



Published in final edited form as:

IEEE Trans Med Imaging. 2020 December ; 39(12): 4186–4197. doi:10.1109/TMI.2020.3014581.

Deep Generalization of Structured Low-Rank Algorithms (Deep-SLR)

Aniket Pramanik, Hemant Aggarwal, Mathews Jacob

Department of Electrical and Computer Engineering at the University of Iowa, Iowa City, IA, 52242, USA

Abstract

Structured low-rank (SLR) algorithms, which exploit annihilation relations between the Fourier samples of a signal resulting from different properties, is a powerful image reconstruction framework in several applications. This scheme relies on low-rank matrix completion to estimate the annihilation relations from the measurements. The main challenge with this strategy is the high computational complexity of matrix completion. We introduce a deep learning (DL) approach to significantly reduce the computational complexity. Specifically, we use a convolutional neural network (CNN)-based filterbank that is trained to estimate the annihilation relations from imperfect (under-sampled and noisy) k-space measurements of Magnetic Resonance Imaging (MRI). The main reason for the computational efficiency is the pre-learning of the parameters of the non-linear CNN from exemplar data, compared to SLR schemes that learn the linear filterbank parameters from the dataset itself. Experimental comparisons show that the proposed scheme can enable calibration-less parallel MRI; it can offer performance similar to SLR schemes while reducing the runtime by around three orders of magnitude. Unlike pre-calibrated and self-calibrated approaches, the proposed uncalibrated approach is insensitive to motion errors and affords higher acceleration. The proposed scheme also incorporates image domain priors that are complementary, thus significantly improving the performance over that of SLR schemes.

Index Terms—

parallel MRI; reconstruction; structured low rank; annihilation; DL

I. Introduction

Parallel MRI was introduced to speed up the traditionally slow MR acquisitions. Specifically, the redundancy between the k-space samples is capitalized to highly undersample k-space data, thereby reducing the scan time. The recovery of images from highly under-sampled multi-channel Fourier measurements is a classical problem in MRI [1]. Pre-calibrated approaches such as SENSE [2] rely on coil sensitivities that are estimated using additional calibration scans. Several image priors [3], including sparsity, have been

used to regularize pre-calibrated image recovery, resulting in improved image recovery at high acceleration factors. Researchers have recently introduced model-based deep-learning (DL) algorithms that use a forward model (capturing the imaging physics) combined with a deep-learned prior [4], [5], [6], [7]. An improvement in image quality as a result of multiple iterations of optimization blocks, sharing weights across the network, and end-to-end training, along with the ability to use multiple learned regularization priors, has been demonstrated in [7]. Since these methods use pre-estimated coil sensitivity maps within the forward model, they suffer from errors in the sensitivity maps resulting from motion or high accelerations. Self-calibrated approaches such as GRAPPA [8], SPIRiT [9] and ESPIRiT [10] estimate coil sensitivities from a fully sampled calibration region in the center of k -space. However, the need for a fully sampled region restricts the achievable acceleration in these settings.

Structured low-rank (SLR) matrix completion approaches [11], [12], [13] were introduced to overcome the challenges with the above calibration based schemes and have been very effective in uncalibrated parallel MRI [10] and multishot acquisitions [14]. In the context of parallel MRI, these methods exploit the annihilation relations between multi-channel Fourier data rather than relying on explicit coil sensitivity estimates. Similar SLR approaches have been used to exploit a variety of other signal properties, including support constraints [12], continuous domain sparsity [15], [16], phase [12], and the exponential structure of an MRI time series [11]. Iterative re-weighted least-squares (IRLS) SLR algorithms make use of the convolutional structure of the matrices [17] to accelerate the algorithms. IRLS methods alternate between estimating an annihilation (null space) filterbank and updating the Fourier coefficients of the signal from the available measurements. Specifically, the missing Fourier coefficients are chosen so they match the measurements while being annihilated by the filterbank; the projection energy of the signal to the signal subspace measured by a residual convolution-deconvolution filterbank, is maximized. While this algorithm is considerably faster than earlier approaches, the iterative estimation of the annihilation filterbank from the under-sampled data is still computationally expensive. The approaches that use calibration information [10], [15], [12] estimate the null space filters from a fully sampled calibration region, resulting in reduced complexity. Since the annihilation filterbank need not be derived from the under-sampled data in an iterative fashion, this approach offers faster reconstructions. However, the challenge with these methods is the need for a calibration region, which restricts the achievable acceleration.

In this paper, we introduce a general DL strategy to reduce the runtime of the SLR algorithms, which is valid for all the signal priors discussed above [15], [16], [12], [11]. Unlike the SLR approach that estimates a specific linear annihilation network for each dataset from the undersampled measured data, we propose to learn a single non-linear CNN from several training datasets. Specifically, the residual convolution-deconvolution linear filterbank in the IRLS-SLR algorithm is replaced with a residual multi-channel CNN. We hypothesize that the pre-learned non-linear CNN behaves as a different linear annihilation filterbank for each specific dataset, annihilating the multi-channel data. The residual CNN behaves as a projection for each dataset, facilitating the *denoising* of the dataset from alias artifacts and noise at each iteration. Similar to MoDL [7], the proposed model unrolls the resulting algorithm and learns the parameters of the non-linear filterbank in an end-to-end

fashion. The proposed work combines this approach with an image domain prior similar to MoDL, which is complementary to the Fourier domain multi-channel relations. This hybrid approach offers improved performance over SLR, while offering around three orders of magnitude reduction in computational complexity. We focus on two representative applications—sparse single-channel recovery and parallel MRI—which use two distinct lifting structures in the SLR approach [11]. Specifically, we show how different lifting structures can be accommodated in the proposed scheme by modifying the data organization of the input and output of the CNN module. This enables the extension of the proposed framework to a range of SLR applications [11], [12], [18] that use one or a combination of the above lifting structures.

This main focus of this work is to introduce a general DL framework for uncalibrated parallel MRI and multishot MRI. This work is related to [19] and [20], which are partially calibrated strategies. Specifically, the MoDL-MUSSELS [19] framework explicitly accounted for the pre-estimated coil sensitivities within the data consistency block, while it performed a calibration-free correction of phase errors between shots. A challenge with these partially calibrated approaches is the potential mismatch between coil sensitivities and the diffusion-weighted acquisition due to motion between the coil sensitivity calibration scan and the diffusion scan. By contrast, the annihilation relation between coils is learned by the non-linear k-space CNN from exemplar data in this work; the application of this framework to the diffusion setting yields a completely uncalibrated algorithm, which jointly accounts for coil sensitivities as well as phase errors between shots. As demonstrated in our experiments in the supplementary material, this approach eliminates errors resulting from a motion-induced mismatch between the calibration scan and the diffusion-weighted one. In addition, the main focus of this work is to show that the proposed approach works well for a range of SLR priors, of which only one was considered in the earlier work [19]. The conference version of this work is presented in the literature [21], [22]. DL methods in k-space were the focus of recent work [23], [24], [25]. The RAKI framework [25] is a calibrated scheme, unlike our calibration-free approach. A direct inversion (model-free) approach was pursued in one study [23]; it differs from the proposed model-based framework, which also combines image domain priors. The KIKI net approach [24] was introduced for a single-channel setting, unlike our uncalibrated multi-channel scheme. The proposed reconstructions are compared against model-free image domain DL [26], k-space DL [23], image domain MoDL [7], and traditional SLR methods. These comparisons reveal the improved performance offered by the Deep-SLR framework.

II. Background

We now briefly describe the background to make the paper self-contained and easily accessible.

A. Forward Model

We model the acquisition of image $\gamma(\mathbf{r})$ as:

$$b_i = \mathcal{S} \left(\underbrace{\mathcal{F}(s_i \gamma)}_{\gamma_i} \right) + \eta_i, \quad i = 1 \dots M, \quad (1)$$

where $s_i, i = 1, \dots, M$ is the coil sensitivity of the i^{th} coil, while b_i is the noisy under-sampled Fourier measurement, and γ_i is the image, corresponding to the i^{th} coil. η_i is the noise term. Here \mathcal{F} is the Fourier transform that maps γ_i onto its k-space samples and \mathcal{S} is the under-sampling operator. We compactly denote the above operation as

$$\mathbf{B} = \mathcal{A}(\mathbf{\Gamma}) + \mathbf{P} \quad (2)$$

where $\widehat{\mathbf{\Gamma}} = [\widehat{\gamma}_1 \dots \widehat{\gamma}_M]$ is the matrix representing multi-channel data in Fourier space, $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_M]$ is the corresponding noisy under-sampled multi-channel Fourier measurement, and $\mathbf{P} = [\boldsymbol{\eta}_1 \dots \boldsymbol{\eta}_M]$ is the multi-channel noise. Note that we denote the image of the i^{th} channel by γ_i , while $\mathbf{\Gamma}$ denotes the concatenation of the channel data in spatial domain. The Fourier domain representation of i^{th} channel image is $\widehat{\gamma}_i$, while $\widehat{\mathbf{\Gamma}}$ is the concatenation of the channel images in Fourier domain.

B. Structured Low-Rank Algorithms

SLR methods rely on different liftings of the Fourier coefficients, designed to exploit specific properties of the signal. We now discuss two representative SLR applications, which illustrate the different types of lifting used in the SLR setting.

