. Seeger et al [30] present a Bayesian approach to optimize k-space sampling trajectories under sparsity constraints. The resulting algorithm is computationally expensive and does not scale well to large data sets. To alleviate this drawback, Liu et al. [20] propose a computationally more efficient strategy to optimize the under-sampling trajectory. However, this method does not consider a sophisticated reconstruction tech-

nique. Instead, they merely optimize for the simple method of inverse Fourier transform with zero-filling.

Below, we describe the proposed method, LOUPE, that computes the optimal probabilistic sub-sampling mask together with a state-of-the-art rapid neural network based reconstruction model. We train LOUPE using an end-to-end unsupervised learning approach with retrospectively sub-sampled images.

## CHAPTER 2

### METHOD

#### 2.1 Learning-based Optimization of the Under-sampling Pattern

In this section, we describe our unsupervised end-to-end optimization technique for the under-sampling pattern, together with the reconstruction model. We provide the details of our novel problem formulation and the approach we implement to solve it. We call our algorithm LOUPE, which stands for Learning-based Optimization of the Under-sampling Pattern. LOUPE employs an end-to-end learning strategy and considers the two fundamental problems of compressed sensing simultaneously: the optimization of the under-sampling pattern and learning a reconstruction model that rapidly solves the ill-posed anti-aliasing problem.

In LOUPE, we seek a “probabilistic mask”  $\mathbf{p}$  that describes an independent Bernoulli (binary) random variable  $\mathcal{B}$  at each k-space (discrete Fourier domain) location on the full-resolution grid. Thus, a probabilistic mask  $\mathbf{p}$  is an image of probability values in k-space. A binary mask  $\mathbf{m}$  has a value of 1 (0) that indicates that a sample is (not) acquired at the corresponding k-space point. We assume  $\mathbf{m}$  is a realization of  $\mathbf{M} \sim \prod_i \mathcal{B}(\mathbf{p}_i)$ , where  $i$  is the k-space location index. Let  $\mathbf{x}_j$  denote a full-resolution (e.g., 2D) MRI slice in the image (spatial) domain, where  $j$  is the scan index. While  $\mathbf{p}$ ,  $\mathbf{M}$ ,  $\mathbf{m}$  and  $\mathbf{x}_j$  are defined on a 2D grid (in k-space or image domain), we vectorize them in our mathematical expressions. Our method is not constrained to 2D images and can be applied 3D sampling grids as well.

LOUPE aims to solve the following optimization problem:

$$\arg \min_{\mathbf{p}, A} \mathbb{E}_{\mathbf{M} \sim \prod_i \mathcal{B}(\mathbf{p}_i)} \left[ \lambda \sum_i \mathbf{M}_i + \sum_j \|A(F^H \text{diag}(\mathbf{M})F \mathbf{x}_j) - \mathbf{x}_j\|_1 \right], \quad (2.1)$$

where  $F$  is the (forward) Fourier transform matrix,  $F^H$  is its inverse (i.e., Hermitian transpose of  $F$ ),  $A(\cdot)$  is an anti-aliasing (de-noising) function that we will parameterize via a neural network,  $\mathbf{M}_i \sim \mathcal{B}(\mathbf{p}_i)$  is an independent Bernoulli,  $\text{diag}(\mathbf{M})$  is a diagonal matrix with diagonal elements set to  $\mathbf{M}$ ,  $\lambda \in \mathbb{R}^+$  is a hyper-parameter, and  $\|\cdot\|_1$  denotes the L1-norm of a vector. While in our experiments  $\mathbf{x}_j$  is real-valued,  $F$  and  $F^H$  are complex valued, and  $A(\cdot)$  accepts a complex-valued input. We design  $A$  to output a real-valued image.

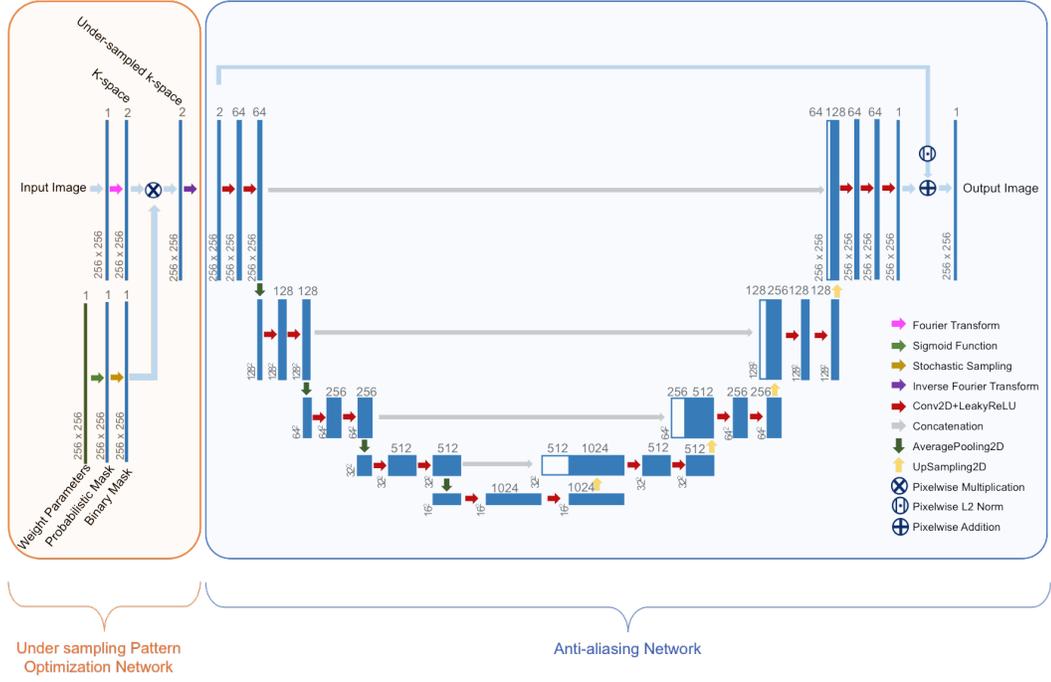
The first term in Eq. 2.1 is a sparsity penalty that encourages the number of k-space points that will be sampled to be small. The hyper-parameter  $\lambda$  controls the influence of the sparsity penalty, where higher values yield a more aggressive sub-sampling factor. We approximate the second term using a Monte Carlo approach. Thus the LOUPE optimization problem becomes:

