



Published in final edited form as:

IEEE J Sel Top Signal Process. 2020 October ; 14(6): 1151–1162. doi:10.1109/jstsp.2020.3004094.

J-MoDL: Joint Model-Based Deep Learning for Optimized Sampling and Reconstruction

Hemant Kumar Aggarwal [Member, IEEE], Mathews Jacob [Senior Member, IEEE]

Department of Electrical and Computer Engineering, University of Iowa, IA, USA, 52242.

Abstract

Modern MRI schemes, which rely on compressed sensing or deep learning algorithms to recover MRI data from undersampled multichannel Fourier measurements, are widely used to reduce the scan time. The image quality of these approaches is heavily dependent on the sampling pattern. We introduce a continuous strategy to optimize the sampling pattern and the network parameters jointly. We use a multichannel forward model, consisting of a non-uniform Fourier transform with continuously defined sampling locations, to realize the data consistency block within a model-based deep learning image reconstruction scheme. This approach facilitates the joint and continuous optimization of the sampling pattern and the CNN parameters to improve image quality. We observe that the joint optimization of the sampling patterns and the reconstruction module significantly improves the performance of most deep learning reconstruction algorithms. The source code is available at <https://github.com/hkaggarwal/J-MoDL>.

Index Terms—

Experiment design; Sampling; Deep learning; Parallel MRI

I. Introduction

MR imaging offers several benefits, including good soft-tissue contrast, non-ionizing radiation, and the availability of multiple tissue contrasts. However, its main limitation is the slow image acquisition rate. The last decade has witnessed several approaches, including parallel MRI and compressed sensing, to recover the images from undersampled k-space measurements. Recently, deep learning methods are emerging as powerful algorithms for the reconstruction of undersampled k-space data; they offer significantly improved computational efficiency and higher image quality than classical methods. Several direct-inversion methods including [1]–[8] use a convolutional neural network (CNN) to recover the images from the undersampled data directly. Another family of methods pose the image recovery as an optimization problem involving a physics-based forward model and a deep-learned regularization prior [9]–[18]. These model-based methods can be thought of as learning based variants of earlier plug-and-play methods [19], [20], which used off-the-shelf denoisers as regularization penalties. In this work, we will focus on our implementation [13], which is termed as model based deep learning (MoDL). We refer the reader to [13] for the

details of MoDL, including its benefits over (a) direct-inversion based methods, (b) similar unrolled architectures and learned plug-and-play priors, (c) the use of conjugate gradients in contrast to steepest descent update to enforce data consistency, (d) as well as its ability to work with smaller CNN modules that allows it to learn from smaller datasets. The image quality offered by all of the above methods heavily depends on the sampling pattern. Early parallel MRI hardware [21] was designed to eliminate the need to sample adjacent k-space samples, making uniform undersampling of k-space a desirable approach. By contrast, compressed sensing [22], [23] advocates for the sampling pattern to be maximally incoherent. Since the k-space center is associated with high energy, variable density schemes that sample the center with a higher density are preferred by practitioners. Many of the current methods rely on the Poisson-disc variable density approach, which is a heuristic that combines the above intuitions [24]. Early empirical studies in the context of deep learning suggest that incoherent sampling patterns, which are widely used in compressed sensing, may not be necessary for good reconstruction performance in this context [14]. Computational methods were introduced as a systematic approach to design the sampling patterns for each setting.

The computational design of sampling patterns has a long history in MRI. Current solutions can be broadly classified as algorithm-dependent and algorithm-agnostic. The algorithm-agnostic approaches such as [25]–[29] consider specific image properties and optimize the sampling patterns to improve the measurement diversity for that class. Image properties, including image support [25], parallel acquisition using sensitivity encoding (SENSE) [26], [28], [30], and sparsity constraints [27], have been introduced. These experiment design strategies often rely on the Cramer-Rao (CR) bound, assuming the knowledge of the image support or location of the sparse coefficients. Algorithm-dependent schemes such as [31], [32] optimize the sampling pattern, assuming specific reconstruction algorithms (e.g., TV or wavelet sparsity). These approaches [31], [32] only consider single-channel settings with undersampled Fourier transform as a forward model. They utilize a subset of discrete sampling locations using greedy or continuous optimization strategies to minimize the reconstruction error. The main challenge with the above computational approaches is the significantly high computational complexity. The main contributor to the complexity is the evaluation of the loss associated with a specific sampling pattern. For instance, algorithm-dependent schemes need to solve the compressed sensing problem for each image in the dataset to evaluate the loss for a specific sampling pattern. The design of sampling pattern thus involves a nested optimization strategy; the optimization of the sampling patterns is performed in an outer loop, while image recovery is performed in the inner loop to evaluate the cost associated with the sampling pattern. The use of deep learning methods for image reconstruction offers an opportunity to speed up the computational design. Specifically, deep learning inference schemes enables the fast evaluation of the loss associated with each sampling pattern. In addition, these methods also facilitates the evaluation of the gradients of the cost with respect the sampling pattern. Unlike classical methods that rely on specific image properties (e.g. sparsity, support-constraints), the non-linear convolutional neural networks (CNN) schemes exploit complex non-linear redundancies that exist in images. This makes it difficult to use the algorithm-agnostic computational optimization algorithms discussed above in this setting. In addition, these learning-based methods often learn

representations that may be strongly coupled to the specific sampling scheme. A joint strategy, which simultaneously optimizes for the acquisition scheme as well as the reconstruction algorithm, is necessary to obtain the best performance.

Most of the current sampling pattern optimization schemes for deep learning relies on a binary sampling mask [33]–[35], which chooses a subset of the Cartesian sampling pattern. For instance, the recent LOUPE algorithm [35] jointly optimizes the sampling density in k-space and the reconstruction algorithm. It assumes each binary sampling location to be an independent random variable. The independence assumption makes it difficult of LOUPE to account for dependencies between sampling locations. We note that the popular Poisson disc sampling strategy [24] assumes the sampling locations to be separated by a minimum distance [36], in addition to following a density. This separation is vital for exploiting the redundancies resulting from multichannel sampling with smooth coil sensitivities as described in [21]. The PILOT approach [33], [34] instead relies on a relaxation of the binary mask to make the cost differentiable. A challenge with this scheme is the large number of trainable parameters, which often translate to convergence issues [34]. In our own settings, a non-parametric strategy that aimed to optimize for all the sampling locations failed to converge, especially when large training datasets are not used. We note that another class of deep learning solutions involve active strategies [37], [38], where a neural network is used to predict the next k-space sample to be acquired based on the image reconstructed from the current samples. We do not focus on such active paradigms in this work. We also note that similar work involving the optimization of the forward model have been also explored in the context of optical imaging [39]–[43].

The main focus of this work is to jointly optimize the sampling pattern and the deep network parameters for parallel MRI reconstruction. We rely on an algorithm-dependent strategy to search for the best sampling pattern. The main contributions of this work are

1. Unlike previous methods [31]–[33], [35] that constrain the sampling pattern to be a subset of the Cartesian sampling pattern, we assume the sampling locations to be continuous variables. The earlier methods [31]–[33], [35] rely on relaxations or approximations of the discrete mask to make the cost function differentiable. The proposed scheme does not need any approximations since the derivatives with respect to the sampling locations are well-defined.
2. Unlike [34], we solve for the sampling pattern rather than the sampling density. Hence, our approach can account for complex dependencies between k-space sampling locations, which may be difficult for a density-based approach.
3. Unlike the previous optimization strategies [31]–[35] that were only restricted to the single-channel setting, we extend the scheme to the multichannel setting where there is the most gain.
4. We introduce a parametric representation of the sampling patterns to reduce the degrees of freedom of the sampling pattern. The reduced search space improves the ability to learn the sampling pattern even from smaller datasets.

The main objective of the proposed work is to optimize the sampling pattern for a specific anatomy (e.g. knee, brain) and protocol, rather than optimizing it for each subject. We note that the earlier optimization strategies in MRI are also designed for similar settings [34], [44]. Our experiments show that most of the deep learning algorithms significantly benefit from sampling pattern optimization, which is a relatively under-explored area compared to reconstruction network architecture and training. Our experiments involving the fastMRI knee dataset [45], acquired from multiple sites and scanners, demonstrate the robustness of the approach.

II. Method

A. Image Formation

We consider the recovery of the complex image $\rho \in \mathbb{C}^{M \times N}$ from its non-Cartesian Fourier samples:

$$b[i, j] = \sum_{\mathbf{m} \in \mathbb{Z}^2} s_j[\mathbf{m}] \rho[\mathbf{m}] e^{-j\mathbf{k}_i^T \mathbf{m}} + n[i, j], \mathbf{k}_i \in \Theta. \quad (1)$$

Here, Θ is a set of sampling locations and $n[i, j]$ is the noise process. $s_j, j = 1, \dots, J$ corresponds to the sensitivity of the j^{th} coil, while \mathbf{k}_i is the i^{th} sampling location. The above mapping can be compactly represented as $\mathbf{b} = \mathcal{A}_\Theta(\rho) + \mathbf{n}$. The measurement operator \mathcal{A}_Θ is often termed to as the forward model. It captures the information about the sampling pattern as well as the receive coil sensitivities. We note that the forward model is often modified to include additional information about the imaging physics, including field inhomogeneity distortions and relaxation effects [46].