1) Continuous domain sparsity: A continuous domain piecewise constant image γ with edges specified by the zero sets of a bandlimited function μ satisfies an image domain annihilation relation, $\nabla \gamma(\mathbf{r}) \cdot \mu(\mathbf{r}) = 0, \forall \mathbf{r}$, where \mathbf{r} represents spatial coordinates. Here, $\nabla \gamma$ denotes the gradient of γ . This relation translates to the following Fourier domain annihilation relations $\widehat{\nabla \gamma}[\mathbf{k}] * n[\mathbf{k}] = 0, \forall \mathbf{k}$, where \mathbf{k} denotes k-space coordinates (Fourier space). Here $\widehat{\nabla \gamma}[\mathbf{k}]$ represents the Fourier coefficients of the gradient of γ and $n[\mathbf{k}]$ is the Fourier transform of $\mu(\mathbf{r})$. We denote the mapping from the Fourier coefficients $\widehat{\gamma}$ to $\widehat{\nabla \gamma}$ by \mathcal{E} :

$$\mathcal{E}(\widehat{\gamma}) = \widehat{\nabla \gamma}[\mathbf{k}] = \begin{bmatrix} j2\pi k_x \widehat{\gamma}[\mathbf{k}] \\ j2\pi k_y \widehat{\gamma}[\mathbf{k}] \end{bmatrix} = \begin{bmatrix} \widehat{\gamma}_x \\ \widehat{\gamma}_y \end{bmatrix}. \quad (3)$$

Note that \mathcal{E} essentially creates two copies of $\widehat{\gamma}$, each with a different Fourier weighting.

The convolution relation $\widehat{\nabla \gamma}[\mathbf{k}] * n[\mathbf{k}] = 0$ can be represented as Hankel matrix multiplication $\mathcal{H}(\widehat{\nabla \gamma})\mathbf{n} = 0$. The number of such null space filters, denoted by V , is often large (see [27])

$$\underbrace{\begin{bmatrix} \mathcal{H}(\widehat{\gamma}_x) \\ \mathcal{H}(\widehat{\gamma}_y) \end{bmatrix}}_{\mathcal{T}(\mathcal{E}(\widehat{\gamma}))} \underbrace{[n_1 \ n_2 \ \dots \ n_V]}_{\mathbf{N}} = 0. \quad (4)$$

resulting in a low-rank matrix $\mathcal{T}(\mathcal{E}(\widehat{\gamma}))$.

Note that the Hankel matrices are vertically stacked to obtain $\mathcal{T}(\mathcal{G}(\hat{\gamma}))$, which is a common approach in SLR [10].

2) Parallel MRI acquisition scheme: Image and Fourier domain multi-channel annihilation relations were shown in two studies [10], [11]. Specifically, each pair of multi-channel images in (1) satisfy a Fourier domain annihilation relation $\hat{\gamma}_i[\mathbf{k}] * \hat{s}_j[\mathbf{k}] - \hat{\gamma}_j[\mathbf{k}] * \hat{s}_i[\mathbf{k}] = 0, \forall \mathbf{k}$, where $\hat{\gamma}_i[\mathbf{k}]$ and $\hat{s}_i[\mathbf{k}]$ are the Fourier coefficients of $\gamma_i(\mathbf{k})$ and $s_i(\mathbf{k})$, respectively. Such annihilation relations exist for every pair of coil images and can be compactly written as

$$\underbrace{[\mathcal{H}(\hat{\gamma}_1) \mathcal{H}(\hat{\gamma}_2) \dots \mathcal{H}(\hat{\gamma}_M)]}_{\mathcal{T}(\hat{\Gamma})} \cdot \mathbf{N} = 0. \quad (5)$$

The columns of \mathbf{N} correspond to the vertical stacking of the filters \hat{s}_i . The large null space \mathbf{N} implies it is low rank. Note that the Hankel matrices are horizontally stacked to obtain $\mathcal{T}(\hat{\Gamma})$. Here $\mathcal{G} = \mathcal{I}$, which is the identity mapping. This is another popular class of lifting used in SLR [11], [12], [16].

C. Calibration-free SLR Methods

In general, SLR schemes aim to recover an image or a series of images $\mathbf{\Gamma}$ from its measurements $\mathcal{A}(\mathbf{\Gamma})$ by solving the optimization problem:

$$\min_{\mathbf{\Gamma}} \text{rank} [\mathcal{T}(\mathcal{G}(\hat{\Gamma}))] \text{ such that } \mathbf{B} = \mathcal{A}(\mathbf{\Gamma}) + \mathbf{P}. \quad (6)$$

Here, $\mathcal{T}(\cdot)$ is a lifting operator that lifts the weighted signal $\mathcal{G}(\hat{\Gamma})$ into a higher dimensional structured matrix. As discussed earlier, the generic weighting matrix \mathcal{G} depends on the specific annihilation relation. The recovery of $\mathbf{\Gamma}$ is often posed as an unconstrained nuclear norm minimization problem

$$\arg \min_{\mathbf{\Gamma}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda \left\| \mathcal{T}(\mathcal{G}(\hat{\Gamma})) \right\|_* \quad (7)$$

where λ is a regularizer to tune the nuclear norm loss term.

D. Iterative Re-weighted Least-Squares (IRLS) Algorithm

The IRLS scheme majorizes the nuclear norm with a weighted Frobenius norm as $\left\| \mathcal{T}(\mathcal{G}(\hat{\Gamma})) \right\|_* \leq \left\| \mathcal{T}(\mathcal{G}(\hat{\Gamma}))\mathbf{Q} \right\|_F^2$ to yield a two-variable optimization problem

$$\arg \min_{\mathbf{\Gamma}, \mathbf{Q}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda \left\| \mathcal{T}(\mathcal{G}(\hat{\Gamma}))\mathbf{Q} \right\|_F^2, \quad (8)$$

which alternates between the null space \mathbf{Q} and image $\mathbf{\Gamma}$,

$$\mathbf{\Gamma}^{(n)} = \arg \min_{\mathbf{\Gamma}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda \left\| \mathcal{T}(\mathcal{G}(\widehat{\mathbf{\Gamma}})) \mathbf{Q}^{(n-1)} \right\|_F^2 \quad (9)$$

$$\mathbf{Q}^{(n)} = \left[\mathcal{T}(\mathcal{G}(\widehat{\mathbf{\Gamma}}^{(n)}))^H \mathcal{T}(\mathcal{G}(\widehat{\mathbf{\Gamma}}^{(n)})) + \epsilon^{(n)} \mathbf{I} \right]^{-1/4} \quad (10)$$

respectively. The matrix \mathbf{Q} can be viewed as a collection of column vectors spanning the null space of $\mathcal{T}(\mathcal{G}(\widehat{\mathbf{\Gamma}}))$.

E. Calibration-based SLR Methods

Several calibration-based MRI schemes (e.g., GRAPPA, SPIRiT [8], [9]) are related to the SLR schemes [10], [11]. These approaches acquire a fully sampled calibration region in the Fourier domain, which corresponds to fully sampled rows of $\mathcal{T}(\mathcal{G}(\widehat{\mathbf{\Gamma}}))$ or, equivalently, the sub-matrix $\mathcal{T}_R(\mathcal{G}(\widehat{\mathbf{\Gamma}}))$. These schemes estimate the null space matrix \mathbf{Q} (or, equivalently, $\mathcal{T}_R(\mathcal{G}(\widehat{\mathbf{\Gamma}}))\mathbf{Q} = 0$ subject to norm constraints on \mathbf{Q} ; see the literature [11] for details.

Once the \mathbf{Q} is pre-estimated from calibration data, the image is recovered from under-sampled Fourier coefficients by minimizing

$$\arg \min_{\mathbf{\Gamma}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda \left\| \mathcal{T}(\mathcal{G}(\widehat{\mathbf{\Gamma}})) \mathbf{Q} \right\|_F^2. \quad (11)$$

The above optimization problem simplifies solving the system of equations $\mathcal{A}(\mathbf{\Gamma}) = \mathbf{B}$; $\mathcal{G}(\widehat{\mathbf{\Gamma}})\mathbf{Q} = 0$ for specific sampling patterns analytically [8], [10]. In other cases [15], [12], (11) is solved iteratively. Both strategies are computationally efficient since \mathbf{Q} is fully known. However, the need for a calibration region restricts the achievable acceleration.

III. Deep Generalization of SLR Methods

The main focus of this work is to introduce a DL solution to improve the computational efficiency of SLR algorithms. We note that calibrated SLR methods, which learn the linear null space projection operator from calibration data, require few iterations for convergence, thus offering fast image recovery. Calibration-free SLR methods by contrast are computationally expensive. Specifically, because the null space matrix \mathbf{Q} is estimated from the data itself, the algorithm requires several iterations to converge.

We propose to pre-learn a CNN-based null space projector from multiple exemplar datasets. The proposed non-linear CNN module learns to estimate the annihilation relations from the under-sampled data based on its training on exemplar data. We view this approach as learning a non-linear filterbank, which behaves like different linear filterbanks for different images. Specifically, the non-linear filterbank can be approximated as a linear filterbank, which projects the data to the null space, thus annihilating the signal but preserving the noise and alias artifacts; the residual block preserves the signal, while suppressing noise.

A. IRLS Algorithm with Variable Splitting

To facilitate the reinterpretation of the reconstruction scheme as an iterative denoising strategy, we introduce an auxiliary variable $\hat{\mathbf{z}}$ in (8) to obtain a three-variable constrained optimization problem,

$$\arg \min_{\mathbf{\Gamma}, \mathbf{Q}, \hat{\mathbf{z}}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda \left\| \mathcal{T}(\hat{\mathbf{z}})\mathbf{Q} \right\|_F^2 \text{ such that } \hat{\mathbf{z}} = \mathcal{G}(\hat{\mathbf{\Gamma}}). \quad (12)$$

We impose the constraint by a penalty term as

$$\arg \min_{\mathbf{\Gamma}, \mathbf{Q}, \hat{\mathbf{z}}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda \left\| \mathcal{T}(\hat{\mathbf{z}})\mathbf{Q} \right\|_F^2 + \beta \left\| \mathcal{G}(\hat{\mathbf{\Gamma}}) - \hat{\mathbf{z}} \right\|_2^2.$$

This formulation is equivalent to (12) when $\beta \rightarrow \infty$. We propose to solve the above problem using the alternating minimization scheme:

$$\mathbf{\Gamma}_{n+1} = \arg \min_{\mathbf{\Gamma}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \beta \left\| \mathcal{G}(\hat{\mathbf{\Gamma}}) - \hat{\mathbf{z}}_n \right\|_2^2 \quad (13)$$

$$\hat{\mathbf{z}}_{n+1} = \arg \min_{\hat{\mathbf{z}}} \beta \left\| \mathcal{G}(\hat{\mathbf{\Gamma}}_{n+1}) - \hat{\mathbf{z}} \right\|_2^2 + \lambda \left\| \mathcal{T}(\hat{\mathbf{z}})\mathbf{Q} \right\|_F^2. \quad (14)$$

At each step, the \mathbf{Q} matrix is updated as in (10).