$$\arg \min_{\mathbf{p}, A} \lambda \sum_i \mathbf{p}_i + \sum_j \frac{1}{K} \sum_{k=1}^K \|A(F^H \text{diag}(\mathbf{m}^{(k)})F \mathbf{x}_j) - \mathbf{x}_j\|_1, \quad (2.2)$$

where  $\mathbf{m}^{(k)}$  is an independent binary mask realization of  $\mathbf{M} \sim \prod_i \mathcal{B}(\mathbf{p}_i)$ , and we use  $K$  samples. We further re-parameterize the second term of Eq. 2.2:

$$\arg \min_{\mathbf{p}, A} \lambda \sum_i \mathbf{p}_i + \sum_j \frac{1}{K} \sum_{k=1}^K \|A(F^H \text{diag}(\mathbf{u}^{(k)} \leq \mathbf{p})F \mathbf{x}_j) - \mathbf{x}_j\|_1, \quad (2.3)$$

where  $\mathbf{u}^{(k)}$  is a realization of a random vector of independent uniform random variables on  $[0, 1]$ , and  $\mathbf{u}^{(k)} \leq \mathbf{p}$  is a binary random vector where each entry is set to 1 if the inequality is satisfied, and 0 otherwise.



**Fig. 2.1.** The neural network architecture for LOUPE. Each vertical blue line represents a 2D image, with the number of channels indicated above and the size shown on the lower left side. The green line represents a 2D real-valued image of *weight* parameters, where one parameter is learned at each location, which is then passed through a sigmoid to yield the probability mask  $\mathbf{p}$ .

## 2.2 Implementation

LOUPE is implemented with deep neural networks, specifically a fully convolutional U-Net for the reconstruction component. The network solves the learning problem via stochastic gradient descent. To avoid vanishing gradient problems and make the loss function differentiable everywhere, we relax the thresholding operation in Eq. 2.3 via a sigmoid:

$$\arg \min_{\mathbf{p}, \theta} \lambda \sum_i \mathbf{p}_i + \sum_j \frac{1}{K} \sum_{k=1}^K \|A_\theta(F^H \text{diag}(\sigma_s(\mathbf{u}^{(k)} - \mathbf{p}))F\mathbf{x}_j) - \mathbf{x}_j\|_1, \quad (2.4)$$

where  $\sigma_s(a) = \frac{1}{1+e^{-sa}}$ , and  $A_\theta$  denotes a neural network parameterized with weights  $\theta$ . We set the slope for this sigmoid to be relatively steep to better approximate the thresholding step function.

The second major component of the model, which is an anti-aliasing and denoising function  $A_\theta$  is a fully-convolutional neural network that builds on the widely used U-Net architecture [29]. The input to  $A_\theta$  is a two-channel 2D image, which correspond to the real and imaginary components of the under sampled image. As in [19], the U-Net estimates the difference between the aliased reconstruction (i.e., the result of applying the inverse Fourier transform to the zero-filled under-sampled k-space measurements), and the fully-sampled ground truth image. Finally, the probabilistic mask  $\mathbf{p}$  is formed by passing an unrestricted real-valued image through a sigmoid. Figure 2.1 illustrates the full architecture that combines all these elements. The red arrows represent 2D convolution layers with a kernel size  $3 \times 3$ , and a Leaky ReLU activation followed by Batch Normalization. The convolutions use zero-padding to match the input and output sizes. The gray arrows indicate skip connections, which correspond to concatenation operations. We also implement a stochastic sampling layer that draws uniform random vectors  $\mathbf{u}^{(k)}$ . This is similar to the Monte Carlo strategy used in variational neural networks [17].

We train our model on a data set of full-resolution images  $\{\mathbf{x}_j\}$ . Thus, LOUPE minimizes the unsupervised loss function 2.4 using an end-to-end learning strategy to obtain the probabilistic mask  $\mathbf{p}$  and the weights  $\theta$  of the anti-aliasing network  $A_\theta$ . The hyper-parameter  $\lambda$  is set empirically to obtain the desired sparsity. We implement our neural network in Keras [4], with TensorFlow [1] as the back-end and using layers from Neuron library [5]. The code is made available

at: <https://github.com/cagladbahadir/LOUPE>. We use the ADAM [16] optimizer with an initial learning rate of 0.001 and terminate learning when validation loss plateaued. Our mini-batch size is 32 for the ABIDE-I data set and 12 for the NYU fastMRI data set and  $K = 1$  for both. The input images are randomly shuffled.

## CHAPTER 3

### EMPIRICAL ANALYSIS

#### 3.1 Data

We conducted our experiments in two distinct data sets. ABIDE-I (Autism Brain Imaging Data Exchange) [6] comprises of reconstructed brain scans and NYU fastMRI ([fastmri.med.nyu.edu](http://fastmri.med.nyu.edu)) [35] comprises of raw k-space values of knee scans. NYU fastMRI investigators (the full list can be found at: [fastmri.med.nyu.edu](http://fastmri.med.nyu.edu)) provided data but did not participate in analysis or writing of this report. The primary goal of NYU fastMRI data set is to evaluate the performance and potential of Machine Learning methods in reconstruction of medical images.

##### 3.1.1 ABIDE-I

In our analyses, we used 3D T1-weighted brain MRI scans from the multi-site ABIDE-1 study [6]. We used 100 high quality volumes, as rated by independent experts via visual assessment, for training LOUPE, while a non-overlapping set of fifty subjects were used for validation. For testing all methods, including LOUPE, we used ten held-out independent test subjects. All our experiments were conducted on 2D axial slices, which consisted of  $1 \times 1\text{mm}^2$  pixels and were of size  $256 \times 256$ . We extracted 175 slices from each 3D volume, which provided full coverage of the brain - our central region of interest, and excluded slices that were mostly background.