B. Regularized Image recovery

Model-based algorithms are widely used for the recovery of images from heavily undersampled measurements, such as (1). These schemes pose the reconstruction as an optimization problem of the form

$$\hat{\rho}_{\{\Theta, \Phi\}} = \arg \min_{\rho} \|\mathbf{b} - \mathcal{A}_\Theta(\rho)\|_2^2 + \mathcal{R}_\Phi(\rho). \quad (2)$$

Here, \mathcal{R}_Φ is a regularization penalty. Regularizers include transform domain sparsity [47], total variation regularization [48], and structured low-rank methods [49]. For instance, in transform domain sparsity, the regularizer is chosen as $\mathcal{R}(\rho) = \lambda \|\mathbf{T}\rho\|_{\ell_1}$, with $\Phi = \{\lambda, \mathbf{T}\}$ denoting the parameters of the regularizer and the transform. We rely on the notation $\hat{\rho}_{\{\Theta, \Phi\}}$ for the solution of (2) to denote its dependence on the regularization parameters as well as the sampling pattern.

C. Deep learning based image recovery

Deep learning methods are increasingly being investigated as alternatives for regularized image reconstruction. Instead of algorithms that rely on the hand-crafted priors discussed

above, these schemes learn the parameters from exemplar data. Hence, these schemes are often termed as data-driven methods.

1) Direct-inversion schemes: Direct-inversion approaches [1], [2] rely on a deep CNN \mathcal{N}_Φ to recover the images from undersampled gridding reconstruction $\mathcal{A}_\Theta^H(\mathbf{b})$ as

$$\rho_{\text{direct}} = \mathcal{N}_\Phi(\mathcal{A}_\Theta^H \mathbf{b}). \quad (3)$$

Here Φ denotes the learnable parameters of the CNN \mathcal{N}_Φ (see Fig. 1(a)).

2) Model-based deep learning: Several unrolled approaches, which combine physics-based priors with learned priors, have been introduced for image recovery [9]–[18]. In this paper, we will focus on the model-based deep learning (MoDL) [13] framework, where image recovery is formulated as

$$\hat{\rho}_{\{\Theta, \Phi\}} = \arg \min_{\rho} \|\mathbf{b} - \mathcal{A}_\Theta(\rho)\|_2^2 + \|\rho - \mathcal{D}_\Phi(\rho)\|_F^2, \quad (4)$$

where \mathcal{D}_Φ is a residual learning-based CNN that is designed to extract the noise and alias terms in ρ . The optimization problem specified by (4) is solved using an iterative algorithm, which alternates between a denoising step and a data consistency step:

$$\rho_{n+1} = (\mathcal{A}_\Theta^H \mathcal{A}_\Theta + \mathcal{F})^{-1} (\mathbf{z}_n + \mathcal{A}_\Theta^H \mathbf{b}) \quad (5)$$

$$\mathbf{z}_{n+1} = \mathcal{D}_\Phi(\rho_{n+1}). \quad (6)$$

Here, (5) is implemented using a conjugate gradient algorithm. This iterative algorithm is unrolled to obtain a deep recursive network $\mathcal{M}_{\Theta, \Phi}$, where the weights of the CNN blocks and data consistency blocks are shared across iterations, as shown in Fig. 1(b). Specifically, the solution to (4) is given by

$$\hat{\rho}_{\{\Theta, \Phi\}} = \mathcal{M}_{\Theta, \Phi}(\mathcal{A}_\Theta(\rho)). \quad (7)$$

Note that once unrolled, the image reconstruction algorithm is essentially a deep network, shown in Fig. 1(b). Thus, the main distinction between MoDL and direct-inversion scheme is the structure of the network $\mathcal{M}_{\Theta, \Phi}$. Please see [13] for details.

D. Optimization of sampling patterns and hyperparameters

The focus of this work is to optimize the sampling pattern specified by Θ in (1) and the parameters Φ of the reconstruction algorithm (2) to improve the quality of the reconstructed images. Conceptually, the regularization priors encourage the solution to be restricted to a family of feasible images (e.g., sparse wavelet representation). The objective is to optimize the sampling pattern to capture information that is maximally complementary to the image representation.

Early approaches that rely on compressed sensing algorithms [31], [32] optimize the sampling pattern Θ such that

$$\{\Theta^*\} = \arg \min_{\Theta} \sum_{i=1}^N \|\hat{\rho}_{i, \{\Theta, \Phi\}} - \rho_i\|_2^2, \quad (8)$$

is minimized. Here ρ_i ; $i = 1, \dots, N$ are the different training images used in the optimization process and $\hat{\rho}_{i, \{\Theta, \Phi\}}$ are the corresponding reconstructed images, recovered using (2).

Greedy [32] or continuous optimization schemes [31] are used to solve (8). However, the main challenge associated with these schemes is the high complexity of the optimization algorithm used to solve (2). Note that the optimization scheme (2) is in the inner loop; for each sampling pattern, the N images have to be reconstructed using computationally expensive CS methods to compute the loss in (8). This makes it challenging to train the pattern using a large batch of training images. In addition, the hyperparameters of the algorithm denoted by Φ are assumed to be fixed during this optimization.

Recent schemes such as LOUPE [35] and PILOT [34] exploit the fast deep-learned reconstruction algorithms to optimize for the sampling pattern. Instead of directly solving for the k -space locations, the LOUPE approach optimizes for the sampling density [35]. Specifically, they assume the k -space sampling locations that are acquired to be binary random variables $\mathbf{k}_i \sim \mathcal{B}(p_i)$ and optimize for the probabilities p_i . They rely on several random realizations of \mathbf{k}_i and the corresponding reconstructions to perform the optimization.

E. Proposed Joint Optimization Strategy

This work proposes a joint model-based deep learning (J-MoDL) framework to jointly optimize both the \mathcal{D}_{Φ} and \mathcal{Q}_{Θ} blocks in the MoDL framework (4) with the goal of improving the reconstruction performance. Specifically, we propose to jointly learn the sampling pattern Θ and the CNN parameters Φ from training data using

$$\{\Theta^*, \Phi^*\} = \arg \min_{\Theta, \Phi} \sum_{i=1}^N \|\mathcal{M}_{\Theta, \Phi}(\mathcal{A}_{\Theta}(\rho_i)) - \mathbf{x}_i\|_2^2. \quad (9)$$

We note that $\mathcal{M}_{\Theta, \Phi}$ denotes a general deep learning network architecture that includes direct-inversion schemes denoted by (3) as well as unrolled architectures denoted by (7).

While the proposed J-MoDL framework (in Fig. 1(b)) can be generalized to other error metrics such as perceptual error, we focus on the ℓ_2 error in this work.

F. Forward model and parametrization of the sampling pattern

We represent the forward model as

$$\mathbf{b}_i = \mathcal{F}_{\Theta}(\mathbf{s}_i \cdot \rho); \quad i = 1, \dots, N_c, \quad (10)$$

where $s_i; i = 1, \dots, N_C$ denotes the coil sensitivities of the i^{th} channel to compactly represent (1). Here, $\mathcal{F}_{\Theta}(\rho)$ denotes the Fourier transform of ρ evaluated at the continuous sampling locations \mathbf{k}_i , whose set is denoted by Θ .

We found it challenging to directly optimize for the large number of trainable sampling locations which would require a huge amount of training data. Hence, we propose to reduce the dimension of the search space by a parametrization of the sampling pattern as illustrated in Fig. 2 for 1-D and 2-D case. Specifically, we assume that the sampling pattern to be the union of transformed versions of a template set Γ :

$$\Theta = \bigcup_{i=1}^P \mathcal{T}_{\theta_i}(\Gamma). \quad (11)$$

Here, \mathcal{T}_{θ_i} is a transformation that is dependent on the trainable parameters θ_i . For example, one may consider the optimization of the phase encoding locations in MRI, while the frequency encoding direction is fully sampled. Specifically, we choose Γ as samples on a line and \mathcal{T}_{θ_i} are translations orthogonal to the line. Here, $\theta_i; i = 1, \dots, P$ are the phase encoding locations. In the 2-D setting, we also consider sampling patterns of the form

$$\Theta = \Theta_v \cap \Theta_h, \quad (12)$$

where Θ_v and Θ_h are 1-D sampling patterns in the vertical and horizontal directions, respectively. Here, we assume that the readout direction is orthogonal to the scan plane and is fully sampled. An example sampling pattern in this setting is shown in Fig. 10. Specifically, the locations $k_{x_i}; i = 1, \dots, h$ and $k_{y_i}; i = 1, \dots, v$ are the unknowns that the algorithm optimizes for. We note that this approach reduces the number of trainable parameters from hv to $h + v$.

In addition to reducing the parameter space, the above approaches also simplifies the implementation. We focus on this setting because the forward model in (10) can be implemented in terms of the 1-D Fourier transforms as

$$\mathbf{B} = \mathbf{F}_h \mathbf{X} \mathbf{F}_v^H. \quad (13)$$

Here, \mathbf{X} is the 2-D image and F_h and F_v are 1-D discrete Fourier transform operators as described in Fig. 2(b). By eliminating the need for non-uniform Fourier transform (NUFT) operators, this approach accelerates the training and inference.