1) Image Update: The first step specified by (13) is a simple Tikhonov regularized optimization problem to recover the multi-channel images γ at the $(n+1)$ -th iteration. When $\mathcal{G} = \mathcal{I}$, the prior reduces to $\left\| \hat{\mathbf{\Gamma}} - \hat{\mathbf{z}}_n \right\|_2^2$. In the general case, the solution to this optimization problem can be determined analytically as

$$\hat{\mathbf{\Gamma}}_n = \left(\mathcal{A}^H \mathcal{A} + \beta \mathcal{G}^H \mathcal{G} \right)^{-1} \left(\mathcal{A}^H \mathbf{B} + \beta \mathcal{G}^H (\hat{\mathbf{z}}_{n-1}) \right), \quad (15)$$

when \mathcal{A} involves a sampling in the Fourier domain. Similar analytical solutions can also be used when \mathcal{G} involves a Fourier domain weighting as in the literature [28].

2) Projection : The sub-problem (14) is essentially a proximal operation. Specifically, the second term of (14) is the energy in projecting $\mathcal{T}(\hat{\mathbf{z}})$ to the subspace \mathbf{Q} . If $\lambda \rightarrow \infty$, we obtain $\hat{\mathbf{z}}$ as the projection of $\hat{\mathbf{\Gamma}}_{n+1}$ onto the signal subspace, orthogonal to \mathbf{Q} .

B. Filterbank Interpretation of the Denoising Subproblem

We will now focus on the denoising sub-problem by showing its linear filterbank structure. We will capitalize on this structure to generalize the algorithm. We will focus on the vertical and horizontal stacking cases separately.

1) Vertical stacking considered in Section II-B1: Consider the term $\mathcal{F}(\hat{\mathbf{Z}})\mathbf{q}_i$, where \mathbf{q}_i is one of the columns of the matrix \mathbf{Q} . When the lifting operation is described by (4), we have

$$\mathcal{F}(\hat{\mathbf{Z}})\mathbf{q}_i = \begin{bmatrix} \mathcal{H}(\hat{\mathbf{z}}_1) \\ \mathcal{H}(\hat{\mathbf{z}}_2) \end{bmatrix} \mathbf{q}_i = \begin{bmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \end{bmatrix} \quad (16)$$

Because $\mathcal{H}(\hat{\mathbf{z}})$ is a Hankel matrix, $\mathcal{H}(\hat{\mathbf{z}})\mathbf{Q}$ corresponds to the linear convolution between $\hat{\mathbf{z}}$ and \mathbf{Q} . Since convolution is commutative, we can rewrite the above expression as

$$\mathcal{F}(\hat{\mathbf{Z}})\mathbf{q}_i = \underbrace{\begin{bmatrix} \hat{\mathbf{z}}_1 \\ \hat{\mathbf{z}}_2 \\ \hat{\mathbf{z}} \end{bmatrix}}_{\hat{\mathbf{Z}}} \mathcal{P}(\mathbf{q}_i), \quad (17)$$

where, $\mathcal{P}(\mathbf{q}_i)$ is a block Hankel matrix constructed from the samples of \mathbf{q}_i . We thus have $\|\mathcal{F}(\hat{\mathbf{Z}})\mathbf{Q}\|^2 = \|\hat{\mathbf{Z}}\mathcal{F}(\mathbf{Q})\|^2$, where $\mathcal{F}(\mathbf{Q})$ is obtained by horizontally stacking the matrices $\mathcal{P}(\mathbf{q}_i)$. We note that $\hat{\mathbf{z}}_1\mathcal{F}(\mathbf{Q})$ corresponds to passing $\hat{\mathbf{z}}_1$ through a single input multiple output (SIMO) filterbank, whose filters are specified by \mathbf{q}_i .

2) Horizontal stacking considered in Section II-B2: Similar to the vertical stacking case, we consider

$$\mathcal{F}(\hat{\mathbf{Z}})\mathbf{q}_i = \underbrace{\begin{bmatrix} \mathcal{H}(\hat{\mathbf{z}}_1) & \dots & \mathcal{H}(\hat{\mathbf{z}}_N) \end{bmatrix}}_{\mathcal{F}(\hat{\mathbf{Z}})} \underbrace{\begin{bmatrix} \mathbf{q}_{i,1} \\ \vdots \\ \mathbf{q}_{i,N} \end{bmatrix}}_{\mathbf{q}_i} \quad (18)$$

$$= \underbrace{\begin{bmatrix} \mathcal{P}(\mathbf{q}_{i,1}) & \dots & \mathcal{P}(\mathbf{q}_{i,N}) \end{bmatrix}}_{\mathcal{F}(\mathbf{Q})} \underbrace{\begin{bmatrix} \hat{\mathbf{z}}_1 \\ \vdots \\ \hat{\mathbf{z}}_N \end{bmatrix}}_{\hat{\mathbf{Z}}} \quad (19)$$

We thus have $\|\mathcal{F}(\hat{\mathbf{Z}})\mathbf{Q}\|^2 = \|\mathcal{F}(\mathbf{Q})\hat{\mathbf{Z}}\|^2$, where

$$\mathcal{F}(\mathbf{Q}) = \begin{bmatrix} \mathcal{P}(\mathbf{q}_{1,1}) & \dots & \mathcal{P}(\mathbf{q}_{1,N}) \\ \vdots & \dots & \vdots \\ \mathcal{P}(\mathbf{q}_{N,1}) & \dots & \mathcal{P}(\mathbf{q}_{N,N}) \end{bmatrix} \quad (20)$$

We note that $\mathcal{F}(\mathbf{Q})\hat{\mathbf{Z}}$ corresponds to passing $\hat{\mathbf{Z}}$ through a multiple input multiple output (MIMO) filterbank, whose filters are specified by \mathbf{q}_i .

C. Approximation of Denoising Sub-problem

We thus rewrite (14) for both lifting approaches as,

$$\widehat{\mathbf{Z}}_{n+1} = \arg \min_{\widehat{\mathbf{Z}}} \beta \|\mathcal{G}(\widehat{\Gamma}_{n+1}) - \widehat{\mathbf{Z}}\|^2 + \lambda \|\mathcal{F}(\mathbf{Q}_n)\widehat{\mathbf{Z}}\|_F^2. \quad (21)$$

which reduces to

$$\widehat{\mathbf{Z}}_{n+1} = \left[\mathbf{I} + \frac{\lambda}{\beta} \mathcal{F}(\mathbf{Q}_n)^H \mathcal{F}(\mathbf{Q}_n) \right]^{-1} \mathcal{G}(\widehat{\Gamma}_{n+1}). \quad (22)$$

We propose to solve the denoising problem approximately. Assuming $\lambda \ll \beta$ and applying first-order Taylor approximation, we obtain an approximate solution for $\widehat{\mathbf{Z}}$ as

$$\widehat{\mathbf{Z}}_{n+1} \approx \underbrace{\left[\mathbf{I} - \frac{\lambda}{\beta} \mathcal{F}(\mathbf{Q}_n)^H \mathcal{F}(\mathbf{Q}_n) \right]}_{\mathcal{L}_n} \underbrace{\mathcal{G}(\widehat{\Gamma}_{n+1})}_{\mathcal{R}_n}. \quad (23)$$

As discussed before, $\mathcal{F}(\mathbf{Q})$ denotes a MIMO or SIMO filterbank, depending on the nature of lifting. The term $\mathcal{F}(\mathbf{Q})^H$ denotes convolution with a flipped version of \mathbf{Q} , often referred to as the deconvolution layer in DL literature. \mathcal{R}_n is a filterbank that projects the signal to the null space, thus killing or annihilating the signal and preserving the noise terms. Thus, the linear operator \mathcal{L}_n is a residual block, which removes the alias or noise terms from the input signal, thus essentially denoising the signal (see Fig. 1).

Note that the filterbank \mathbf{Q}_n has a subscript n since it is updated at each iteration. The joint estimation of \mathbf{Q}_n and reconstruction $\widehat{\Gamma}_n$ results in high computational complexity. On the other hand, calibration-based methods pre-estimate \mathbf{Q} and hence the residual filterbank \mathcal{L} , thus resulting in significantly reduced computational complexity.

D. SLR-inspired Model-based k-space DL

The main disadvantage of the IRLS strategy discussed above is the high computational complexity. Specifically, this iterative approach requires an singular value decomposition (SVD) at each iteration, and thus results in a computationally expensive algorithm. To improve the computational efficiency, we propose to pre-learn a non-linear CNN annihilation filterbank \mathcal{N}_k from exemplar data. The subscript k indicates that the network performs convolutions in k -space. We pose a reconstruction similar to (11):

$$\arg \min_{\Gamma} \left\| \mathcal{A}(\Gamma) - \mathbf{B} \right\|_2^2 + \lambda_1 \left\| \underbrace{(\mathcal{F} - \mathcal{D}_k)}_{\mathcal{N}_k} (\mathcal{G}(\widehat{\Gamma})) \right\|_2^2. \quad (24)$$

Here, \mathcal{N}_k is a CNN that kills or annihilates the signal while preserving the noise or alias terms, which is conceptually similar to \mathcal{R}_n in (23). Thus, the operator $\mathcal{D}_k = \mathcal{F} - \mathcal{N}_k$ can be viewed as a denoiser similar to \mathcal{L}_n in (23).