### 3.1.2 NYU fastMRI

We also conducted experiments on NYU fastMRI dataset [35]; which is an openly available, large-scale, public data set of raw k-space data of knee scans. The data set originally comprises of 2D coronal knee scans acquired with two pulse sequences that result with Proton Density (PD) and Proton Density Fat Suppressed (PDFS) weighted images. The multi-coil knee scans were collected from four different MRI scanners: Skyra (3T), Prisma (3T), Biograph mMR (3T) and Aera (1.5T). We used the provided emulated single-coil (ESC) k-space data of scans, which is derived from raw 15-channel multi-coil data and linearly combined to fit the ground truth of root mean squares solution in a least-squares sense. Amongst the four scanners, we picked Biograph mMR (3T), due to appropriate number and quality of the scans. We used 100 volumes for training from the official provided training file, and split the validation file as: 10 volumes for validation and 10 volumes for test, due to original test file not having the fully sampled ground truth scans for challenge purposes. The provided sequence parameters were: Echo train length of 4, matrix size of 320x320, in plane resolution of 0.5mm x 0.5 mm, slice thickness of 3mm and no gap between slices. The time of repetition (TR) varied between 2200 and 3000ms and the Echo Time (TE) ranged between 27 and 34ms. Training volumes had  $38 \pm 4$  slices, where the validation volumes had  $37 \pm 3$  and test volumes had  $38 \pm 4$ .

Each set (training, validation and test) had differing slice sizes amongst volumes. After taking the Inverse Fourier Transform (IFFT) of the ESC k-space data, and rotating and flipping the images to match the orientation in the fastMRI paper, we cropped the central 320x320 as suggested and normalized the magnitude within each volume.

The U-Net configuration was also used for NYU fastMRI and only differed in terms of image size and two channel input layer that accepts complex valued images. The model was trained with the Adam optimizer, with an initial learning rate of 0.001 and a batch size of 12. A constraint of 32x32 central calibration region was imposed on the mask. Furthermore, a linearly increasing slope after each epoch was used for the sigmoid in Eq. 2.4 for a facilitated optimization and better approximation of the binary mask.

## 3.2 Evaluation

PSNR is a standard and well-established metric of reconstruction quality used in compressed sensing MRI [31]. During testing, we computed peak signal to noise ratio (PSNR) between the reconstructions of the different models and the full-resolution ground truth images for each volume. Our quantitative results with other metrics (not included) were also consistent.

## 3.3 Benchmark Reconstruction Methods

We used several benchmark reconstruction techniques to evaluate the performance of the optimized mask compared to other mask configurations, in widely used and already established frameworks. The first benchmark method is ALOHA [18], which uses a low-rank Hankel matrix to assign missing k-space values. We employed the code distributed by the authors<sup>1</sup> for the ABIDE-I experiments. We optimized the input parameters to minimize the MAE on a

---

<sup>1</sup><https://bispl.weebly.com/aloha-for-mr-recon.html>

training subject, since the default settings of the model did not yield acceptable reconstructions for our data set.

The second benchmark reconstruction method employs a novel regularized regression technique that combines total generalized variation (TGV) and the shearlet transform. This method has been demonstrated to preserve the edge information and be robust against artifacts and overall yield excellent accuracy in compressed sensing MRI [10]. We used the code provided by the authors<sup>2</sup>.

Our third benchmark method is based on the Block Matching 3D (BM3D) method, which was recently shown to offer high quality reconstructions for under-sampled MRI data [7]. BM3D is an iterative method that alternates between a de-noising step and a reconstruction step. We employed the open source code<sup>3</sup>.

The fourth benchmark method is based on low-rank modeling of local k-space neighborhoods (LORAKS) [11, 14], which we used for the NYU fastMRI experiments. The model uses a novel regularization term and algorithm to encourage low rank during reconstruction. The method is shown to yield successful reconstructions and denoising for under-sampled cases. We used the open source code provided by the authors<sup>4</sup>.

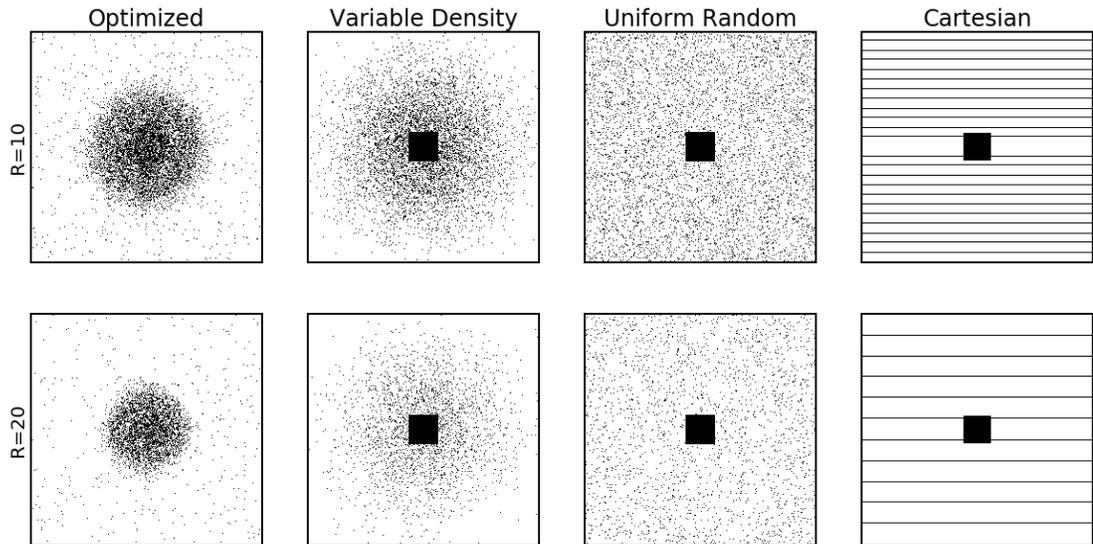
Finally, we consider a U-Net based reconstruction method, similar to the recently proposed deep residual learning for anti-aliasing technique of [19]. This reconstruction model is the one we used in LOUPE, with an important difference: in the benchmark implementation, the anti-aliasing model is trained from scratch, for each sub-sampling mask, separately. In LOUPE, this model is