G. Architecture of the networks used in joint optimization

Figure 1(b) shows the proposed J-MoDL framework. The framework alternates between data consistency blocks \mathcal{Q}_{Θ} , that depend only on the sampling pattern, and the CNN blocks \mathcal{D}_{Φ} . We unrolled the MoDL algorithm in Fig. 1(b) for $K=5$ iterations (i.e., five iterations of alternating minimization) to solve Eq. (4). The forward operator \mathcal{A}_{Θ} is implemented as a 1-D discrete Fourier transform to map the spatial locations to the continuous domain Fourier samples specified by Θ , following the weighting by the coil sensitivities, as described by (1)

and Fig. 2. The data consistency block \mathcal{Q}_Θ is implemented using 10 iterations of the conjugate gradient algorithm. The CNN block \mathcal{D}_Φ is implemented as a UNET with four pooling and unpooling layers with 3×3 trainable filters as in the UNET model [50]. The parameters of the blocks \mathcal{D}_Φ and \mathcal{Q}_Θ are optimized to minimize (9). We relied on the automatic differentiation capability of TensorFlow to evaluate the gradient of the cost function with respect to Θ and Φ .

We also study the optimization of the sampling pattern in the context of direct-inversion (i.e., when a UNET is used for image inversion). A UNET with the same number of parameters as the MoDL network considered above was used to facilitate fair comparison. This optimization scheme, where both sampling parameters and the UNET parameters are learned jointly, is termed as J-UNET (Fig. 1(a)). Since MR images are inherently complex valued, all the networks were trained using complex k-space data as input and the training loss was calculated on the complex images. The complex data was split into real and imaginary parts, which were fed into the neural networks. The data consistency steps explicitly worked with the complex data type.

H. Proposed continuous optimization training strategy

We first consider a collection of variable density random sampling patterns with 4% fully sampled locations in the center of the k-space, and train only the network parameters Φ . This training strategy is referred to as Φ -alone optimization. Once this training is completed, we fixed the trained network parameters and optimize for the sampling locations alone. Specifically, we consider the sampling operator \mathcal{A}_Θ and its adjoint as layers of the corresponding networks. The parameters of these layers are the location of the samples, denoted by Θ . We optimize for the parameters using stochastic gradient descent, starting with random initialization of the sampling locations Θ . The gradients of the variables are evaluated using the automatic differentiation capability of TensorFlow. This strategy, where only the sampling patterns are optimized, is referred to as the Θ -alone optimization; the parameters of the network derived from the Φ -alone optimization are held constant. The third strategy, we refer as Θ , Φ -Joint or just Joint, simultaneously optimizes for both, the sampling parameter Θ as well as the network parameters Φ . The Φ -alone optimization strategy take 5.5 hours to train in single-channel settings as described in section III-B. The Θ -alone and Θ , Φ -joint strategies only take 1 hour to train with an initialization from Φ -alone model.

III. Experiments and Results

A. Datasets

We relied on three datasets for comparison.

1) Single-channel knee data from fastMRI database: We used the data from the NYU fastMRI Initiative database [45] (fastmri.med.nyu.edu) in this section. As such, NYU fastMRI investigators provided data, but did not participate in analysis or writing of this article. The primary goal of fastMRI is to test whether machine learning can aid in the reconstruction of medical images. We relied on a PCA-based complex combination of the

multichannel images from the database to obtain single-coil images. The k-space data of these images, computed using the forward model in (1) with $J=1$ and $s_1(x) = 1$, are the input to the networks, while the corresponding complex images are used as the ground truth for training. We chose three subsets of the fastMRI dataset, consists of 100 training, 50 validation, and 100 test subjects. Unlike the other datasets considered in this work, this data was acquired on multiple scanners at different institutions, thus exhibiting significant diversity in the measurement settings. This dataset thus enables the evaluation of the scheme in a multi-site setting.

2) Multichannel knee dataset: We used a publicly available parallel MRI knee dataset as in [14]. The training data constituted of 381 slices from ten subjects, whereas test data had 80 slices from two subjects. Each slice in the training and test dataset had different coil sensitivity maps that were estimated using the ESPIRIT [51] algorithm. Since the data was acquired by using a 2-D Cartesian sampling scheme, we relied on a 1-D undersampling of this data.

3) Multichannel brain dataset: We consider a parallel MRI brain data using a 3-D T2 CUBE sequence with Cartesian readouts using a 12-channel head coil at the University of Iowa on a 3T GE MR750w scanner. The data was acquired according to the approved IRB protocol. Written consent was obtained from all subjects prior to the scan. The matrix dimensions were $256 \times 232 \times 208$ with a 1 mm isotropic resolution. Fully sampled multichannel brain images of nine volunteers were collected, out of which data from five subjects were used for training, while the data from two subjects were used for testing and the remaining two for validation. Since the data was acquired with a 3-D sequence, we used this data to determine the utility of 1-D and 2-D sampling in parallel MRI settings. Specifically, we performed a 1-D inverse Fourier transform along the readout direction and considered the recovery of each slice in the volume. Since the undersampling was performed on the phase encoding directions, these simulation studies are realistic. Following the image formation model in (1), additive white Gaussian noise of standard deviation $\sigma = 0.01$ was added in k-space in all the experiments.

B. Single-Channel Results

We first consider the single-channel experiments using the fastMRI data, as described in Section III-A1. We note that almost all sampling pattern optimization schemes have considered the single-channel settings [31]–[35], where an undersampled Fourier sampling forward operator is considered. Unlike the discrete optimization schemes that rely on relaxations of a discrete sampling mask [31]–[35], we consider the optimization of the continuous values of the phase encoding locations k_1, \dots, k_M , as shown in Fig. 2(a). We consider an undersampling factor of four in this experiment.

Table I reports the average PSNR and SSIM values obtained on the test data from 100 subjects. The top row corresponds to the optimization of the network parameters Φ -alone, assuming the random variable density undersampling patterns with 4% fully sampled center of the k-space. Each training slice had a different sampling pattern, whereas during testing each subject had a different sampling pattern; all slices of a subject had same sampling

pattern. This approach made the network relatively insensitive to the specific sampling pattern, compared to the learning with a single pattern. We note that the higher complexity of the MoDL framework translated to an approximate 3.5 dB improvement in performance over a UNET scheme in the Φ -alone setting, even though both methods had the same number of parameters. This observation is in line with the experiments in [13]. The second row in Table I reports the result of only optimizing the sampling parameters Θ , while keeping the reconstruction network fixed as the one trained in the first row (Φ -alone). The last row of Table I corresponds to the joint optimization scheme, where both Θ and Φ are trained with the initial sampling pattern used in the top row. The resulting J-MoDL scheme offers a 3.13 dB improvement in performance over the case where only the network is trained. The J-MoDL scheme is also better by 1.32 dB compared to only optimizing the sampling pattern. By contrast, the J-UNET approach provided only a 1.05 dB improvement over the initialization. The results demonstrate the benefit of the decoupling of the sampling pattern and CNN parameters offered by MoDL.

The visual comparisons of these strategies are shown in Fig. 3. The proposed J-MoDL method provides significantly improved results over the MoDL scheme, as highlighted by the zoomed region. The red arrows clearly show that the proposed J-MoDL architecture preserves the high-frequency details better than the MoDL architecture. The optimization of the sampling patterns also improved the UNET performance.

C. Parallel Imaging (Multichannel) with 1-D sampling

Table II summarizes the results in the 1-D parallel MRI setting on knee images, described in Section III-A2. The first row denoted as Φ -alone in Table II corresponds to optimizing the network parameters alone without optimizing the sampling mask. Unlike the setting in Section III-B, the network was not trained with different sampling masks. We choose the sampling mask as a single pseudo-random pattern for all the slices. In the second row, denoted as Θ -alone optimization, only the sampling mask is optimized, while keeping the reconstruction parameters fixed to optimal values as derived in the first row. Here, the network parameters were initialized with the ones derived from the Φ -alone optimization. Unlike the trend in Table I, we observe that the performance of the UNET scheme dropped slightly, while the MoDL scheme that was trained with the same setting provided improved results. The last row compares joint optimization using direct-inversion and model-based techniques. We observe that both methods improved in this case. The J-MoDL provides around 7 dB improvement over Φ -alone in the 4x setting and 3.5 dB in the 6x settings.

Figure 4.(a) shows an example slice from the test dataset that illustrates the benefit of jointly optimizing both the sampling pattern and the network parameters (Fig. 4.(c)), compared to the network-alone in the model-based deep learning framework (Fig. 4.(b)). The zoomed image portion shows that joint learning using J-MoDL better preserves the soft tissues in the knee at the four-fold acceleration case in parallel MRI settings.

To understand the drop in performance of the UNET scheme during the Θ -alone optimization, we compare the optimization landscape of the two schemes (MoDL and UNET in Θ -alone settings) in Fig. 5. Since this is a large dimensional problem, we plot the variation in MSE with respect to two variables (sampling locations) at a time.

As described above, a single sampling pattern, shown in Fig. 5, was used to train UNET and MoDL architectures on the parallel imaging knee dataset. We then computed the loss of the networks for perturbations of the sampling locations around the sampling pattern shown in Fig. 5. Specifically, the trained models were used to reconstruct the test dataset, while two of the original sampling locations (denoted by the green and red lines in Fig. 5) are perturbed from their original values. The loss evaluated for each of the perturbations are plotted in (a)-(d) and (e)-(h), respectively. Specifically, (a)-(d) corresponds to perturbations around the green locations 6 and 15 from the high-frequency samples, while (e)-(h) correspond to the samples 135 and 167, closer to the k-space center. The losses of the networks are plotted in Fig. 5. The $n1$ and $n2$ axes on these four plots correspond to the sampling locations, while the vertical z-axis shows the mean squared error ($\text{MSE} \times 1000$) between the predicted and original test image. Each of the $n1$ and $n2$ axis were varied for 100 points around them, thus resulting in a total of 10,000 MSE evaluations on each of the plots. The plots show that the MoDL network exhibits a smoother cost landscape around its minimum, while the UNET, which was trained using the same settings and initialization, resulted in highly oscillatory landscape.

We note that the proposed sampling pattern optimization scheme relies on stochastic gradient descent. The high gradients resulting from the oscillatory landscape, as well as the randomness in the gradient updates, likely resulted in the UNET converging to a bad local minimum. As shown in Table I, this problem can be mitigated by sampling pattern augmentation. However, we note that the MoDL scheme does not require the network to be trained with multiple sampling patterns to have a smoother optimization landscape, which explains its improved performance in this setting.

D. Parallel Imaging (Multichannel) with 2-D sampling

Table III summarizes the comparison results in the multichannel setting with 2-D sampling patterns, as described by (12) on the brain data described in Section III-D. Both the direct-inversion based framework (UNET) and the model-based framework (MoDL) are compared in Table III at three different optimization strategies at 6x acceleration factor. As described in Section III-B, the Φ -alone network was trained with multiple sampling patterns to reduce its sensitivity to sampling patterns. The trends of the different methods continue to be the same as in Section III-B.

The improved performance offered by the optimization of the sampling pattern in 2-D parallel imaging settings can be appreciated from Fig. 6. The zoomed portion in Fig. 6 shows the cerebellum region in which all the fine features are reconstructed well by the proposed J-MoDL approach at 6x acceleration. The red arrows are pointing to a high-frequency feature that is not recovered by the joint learning in the direct-inversion framework (J-UNET). This feature is also not recovered by the fixed model-based deep learning framework without joint optimization (see Fig. 6(b) and (c)). Fig. 6 also shows that proposed method can reconstruct the phase of the MR images. The error maps in Fig. 6 shows that the proposed J-MoDL approach has the least error in reconstruction among competing methods.

E. Comparison with classical sampling patterns

Figure 7 shows the visual comparison of reconstruction quality obtained with three different sampling patterns at the same 10x acceleration in 2-D parallel MRI settings with the model-based deep learning framework. Figure 7(a) and (b) are showing pseudo-random and variable density (VD) masks, respectively, while Fig. 7(c) shows the 2-D mask learned using joint learning with J-MoDL. These masks result in gridding reconstructions, as shown in Fig. 7(d), (e), and (f). It can be observed from Fig. 7(f) that learned mask results in a gridding reconstruction with comparatively fewer artifacts. Figures 7(g) and 7(h) are the reconstructed images using the Φ -alone optimization, whereas Fig. 7(i) corresponds to the reconstruction using joint learning.

F. Impact of noise

We study the impact of noise on the learned optimal sampling pattern and the reconstruction performance in Fig. 8. We note that the data was already corrupted by noise. We further added complex Gaussian noise with different standard deviations to the 8x undersampled k-space measurements. The results show that as the noise standard deviation increases, the optimal sampling patterns get concentrated to the center of k-space. This is expected since the energy of the Fourier coefficients in the center of k-space is higher. As the standard deviation of the noise increases, the outer k-space regions become highly corrupted with noise and hence sampling them does not aid the reconstruction performance. As expected, the restriction of the sampling pattern to the center of k-space results in image blurring. It can be noted that during training with different noise levels no extra constraints were imposed to promote a low-frequency mask. This behavior can be attributed to the explicit data consistency step in the model-based deep learning framework. This experiment empirically shows that the proposed J-MoDL technique indeed conforms with classical model-based techniques while retaining the benefits of deep learning methods.

G. Convergence of the sampling pattern optimization scheme

We empirically study the convergence of the sampling pattern optimization schemes in Fig. 9. We consider three different initial pseudo-random sampling patterns, each with fully sampled center having 4% lines as initialization. In each experiment, training is performed for 50 epochs. Figure 9(a) shows the decay of training loss with the proposed J-MoDL scheme, while Fig. 9(b) correspond to the J-UNET scheme. We observe that despite the highly non-convex optimization scheme the J-MoDL network was able to converge to solutions with almost the same cost. We also note that the images reconstructed with the final network are similar, even though the sampling patterns are different. We observe that the convergence of the J-UNET scheme was relatively less smooth, likely because of the non-smooth optimization landscape. The image quality of the final reconstructions are also more variable in this case. By contrast, the J-UNET scheme exhibit more variability in the final results.

H. Comparison with other sample optimization schemes

We compare the proposed J-MoDL scheme against other sampling pattern optimization methods. In particular, we compare J-MoDL against LOUPE [35], which is a discrete

optimization strategy. We extended the original LOUPE algorithm to the multichannel setting to compare with the proposed scheme. We also study the proposed continuous sampling pattern optimization scheme for a range of network architectures, including ISTANet, UNET, and MoDL with one iteration. These three methods were trained with identical initialization of the sampling mask as shown in Fig. 10(a). See Fig. 10 for the visual comparison of reconstructed images and learned masks.

Table IV summarizes the quantitative comparative results at 8x acceleration in the multichannel settings. The columns Φ -alone denotes the network-alone optimization, when a variable density 2D pattern is used. The results of the joint optimization of the sampling pattern and network is shown in $\Theta \Phi$ Joint columns. Since the LOUPE implementation available from the authors cannot be run without joint optimization, the results for the Φ -alone case are not reported. We observe that the image quality of all methods improved significantly with the optimization of the sampling pattern.

We note that all of the architectures in the above study have roughly the same number of trainable parameters. The ISTANet and MoDL ($K = 5$) approaches repeat the UNET five times, and hence have higher computational complexity over the UNET network and MoDL ($K = 1$) network. We observe that MoDL ($K = 1$) and UNET differ mostly in the addition of a data consistency step at the end. The results show that this gives around 3 dB improvement in performance during the network-alone training. We observe that the MoDL ($K = 5$) network provided an additional 2dB improvement in performance over the one-iteration MoDL with a fixed sampling pattern, which is consistent with our earlier findings [13]. However, the performance improvement offered by this scheme with joint optimization is even more significant.

IV. Discussion and Conclusion

We introduced an approach for the joint optimization of the continuous sampling locations and the reconstruction network for parallel MRI reconstruction. Unlike past schemes, we consider a Fourier operator with continuously defined sampling locations, which facilitated the optimization of the sampling pattern without approximations. Our experiments show the benefit of the joint optimization strategy. We relied on a parametric sampling pattern with few parameters, which improved the convergence of the network with limited data. The experimental results demonstrate the significant benefits in the joint optimization of the sampling pattern in the proposed model-based framework.

We note that the continuous optimization problem is highly non-convex with potentially many local minima and global minima. Specifically, any permutation of the optimal sampling pattern would be associated with the minimal cost. We note that similar symmetries do exist in the weights of neural networks. Fortunately, the stochastic gradient descent scheme is able to provide good solutions, despite the challenges in optimization.

We note that the MoDL scheme relies on end-to-end training of the network parameters. This training approach is different from plug-and-play methods, where the network parameters are pre-trained. We refer the interested readers to [13], where the benefit of end-

to-end training over pre-training is demonstrated. Similarly, [13] also shows the benefit of using conjugate gradients in the data consistency blocks over steepest descent updates as in ISTANet. Further, a detailed comparisons between direct-inversion schemes and model-based schemes are covered in [13]. We omit such comparisons in this work for brevity. We note that the proposed sampling pattern optimization framework can also be utilized along with GAN-based reconstruction networks such as [3], [4].

In this work, we used 10 iterations of the conjugate gradient in the data consistency step. Both the UNET and the J-MoDL unrolled for 10 iterations are trained on a 12 GB TitanV or any similar graphics card. The offline training time for MoDL is almost 5 times longer than that of basic UNET. During inference, the basic UNET reconstructs 110 slices per second whereas MoDL reconstructs only 18 slices per second. We note that the UNET is around six times faster than the MoDL framework. However, we believe that the MoDL scheme is considerably faster than compressed sensing methods, and the improved image quality justifies its use in many applications.

We note that the MSE between the final reconstructions and original images was chosen as the loss function for both MoDL and UNET. However, we note that the final images in MoDL are obtained as the minimization of the cost function in Eq. (4). Thus, one may view the MoDL training of the network parameters as consisting of two loss terms, one corresponding to the comparison in the image domain, and one corresponding to comparison with the measured noisy samples. Since this training is more fine-tuned to the measurement process, the optimization is expected to yield improved results than the UNET approach, which is confirmed by our experimental findings. We have reported the standard deviation across slices in our multichannel experiments. We understand that this might be an underestimate since the slices across a single subject may be correlated. A larger study involving more testing subjects will be needed to address this issue

As discussed previously, the joint optimization scheme with a 2-D non-parametric sampling pattern did not converge, possibly due to limited training data. In our future work, we will study the possibility of 2-D non-parametric sampling with more training data. In this work, we constrained the sampling pattern as the tensor product of two 1-D sampling patterns and optimized for the encoding locations. This approach reduced the number of trainable parameters, thus significantly improving the convergence of the algorithm over non-parametric strategies. We note that several alternate approaches may be used to achieve similar goals. For instance, one may search and pick a variable density pattern that yields the best performance from several randomly selected variable patterns. However, since the network has to be trained for each pattern, the anticipated training time is expected to be high. We note that constraining the sampling pattern as the tensor product is a limitation of this work. Optimizing the parameters of a truly 2-D parametric sampling pattern such as [52] may improve the performance. We will explore these ideas systematically in our future work.

Acknowledgments


This work is supported by 1R01EB019961-01A1. This work was conducted on an MRI instrument funded by 1S10OD025025-01

References

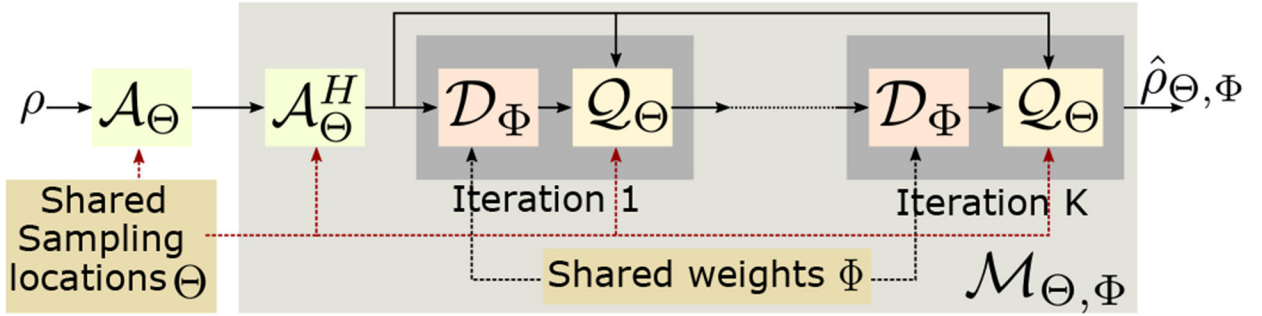
- [1]. Chen H, Zhang Y et al., “Low-Dose CT with a Residual Encoder-Decoder Convolutional Neural Network,” *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2524–2535, 2017.
- [2]. Han Y, Sunwoo L, and Ye JC, “k-space deep learning for accelerated MRI,” *IEEE Trans. Med. Imag.*, 2019.
- [3]. Yang G, Yu S et al., “DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction,” *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1310–1321, 2017.
- [4]. Quan TM, Nguyen-Duc T, and Jeong W-K, “Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss,” *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1488–1497, 2018.
- [5]. Hammernik K, Schlemper J et al., “Sigma-Net: Systematic evaluation of iterative deep neural networks for fast parallel MR image reconstruction,” *arXiv preprint arXiv:1912.09278*, 2019.
- [6]. Dar SUH, Yurt M et al., “Synergistic reconstruction and synthesis via generative adversarial networks for accelerated multi-contrast MRI,” *arXiv preprint arXiv:1805.10704*, 2018.
- [7]. Dar SUH, Özbey M et al., “A transfer-learning approach for accelerated MRI using deep neural networks,” *Magnetic Resonance in Medicine*, 2017.
- [8]. Zhu B, Liu JZ et al., “Image reconstruction by domain-transform manifold learning,” *Nature*, vol. 555, no. 7697, pp. 487–492, 2018. [PubMed: 29565357]
- [9]. Schmidt U and Roth S, “Shrinkage fields for effective image restoration,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 2774–2781.
- [10]. yang y., Sun J et al., “Deep ADMM-Net for compressive sensing MRI,” in *Advances in Neural Information Processing Systems 29*, 2016, pp. 10–18.
- [11]. Zhang J and Ghanem B, “ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1828–1837.
- [12]. Schlemper J, Caballero J et al., “A deep cascade of convolutional neural networks for dynamic MR image reconstruction,” *IEEE Trans. Med. Imag.*, vol. 37, no. pp. 491–503, 2018.
- [13]. Aggarwal HK, Mani MP, and Jacob M, “MoDL: Model based deep learning architecture for inverse problems,” *IEEE Trans. Med. Imag.*, vol. 38, no. pp. 394–405, 2019.
- [14]. Hammernik K, Klatzer T et al., “Learning a Variational Network for Reconstruction of Accelerated MRI Data,” *Magnetic resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2017. [PubMed: 29115689]
- [15]. Zhang L and Zuo W, “Image Restoration: From Sparse and Low-Rank Priors to Deep Priors,” *IEEE Signal Process. Mag.*, vol. 34, no. 5, pp. 172–179, 2017.
- [16]. Pramanik A, Aggarwal HK, and Jacob M, “Off-the-grid model based deep learning O-MoDL,” in *IEEE 16th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2019, pp. 1395–1398.
- [17]. Aggarwal HK, Mani MP, and Jacob M, “MoDL-MUSSELS: Model-based deep learning for multishot sensitivity-encoded diffusion MRI,” *IEEE Trans. Med. Imag.*, 2019.
- [18]. Mardani M, Gong E et al., “Deep generative adversarial neural networks for compressive sensing MRI,” *IEEE Trans. Med. Imag.*, vol. 38, no. 1, pp. 167–179, 2018.
- [19]. Venkatakrisnan SV, Bouman CA, and Wohlberg B, “Plug-and-play priors for model based reconstruction,” in *2013 IEEE Global Conference on Signal and Information Processing*. IEEE, 2013, pp. 945–948.
- [20]. Ahmad R, Bouman CA et al., “Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery,” *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 105–116, 2020.
- [21]. Sodickson DK and Manning WJ, “Simultaneous acquisition of spatial harmonics SMASH: fast imaging with radiofrequency coil arrays,” *Magnetic resonance in medicine*, vol. 38, no. 4, pp. 591–603, 1997. [PubMed: 9324327]
- [22]. Candes E and Romberg J, “Sparsity and incoherence in compressive sampling,” *Inverse problems*, vol. 23, no. 3, p. 969, 2007.

- [23]. Lustig M, Donoho DL et al., “Compressed sensing MRI,” *IEEE signal processing magazine*, vol. 25, no. p. 72, 2008.
- [24]. Levine E, Daniel B et al., “3D Cartesian MRI with compressed sensing and variable view sharing using complementary Poisson-disc sampling,” *Magnetic resonance in medicine*, vol. 77, no. 5, pp. 1774–1785, 2017. [PubMed: 27097596]
- [25]. Gao Y and Reeves SJ, “Optimal k-space sampling in MRSI for images with a limited region of support,” *IEEE Trans. Med. Imag.*, vol. 19, no. 12, pp. 1168–1178, 2000.
- [26]. Xu D, Jacob M, and Liang Z, “Optimal sampling of k-space with Cartesian grids for parallel MR imaging,” in *Proc Int Soc Magn Reson Med*, vol. 13, 2005, p. 2450.
- [27]. Haldar JP and Kim D, “OEDIPUS: An experiment design framework for sparsity-constrained MRI,” *IEEE Trans. Med. Imag.*, 2019.
- [28]. Levine E and Hargreaves B, “On-the-fly adaptive k-space sampling for linear MRI reconstruction using moment-based spectral analysis,” *IEEE Trans. Med. Imag.*, vol. 37, no. pp. 557–567, 2017.
- [29]. Senel LK, Kilic T et al., “Statistically segregated k-space sampling for accelerating multiple-acquisition MRI,” *IEEE Trans. Med. Imag.*, vol. 38, no. 7, pp. 1701–1714, 2019.
- [30]. Liu F, Samsonov A et al., “SANTIS: Sampling-augmented neural network with incoherent structure for MR image reconstruction,” *Magnetic resonance in medicine*, 2019.
- [31]. Sherry F, Benning M et al., “Learning the sampling pattern for MRI,” arXiv preprint arXiv:1906.08754, 2019.
- [32]. Gözcü B, Mahabadi RK et al., “Learning-based compressive MRI,” *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1394–1406, 2018.
- [33]. Weiss T, Vedula S et al., “Learning fast magnetic resonance imaging,” arXiv preprint arXiv:1905.09324, 2019.
- [34]. Weiss T, Senouf O et al., “PILOT: Physics-informed learned optimal trajectories for accelerated MRI,” arXiv preprint arXiv:1909.05773, 2019.
- [35]. Bahadir CD, Dalca AV, and Sabuncu MR, “Learning-based optimization of the under-sampling pattern in MRI,” in *International Conference on Information Processing in Medical Imaging*. Springer, 2019, pp. 780–792.
- [36]. Lustig M and Pauly JM, “SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space,” *Magnetic resonance in medicine*, vol. 64, no. pp. 457–471, 2010. [PubMed: 20665790]
- [37]. Jin KH, Unser M, and Yi KM, “Self-supervised deep active accelerated MRI,” arXiv preprint arXiv:1901.04547, 2019.
- [38]. Zhang Z, Romero A et al., “Reducing uncertainty in undersampled MRI reconstruction with active acquisition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2049–2058.
- [39]. Metzler CA, Ikoma H et al., “Deep optics for single-shot high-dynamic-range imaging,” arXiv preprint arXiv:1908.00620, 2019.
- [40]. Muthumbi A, Chaware A et al., “Learned sensing: jointly optimized microscope hardware for accurate image classification,” *Biomed. Opt. Express*, vol. 10, no. 12, pp. 6351–6369, 2019. [PubMed: 31853404]
- [41]. Horstmeyer R, Chen RY et al., “Convolutional neural networks that teach microscopes how to image,” 2017.
- [42]. Cheng YF, Strachan M et al., “Illumination pattern design with deep learning for single-shot fourier ptychographic microscopy,” *Opt. Express*, vol. 27, no. pp. 644–656, 2019. [PubMed: 30696147]
- [43]. Chakrabarti A, “Learning sensor multiplexing design through back-propagation,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3081–3089.
- [44]. Knoll F, Clason C et al., “Adapted random sampling patterns for accelerated mri,” *Magnetic resonance materials in physics, biology and medicine*, vol. 24, no. 1, pp. 43–50, 2011.
- [45]. Zbontar J, Knoll F et al., “fastMRI: An open dataset and benchmarks for accelerated MRI,” arXiv preprint arXiv:1811.08839, 2018.

- [46]. Doneva M, Börnert P et al., “Compressed sensing reconstruction for magnetic resonance parameter mapping,” *Magnetic Resonance in Medicine*, vol. 64, no. 4, pp. 1114–1120, 2010. [PubMed: 20564599]
- [47]. Figueiredo MAT, Nowak RD et al., “An EM Algorithm for Wavelet-Based Image Restoration,” *IEEE Trans. Image Process*, vol. 12, no. 8, pp. 906–916, 2003. [PubMed: 18237964]
- [48]. Ma S, Yin W et al., “An Efficient Algorithm for Compressed MR Imaging using Total Variation and Wavelets,” in *Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
- [49]. Jacob M, Mani MP, and Ye JC, “Structured low-rank algorithms: Theory, MR applications, and links to machine learning,” *IEEE Signal Processing Magazine*, pp. 1–12, 2019.
- [50]. Ronneberger O, Fischer P, and Brox T, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2015, pp. 234–241.
- [51]. Uecker M, Lai P et al., “ESPIRiT - An eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA,” *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014. [PubMed: 23649942]
- [52]. Lazarus C, Weiss P et al., “SPARKLING: variable-density k-space filling curves for accelerated T2*-weighted MRI,” *Magnetic resonance in medicine*, vol. 81, no. 6, pp. 3643–3661, 2019. [PubMed: 30773679]



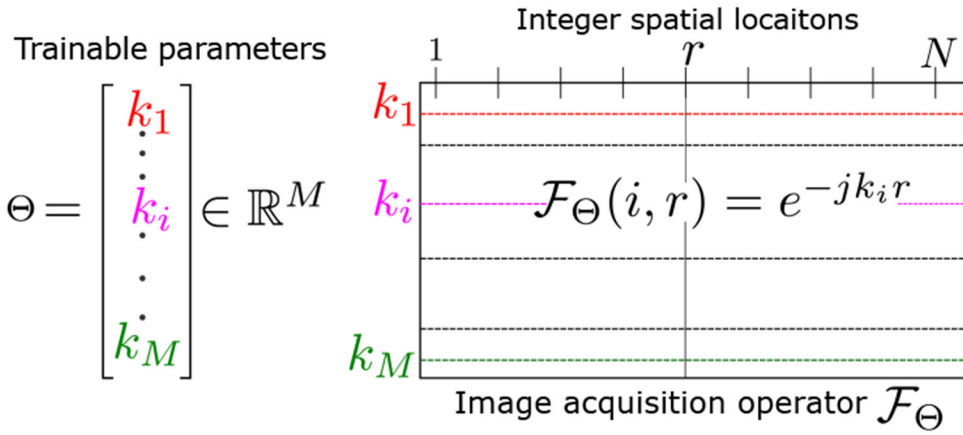
(a) J-UNET architecture



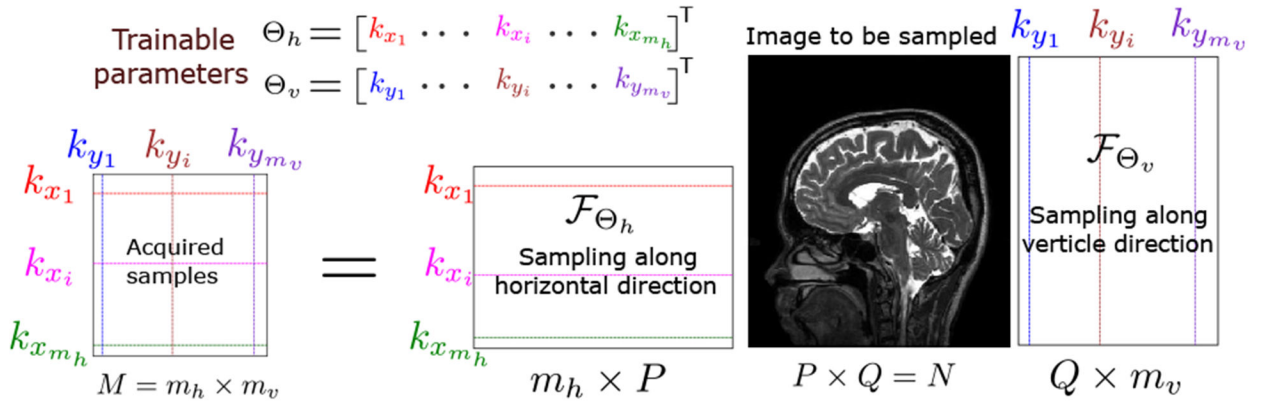
(b) proposed J-MoDL architecture

Fig. 1.

Illustration of the simultaneous sampling and reconstruction architectures. (a) The direct-inversion (J-UNET) architecture described by (3), where a CNN \mathcal{N}_Φ is used to recover the images from $\mathcal{A}_\Theta^H \mathbf{b}$. As discussed previously, the CNN parameters are closely coupled with the specific sampling pattern, making joint optimization challenging. (b) corresponds to the J-MoDL architecture, described by (5) and (6). Each iteration alternates between the CNN denoiser \mathcal{D}_Φ and the data consistency block \mathcal{Q}_Θ . The data consistency block \mathcal{Q}_Θ inverts the measured Fourier samples assuming \mathbf{z}_n , while \mathcal{D}_Φ acts as a denoiser of the current iterate. The blocks \mathcal{D}_Φ and \mathcal{Q}_Θ are relatively independent of Θ and Φ , respectively.



(a) sampling 1-D



(b) sampling 2-D

Fig. 2.

Illustration of the proposed sampling parameterization to acquire M samples of an N dimensional signal such that the acceleration factor = N/M . (a) The sampling operator \mathcal{F}_Θ is an $M \times N$ matrix that can capture M samples from possibly non-integer locations k_1, \dots, k_M . These M locations are real-valued trainable parameters constrained between $[0, 1]$. (b) In the 2-D case, we utilize two sampling operators, \mathcal{F}_{Θ_h} and \mathcal{F}_{Θ_v} , in the horizontal and the vertical directions, respectively. \mathcal{F}_{Θ_h} acquires m_h samples, whereas \mathcal{F}_{Θ_v} acquires m_v samples such that total $M = m_h \times m_v$ samples are acquired from $N = P \times Q$ dimensional image.

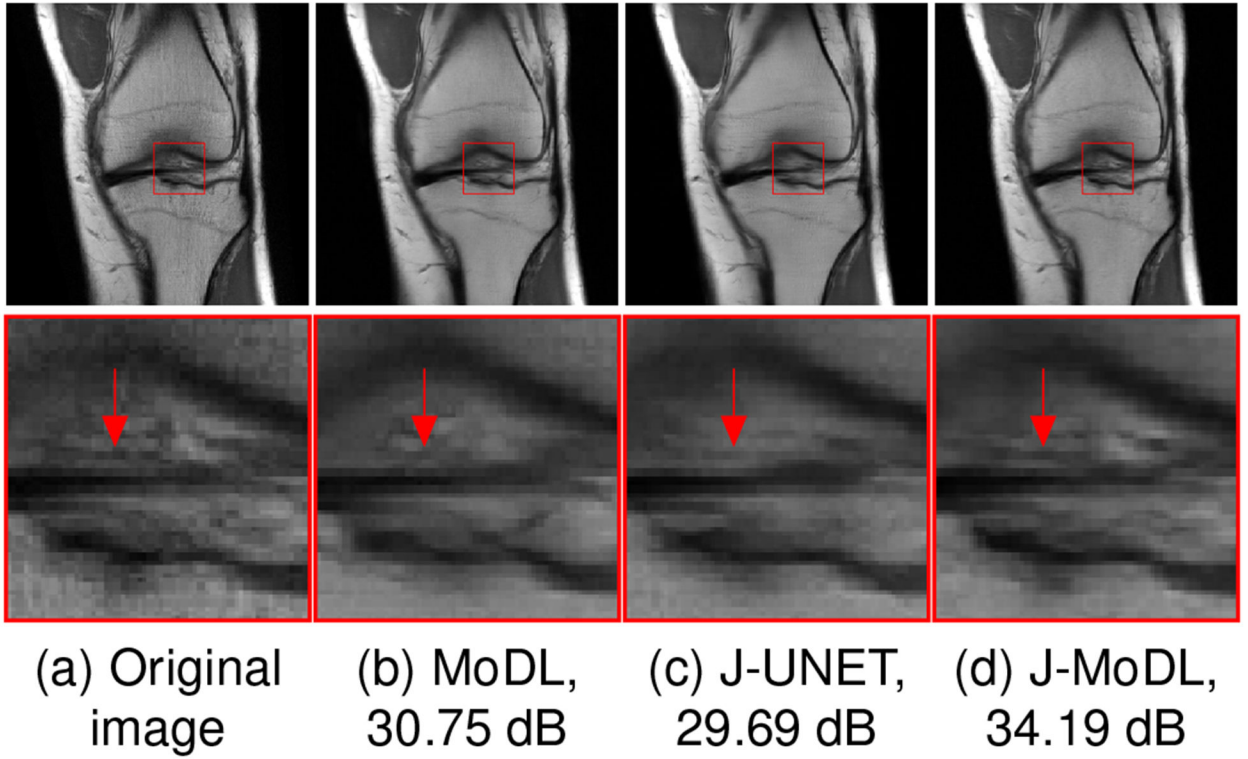


Fig. 3.

The visual comparisons of different optimization strategies, described in section III-B, on a test slice. The numbers in the subcaption show the SSIM values. The green arrow in the zoomed area points to a feature not captured by J-UNET despite having higher SSIM as compare to MoDL. The red arrow points to thin vertical features not captured by the Greedy approach.

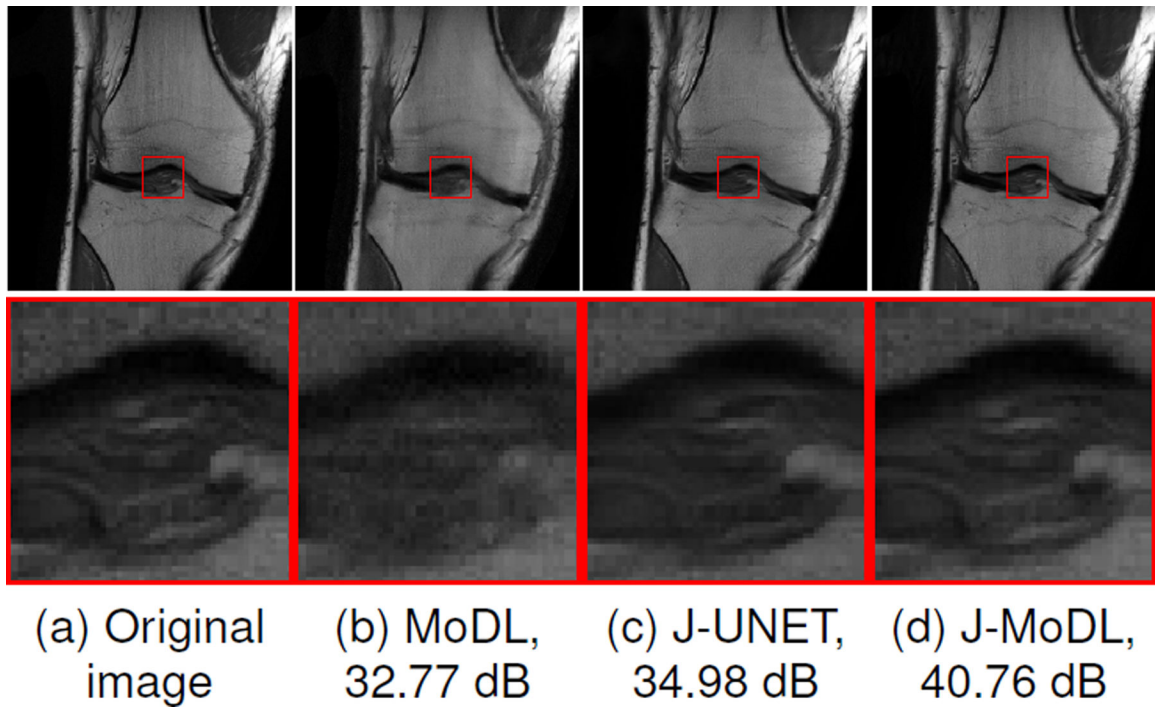


Fig. 4.

Comparison of joint and network-alone optimization in parallel imaging settings, described in section III-C, with a 1-D sampling mask. The numbers in subcaptions are showing the PSNR (dB) values. (a) shows a fully sampled image from the test dataset. (b) shows the reconstructed image with a pseudo-random 4x acceleration mask using the MoDL approach. (c,d) shows joint optimization of sampling as well as network parameters using direct-inversion and model-based techniques, respectively. The zoomed areas clearly show that joint learning better preserves the fine details.

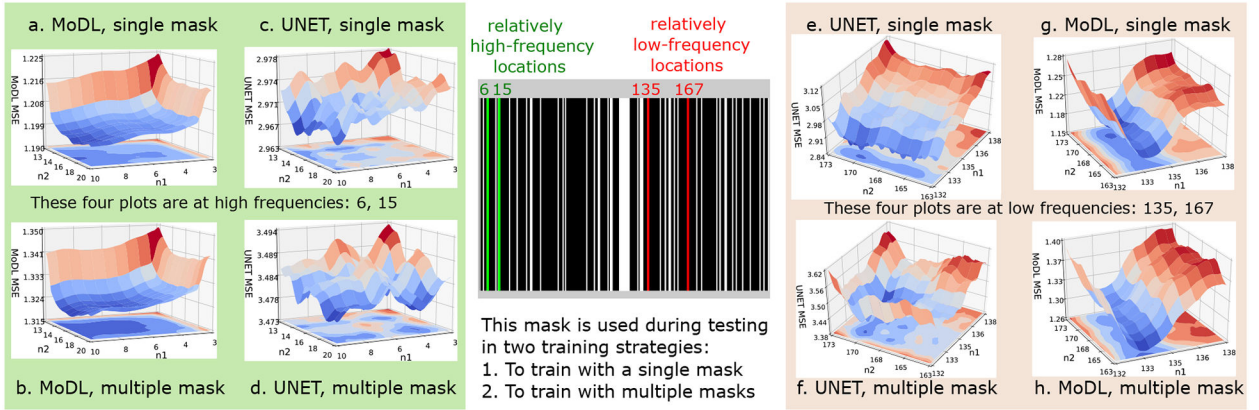


Fig. 5. This figure compares the optimization landscape of the MoDL ($K=1$) and UNET architecture for 1-D multichannel data. These plots show the mean squared error ($MSE \times 1000$) between the reconstructions and the corresponding original images. The n_1 and n_2 axes represent continuous valued sampling locations around the ones marked on the mask. (a) and (b) show the landscape plot for MoDL architecture trained with a single sampling pattern and multiple sampling patterns, respectively. Similarly (c) and (d) shows corresponding plots for the UNET architecture. These plots (a)-(d) are plotted at high-frequency values around locations 6 and 15, as marked with green in the mask. Similarly, (e)-(h) show landscape plots at relatively low frequencies around locations 135 and 167. From this controlled experiment, we observe that MoDL results in a smoother landscape as compared to UNET both at low and high frequencies. In addition, the UNET landscapes become comparatively smoother with the sampling pattern augmentation strategy, which makes the approach relatively insensitive to small differences in sampling pattern, as seen from (c) to (d) and (e) to (f).




Fig. 6.

This figure compares the different schemes in parallel imaging settings with a 2-D sampling mask at 6x acceleration, as described in section III-D. Row 1 and row 2 show magnitude images and a zoomed cerebellum region, respectively. Row 3 shows reconstructed phase images, while row 4 shows error maps of the reconstructed images with respect to the original image. Φ -UNet (a) and Φ -MoDL (b) optimize only the network parameters Φ . Θ -UNet and Θ -MoDL optimize only the sampling parameters Θ with initialization from respective Φ -alone models. The sub-captions denote the SSIM values. Finally, J-UNET and J-MoDL are the proposed joint optimization models. The J-MoDL approach preserves the fine features in the cerebellum region, as shown by the zoomed area. The red arrows in the zoomed area point to a feature that is well preserved by the joint optimization techniques versus the results of the networks that optimized only network parameters.


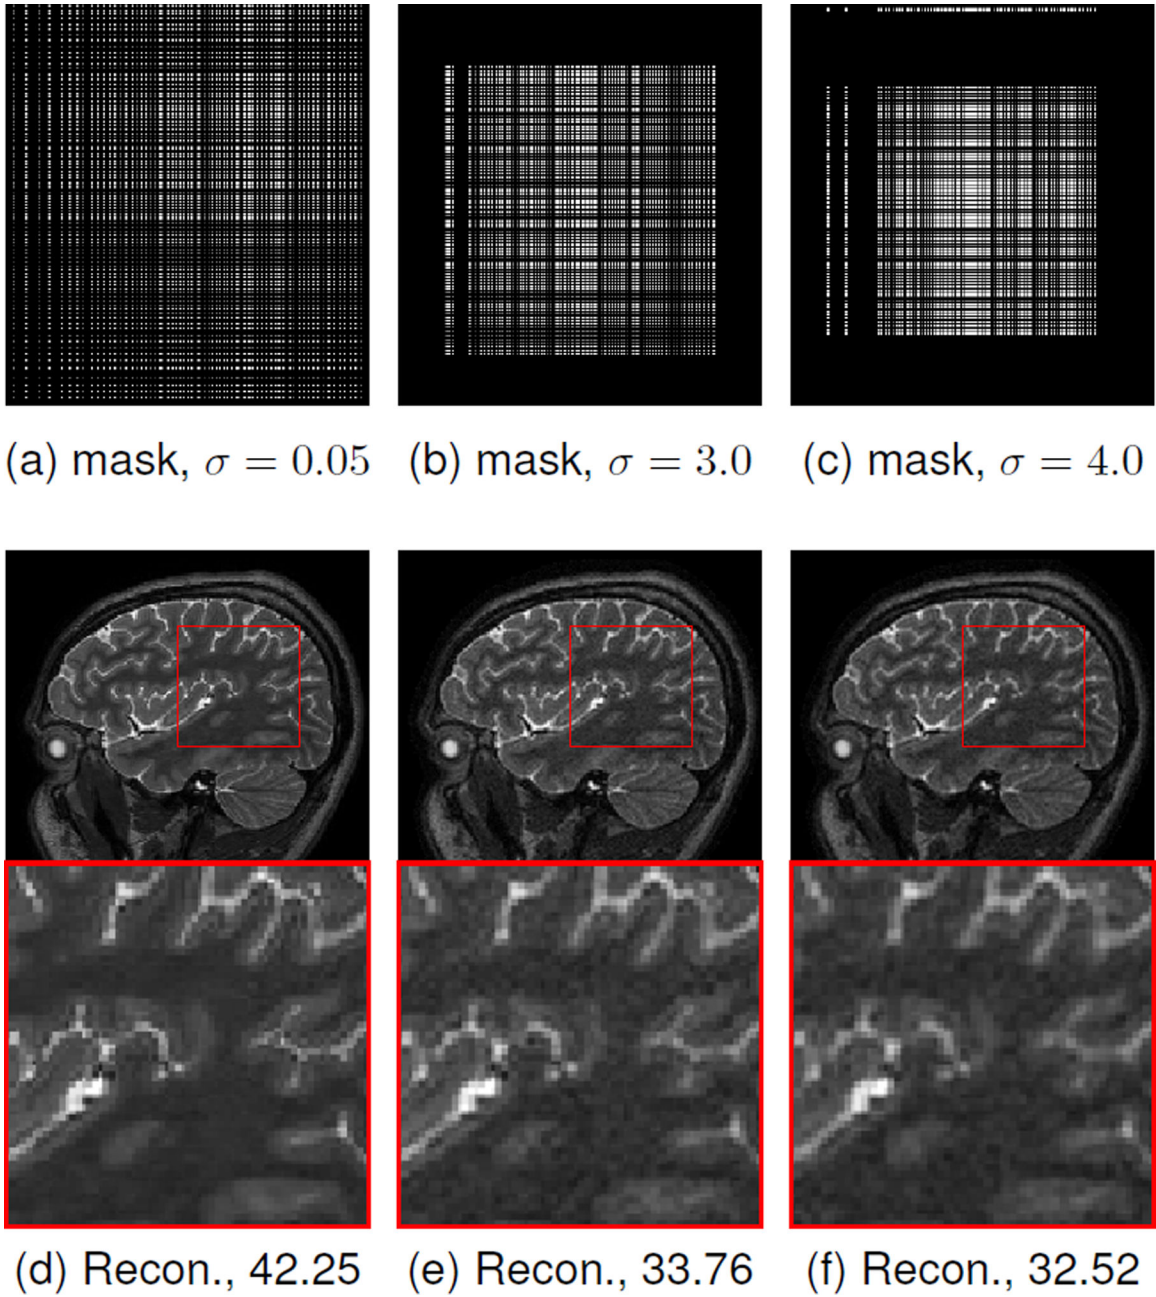


Fig. 7.

This figure compares the reconstruction quality obtained by different sampling masks in 2-D parallel imaging settings at 10x acceleration as described in section III-D. Rows one, two, and three shows masks, A^{Hb} , and reconstruction outputs, respectively. Two Φ -alone models using the MoDL approach were trained with random masks as well as random variable-density (VD) masks. It can be observed that the 2-D mask learned using the J-MoDL approach outperforms the reconstruction using fixed random and variable-density masks. The numbers in sub-captions are showing PSNR (dB) values.

**Fig. 8.**

This figure demonstrates the impact of adding a high amount of noise in the k-space samples in 2-D parallel MRI settings at 8x acceleration, as described in section III-D. The first row shows different masks learned with the J-MoDL approach when the Gaussian noise of standard deviation σ is added in the k-space samples. The second row shows corresponding reconstructions (Recon.). Subcaptions of (d), (e), and (f) are showing PSNR (dB) values. As expected, higher noise levels promote the algorithm to learn the sampling parameters that sample more of the low-frequency components from the center of k-space, leading to low-resolution reconstructions.

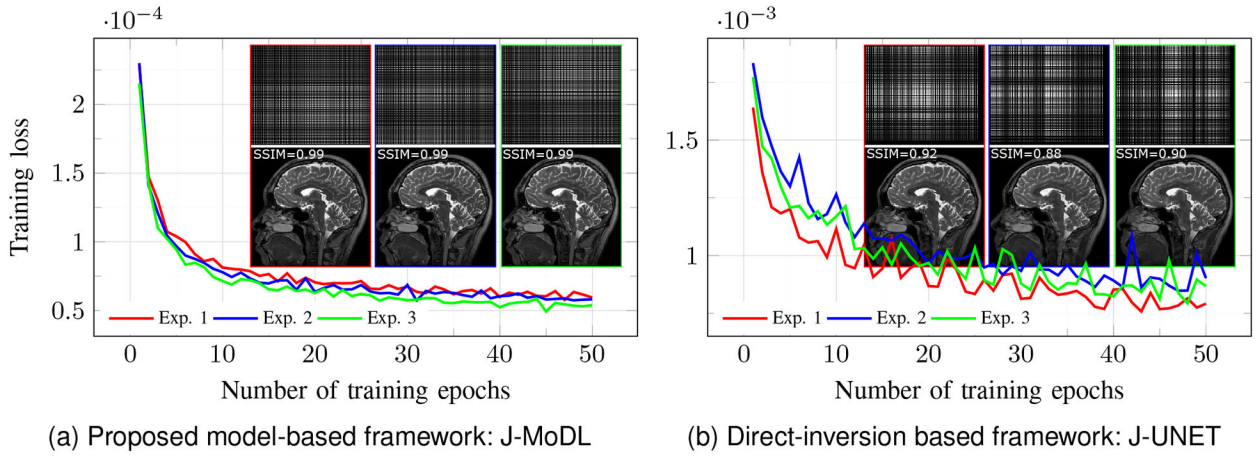



Fig. 9.

This figure compares the convergence of training loss in joint optimization with direct-inversion and model-based deep learning frameworks. We performed three independent experiments for each of the two frameworks. The experimental setup for all the three experiments was identical except for the initialization of sampling parameters. The learned masks, reconstructed images, as well as the training loss, are plotted for each of the three experiments. We observe that J-MoDL convergence is relatively smoother compared to the J-UNET approach; all three initializations resulted in results with somewhat similar image quality. We note that the cost function may have multiple global minima. The J-UNET convergence was slower compared to the J-MoDL setting, likely because of the non-smooth optimization landscape as seen from Fig. 5.

**Fig. 10.**

This figure shows comparative results between the proposed J-MoDL approach and existing algorithms at 8x acceleration in multichannel (MC) settings. (a) shows the 8x sampling mask used for the initialization of the J-ISTANet, J-UNET, and J-MoDL approaches. Row 1 shows the learned masks by the respective algorithms. The learned masks by both J-MoDL approaches takes samples from all the locations in the k-space that help in preserving the high-frequency details in the reconstructed images as shown by the zoomed area near cerebellum.

TABLE I

Single-channel settings: The mean \pm std values of PSNR (dB) and SSIM over the test data of hundred subjects using different optimization strategies at 4x acceleration.

Optimize	PSNR		SSIM	
	UNET	MoDL	UNET	MoDL
Φ -alone	28.65 ± 1.14	30.65 ± 1.43	0.80 ± 0.03	0.82 ± 0.04
Θ -alone	29.02 ± 1.03	32.46 ± 1.07	0.80 ± 0.03	0.84 ± 0.03
Joint	29.70 ± 1.06	33.78 ± 1.13	0.82 ± 0.03	0.87 ± 0.03

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE II

Impact of optimization strategies for parallel MRI recovery of knee images using 1-D sampling. The results correspond to two subjects with a total of 80 slices.

	PSNR		SSIM	
	UNET	MoDL	UNET	MoDL
Optimize	4x acceleration			
Φ -alone	29.95 ± 3.76	34.21 ± 3.14	0.83 ± 0.13	0.91 ± 0.04
Θ -alone	28.85 ± 3.94	37.66 ± 3.30	0.86 ± 0.04	0.96 ± 0.03
Joint	34.02 ± 3.31	41.28 ± 3.07	0.93 ± 0.04	0.96 ± 0.02
	6x acceleration			
Φ -alone	29.24 ± 3.94	32.40 ± 3.00	0.82 ± 0.13	0.89 ± 0.04
Θ -alone	24.45 ± 3.65	33.31 ± 3.17	0.78 ± 0.09	0.93 ± 0.03
Joint	29.62 ± 2.54	35.93 ± 2.74	0.89 ± 0.05	0.93 ± 0.03

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE III

Impact of optimization strategies for parallel MRI recovery of the brain images using 2-D sampling. The PSNR and SSIM values are reported for the average \pm std. of 200 test slices at 6x acceleration factor.

Optimize	PSNR		SSIM	
	UNET	MoDL	UNET	MoDL
Φ -alone	27.34 ± 1.14	34.19 ± 1.03	0.82 ± 0.02	0.94 ± 0.01
Θ -alone	28.56 ± 0.93	37.47 ± 0.57	0.85 ± 0.02	0.94 ± 0.01
Joint	34.31 ± 0.81	37.60 ± 0.56	0.94 ± 0.01	0.96 ± 0.01

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE IV

The average PSNR and SSIM values along with standard deviation are shown for different algorithms at 8x acceleration in the multichannel settings. The LOUPE algorithm is for joint optimization therefore its results are not available (NA) for the network alone (Φ -alone) case.

Algorithm	PNSR		SSIM	
	Φ alone	$\Theta \Phi$ Joint	Φ alone	$\Theta \Phi$ Joint
MC-LOUPE	NA	33.68 ± 3.23	NA	0.92 ± 0.02
ISTA K=5	28.66 ± 1.58	34.38 ± 1.38	0.86 ± 0.03	0.94 ± 0.01
UNET	26.27 ± 1.33	30.63 ± 1.22	0.79 ± 0.03	0.91 ± 0.02
MoDL K=1	30.68 ± 1.53	31.85 ± 0.84	0.89 ± 0.02	0.92 ± 0.01
MoDL K=5	32.72 ± 1.34	36.38 ± 0.54	0.92 ± 0.02	0.95 ± 0.01