We propose to pre-learn the parameters of the network from exemplar data. Unlike calibrated schemes that learn a small linear network from a small subset of Fourier data (calibration region), the CNN parameters are learned from several fully sampled exemplar datasets. This approach enables us to learn a larger CNN, which can generalize to other datasets. We hypothesize that this pre-learned non-linear network can behave like a linear projection for each dataset, thereby facilitating their recovery from under-sampled data. Since the parameters of the network do not need to be self-learned, this approach is significantly faster than uncalibrated SLR approaches.

We use an alternating minimization strategy similar to (13) and (14) to minimize (24). The resulting algorithm translates to a recursive network, which alternates between the denoising network \mathcal{D}_k , which removes the noise and alias terms, and data consistency (DC) blocks:

$$\hat{\mathbf{Z}}_n = \mathcal{D}_k(\mathcal{E}(\hat{\mathbf{\Gamma}}_n)) \quad (25)$$

$$\hat{\mathbf{\Gamma}}_{n+1} = (\mathcal{A}^H \mathcal{A} + \lambda_1 \mathcal{E}^H \mathcal{E})^{-1} (\mathcal{A}^H \mathbf{B} + \lambda_1 \mathcal{E}^H \hat{\mathbf{Z}}_n) \quad (26)$$

Similar to [7], we consider K iterations of the above algorithm and unroll the above iterative scheme to obtain a deep network. The unrolled network consists of K number of repetitions of both \mathcal{D}_k and DC blocks with parameters of \mathcal{D}_k being shared across iterations. At each iteration, the *noisy* input $\hat{\mathbf{\Gamma}}_n$ is projected to the signal subspace and hence denoised. The output of \mathcal{D}_k is given by $\mathcal{D}_k(\mathcal{E}(\hat{\mathbf{\Gamma}}_n)) = \mathcal{E}(\hat{\mathbf{\Gamma}}_n) - \mathcal{N}_k(\mathcal{E}(\hat{\mathbf{\Gamma}}_n))$. The output is then fed into the DC block as shown in Fig 2. As discussed previously, this iterative algorithm is similar to an alternating scheme to solve (11), with the distinction that the linear convolution-deconvolution block is replaced by a non-linear CNN. Unlike the setting in (11), where the filter parameters are learned from the calibration data of each dataset, we propose to pre-learn a CNN from exemplary data.

E. Hybrid Regularized DL

The SLR methods exploit the redundancies in k-space resulting from specific structures in the signal. However, the image patches in MR images often exhibit extensive redundancy, which is exploited in our MoDL scheme [7] as well as other image domain methods [29], [5], [6]. These priors are complementary to the SLR priors discussed in the previous section. We propose to modify the cost function in (24) as

$$\arg \min_{\mathbf{\Gamma}} \left\| \mathcal{A}(\mathbf{\Gamma}) - \mathbf{B} \right\|_2^2 + \lambda_1 \left\| \mathcal{N}_k(\mathcal{E}(\hat{\mathbf{\Gamma}})) \right\|_2^2 + \lambda_2 \left\| \mathcal{N}_I(\mathbf{\Gamma}) \right\|_2^2. \quad (27)$$

Here, \mathcal{N}_I and \mathcal{N}_k are two residual CNNs. The alternating minimization of this scheme results in the following steps:

$$\Theta_n = \mathcal{D}_k(\mathcal{E}(\hat{\mathbf{\Gamma}}_n)) \quad (28)$$

$$\Phi_n = \mathcal{D}_1(\Gamma_n) \quad (29)$$

$$\widehat{\Gamma}_{n+1} = (\mathcal{A}^H \mathcal{A} + \lambda_1 \mathcal{G}^H \mathcal{G} + \lambda_2 \mathcal{F})^{-1} (\mathcal{A}^H \mathbf{B} + \lambda_1 \mathcal{G}^H \Theta_n + \lambda_2 \Phi_n) \quad (30)$$

as shown in Fig 3. The \mathcal{D}_k relies on annihilation relations in k-space, while \mathcal{D}_1 exploits the image domain priors. We propose to learn the parameters of the CNNs \mathcal{D}_k and \mathcal{D}_1 using exemplary data.

F. Special Cases

We show applications of our proposed methods in both single-channel and multi-channel settings, and show that the image sub-problem can be solved analytically in both cases. This approach will accelerate the training and testing procedures.

1) Piecewise Constant Image Structure: The GIRAF [17] algorithm is an SLR scheme that exploits the piecewise constant nature of images, as described in Section II-B1. Here, the operator $\mathcal{G}(\hat{\gamma}) = \widehat{\nabla} \gamma$ as defined in (3). In this case, we have

$$\mathcal{G}^H \begin{pmatrix} \hat{\mathbf{z}}_1 \\ \hat{\mathbf{z}}_2 \end{pmatrix} [\mathbf{k}] = - (j2\pi k_x \hat{\mathbf{z}}_1[\mathbf{k}] + j2\pi k_y \hat{\mathbf{z}}_2[\mathbf{k}]) \quad (31)$$

$$\mathcal{G}^H \mathcal{G}(\hat{\gamma})[\mathbf{k}] = 4\pi^2 (k_x^2 + k_y^2) \hat{\gamma}[\mathbf{k}] \quad (32)$$

Note that the matrix $(\mathcal{A}^H \mathcal{A} + \lambda_1 \mathcal{G}^H \mathcal{G} + \lambda_2 \mathcal{F})$ in (30) can be viewed as a weighting operator in the Fourier domain in the single-channel setting. We can thus solve (30) analytically.

2) Parallel MRI Acquisition: In a parallel MRI setting, $\mathcal{G} = \mathcal{F}$ and hence the data consistency term simplifies to $(\mathcal{A}^H \mathcal{A} + (\lambda_1 + \lambda_2) \mathcal{F})^{-1} (\mathcal{A}^H \mathbf{B} + \lambda_1 \Theta_n + \lambda_2 \Phi_n)$. The term $(\mathcal{A}^H \mathcal{A} + (\lambda_1 + \lambda_2) \mathcal{F})$ is separable across the channels. Hence, one can independently solve for each channel of $\widehat{\Gamma}_{n+1}$ in the Fourier domain in an analytical fashion.

IV. Implementation Details

A. Datasets

The datasets used for single channel experiments were multi-coil k-space of knee from (www.mridata.org) and multi-coil brain from the Calgary-Campinas Public (CCP) dataset [30] in (<https://sites.google.com/view/calgary-campinas-dataset>). The CCP consists of 12-channel T1-weighted brain MR datasets acquired on a 3T scanner. It is a 3D acquisition that allows undersampling along two directions (phase and slice encoding). We used the single channel complex valued images provided by the organizers, which were generated by multi-channel coil combination. The Fourier Transform (FT) was applied to obtain k-space samples from the coil combined images. Since the frequency encoding dimension is fully

sampled, we performed an IFFT along this dimension and considered the recovery of each 2D slice independently. We chose twenty subjects for training, five for validation, and ten for testing. The other dataset consisting of multi-channel knee data of twenty subjects was acquired with a 3D fast spin echo (FSE) sequence on a 3T scanner. The parameters set for the scan were: repetition time $TR = 1550$ ms, echo time $TE = 25$ ms, and a flip angle of 90° . There are 256 sagittal slices and 320 coronal slices per subject with matrix sizes of 320×320 and 320×256 , respectively, at a slice thickness of 0.5 mm. A coil combination of the 8-channel knee k-space data was performed using principal component analysis (PCA). Specifically, we performed a PCA along the coil dimension and picked the first component along the coil dimension as the single-channel complex image. This coil compression preserved on average about 90% of the energy from multi-channel data. Fifteen subjects were used for training, two for validation, and the remaining three for testing.

A set of experiments were done to study the workings of the k-space Deep-SLR scheme for single-channel MRI recovery. The NIfTI formatted T2-weighted brain datasets from the Human Connectome Project (HCP) [31] were used. The T2-weighted brain images were acquired by a Siemens 3T MR scanner using a 3D Cartesian spin-echo sequence. The TR and TE parameters were 3200 ms and 565 ms, respectively, while the matrix size was 320×256 with a field of view (FOV) of 224×224 mm².

Parallel MRI experiments were performed on multi-channel brain and knee datasets. The knee dataset [6] is a multi-slice 2D dataset consisting of 15-channel slices from 20 subjects with roughly 40 slices per subject. The slices are of dimension $640 \times 368 \times 15$. Twelve subjects were used for training, one for validation, and the remaining seven for testing. The data was under-sampled by varying density along the phase encodes. Brain MRI was collected from nine subjects at University of Iowa Hospitals and Clinics using a 3D T2 CUBE sequence with Cartesian readouts using a 12-channel head coil. There are 140 3D slices per subject with dimensions $12 \times 256 \times 232$. We used five subjects for training, one for validation, and the remaining three for testing.

In both cases, the fully sampled complex k-space data was under-sampled and used for training. The complex image obtained by evaluating the IFFT of the individual coil data was used as ground truth in training and testing.

B. Quality Evaluation Metric

We quantitatively evaluate the recovered images in terms of signal-to-noise ratio (SNR) and structural similarity (SSIM) index. The SNR of an image is computed as

$$\text{SNR} = 20 \cdot \log_{10} \left(\frac{\|\mathbf{x}_{\text{rec}}\|_2}{\|\mathbf{x}_{\text{org}} - \mathbf{x}_{\text{rec}}\|_2} \right),$$

where \mathbf{x}_{org} and \mathbf{x}_{rec} are original ground truth and reconstructed images, respectively.