---

<sup>2</sup>[http://www.math.ucla.edu/~wotaoyin/papers/tgv\\_shearlet.html](http://www.math.ucla.edu/~wotaoyin/papers/tgv_shearlet.html)

<sup>3</sup>[http://web.itu.edu.tr/eksioglu/pubs/BM3D\\_MRI.htm](http://web.itu.edu.tr/eksioglu/pubs/BM3D_MRI.htm)

<sup>4</sup><https://mr.usc.edu/download/loraks2/>

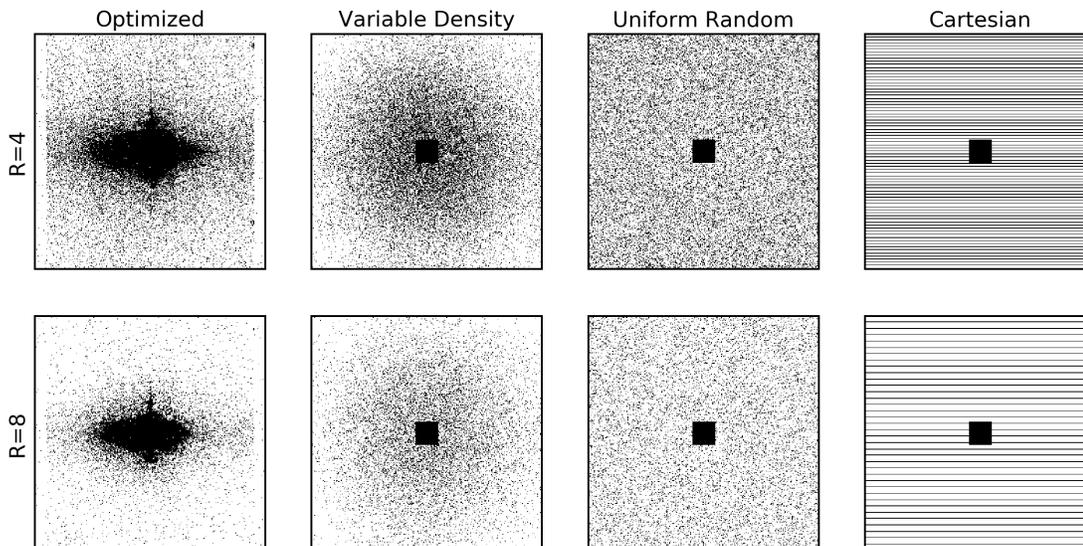


**Fig. 3.1.** Optimized and benchmark masks for two levels of sub-sampling rates for ABIDE-I data set:  $R = 10$  and  $R = 20$ . Figures are in 2D k-space and black dots indicate the points at which a sample is acquired. Representative instantiations are visualized for the random masks.

trained *jointly* with the optimization of the sub-sampling mask.

### 3.4 Sub-sampling Masks

In this study, we consider three different sub-sampling patterns that are widely used in the literature: Random Uniform [8], Random Variable Density [33] and equispaced Cartesian [12] - all with a fixed  $32 \times 32$  so-called “calibration region” in the center of the k-space. The calibration region is a fully sampled rectangular region around the origin, and has been demonstrated to yield better reconstruction performance [32]. We experimented with excluding the calibration region and sub-sampling over the entire k-space. However, reconstruction performance was no better than including the calibration region, so we omit these



**Fig. 3.2.** Optimized and benchmark masks for two levels of sub-sampling rates for NYU fastMRI data set:  $R = 4$  and  $R = 8$ . Figures are in 2D k-space and black dots indicate the points at which a sample is acquired. Representative instantiations are visualized for the random masks.

results.

The Uniform and Variable Density patterns were randomly generated by drawing independent Bernoulli samples. For Uniform, the probability value at each k-space point was the same and equal to the desired sparsity level. For Variable Density, the probability value at each k-space point was chosen from a Gaussian distribution, centered at the k-space origin. The proportionality constant was set to achieve the desired sparsity level. The Cartesian sub-sampling pattern is deterministic, and yields a k-space trajectory that is straightforward to implement. Figure 3.1 visualizes these masks for ABIDE-I and figure 3.2 visualizes the masks for NYU fastMRI. We consider two sparsity levels for the ABIDE-I data set: 10% and 5%, which correspond to  $R = 10$  and  $R = 20$  sub-sampling rates. We also used two sparsity levels for the NYU fastMRI data set:

25% and 12.5%, which correspond to  $R = 4$  and  $R = 8$  sub-sampling rates.

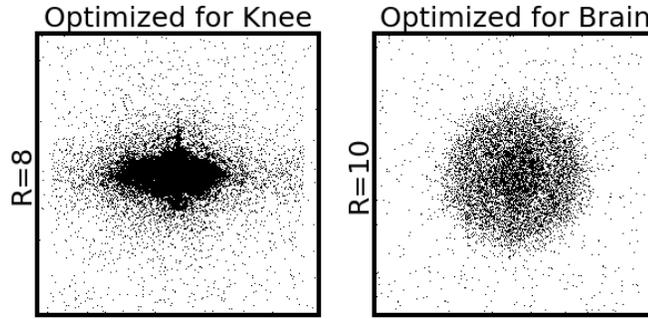
## CHAPTER 4

### RESULTS

The models have been evaluated in terms of reconstruction quality and time efficiency. Table 4.1 lists run time statistics for the different reconstruction methods, computed on the test subjects. For the U-Net, we provide run-times for both GPU (NVidia Titan Xp) and CPU. The U-Net model is significantly faster than the other reconstruction methods, which are all iterative. This speed, combined with the fact that the neural network model is differentiable, enabled us to use the U-Net in the end-to-end learning of LOUPE, and optimize the sub-sampling pattern.

Figure 3.1 shows the optimized sub-sampling mask that was computed by LOUPE on T1-weighted brain MRI scans from 100 training subjects. The resulting mask has similarities with to the Variable Density mask. While it does not include a calibration region, it exhibits a denser sampling pattern closer to the origin of k-space. However, at high frequency values, the relative density of the optimized mask is much smaller than the Variable Density mask.

Figure 3.2 shows the optimized sub-sampling mask that was computed by LOUPE on PD and PDFS weighted images of knee MRI scans. The mask depicts a non-symmetrical structure, compared to the optimized mask for the brain data set. Although, the mask is comparable to Variable Density in terms of emphasizing the lower frequencies, it is apparent that the horizontal frequencies are favored significantly more than the vertical frequencies. This apparent symmetry difference between the optimized masks for two distinct data sets shows the sensitivity of the model to the anatomical differences and supports the idea of optimal masks being data driven. Figure 4.1 shows a side by side compar-



**Fig. 4.1.** Optimized masks compared side to side for the knee anatomy and brain anatomy for sub-sampling rate of  $R = 8$  and  $R = 10$  respectively.

ison of the two optimized masks for the knee and brain anatomy respectively. The masks share the behavior of sudden drop of density from lower to higher frequencies, however differ significantly as the knee mask favors the horizontal frequencies more due to knee anatomy and tissue being more continuous in vertical alignment and having the majority of the contrast change in the horizontal alignment.

Figure 4.2 includes box plots for subject-level PSNR values of reconstructions obtained with four reconstruction methods, four different masks, and two sub-sampling rates. The Cartesian and Uniform masks overall yielded worse reconstructions than the Variable Density and Optimized masks. In all except a single scenario, the Optimized mask significantly outperformed other masks (FDR corrected  $q < 0.01$  on paired t-tests). The only case where the Optimized mask was not the best performer was for the 10% sub-sampling rate, coupled with the BM3D reconstruction method [7]. Here, the PSNR values were slightly worse than the best-performing mask, that of Variable Density.

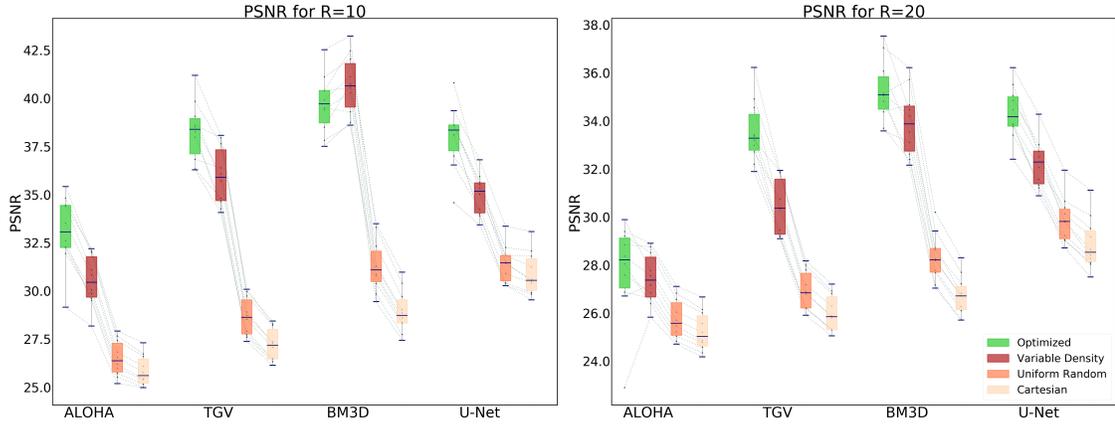
Figure 4.3 is also structured similar to the figure 4.2. The box-plots show that the U-Net configuration with the optimized mask yields the best PSNR values,

**Table 4.1.** Average per volume run times (in sec) for different reconstruction methods. All except U-Net (GPU) were evaluated on a CPU - a dual Intel Xeon (E5-2640, 2.4GHz).

ALOHA [18]	TGV [10]	BM3D [7]	U-Net [19] (CPU)	U-Net [19] (GPU)
$498 \pm 43.9$	$492 \pm 33.8$	$1691.1 \pm 216.4$	$55.9 \pm 0.3$	$1.6 \pm 0.4$

in both sub-sampling rates compared to other mask configurations and reconstruction methods, for each subject. Variable density has the second best reconstruction quality after the optimized mask, similar to the results for the ABIDE-I data set. The PSNR results within each reconstruction method and mask configuration show a division between the data set. The test data set comprises of 6 subjects with PD weighted images and 4 subjects with PDFS weighted images. The fat suppression operation inherently lowers the signal level, as fat has the highest levels of signal in an MRI scans, and yields a noisy image, where small details are more apparent. Whereas, the PD weighted scans that still have the fat tissue has inherently higher SNR. Therefore, the PD scans yield higher PSNR values in reconstruction compared to the PDFS scans. Overall, the optimized mask yields the highest PSNR in each reconstruction method compared to other mask configurations both for PD and PDFS scans, while the difference in PSNR for PD scans are more apparent.

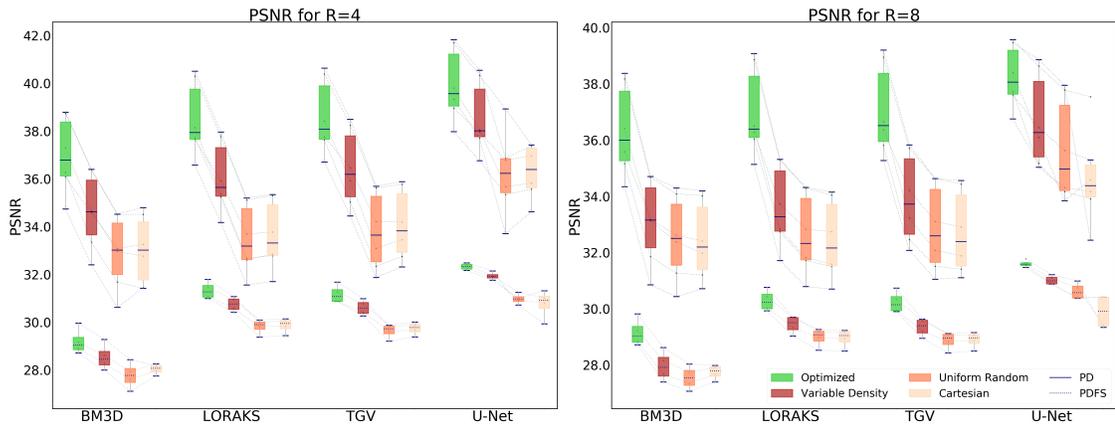
While the quantitative results give us a sense of overall quality, we found it very informative to visually inspect the reconstructions. Figures 4.4 and 4.5 show typical examples of reconstructed images for the ABIDE-I data set. We observe that our optimized mask yielded reconstructions that capture much more anatomical detail than what competing masks yielded (highlighted with red arrows in the pictures). In particular, the cortical folding pattern and the boundary of the putamen – a subcortical structure – were much better discernible for our



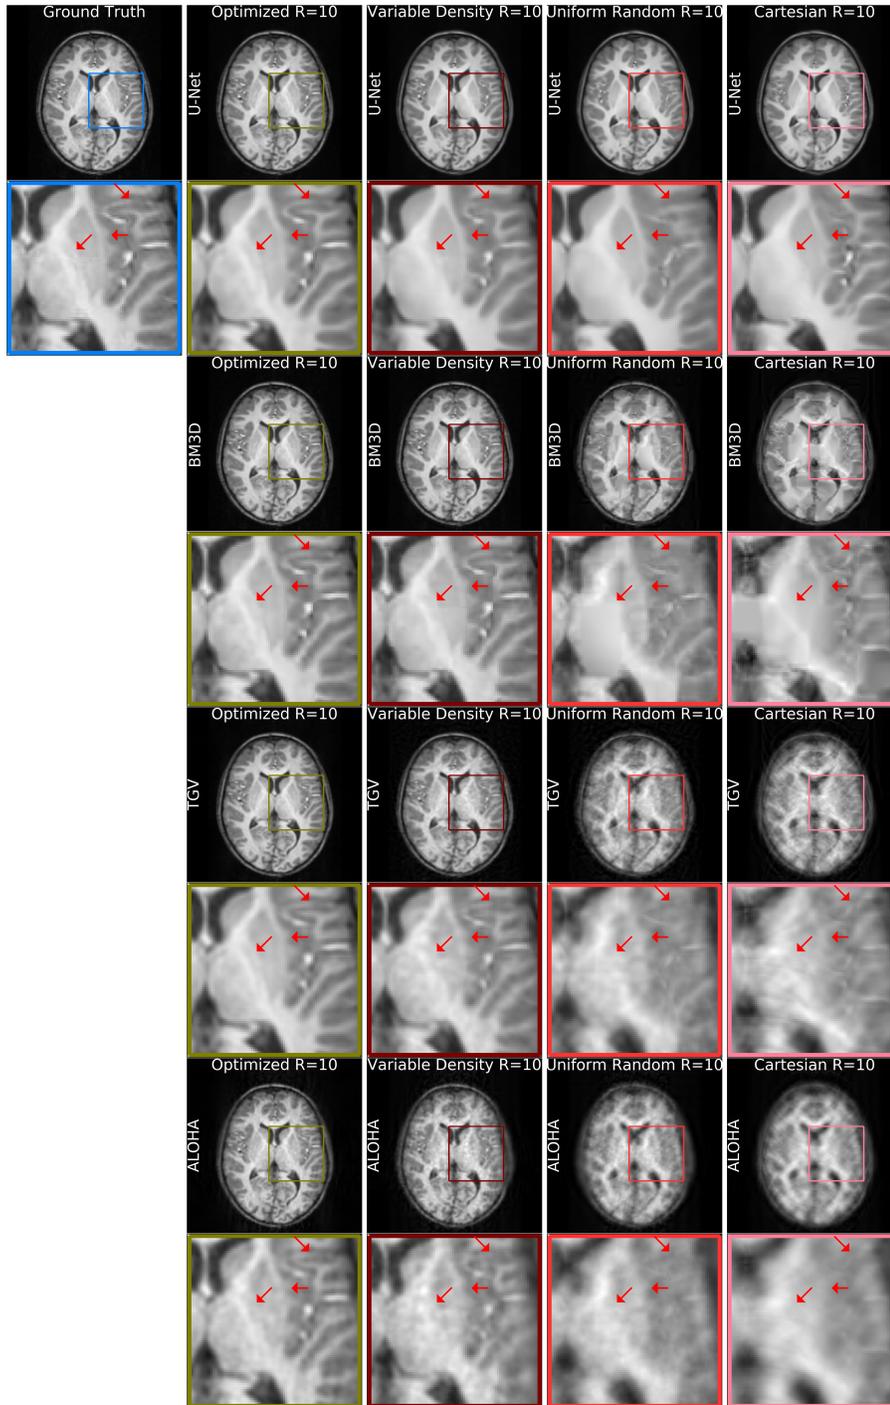
**Fig. 4.2.** Quantitative evaluation of reconstruction quality for ABIDE-I data set. For each plot, we show four reconstruction methods using four acquisition masks, including the Optimized Mask obtained using LOUPE in green. Each dot is the PSNR value for a single test subject across slices. For each box, the blue line shows the median value, and the whiskers indicate the the most extreme (non-outlier) data points.

optimized mask. The difference in reconstruction quality between the different methods can also be appreciated. Overall, U-Net and BM3D offer more faithful reconstructions that can be recognized in the zoomed-in views.

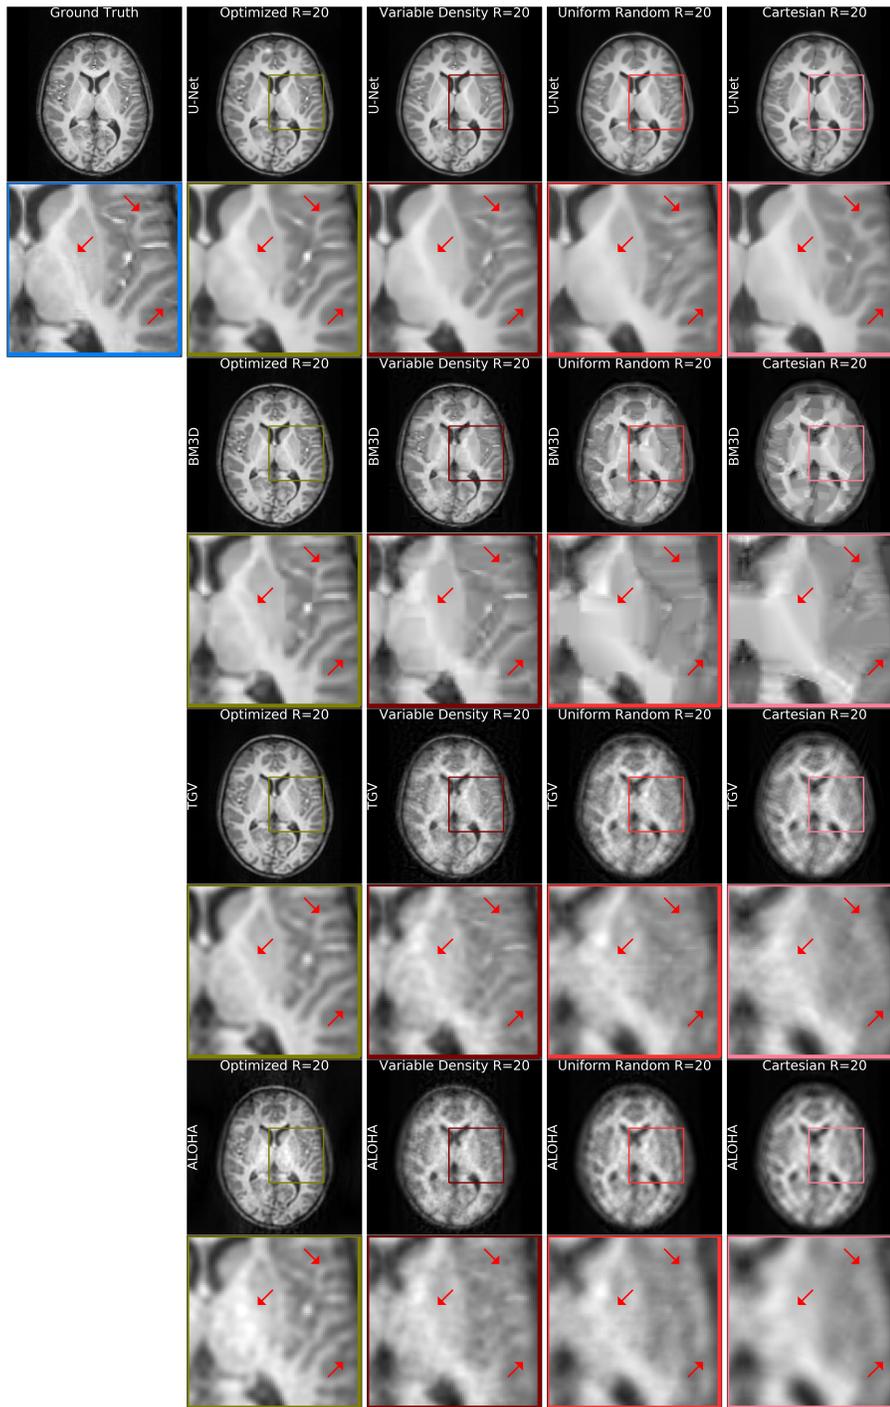
Furthermore, the superiority of the reconstruction quality and faithfulness to the details was also apparent in the visual inspections on the NYU fastMRI data-set. Figures 4.6 and 4.7 depict a representative PD weighted image of a subject from the test data set. The optimized mask shows robustness against blurring effects and phantom artifacts compared to benchmark mask configurations in each of the reconstruction method. The U-Net method with the optimized mask yields the best reconstruction results overall and the method appears less prone to artifacts even for other benchmark mask configurations.



**Fig. 4.3.** Quantitative evaluation of reconstruction quality for the NYU fastMRI data set. For each plot, we show four reconstruction methods using four acquisition masks, including the Optimized Mask obtained using LOUPE in green. Each dot is the PSNR value for a single test subject across slices. For each box, the blue straight line shows the median value for patients with Proton Density images while the blue dashed line shows the median value for patients with Proton Density Fat Suppressed images. The whiskers indicate the the most extreme (non-outlier) data points..



**Fig. 4.4.** Reconstructions for a representative slice from ABIDE-I experiments with  $R = 10$  sub-sampling rate. Each row is a reconstruction method. Each column corresponds to a sub-sampling mask. We observe that our optimized mask yields reconstructions that capture more anatomical detail. Red arrows highlight some nuanced features that were often missed in reconstructions.



**Fig. 4.5.** Reconstructions for a representative slice from ABIDE-I experiments with  $R = 20$  sub-sampling rate. See caption of Figure 4.4 and text for more detail.



**Fig. 4.6.** Reconstructions for a representative slice from NYU fastMRI experiments with  $R = 4$  sub-sampling rate. Each row is a reconstruction method. Each column corresponds to a sub-sampling mask. We observe that our optimized mask yields reconstructions that capture more anatomical detail.



**Fig. 4.7.** Reconstructions for a representative slice from NYU fastMRI experiments with  $R = 8$  sub-sampling rate. See caption of Figure 4.6 and text for more detail.

## CHAPTER 5

### DISCUSSION

We presented a novel problem formulation and a learning-based approach to simultaneously optimize the sub-sampling pattern and reconstruction model. Our experiments on retrospectively under-sampled brain and knee MRI scans suggest that our optimized mask can yield reconstructions that are of higher quality than those computed from other widely-used under-sampling masks. The model also provides undeniable proof to the optimal mask being data driven due to apparent difference in the shapes of optimized masks for two distinct anatomies.

This model creates an opportunity and a strong base for future directions that we would like to explore. First, sampling associated cost is captured with an L1 penalty in our formulation. We are interested in incorporating hardware limitations into our cost function to better capture true cost of a k-space trajectory. Second, LOUPE we used L1 norm for reconstruction loss, which is one of the many evaluation metrics used for determining the reconstruction quality. This can also be replaced with alternate metrics, such as those based on adversarial learning or emphasizing subtle yet important anatomical details and/or pathology. Third, we will consider combining LOUPE with a multi-coil parallel imaging approach to obtain even higher levels of acceleration. Fourth, we plan to explore optimizing sub-sampling patterns for other MRI sequences and organ domains on top of the brain and the knee anatomies. Overall, we believe that the proposed framework can be used in other compressed sensing applications beyond medical imaging and further extend to communication and other signal processing related applications.

## BIBLIOGRAPHY

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016.
- [2] Cagla Deniz Bahadir, Adrian V Dalca, and Mert R Sabuncu. Learning-based optimization of the under-sampling pattern in mri. *arXiv preprint arXiv:1901.01960*, 2019.
- [3] Luca Baldassarre, Yen-Huan Li, Jonathan Scarlett, Baran Gözcü, Ilija Bogunovic, and Volkan Cevher. Learning-based compressive subsampling. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):809–822, 2016.
- [4] François Chollet et al. Keras, 2015.
- [5] Adrian V Dalca, John Guttag, and Mert R Sabuncu. Anatomical priors in convolutional networks for unsupervised biomedical segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9290–9299, 2018.
- [6] Adriana Di Martino, Chao-Gan Yan, Qingyang Li, Erin Denio, Francisco X Castellanos, Kaat Alaerts, Jeffrey S Anderson, Michal Assaf, Susan Y Bookheimer, Mirella Dapretto, et al. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry*, 19(6):659, 2014.
- [7] Ender M Eksioğlu. Decoupled algorithm for mri reconstruction using non-local block matching model: Bm3d-mri. *Journal of Mathematical Imaging and Vision*, 56(3):430–440, 2016.
- [8] Urs Gamper, Peter Boesiger, and Sebastian Kozerke. Compressed sensing in dynamic mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 59(2):365–373, 2008.
- [9] Baran Gözcü, Rabeeh Karimi Mahabadi, Yen-Huan Li, Efe Ilıcak, Tolga Cukur, Jonathan Scarlett, and Volkan Cevher. Learning-based compressive mri. *IEEE transactions on medical imaging*, 37(6):1394–1406, 2018.

- [10] Weihong Guo, Jing Qin, and Wotao Yin. A new detail-preserving regularization scheme. *SIAM journal on imaging sciences*, 7(2):1309–1334, 2014.
- [11] Justin P Haldar. Low-rank modeling of local  $k$ -space neighborhoods (loraks) for constrained mri. *IEEE transactions on medical imaging*, 33(3):668–681, 2014.
- [12] Justin P Haldar, Diego Hernando, and Zhi-Pei Liang. Compressed-sensing mri with random encoding. *IEEE transactions on Medical Imaging*, 30(4):893–903, 2011.
- [13] Justin P Haldar and Daeun Kim. Oedipus: An experiment design framework for sparsity-constrained mri. *IEEE transactions on medical imaging*, 2019.
- [14] Justin P Haldar and Jingwei Zhuo. P-loraks: Low-rank modeling of local  $k$ -space neighborhoods with parallel imaging data. *Magnetic resonance in medicine*, 75(4):1499–1514, 2016.
- [15] Yue Huang, John Paisley, Qin Lin, Xinghao Ding, Xueyang Fu, and Xiaoping Zhang. Bayesian nonparametric dictionary learning for compressed sensing mri. *IEEE Transactions on Image Processing*, 23(12):5007–5019, 2014.
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [18] Dongwook Lee, Kyong Hwan Jin, Eung Yeop Kim, Sung-Hong Park, and Jong Chul Ye. Acceleration of mr parameter mapping using annihilating filter-based low rank hankel matrix (aloha). *Magnetic resonance in medicine*, 76(6):1848–1864, 2016.
- [19] 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.
- [20] Duan-duan Liu, Dong Liang, Xin Liu, and Yuan-ting Zhang. Under-sampling trajectory design for compressed sensing mri. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 73–76. IEEE, 2012.

- [21] Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72, 2008.
- [22] Shiqian Ma, Wotao Yin, Yin Zhang, and Amit 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. IEEE, 2008.
- [23] Rabeeh Karimi Mahabadi, Cosimo Aprile, and Volkan Cevher. Real-time dct learning-based reconstruction of neural signals. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 1925–1929. IEEE, 2018.
- [24] Rabeeh Karimi Mahabadi, Junhong Lin, and Volkan Cevher. A learning-based framework for quantized compressed sensing. *IEEE Signal Processing Letters*, 2019.
- [25] Morteza Mardani, Enhao Gong, Joseph Y Cheng, Shreyas Vasanawala, Greg Zaharchuk, Marcus Alley, Neil Thakur, Song Han, William Dally, John M Pauly, et al. Deep generative adversarial networks for compressed sensing automates mri. *arXiv preprint arXiv:1706.00051*, 2017.
- [26] Xiaobo Qu, Yingkun Hou, Fan Lam, Di Guo, Jianhui Zhong, and Zhong Chen. Magnetic resonance image reconstruction from undersampled measurements using a patch-based nonlocal operator. *Medical image analysis*, 18(6):843–856, 2014.
- [27] Tran Minh Quan, Thanh Nguyen-Duc, and Won-Ki Jeong. Compressed sensing mri reconstruction using a generative adversarial network with a cyclic loss. *IEEE transactions on medical imaging*, 37(6):1488–1497, 2018.
- [28] Saiprasad Ravishankar and Yoram Bresler. Mr image reconstruction from highly undersampled k-space data by dictionary learning. *IEEE transactions on medical imaging*, 30(5):1028–1041, 2011.
- [29] 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.
- [30] Matthias Seeger, Hannes Nickisch, Rolf Pohmann, and Bernhard Schölkopf. Optimization of k-space trajectories for compressed sensing

by bayesian experimental design. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 63(1):116–126, 2010.

- [31] Jian Sun, Huibin Li, Zongben Xu, et al. Deep adm-net for compressive sensing mri. In *Advances in neural information processing systems*, pages 10–18, 2016.
- [32] 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.
- [33] Zhongmin Wang and Gonzalo R Arce. Variable density compressed image sampling. *IEEE Transactions on image processing*, 19(1):264–270, 2010.
- [34] 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, 2018.
- [35] 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 Drozdal, 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. *arXiv preprint arXiv:1811.08839*, 2018.