C. Architecture of the CNNs

The modular nature of the proposed scheme allows us to use any residual CNN architecture to define the prior. A key difference with the approach in Fig. 1 is that the CNN parameters are fixed and do not change with iterations as in Fig. 1.(b). The pre-learning of the CNN

parameters using exemplar data allows us to significantly reduce the number of alternating steps compared to the self-learning strategy in Fig. 1. Image domain CNN \mathcal{N}_I is structurally identical to the Fourier domain CNN \mathcal{N}_k , with an equal number of parameters. The residual block \mathcal{D}_I performs an IFFT that feeds spatial domain input $\mathbf{\Gamma}$ to the CNN \mathcal{N}_I and transforms the residual output back to k-space by a fast Fourier transform (FFT) operation. For implementation purposes, we split the real and imaginary parts of the input k-space data into real and imaginary components, which are fed as two channels. The two output channels are combined to recreate the complex output k-space data. The models were implemented

1) Single-channel Case: We use a residual UNET as \mathcal{N}_k in the single-channel setting for the proposed k-space Deep-SLR (K-DSLRL) scheme. We use its modified version of UNET with only 12 layers (two pooling and unpooling operations). The number of filters per layer grows from 64 to a maximum of 256. The UNET operates on single-channel Fourier data ($M=1$). For the proposed hybrid scheme H-DSLRL, the number of parameters in both the UNETs were halved layer by layer to keep the total similar to the K-DSLRL network for fair comparisons.

2) Parallel MRI (Multi-channel Case): A residual five-layer MIMO CNN \mathcal{N}_k as shown in Fig. 2.(b) is used as the k-space network in K-DSLRL. The input and output channels of the network are adjusted according to the dataset. For example, $M=12$ and $M=15$ channels are set for multi-channel brain and knee data, respectively. Each convolution layer consists of $64 \ 3 \times 3$ filters, followed by ReLU non-linearity. The number of filters per layer was halved to 32 for both the CNNs in H-DSLRL compared to 64 in K-DSLRL for fair comparison.

We trained the unrolled recursive network for different iterations of K . $K=10$ was found to be the best-performing model on test data for both cases and the performance saturated afterwards. We were constrained by 16 GB GPU memory, which restricted us from going beyond 15 iterations. The regularization parameters were fixed at $\lambda_1 = \lambda_2 = 1$ for all the experiments. The weights were Xavier initialized and trained for 500 epochs with an Adam optimizer to reduce the mean square error (MSE) at a learning rate of 10^{-4} . All the DL models were implemented using Tensorflow version 1.15. The proposed $K=10$ iteration models for single and multi-channel cases took 5 and 10 hours for training respectively. The source code for the proposed H-DSLRL scheme on multi-channel MRI datasets can be viewed and downloaded from the github link: <https://github.com/anikpram/Deep-SLR>.

D. State-of-the-art Methods for Comparison

We compare our scheme for single-channel recovery against the SLR algorithm (GIRAF) [17], a k-space UNET (K-UNET) [23], and an image domain UNET (I-UNET). The K-UNET is a direct DL approach with a 20-layer 2D UNET in k-space without a DC step. It accepts a real image formed by concatenation of real and imaginary parts of 2D complex k-space. The I-UNET is the spatial version of K-UNET where learning is performed in spatial domain. The I-UNET structure and its number of parameters are exactly the same as in K-

UNET. These networks were trained and tested on single-channel knee and brain datasets described in Section IV-A.

In the parallel MRI setting, we compare the proposed scheme with MoDL [7], K-UNET [23], and the calibration-less parallel SLR algorithm, which motivated our proposed scheme. K-UNET is also a multi-channel calibration-less direct DL approach in k-space without a DC step [23]. Its structure is similar to single-channel K-UNET, with the only difference being the multi-channel input and output. MoDL [7] is a pre-calibrated approach that uses coil sensitivity information and spatial domain regularization. The coil sensitivities for MoDL were estimated using ESPiRiT [10]. All the parallel MRI methods were evaluated on the brain and knee datasets mentioned in Section IV-A.

V. Experiments and Results

Experiments were done on multiple datasets for both single-channel sparse MRI and parallel MRI recovery. Additional experiments were also done on diffusion MRI recovery and are discussed in the supplementary material.

A. Single-channel Signal Recovery

Comparisons of the proposed single-channel schemes against state-of-the-art methods are shown in Fig. 4 for the CCP dataset described in Section IV-A. We observe that the k-space Deep-SLR (K-DSLRL) approach in Fig. 4.(d) provides results that are comparable to the model-based GIRAF [17] method in Fig. 4.(b). By contrast, the direct inversion based I-UNET and K-UNET provides lower performance, even though the number of trainable parameters are larger. Of these, the K-UNET provides slightly lower errors. The improved performance of K-DSLRL over K-UNET may be attributed to the model-based approach, which repeatedly enforces DC. Fig. 4.(f) corresponds to H-DSLRL, which uses both k-space and image domain priors. The H-DSLRL scheme significantly reduces errors. The number of parameters in this model is similar to the one in Fig. 4.(d) since the number of output channels of each intermediate layer is halved. However, the addition of the complementary prior significantly reduces the errors. A similar set of experiments were also performed on the single-channel coronal and sagittal view of knee images described in Section IV-A. The comparisons are shown and discussed in the supplementary paper. The quantitative results of all the experiments are recorded in Table S1 in the supplementary section.

B. Annihilation Operators on Piecewise Constant Images

Single-channel Deep-SLR scheme is used to study its inner workings and its similarity to classical SLR methods in Fig. 5. Note that k-space SLR methods for single-channel MRI schemes [15], [27] learn linear annihilation relations in k-space. As shown in the literature [15], [27], the SLR penalty in (11) is a weighted ℓ_2 norm of the gradients of the image, where the weights correspond to the sum of squares (SOS) of the estimated null space filters. Specifically, the linear annihilation operator has several linearly independent null space vectors; the sum of squares of the IFFT of the null space vectors yield zeros in the location of the gradients in the single-channel setting as shown in the literature [17]. The SLR scheme estimates annihilation relations from under-sampled data using an optimization

strategy. By contrast, the proposed scheme learns to estimate the annihilation relations from under-sampled measurements based on its training on exemplar data.

The solutions provided by the unconstrained setting considered in this paper and [7] are similar to constrained setting in [32], where the formulation in (24) is replaced by

$$\gamma^* = \arg \min \|\mathcal{N}_k(\mathcal{G}(\hat{\gamma}))\|^2 \text{ such that } \|\mathcal{A}(\gamma) - \mathbf{b}\|^2 \leq \sigma^2, \quad (33)$$

where σ is the noise variance; see [32] for detailed performance comparisons of the constrained and unconstrained formulations. In this case, the data consistency layer specified by (30) gets modified as

$$\gamma_{n+1} = \arg \min \|\mathcal{G}(\hat{\gamma}) - \hat{\mathbf{z}}_n\|^2 \text{ such that } \|\mathcal{A}(\gamma) - \mathbf{b}\|^2 \leq \sigma^2. \quad (34)$$

When \mathcal{A} satisfies the restricted isometry conditions [33], then

$\epsilon \|\gamma - \gamma^*\|^2 \leq \mathcal{A}(\gamma - \gamma^*) \leq \delta \|\gamma - \gamma^*\|^2$, where ϵ and δ are the restricted isometry property (RIP) constants. Thus,

$$\|\gamma_{n+1} - \gamma^*\|^2 \leq \frac{\sigma^2}{\epsilon}, \quad (35)$$

where γ^* is the true solution. This relation implies that at each iteration n , the input to the network γ_n is within a σ^2/ϵ ball of the true solution γ^* .

We note that an arbitrary non-linear function can be approximated by its first-order Taylor series representation in a small neighborhood. Our hypothesis is that the first order Taylor series approximation of the non-linear annihilation block $\mathcal{N}_k(\mathcal{G}(\hat{\gamma}))$ within the σ^2/ϵ ball around γ^* closely matches the linear annihilation relations in SLR schemes. Specifically, the annihilation filters would kill the high gradients, while preserving the noise. The use of this annihilation filterbank within the residual block, results in preserving the true signal while suppressing the noise-like perturbations.

In order to test this hypothesis, small random perturbations of variance $\sigma = 0.01$ are added to a given image γ^* and the corresponding output perturbations are analyzed; the sum of squares of the corresponding perturbations is an indicator of the response of the annihilation operator. We consider a piecewise constant image in Fig. 5, which was derived from an image from the HCP dataset (described in Section IV-A), by thresholding. The CNN network with the same architecture as above (12-layer UNET as \mathcal{N}_k) are trained using piecewise constant brain images from 10 training subjects, also obtained by thresholding the HCP data. Following training, random perturbations are added to a new dataset and its k-space data is passed through the network.

The sum of squares of the IFFT of the outputs for 1000 realizations are evaluated, which are shown in Fig. 5.(d). Note that the zeros of the SOS output function closely mimics the SOS function in Fig. 5.(c). This behaviour is observed across a wide variety of testing slices unseen by the trained network, as shown in Fig. 5 and also in Fig. S1 in the supplemental

document, which justifies its generalizability. The experiment strengthens our hypothesis that the proposed network behaves as a linear projector similar to classical SLR schemes for each image γ^* . While similar results are observed for natural images, it is difficult to visualize this due to the large dynamic range. We note that without additional constraints on the weights, one cannot guarantee that the Lipschitz constants of the network is bounded for all inputs, including adversarial perturbations.

C. Parallel MRI Recovery

Proposed multi-channel schemes are compared against state-of-the-art calibration-less and calibrated schemes in Figures 6 and 7, and Table I. The methods have been tested on 420 3D brain slices collected from three subjects. The same set of methods have also been tested on approximately 300 (seven subjects) 3D knee slices. Similar to the single-channel case, the performance of the multi-channel K-DSLRL is comparable to the parallel SLR (PSLR) scheme. The k-space network exhibits some residual aliasing in the knee example in Fig. 6. (d), which can be attributed to the highly structured/uniform nature of sampling. Note that the data was acquired with a calibration region, which the iterative PSLR scheme seems to have benefited from, even though we did not explicitly rely on a calibrated approach. The table reveals that the proposed H-DSLRL outperformed the multi-channel PSLR and K-UNET [23] and that it is slightly better than the pre-calibrated approach MoDL [7]. Note that MoDL is a calibrated scheme, which requires the explicit knowledge of the coil sensitivities. The coil sensitivities were estimated from the fully sampled region in the center of k-space for both knee (Fig. 6) and brain (Fig. 7) experiments using ESPIRiT [10]. The calibration-less methods compared here (PSLR, the proposed method and K-UNET) perform an interpolation in k-space without explicit knowledge of the coil sensitivities. The addition of the image domain prior (H-DSLRL in Fig. 6.(f)) is found to suppress the artifacts and provide reconstructions that are comparable to the MoDL scheme. The proposed Deep-SLR scheme facilitates the recovery of the images without the knowledge of the coil sensitivities. This approach thus eliminates the potential mismatch between the calibration scans for the estimation of the coil sensitivities and the main scan in approaches that rely on an extra calibration scans. By removing the need for an explicit calibration region, this approach enables higher acceleration factors. An additional study of the robustness of our proposed approach to acceleration factors for both knee and brain datasets is presented in the supplementary section of this paper.

D. Benefits over Calibrated Approaches

Pre-calibrated approaches, which estimate coil sensitivities from calibration scans, suffer from motion-induced mismatch between the calibration and main scans, resulting in artifacts. We study the benefit of the uncalibrated deep SLR methods using a simulation. Specifically, we simulate a mismatch by modulating the k-space data of the accelerated scan with a linearly varying phase term, which corresponds to a shift in image domain. A phase shift of 5 pixels along horizontal as well as vertical direction was applied on the 2D slices, assuming a minor physical motion during scan would lead to a similar amount of shift in either direction. We compare the pre-trained MoDL and H-DSLRL framework on this data, whose results are shown in Fig. 8(a)–(d). Due to the mismatch between coil images and the corresponding sensitivities, there are visible striped artifacts in the MoDL reconstruction. By

contrast, we observe that the the proposed hybrid DSLR framework remains unaffected. This simulation study shows the benefit of our proposed method over calibrated setting during motion.

Self-calibrated approaches do not require an additional calibration scan and hence are not sensitive to the above motion errors. They instead leverage a fully sampled calibration region (center of k-space) to estimate the coil sensitivities. However, this approach restricts the achievable acceleration rates. While the acceleration rate can be increased by reducing the size of the calibration region, the smaller calibration region results in inaccurate sensitivity estimates. The sensitivities were estimated from a 24×24 region in our MoDL scheme. We now estimate the sensitivities using ESPIRiT [10] from a calibration window of 16×16 . The pre-trained MoDL was tested on the dataset using those estimated sensitivities. As seen in our experimental results in Fig. 8(e)–(h), the inaccurate sensitivities resulted in several visible artifacts in the MoDL reconstructions. The proposed method does not suffer from these artifacts since it is an uncalibrated scheme and does not rely on the central k-space region to estimate the sensitivities.

E. Comparison of the Computational Complexity

A key benefit of the proposed Deep-SLR scheme over SLR methods is the quite significant reduction in runtime, along with the improved performance offered by the combination of the image domain prior. The recorded runtimes are shown in Table II. We report runtimes for 10 iterations ($K = 10$) of our proposed k-space and hybrid Deep-SLR algorithms, and MoDL. We note that the DL approaches are roughly a few thousands-fold faster than the IRLS-SLR schemes in both cases. As discussed previously, SLR methods estimate the linear projection operator on the fly and require at least 50 iterations to converge. The high complexity of the SVD and the evaluation of the Gram matrix, along with the large number of iterations, is the main reason for the long runtime of the SLR methods. By contrast, the Deep-SLR approaches pre-learn the CNNs from exemplar data, which eliminates the need for (10). The hybrid Deep-SLR approach is slightly slower than k-space Deep-SLR in both the cases since the former uses two CNNs compared to one by the latter even if the effective number of parameters are the same. In a single-channel setting, although K-UNET and I-UNET have more learnable parameters, these approaches are faster by virtue of a single iteration rather than multiple iterations in proposed schemes. Note that the iterative approach brings improved performance as discussed in the previous sections. In the parallel MRI setting, the Deep-SLR schemes use five-layer CNNs that make them faster than K-UNET even after multiple iterations. We note that the MoDL scheme uses a multi-channel forward model that requires a conjugate gradient (CG) algorithm to enforce DC, which makes it slower than the Deep-SLR schemes. By contrast, the proposed scheme recovers the coil images; the forward model only includes Fourier sampling, which makes these schemes faster in training and testing.

VI. Discussion and Conclusion

We introduced a general model-based DL framework to significantly accelerate SLR matrix-completion algorithms. The key distinction with SLR methods is the pre-learning of the

CNN parameters from exemplar data. Since the parameters need not be estimated from the measured data itself, the proposed algorithm is faster by several orders of magnitude. In addition, an additional image domain prior helps to further improve performance. We showed the utility of the proposed scheme in two representative applications with drastically different lifting structure.

In most cases considered in this work, the performance of the k-space network is comparable or better than the corresponding PSLR scheme. The addition of the image domain network further improved performance. The hybrid DSLR outperforms the existing pre-calibrated MoDL scheme in the parallel MRI setting. However, the performance of the k-space DSLR scheme is marginally lower than the corresponding SLR scheme in the single-channel brain case. Additional experiments on larger datasets are needed to understand whether this is a consistent observation. The proposed framework is applicable in theory to a wide range of SLR priors described in earlier work [11]. In this study, we restricted our attention to three representative applications. The applicability of the proposed framework to other problem settings is beyond the scope of this work and will be considered elsewhere. The MSE was used as the loss to train the networks. Since perceptual metrics such as SSIM are related to the MSE in a non-linear fashion, the performance of the proposed networks with respect to the SSIM may be better or worse with respect to others. The training can be changed to use arbitrary loss metrics, including SSIM, which may yield more visually pleasing images than the ones trained using MSE loss. Most of the experiments in this paper were restricted to scans on the same scanners. More work is needed to determine its utility in a multi-scanner and multi-center setting. We have not addressed the design of the sampling scheme that is optimal for the problem in this work. We refer the readers to our recent work that focuses on this aspect [34].

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work is supported by grant NIH 1R01EB019961-01A1.

References

- [1]. Lustig M, Donoho D, and Pauly JM, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007. [PubMed: 17969013]
- [2]. Pruessmann KP, Weiger M, Scheidegger MB, and Boesiger P, "SENSE: sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999. [PubMed: 10542355]
- [3]. Doneva M, "Mathematical models for magnetic resonance imaging reconstruction: An overview of the approaches, problems, and future research areas," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 24–32, 2020.
- [4]. Sun J, Li H, Xu Z et al., "Deep ADMM-Net for compressive sensing MRI," in *Advances in neural information processing systems*, 2016, pp. 10–18.

- [5]. Schlemper J, Caballero J, Hajnal JV, Price AN, and Rueckert D, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE transactions on Medical Imaging*, vol. 37, no. 2, pp. 491–503, 2017. [PubMed: 29035212]
- [6]. Hammernik et al., "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018. [PubMed: 29115689]
- [7]. Aggarwal HK, Mani MP, and Jacob M, "MoDL: Model-based deep learning architecture for inverse problems," *IEEE transactions on medical imaging*, vol. 38, no. 2, pp. 394–405, 2018. [PubMed: 30106719]
- [8]. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, and Haase A, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002. [PubMed: 12111967]
- [9]. Lustig M and Pauly JM, "SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space," *Magnetic Resonance in Medicine*, vol. 64, no. 2, pp. 457–471, 2010. [PubMed: 20665790]
- [10]. Uecker M, Lai P, Murphy MJ, Virtue P, Elad M, Pauly JM, Vasanawala SS, and Lustig M, "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]
- [11]. Jacob M, Mani MP, and Ye JC, "Structured Low-Rank Algorithms: Theory, magnetic resonance applications, and links to machine learning," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 54–68, 2020.
- [12]. Haldar JP and Setsompop K, "Linear Predictability in Magnetic Resonance Imaging Reconstruction: Leveraging shift-invariant fourier structure for faster and better imaging," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 69–82, 2020.
- [13]. Haldar JP, "Low-rank modeling of local k-space neighborhoods (loraks) for constrained mri," *IEEE transactions on medical imaging*, vol. 33, no. 3, pp. 668–681, 2014. [PubMed: 24595341]
- [14]. Mani M, Jacob M, Kelley D, and Magnotta V, "Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS)," *Magnetic Resonance in Medicine*, vol. 78, no. 2, pp. 494–507, 2017. [PubMed: 27550212]
- [15]. Ongie G and Jacob M, "Super-resolution MRI using finite rate of innovation curves," in *IEEE Int. Symp. Bio. Imag IEEE*, 2015, pp. 1248–1251.
- [16]. Lee D, Jin KH, Kim EY, Park S-H, and Ye JC, "Acceleration of MR parameter mapping using annihilating filter-based low rank hankel matrix (ALOHA)," *Magnetic Resonance in Medicine*, vol. 76, no. 6, pp. 1848–1864, 2016. [PubMed: 26728777]
- [17]. Ongie G and Jacob M, "A fast algorithm for convolutional structured low-rank matrix recovery," *IEEE transactions on computational imaging*, vol. 3, no. 4, pp. 535–550, 2017. [PubMed: 29911129]
- [18]. Shin et al., "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," *Magnetic Resonance in Medicine*, vol. 72, no. 4, pp. 959–970, 2014. [PubMed: 24248734]
- [19]. Aggarwal HK, Mani MP, and Jacob M, "MoDL-MUSSELS: Model-based deep learning for multishot sensitivity-encoded diffusion mri," *IEEE transactions on medical imaging*, 2019.
- [20]. Hu Y, Shi X, Tian Q, Guo H, Deng M, Yu M, Moran C, Yang G, McNab J, Daniel B et al., "Reconstruction of multi-shot diffusion-weighted MRI using unrolled network with U-nets as priors," in *Proceedings of the 27th Annual Meeting of ISMRM*, 2019.
- [21]. Pramanik A, Aggarwal H, and Jacob M, "Off-the-grid model based deep learning (O-MODL)," in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019) IEEE*, 2019, pp. 1395–1398.
- [22]. Pramanik A, Aggarwal H, and Jacob M, "Calibrationless parallel mri using model based deep learning (c-modl)," in *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI) IEEE*, 2020, pp. 1428–1431.
- [23]. Han Y, Sunwoo L, and Ye JC, "k-space deep learning for accelerated MRI," *IEEE transactions on medical imaging*, 2019.

- [24]. Eo T, Jun Y, Kim T, Jang J, Lee H-J, and Hwang D, "KIKI-net: cross-domain convolutional neural networks for reconstructing under-sampled magnetic resonance images," *Magnetic Resonance in Medicine*, vol. 80, no. 5, pp. 2188–2201, 2018. [PubMed: 29624729]
- [25]. Akçakaya M, Moeller S, Weingärtner S, and Urbil K, "Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging," *Magnetic Resonance in Medicine*, vol. 81, no. 1, pp. 439–453, 2019. [PubMed: 30277269]
- [26]. Ye JC, Han Y, and Cha E, "Deep convolutional framelets: A general deep learning framework for inverse problems," *SIAM Journal on Imaging Sciences*, vol. 11, no. 2, pp. 991–1048, 2018.
- [27]. Ongie G, Biswas S, and Jacob M, "Convex recovery of continuous domain piecewise constant images from nonuniform fourier samples," *IEEE Transactions on Signal Processing*, vol. 66, no. 1, pp. 236–250, 2017. [PubMed: 30140146]
- [28]. Ongie G and Jacob M, "Off-the-Grid Recovery of Piecewise Constant Images from Few Fourier Samples," *SIAM on Imag. Sci*, vol. 9, no. 3, pp. 1004–1041, 2016.
- [29]. Lee D, Yoo J, Tak S, and Ye JC, "Deep residual learning for accelerated MRI using magnitude and phase networks," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 9, pp. 1985–1995, 2018. [PubMed: 29993390]
- [30]. Souza R, Lucena O, Garrafa J, Gobbi D, Saluzzi M, Appenzeller S, Rittner L, Frayne R, and Lotufo R, "An open, multi-vendor, multfield-strength brain MR dataset and analysis of publicly available skull stripping methods agreement," *NeuroImage*, vol. 170, pp. 482–494, 2018. [PubMed: 28807870]
- [31]. Essen et al., "The Human Connectome Project: a data acquisition perspective," *Neuroimage*, vol. 62, no. 4, pp. 2222–2231, 2012. [PubMed: 22366334]
- [32]. Tamir JI, Stella XY, and Lustig M, "Unsupervised deep basis pursuit: Learning reconstruction without ground-truth data," in *Proc. Intl. Soc. Mag. Reson. Med*, vol. 27, 2019, p. 0660.
- [33]. Candes EJ et al., "The restricted isometry property and its implications for compressed sensing," *Comptes rendus mathématique*, vol. 346, no. 9–10, pp. 589–592, 2008.
- [34]. Aggarwal HK and Jacob M, "J-MoDL: Joint model-based deep learning for optimized sampling and reconstruction," *IEEE Journal of Selected Topics in Signal Processing*, 2020.




Fig. 1. Illustration of the network structure of the IRLS algorithm used in structured low-rank algorithms: (a) shows the linear residual convolutional-deconvolutional block, which projects the signal at the n^{th} iteration to the signal subspace; (b) illustrates the network structure of the SLR algorithm, which alternates between the projection and the data consistency block.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

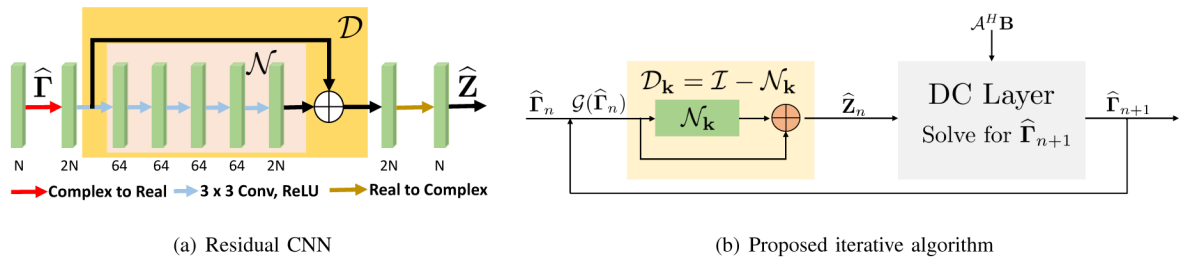


Fig. 2. Network structure of the proposed recursive CNN in k-space, described in Section III-D. The main difference of the proposed scheme with the approach in Fig. 1 is the use of the deep residual CNN in (b), instead of the linear convolution-deconvolution block in Fig. 1.(a).




Fig. 3. Hybrid network: It consists of two identically structured residual CNNs $\mathcal{D}_k, \mathcal{D}_I$ for k-space and image domain learning, respectively. The \mathcal{D}_I block learns redundancies in patches, while \mathcal{D}_k block exploits k-space annihilation relations. The \mathcal{D}_I block does an inverse fast Fourier transform (IFFT) on input k-space Γ and passes it to the residual CNN. The residual image output is transformed back to k-space Γ by a fast Fourier transform (FFT). The output of \mathcal{D}_k and \mathcal{D}_I at the n^{th} iteration are denoted by Θ_n and Φ_n according to (28), (29) and (30). The parameters are not shared between \mathcal{D}_k and \mathcal{D}_I . Both \mathcal{D}_k and \mathcal{D}_I in hybrid Deep-SLR have half the number of feature maps per layer compared to \mathcal{D}_k in k-space Deep-SLR to keep number of trainable parameters the same in both networks for fair comparison. The network parameters are shared across iterations similar to the MoDL [7] framework.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Fig. 4. Reconstruction results of 4x accelerated single-channel brain data. SNR (dB)/PSNR (dB)/SSIM values are reported for each case. The data was under-sampled using a cartesian 2D non-uniform variable-density mask. The top row shows reconstructions (magnitude images), while the bottom row shows corresponding error images. The additional image domain prior in **H-DSLRL** ensures significant improvement in performance over other schemes.

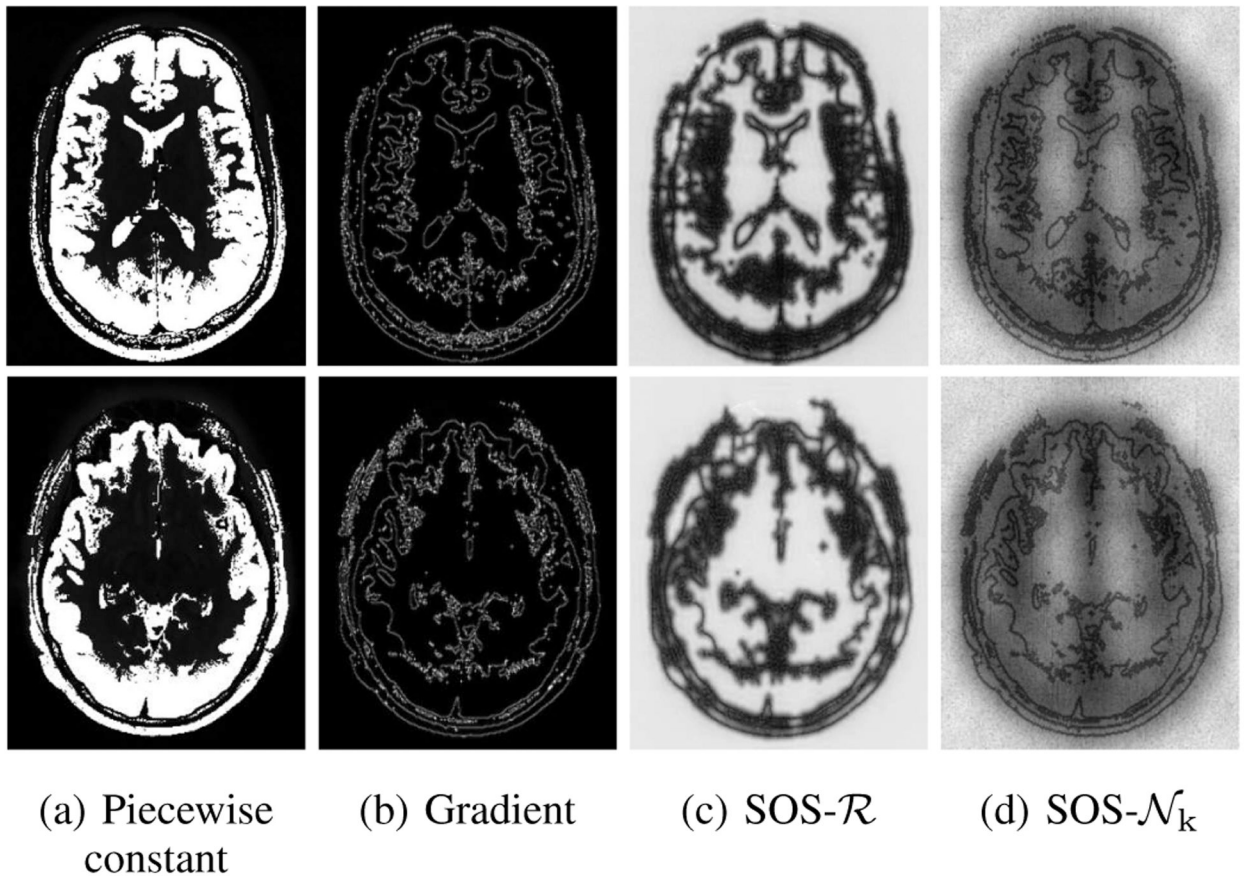
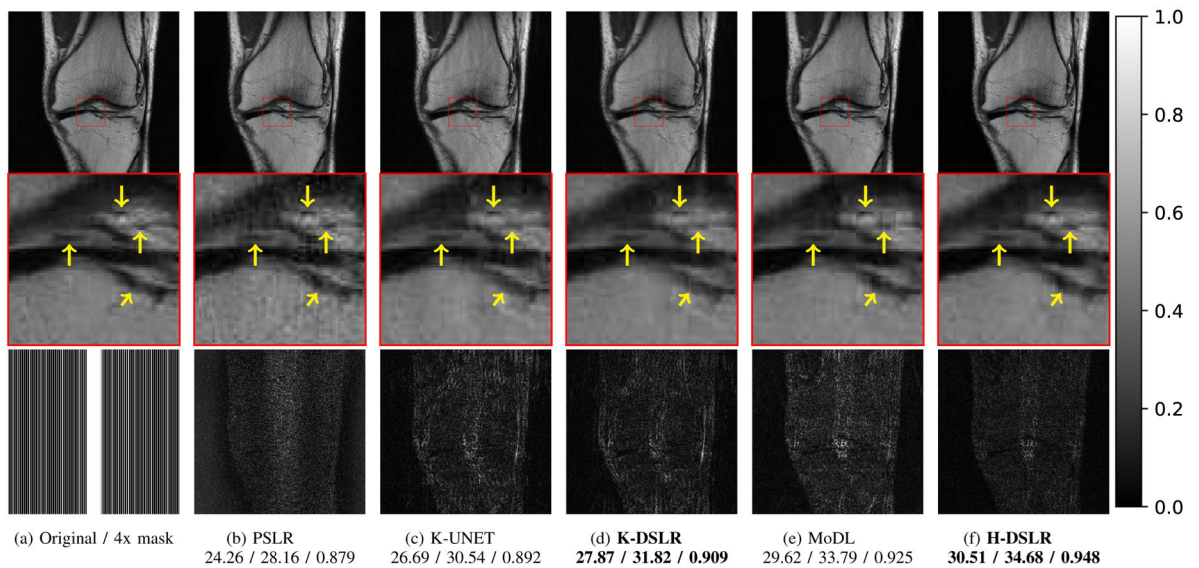


Fig. 5. Illustration of the non-linear and linear annihilation operators. Piecewise constant images and their gradients are shown in (a) and (b), respectively. The non-linear block \mathcal{N}_k behaves like a linear projector to the null space for each image. Pseudo-random perturbations of small magnitude are added to the gradients of the image and fed to \mathcal{N}_k . The SOS of the output perturbations are shown in (d). The SOS function on \mathcal{N}_k closely mimics the linear operator \mathcal{R} in (c). Specifically, it annihilates or kills the gradient components close to the edge locations while preserving the noise components far from the edges. We show more results on different slices in the supplementary material (see Fig. S1).

**Fig. 6.**

Reconstruction results of 4x accelerated 15-channel knee data. A 2D Cartesian structured under-sampling along phase encodes was done. The top row displays reconstructions (SOS), and the bottom row shows corresponding error images. The yellow arrows in the zoomed cartilage region show minute details better preserved by the proposed scheme over other state-of-the-art methods. The numbers are SNR reported in dB. The k-space Deep-SLR scheme **K-DSLRL** yields comparable results to the parallel SLR scheme. The addition of the image domain prior further improves performance. We show H-DSLRL reconstructions of different slices with different accelerations factors in the supplementary material (see Fig. S5).

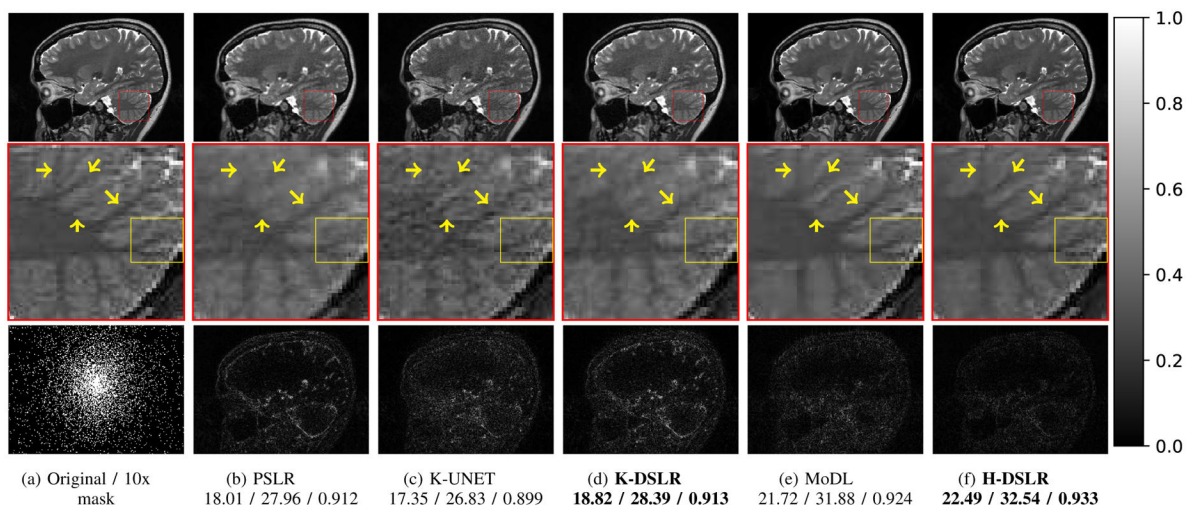



Fig. 7. Reconstruction results of 10x accelerated 12-channel brain data. SNR (dB)/PSNR (dB)/SSIM values are reported for each case. The under-sampling pattern was chosen to be a 2D Cartesian non-uniform variable density. The top row images are reconstructions (SOS), while the bottom row shows corresponding error images. The yellow arrows in the zoomed cerebellum region show minute details better preserved by the proposed scheme than by other state-of-the-art methods. The **K-DSLRL** scheme has errors of lower magnitude than the calibration-less k-space methods PSLR and K-UNET. The proposed hybrid scheme **H-DSLRL** performs comparably to the pre-calibrated approach MoDL. We show H-DSLRL reconstructions of different slices with different accelerations factors in the supplementary material (see Fig. S4).

**Fig. 8.**

The top row of images (a)-(d) show comparisons of pre-calibrated MoDL with the proposed calibration-less approach during mismatches in scans. A cartesian 2D 6-fold under-sampling mask in (b) was used for under-sampling the k-space. The acquired k-space measurements were translated in spatial domain to emulate motion. The MoDL reconstruction shows diagonally striped motion artifacts due to mismatch. Our proposed scheme remains unaffected. The bottom row of images (e)-(h) display comparisons of the proposed approach with self-calibrated MoDL. The mask in (f) is used for under-sampling the k-space data and subsequent reconstruction. It samples 16 fully sampled lines in the center for calibration purposes. The coil sensitivities for MoDL are estimated using ESPIRiT [10] from the calibration window of 16×16 at the center of k-space. The performance of self-calibrated MoDL breaks down due to inaccurate sensitivities estimated from a smaller calibration region. Thus, the requirement of a larger calibration region limits acceleration. Our proposed scheme is robust to acceleration in the calibration region, thus pushing it further.

TABLE I

Quantitative comparison of PSLR, MoDL, proposed, and UNET reconstructions in terms of SNR (dB), PSNR (dB) and SSIM. The bold-faced methods are the proposed ones.

Signal-to-Noise Ratio (SNR)			
Brain			Knee
Acceleration	6x	10x	4x
Methods	SNR	SNR	SNR
PSLR	21.02 ± 2.33	18.12 ± 2.58	24.26 ± 2.12
K-UNET	19.58 ± 2.01	17.28 ± 1.98	26.81 ± 2.05
K-DSLRL	21.58 ± 1.74	18.71 ± 1.83	27.87 ± 1.36
MoDL	23.30 ± 1.53	21.63 ± 1.62	29.77 ± 1.19
MoDL	23.30 ± 1.53	21.63 ± 1.62	29.77 ± 1.19
H-DSLRL	24.34 ± 1.15	22.20 ± 1.23	30.57 ± 0.96
Peak Signal-to-Noise Ratio (PSNR)			
Brain			Knee
Acceleration	6x	10x	4x
Methods	PSNR	PSNR	PSNR
PSLR	31.17 ± 2.30	28.21 ± 2.61	28.19 ± 2.03
K-UNET	29.49 ± 1.96	27.14 ± 1.91	30.94 ± 2.14
K-DSLRL	31.66 ± 1.77	28.55 ± 1.84	31.67 ± 1.33
MoDL	33.43 ± 1.51	31.53 ± 1.57	33.77 ± 1.19
H-DSLRL	34.46 ± 1.22	32.31 ± 1.23	34.69 ± 0.98
Structural Similarity (SSIM)			
Brain			Knee
Acceleration	6x	10x	4x
Methods	SSIM	SSIM	SSIM
PSLR	0.942 ± 0.035	0.918 ± 0.041	0.873 ± 0.027
K-UNET	0.920 ± 0.029	0.883 ± 0.030	0.887 ± 0.021
K-DSLRL	0.938 ± 0.018	0.913 ± 0.023	0.904 ± 0.011
MoDL	0.951 ± 0.020	0.921 ± 0.026	0.928 ± 0.015
H-DSLRL	0.958 ± 0.011	0.935 ± 0.013	0.944 ± 0.008

TABLE II

Comparison of Single-channel and Parallel MRI reconstruction times. The reported values are average reconstruction times per subject in minutes. The bold-faced methods are the proposed ones.

Single-channel recovery (minutes per subject)					
Organ	GIRAF	K-UNET	K-DSL R	I-UNET	H-DSL R
Knee/Brain	197.33	0.07	0.32	0.07	0.37
Parallel MRI recovery (minutes per subject)					
Organ	PSLR	K-UNET	K-DSL R	MoDL	H-DSL R
Brain	1223	0.7	0.17	0.83	0.19
Knee	3106.67	2.83	0.63	4.40	0.75

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript