



Published in final edited form as:

IEEE Signal Process Mag. 2020 January ; 37(1): 54–68. doi:10.1109/msp.2019.2950432.

Structured Low-Rank Algorithms: Theory, Magnetic Resonance Applications, and Links to Machine Learning

Mathews Jacob [Senior Member, IEEE],

University of Iowa, Iowa City, IA 52242

Merry P. Mani,

University of Iowa, Iowa City, IA 52242

Jong Chul Ye [Senior Member, IEEE]

Department of Bio and Brain Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 34141, Republic of Korea

Abstract

In this survey, we provide a detailed review of recent advances in the recovery of continuous domain multidimensional signals from their few non-uniform (multichannel) measurements using structured low-rank matrix completion formulation. This framework is centered on the fundamental duality between the compactness (e.g., sparsity) of the continuous signal and the rank of a structured matrix, whose entries are functions of the signal. This property enables the reformulation of the signal recovery as a low-rank structured matrix completion, which comes with performance guarantees. We will also review fast algorithms that are comparable in complexity to current compressed sensing methods, which enables the application of the framework to large-scale magnetic resonance (MR) recovery problems. The remarkable flexibility of the formulation can be used to exploit signal properties that are difficult to capture by current sparse and low-rank optimization strategies. We demonstrate the utility of the framework in a wide range of MR imaging (MRI) applications, including highly accelerated imaging, calibration-free acquisition, MR artifact correction, and ungated dynamic MRI.

Keywords

Compressed sensing; matrix completion; accelerated MRI; structured low-rank matrix

I. INTRODUCTION

The slow nature of signal acquisition in magnetic resonance imaging (MRI), where the image is formed from a sequence of Fourier samples, often restricts the achievable spatial and temporal resolution in multi-dimensional static and dynamic imaging applications. Discrete compressed sensing (CS) methods provided a major breakthrough to accelerate the magnetic resonance (MR) signal acquisition by reducing the sampling burden. As described

in an introductory article in this special issue [1] these algorithms exploited the sparsity of the discrete signal in a transform domain to recover the images from a few measurements.

In this paper, we review a continuous domain extension of CS using a structured low-rank (SLR) framework for the recovery of an image or a series of images from a few measurements using various compactness assumptions [2]–[22]. The general strategy of the SLR framework starts with defining a *lifting* operation to construct a structured matrix, whose entries are functions of the signal samples. The SLR algorithms exploit the dual relationships between the signal compactness properties (e.g. sparsity, smoothness) and the rank of the lifted matrix. This dual relationship allows recovery of the signal from a few samples in the measurement domain as an SLR optimization problem.

While this strategy may seem contrived, the main benefit of this framework is its remarkable flexibility in exploiting the compactness properties of a variety of signal structures, including continuous-domain real-world signals that classical approaches have difficulty capturing. For example, SLR approaches can recover continuous domain images from a few Fourier measurements with minimal discretization errors, unlike compressed sensing approaches. The SLR methods can account for signals with an infinite number of signal discontinuities that are localized to a curve/surface, which is more general than signals with a finite number of isolated signal discontinuities considered in recent super-resolution methods [23]. In addition, the SLR schemes come with fast algorithms that are readily applicable to large-scale imaging applications such as static and dynamic MRI, unlike super-resolution methods [23] that rely on semi-definite programming. Another example of SLR methods is the recovery of images from their multichannel measurements with unknown sensitivities [5], [9], [13], [18]. These schemes rely on the low-rank structure of a structured matrix, obtained by concatenating block Hankel matrices, which are obtained from each channel image. It is difficult for classical convex compressed sensing algorithms to exploit such complex relations between multichannel measurements. The SLR schemes have also been recently extended to recover an ensemble of images that lie on a smooth surface in high-dimensional space. Current subspace model or a union of subspace model is not efficient in capturing the above property; inspired by kernel methods [24], one can define a non-linear mapping that will transform the smooth surface to a subspace. In SLR approaches, this structure can be exploited by the construction of a structured matrix, whose columns are non-linearly transformed signal samples.

The SLR framework is closely related to the extensive work on linear prediction in MRI, which was often formulated using structure low-rank matrix [2]–[4], [6], [7]. Early work in the context of MRI dates back to 1985 in the context of MR spectroscopy [2], followed by its generalization to MRI by modeling the images as piecewise 1D polynomials by Liang et al. [3]. The finite rate of innovation (FRI) theory [25] also considers the recovery of piecewise polynomials and similar models from a finite number of samples. Multichannel linear predictive relations were introduced in the parallel MRI context in [4], [6], while linear-predictive k-space relations resulting from support and phase constraints were introduced in [7]. The readers are referred to [22] in the same special issue, which is focused on a more elaborate review of linear prediction approaches and their history

in the MRI context. We will review recent development in the low-rank structured matrix approaches for MRI.

1. The early 1-D approaches [3], [25] are generalized to multidimensional continuous domain signals [7], [9], [10]. Theoretical guarantees are also available for multidimensional signals, whose discontinuities are localized to sets of infinite size, but of zero Lebesgue measure (e.g., curves in 2D and surfaces in 3D) [10].
2. Low-rank matrix recovery algorithms with recovery guarantees are used to recover signals from nonuniform measurements [7], [9], [15], as opposed to the explicit annihilating filter estimation from the uniform sampling setting in the classical FRI [3], [25]. Note that an earlier work in this respect goes back to the work by Dologlou et al [26].
3. The framework is generalized to the recovery of signals from non-uniform multichannel measurements with unknown channel responses, facilitating the reduction in samples and calibration-free recovery in a unified matrix completion framework [5], [6], [9], [13], [18].
4. The non-linear generalization of FRI theory enables the recovery of points on smooth surfaces in high-dimensional spaces, which facilitates the joint reconstruction of an ensemble of images from their few measurements [21].
5. The main practical benefit of SLR schemes is their ability to capture a broad range of signal priors, resulting in a wide range of applications, including static MRI [5], [7], [9], [15], [16], dynamic MRI, diffusion MRI [18], MR artifact correction [17], parallel MRI [5], [6], [9], multi-contrast image recovery [19], spectroscopic imaging [11], [27], and field inhomogeneity compensation [28], [29].

The above generalizations come with theoretical recovery guarantees [14], [16], and fast algorithms [9], [15]. In addition, this framework provides rich insights in the deep links between 1-D FRI sampling theory [25], CS [30], low-rank matrix completion, and super-resolution theory [23].

II. OVERVIEW

A. Image acquisition in MRI

In this section, we will briefly describe the image formation in MRI and introduce notations and terminologies used throughout the paper.

1) Single-channel MRI measurements: The image acquisition in MRI constitutes the sampling of the Fourier transform of the image $f(\mathbf{x})$. The measurements in the Fourier domain (also referred to as k-space) are denoted by

$$\hat{f}(\mathbf{k}) = [\mathcal{F}f](\mathbf{k}) = \int f(\mathbf{x})e^{-i2\pi\mathbf{k}^T\mathbf{x}}d\mathbf{x}. \quad (1)$$

Here, $\mathbf{x} \in \mathbb{R}^d$, $d = 2, \dots, 4$ and $\mathbf{k} \in \mathbb{Z}^d$ denote the image domain and k-space coordinates, respectively. The goal of MR image recovery is then to reconstruct $f(\mathbf{x})$ from the above measurements, which are measured on a *sparse* subset of the Fourier domain.

2) Multichannel measurements : Modern scanners acquire the Fourier domain data using multiple receive coils to accelerate the acquisition. These receive coils have different spatial sensitivity patterns, thus providing complementary information. The measurement from the i -th coil is the Fourier transform of the coil-sensitivity weighted image and is denoted by

$$\hat{f}_i(\mathbf{k}) = [\mathcal{F} f_i](\mathbf{k}), \quad \text{where } f_i(\mathbf{x}) = s_i(\mathbf{x})f(\mathbf{x}); \quad i = 1, \dots, N_c. \quad (2)$$

Here, $s_i(\mathbf{x})$ denotes coil sensitivities, $f_i(\mathbf{x})$ denotes the coil-sensitivity weighted image, and N_c is the number of receive coils. The goal of parallel MR image (pMRI) recovery is then to reconstruct $f(\mathbf{x})$ from a few measurements with or without the knowledge of the coil sensitivities.

While the term *multichannel measurements* typically implies the multi-coil measurements described above, we will use the term in a broader sense. As discussed later, different k-space regions are often acquired with slightly different acquisition conditions in MRI. The distortions can often be modeled as spatial weighting terms, similar to coil sensitivities. For example, the signals corresponding to the odd-only and the even-only lines of k-space often differ in phase errors; we consider it a two-channel acquisition with unknown spatial weighting terms.

B. Structured Low-Rank methods: Bird's-eye view

The SLR framework offers a versatile toolbox that exploits a variety of properties of continuous domain multichannel signals without the need for discretization. The SLR algorithms rely on a *lifting* of the original signal to a matrix (see Fig. 1); the structure of the matrix depends on the specific signal properties (e.g., continuous domain sparsity, multichannel relations). The framework relies on the duality relations between the *compactness* of the signal (e.g., sparsity) and the *rank* of the lifted matrix.

Interpolation via structured low-rank matrix completion: In several accelerated MRI acquisitions, the measurements from the full k-space locations are not available. When the signal samples are acquired in a non-uniform fashion, one can rely on a structured low-rank matrix completion to interpolate the missing entries in the lifted matrix. Specifically, this entails determining a matrix with the lowest rank that preserves the structure of the lifted matrix and is consistent with the measured matrix entries. Once the matrix is completed, one can perform inverse lifting to recover the samples of the continuous domain signal (see Fig. 1.(a)). We term this class of methods as interpolation schemes [5], [7], [9], [15].

Extrapolation using two-step SLR algorithms : In some applications, a certain limited region of k-space is acquired without any under-sampling (usually referred to as the

calibration region, typically in the low-frequency regions). The super-resolution approaches in signal processing [23] aim to extrapolate these Fourier coefficients to high-frequency regions, thus recovering the images at higher resolution. The SLR extrapolation strategy is to estimate the null space of the lifted matrix from the fully sampled rows (indicated by the red boxes in Fig. 1.(b)). Given the null space filters, one can (i) estimate a signal model based on the roots of the null-space filters, followed by linear estimation of the unknown signal model parameters, or (ii) perform linear prediction to extrapolate the high-frequency samples from the measured ones subject to the null-space constraints [3], [4], [10].

III. SIGNAL EXTRAPOLATION USING SLR

We will now focus on specific continuous domain signal models and reveal their connection to the lifted matrices that facilitates the recovery of the continuous domain signal. In this section, we will focus on SLR extrapolation where the low-frequency Fourier region is fully sampled as in Fig. 1.(b).

A. FRI theory for piecewise smooth 1D signals

This approach [2], [3], [25] is a generalization of Prony's method to continuous domain 1-D signals. Consider the recovery of r -stream of Diracs f (see Fig. 2.(a)) at locations $x_j, j = 1, \dots, r$ with weights w_j :

$$f(x) = \sum_{j=1}^r w_j \delta(x - x_j), \quad x_j \in [0, 1]. \quad (3)$$

The discrete Fourier transform of the above continuous domain signal is a linear combination of complex exponentials with frequencies $\alpha_j = e^{-i2\pi x_j}$, specified by

$$\hat{f}[k] = \sum_{j=1}^r w_j (\alpha_j)^k, \quad \forall k. \quad (4)$$

The classical Prony's results [25] rely on the annihilation of such exponential signals by

$$\hat{h}(z) = \prod_{j=1}^r (1 - \alpha_j z^{-1}) = \sum_{n=0}^r h[n] z^{-n}, \quad (5)$$

such that the associated $(r+1)$ -tap filter $\hat{h}[n]$ annihilates the Fourier samples \hat{f} :

$$(h * \hat{f})[k] = 0, \quad \forall k \quad (6)$$

This linear convolution relation (6) can be re-expressed in the matrix form (see Fig. 1) to solve for \hat{h} :

$$\underbrace{\begin{bmatrix} \hat{f}[0] & \hat{f}[1] & \cdots & \hat{f}[r] \\ \hat{f}[1] & \hat{f}[2] & \cdots & \hat{f}[r+1] \\ \vdots & \vdots & \ddots & \vdots \\ \hat{f}[r-1] & \cdots & \cdots & \hat{f}[2r-1] \end{bmatrix}}_{\mathbf{H}_{[r+1]}^{[2r]}(\hat{f})} \underbrace{\begin{bmatrix} h[r] \\ \vdots \\ h[0] \end{bmatrix}}_{\mathbf{h}} = 0. \quad (7)$$

Here, the notation $\mathbf{H}_{[r+1]}^{[2r]}(\hat{f})$ represents the linear convolution matrix, which is defined for the $2r$ samples of \hat{f} to be convolved with a filter of size $r+1$; here $[r]$ denotes the set $0, 1, \dots, r$. Note that $\mathbf{H}_{[r+1]}^{[2r]}(\hat{f})$ is a Hankel structured matrix and that the *matrix lifting* thus originates from the linear convolution embedded in (6)-(7) and is illustrated in Fig. 1. The relation (7) can be used to identify \mathbf{h} . Once this linear-prediction filter \mathbf{h} is available, the linear predictive relations (6) can be used to extrapolate the signal samples $\hat{f}[k]$; $k = 0, \dots, 2r - 1$ to any extent in Fourier space.

1) Theoretical Guarantees: Based on (7), the filter \mathbf{h} can be estimated as a null-space vector of the *lifted* matrix as shown in Fig. 1(b). If $2r$ signal consecutive samples are available, the null-space vector is unique [3], [25]. Note from (5) that the roots/zeros of $\hat{h}(z)$ will specify a_j in (4). One can then solve a system of equations with r unknowns to recover the w_j in (3); this estimation is unique with $2r$ measurements.

2) Spatial annihilation relations: The above results can be generalized to recover piecewise constant signal f (see Fig. 2.(a)), with r discontinuities located at t_j ; $j = 1, \dots, r$. The derivative of the signal, denoted by ℓ , is a periodic stream of Diracs as in (3). The above theory can be adapted to this setting by considering the Fourier coefficients of the derivative of $\ell(x)$ as illustrated in Fig. 2(a). The Fourier domain convolution based annihilation relations can also be viewed as space domain multiplication-based annihilation relation

$$\partial f(x) \cdot \hat{h}(x) = 0 \quad (8)$$

as shown in Fig. 2(a). Here, the space domain function $\hat{h}(x)$ is the Fourier transform of the filter $h[n]$ that annihilates the Fourier coefficients of f . The zeros of \hat{h} overlap with the location of Diracs or the location of the discontinuities of the piecewise constant signal [3]. We will now use this spatial domain annihilation relations to extend the 1-D results to higher dimensions.

B. Piecewise smooth signals in higher dimensions

We now review how to extend the classical 1-D FRI theory to recover the piecewise constant images $f(\mathbf{x})$ shown in Fig. 2(b). We assume that the Fourier samples are available in a rectangular region $\Gamma \subset \mathbb{Z}^2$ [10]. Note that the edges of f consist of a set of curves, denoted by \mathcal{C} . Extending the space domain annihilation relation in (8), we assume the edge locations of the image to be represented by the zero level-sets of a 2-D bandlimited function \hat{h} [10]:

$$\mathcal{E} = \{\mathbf{x} | \hat{h}(\mathbf{x}) = 0\} \quad (9)$$

where

$$\hat{h}(\mathbf{x}) = \sum_{\mathbf{k} \in \Lambda} h[\mathbf{k}] \exp(i\mathbf{k}^T \mathbf{x}) \quad (10)$$

is bandlimited to a rectangular region $\Lambda \subset \mathbb{Z}^2$. The bandwidth of \hat{h} , denoted by $|\Lambda|$, is a measure of the complexity of the edge set [10], [31], which generalizes the notion of sparsity in compressive sensing.

Similar to (8), we have space domain annihilation relation (see Fig. 2(b)) specified by $\nabla f \cdot \hat{h} = 0$; the Fourier transform of this relation provides the 2-D Fourier domain annihilation relations similar to (6):

$$\sum_{\mathbf{k} \in \Lambda} \widehat{\nabla f}[\ell - \mathbf{k}] h[\mathbf{k}] = \mathbf{0}, \quad \forall \ell \in \Gamma \ominus \Lambda. \quad (11)$$

Here, Γ is the set of Fourier measurements of f . Note that the convolution of $\widehat{\nabla f}$ with h at a location \mathbf{k} requires the samples of $\widehat{\nabla f}$ within the set $\mathbf{k} + \Lambda$, where the addition amounts to the translation of the set Λ to the location \mathbf{k} . Since we only have the Fourier samples of f within Γ , the convolutions specified by (6) can only be evaluated within the set $\Gamma \ominus \Lambda$, which is the morphological erosion of the set Γ by Λ . Similar to (7), the convolution relation in (11) can be rewritten in the matrix form similar to (7) as

$$\underbrace{\begin{bmatrix} \mathbf{H}_{\Lambda}^{\Gamma}(i\omega_x \hat{f}) \\ \mathbf{H}_{\Lambda}^{\Gamma}(i\omega_y \hat{f}) \end{bmatrix}}_{\mathbf{V}_{\Lambda}^{\Gamma}(\hat{f})} \mathbf{h} = \mathbf{0}, \quad (12)$$

where \mathbf{h} represents the vectorized filter coefficients. Here, $\mathbf{H}_{\Lambda}^{\Gamma}(i\omega_x \hat{f})$ is the matrix corresponding to the 2-D convolution of $i\omega_x \hat{f}$, and $\mathbf{V}_{\Lambda}^{\Gamma}(\hat{f})$ is a composite matrix obtained by stacking the block Hankel matrices vertically as shown in Fig. 3.(a).

1) Theoretical Guarantees: The filter \mathbf{h} , and equivalently the edges of f can be estimated from (12) as the null-space vector of $\mathbf{V}_{\Lambda}^{\Gamma}(\hat{f})$. It is shown in [10, Theorem 1] that $\mathbf{V}_{\Lambda}^{\Gamma}(\hat{f})$ will have a unique null-space vector when the uniform Fourier samples of \hat{f} are available within 3Λ , which is the three-fold dilation of Λ . The piecewise constant signal will then be uniquely identified from its Fourier measurements within 3Λ [10, Theorem 4].

Note that the piecewise constant signal model in this section assumes a finite or infinite number of signal discontinuities, supported on a set of zero measure. This framework is more general than [23], which assumes a finite number of isolated discontinuities. We note that even more general continuous domain models [7], which assume the signal support

to be a region of non-zero measure, do exist. Since such models have infinite degrees of freedom, it is difficult to guarantee the recovery of such a signal from finite measurements.

C. Exponential 1-D signals

Recovery of exponentially decaying signals are of interest in MRI since the temporal evolution of MRI signals in the presence of chemical shifts, field inhomogeneity and multiple relaxation mechanisms can be modeled as a linear combination of exponentials [2], [11], [27]. Specifically, the 1-D exponential signal at the spatial location \mathbf{x} evolving as a function of time can be expressed as

$$f(\mathbf{x}, t) = \sum_{j=1}^r a_j(\mathbf{x})(\beta_j(\mathbf{x}))^t; \quad t = 1, \dots, T, \quad (13)$$

where the goal is to estimate the exponential parameters from their time series. Since the relation in Eq. (5) is true for the above signal model, one can build Hankel matrices $\mathbf{H}_{[r]}^{[n]}(\hat{f}_{\mathbf{x}})$ of the form (7) of each pixel at location \mathbf{x} , analogous to the description in Section III-A. Following (4)-(7), these matrices will have a null space vector parameterized by the exponential parameters. The roots of the null space vector are estimated using singular value decomposition (SVD), which are used to identify the exponential parameters. These subspace strategies based on SVD of $\mathbf{H}_{[r]}^{[n]}(\hat{f}_{\mathbf{x}})$ is widely used in the context of MR spectroscopy [2], [11], [27].

IV. SIGNAL INTERPOLATION USING SLR

This section shows how the extrapolation-based approaches in the previous section can be extended to the non-uniform setting. As discussed in Fig. 1, we will exploit the duality between the rank of the lifted matrices and the compactness priors to fill in the missing entries. Note that such sampling patterns are of high relevance in CS applications. The generalization of SLR can be achieved with the theoretical performance guarantees that is similar to that of the standard CS approaches [30].

A. Piecewise smooth 1-D signals

Consider the signal model in Section III-A with r Diracs. Consider a Hankel structured matrix $\mathbf{H}_{[d]}^{[n]}(\hat{f}) \in \mathbb{C}^{(n-d) \times d}$ with $d > r$. Then, the authors in [14] showed that

$$\text{rank}(\mathbf{H}_{[d]}^{[n]}(\hat{f})) = r. \quad (14)$$

Similar observation was made in early work [3]. The authors of [14] further showed that (14) holds for general FRI signals with the minimum annihilating filter size of $r+1$ on its spectral domain. This property is useful for recovering sparse signals from Fourier measurements. Specifically, let χ be a multi-set consisting of random indices from $\{0, \dots, n-1\}$ such that $|\chi| = m < n$. Then, as shown in Fig. 1.(a), we can interpolate $\hat{f}[k]$ for all $k \in \{0, \dots, n-1\}$ from the sparse Fourier samples exploiting (14):

$$\underset{\hat{g} \in \mathbb{C}^n}{\text{minimize}} \|\mathbf{H}_{[d]}^{[n]}(\hat{g})\|_* \text{ subject to } P_{\chi}(\hat{g}) = P_{\chi}(\hat{f}). \quad (15)$$

where $\|\cdot\|_*$ denotes the matrix nuclear norm. Here, P_{χ} denotes a projection operation to the set χ . Instead of using the nuclear norm, a similar problem was solved by nonconvex optimization method [7], [26].

1) Theoretical Guarantees: If the sampling pattern satisfies the incoherence condition with the parameter ρ , then there exists a constant c such that \hat{f} is the unique minimizer to (15) with probability $1 - 1/n^2$ [14], provided

$$|\chi| \geq c \rho r \log^{\alpha} n, \quad (16)$$

where $\alpha = 2$ if the discontinuities are located on uniform grid; $\alpha = 4$ for general continuous domain signals. This suggests a near-optimal performance similar to the discrete-domain CS approaches [30].

B. Piecewise smooth multidimensional signals

We now consider the recovery of the Fourier samples of the piecewise constant image described in Section III-B on $\Gamma \subset \mathbb{Z}^2$. We consider matrices $\mathbf{V}_{\Lambda_1}^{\Gamma}(\hat{f})$ shown in (12), where $\Lambda \subset \Lambda_1$. It is shown in [10] that

$$r = \text{rank}(\mathbf{V}_{\Lambda_1}^{\Gamma}(\hat{f})) = |\Lambda_1| - |\Lambda_1 \ominus \Lambda|. \quad (17)$$

It is also shown in [10] that there is a *duality* relation between the above rank and the smoothness of the edge set of the images. The low-rank structure of $\mathbf{V}_{\Lambda_1}^{\Gamma}$ enables the recovery of \hat{f} on Γ using SLR.

1) Theoretical Guarantees: The results in [16, Theorem 4] show that the SLR algorithm recovers the true Fourier samples $\hat{f}[\mathbf{k}]; \mathbf{k} \in \Gamma$ from its uniform random measurements on $\chi \subset \Gamma$ with probability exceeding $1 - |\Gamma|^{-2}$, provided

$$|\chi| \geq c \rho c_s r \log^4 |\Gamma|, \quad (18)$$

where r is defined as in (17), $c_s = |\Gamma|/|\Lambda_1|$, c is a universal constant, and $\rho > 1$ is an incoherence measure depending on the geometry of the edge set defined in [16]. Note that the result bears remarkable similarity to (16), in the case of continuous domain signals not on the grid.

C. Exponential signals with spatially smooth parameters

The approaches in Section III-C considered the pixel-by-pixel recovery of exponential parameters at each spatial location \mathbf{x} from their uniform samples. The recovery of

exponentials from non-uniform $k-t$ domain samples is key to accelerating the acquisition in several MR applications. Note that the coefficients of the annihilation filters $h_{\mathbf{x}}[z]$ depend on the exponential parameters at the specific pixel location \mathbf{x} . In imaging applications, the exponential parameters are often spatially smooth. The spatial Fourier transform of the filter coefficients can be assumed to be support limited to a 3-D cube of size Λ . The annihilation relations thus can be compactly expressed as

$$\mathbf{H}_{\Lambda}^{\Gamma}(\hat{f})\mathbf{h} = 0, \quad (19)$$

where $\mathbf{H}_{\Lambda}^{\Gamma}(\hat{f})$ is the block Hankel matrix corresponding to the 3-D convolution of $\hat{f}[\mathbf{k}, n]$ with the filter \mathbf{h} . Similar to the discussions in Sections IV-A and IV-B, the matrix $\mathbf{H}_{\Lambda_1}^{\Gamma}(\hat{f})$ is low-rank if $\Lambda \subset \Lambda_1$. Thus, the recovery of the 3-D data set from its undersampled $k-t$ space measurements can be posed as a structured low-rank problem.

V. RECOVERY OF MULTICHANNEL MRI DATA

We now consider the recovery of multichannel data using calibration-based and calibration-free strategies, which are analogous to the extrapolation and interpolation methods discussed above.

A. Blind multichannel deconvolution

In the blind multichannel deconvolution problems in signal processing, subspace techniques such as the eigenvector-based algorithm for multichannel blind deconvolution (EVAM) [32] rely on multichannel annihilation relations. Specifically, the multichannel measurement for blind multichannel deconvolution problem is given by $\hat{f}_i = \hat{s}_i * \hat{f}$, $i = 1, \dots, N_c$, where N_c denotes the number of channels, \hat{s}_i denotes the unknown convolution kernels, and \hat{f} is the unknown underlying signal. The goal of the blind multichannel deconvolution problem is to estimate \hat{f} and \hat{s}_i from \hat{f}_i .

EVAM [32] relies on the cross-channel annihilating filter relations:

$$\hat{f}_i * \hat{s}_j - \hat{f}_j * \hat{s}_i = 0, \quad \forall i \neq j. \quad (20)$$

Assuming that the convolution kernel \hat{s}_j can be represented by the d -tap filter, the above convolution relations can be compactly represented in the matrix form as

$$\left[\mathbf{H}_{\Lambda}^{\Gamma}(\hat{f}_i) \mid \mathbf{H}_{\Lambda}^{\Gamma}(\hat{f}_j) \right] \begin{bmatrix} \hat{s}_j \\ -\hat{s}_i \end{bmatrix} = 0. \quad (21)$$

We now show that the extension of this work to the MRI setting is similar to the popular parallel MRI schemes [5], [33], which were discovered independently without realizing this connection.

B. Calibration-based parallel MRI recovery

The space domain annihilation relations $f_i(\mathbf{x}) \cdot s_j(\mathbf{x}) - f_j(\mathbf{x}) \cdot s_i(\mathbf{x}) = 0$ corresponding to (20) were used in [34].

The Fourier domain annihilation relations between each pair of channels as in (21) implies that the matrix formed by linearly stacking the Hankel matrices of f_i , $i = 1, \dots, N_c$

$$\mathbf{L}_\Lambda^\Gamma(f_1, \dots, f_{N_c}) = [\mathbf{H}_\Lambda^\Gamma(\widehat{f}_1) | \mathbf{H}_\Lambda^\Gamma(\widehat{f}_2) | \dots | \mathbf{H}_\Lambda^\Gamma(\widehat{f}_{N_c})] \quad (22)$$

is low-rank [4]–[6], [9], [13], [18], [26].

In calibration-based parallel MRI, one often acquires the central Fourier regions in a fully sampled fashion. In this case, one can estimate the null space \mathbf{V} from the central regions similar to the approach in Fig. 1.(b). Once the null space filters are available, the relation $\mathbf{L}_\Lambda^\Gamma \mathbf{V} = 0$ can be used to complete the full matrix (22) from its partial entries. Consider one column of \mathbf{V} , denoted by $\mathbf{v} = [\mathbf{w}_1^T | \mathbf{w}_2^T | \dots | \mathbf{w}_{N_c}^T]^T$; each vector \mathbf{w}_i is of dimension $|\Lambda|$. The relation $\mathbf{L}_\Lambda^\Gamma(f_1, \dots, f_{N_c})\mathbf{v} = 0$ can thus be written as

$$\widehat{f}_1 * w_1 + \widehat{f}_2 * w_2 + \dots + \widehat{f}_{N_c} * w_{N_c} = 0, \quad (23)$$

where $*$ denotes linear convolution. The above equation implies that the k-space sample of the first coil can be predicted as the linear combination of the nearby samples in all the channels, which is essentially the simultaneous auto calibrating and k-space estimation (SAKE) model [5], [33]. The generalized auto calibrating partially parallel acquisitions GRAPPA approach, which preceded SAKE [5], [6], can be viewed as an approximation of this linear prediction relation. See [13] for a more detailed review of similar multichannel methods.

C. Calibration-less multichannel recovery

The above calibration-based framework is extended to non-uniformly sampled k-space data in [5], [6] using the SLR framework, analogous to the interpolation strategies discussed in Section IV. The recovery can be interpreted as the rank-deficiency of $\mathbf{L}_\Lambda^\Gamma(f_1, \dots, f_c)$.

Accordingly, [5], [9], [13] use the the following low-rank matrix completion to estimate the missing k-space data without any calibration data:

$$\min_{[g_1 \dots g_{N_c}]} \text{rank}(\mathbf{L}_\Lambda^\Gamma(g_1, \dots, g_{N_c})) \text{ subject to } P_{\mathcal{X}}(\widehat{g}_i) = P_{\mathcal{X}}(\widehat{f}_i); \quad i = 1, \dots, N_c.$$

To further exploit the annihilation relationship from the FRI model in addition to the multichannel annihilating relationship, the authors in [9] solve the optimization problem (24) after the k-space weighting.

VI. SLR RECOVERY: FAST ALGORITHMS

We will explain the algorithms in the context of a simple Hankel matrix lifting. Earlier algorithms [5]–[7] relied on non-convex rank constrained formulations:

$$\min_{\hat{g} \in \mathbb{C}^n} \left\| P_{\chi}(\hat{g}) - P_{\chi}(\hat{f}) \right\|_2^2 \text{ such that } \text{rank} \left(\mathbf{H}_{[d]}^{[n]}(\hat{g}) \right) \leq r, \quad (24)$$

where r is a pre-defined parameter. Here, P_{χ} denotes a projection operation to the set χ . Convex unconstrained formulations of the form

$$\min_{\hat{g} \in \mathbb{C}^n} \lambda \left\| \mathbf{H}_{[d]}^{[n]}(\hat{g}) \right\|_* + \left\| P_{\chi}(\hat{g}) - P_{\chi}(\hat{f}) \right\|_2^2. \quad (25)$$

where $\| \cdot \|_*$ denotes the matrix nuclear norm and λ denotes the regularization parameter are recently being used in applications due to their ability to converge to global minima.

A. Iterative singular value shrinkage

The formulation in (25) can be solved using the singular value thresholding scheme [8]. Specifically, the algorithm alternates between lifting the original signal \hat{f} to form $\mathbf{H}(\hat{f})$, followed by singular value soft shrinkage of $\mathbf{H}(\hat{f})$ to obtain \mathbf{Q} and un-lifting \mathbf{Q} to impose the Hankel structure and enforce data consistency. A challenge with this scheme is the high computational complexity and memory demand, especially for multichannel recovery problems.

In [5], [7], the authors used the hard thresholding as a singular value shrinkage operation within a similar iterative optimization framework.

B. UV factorization

To reduce the computational cost, the authors in [9], [35] used the following observation:

$$\| \mathbf{A} \|_* = \min_{\mathbf{U}, \mathbf{V}: \mathbf{A} = \mathbf{U}\mathbf{V}^H} \left(\left\| \mathbf{U} \right\|_F^2 + \left\| \mathbf{V} \right\|_F^2 \right), \quad (26)$$

to realize an SVD-free structured rank minimization algorithm. The nuclear norm minimization problem can be solved by replacing the nuclear norm term by the above relation. The constraints are enforced using an alternating direction method of multipliers (ADMM) algorithm, whose complexity is determined by the rank. So for sparse signals, a significant computation gain can be obtained.

C. Generic iterative reweighted annihilating filter (GIRAF) algorithm

Both of the above strategies rely on the explicit lifting of the signal to the large Hankel matrix, which makes the memory demand and computational complexity of the algorithms high. The GIRAF algorithm avoids the lifting steps altogether and reduces the number of SVD steps; it solves the problem in the original signal domain, thus realizing an algorithm that is comparable in computational complexity to compressed sensing methods [15]. The

GIRAF algorithm uses the iterative re-weighted least-squares [36] strategy, which majorizes the nuclear norm penalty as $\|\mathbf{A}\|_* \leq \|\mathbf{A}\mathbf{Q}\frac{1}{2}\|_F$, where $\mathbf{Q} = (\mathbf{A}^*\mathbf{A})^{-\frac{1}{4}}$. The i^{th} iteration of this algorithm updates \mathbf{Q} as $\mathbf{Q}_{i+1} = (\mathbf{H}_{[d]}^{[n]}(\hat{g})^* \mathbf{H}_{[d]}^{[n]}(\hat{g}_i))^{-\frac{1}{4}}$, followed by the minimization of a quadratic cost function to solve for g_{i+1} obtained by replacing the nuclear norm in (25) by $\|\mathbf{H}_{[d]}^{[n]}(\hat{g})\mathbf{Q}_n\|_F^2$. The GIRAF algorithm relies on fast Fourier transforms to evaluate the matrix vector product between $\mathbf{H}_{[d]}^{[n]}(\hat{g})$ and each column of \mathbf{Q} ; this approach exploits the convolutional structure of $\mathbf{H}_{[d]}^{[n]}(\hat{g})$. We note that this approach does not require the creation or storage of the large structured matrix, but works directly with the signal samples of g . The MATLAB implementation of the GIRAF algorithm is available at <https://github.com/cbig-iowa/giraf>.

D. Algorithms for non-convex SLR formulations

All of the above algorithms are designed for convex formulations (25). The early algorithms [5], [6] relied on successively solving

$$g_{n+1} = \min_{\hat{g} \in \mathbb{C}^n} \left\| P_{\mathcal{X}}(\hat{g}) - P_{\mathcal{X}}(\hat{f}) \right\|_2^2 + \lambda \|\hat{g} - M_r(\hat{g}_n)\|^2, \quad (27)$$

where $M_r(\hat{g}_n)$ is the approximation of the n^{th} iterate \hat{g}_n such that $\mathbf{H}_{[d]}^{[n]}(\hat{g})$ is of rank r .

Recently, the multiplicative majorization strategy in the GIRAF algorithm [15] has been adapted into the LORAKS setting to realize faster algorithms.

VII. GENERALIZATION TO MACHINE LEARNING

A. Recovery of point clouds on surfaces: non-linear generalization of union of subspaces model

We will now consider a non-linear generalization of the FRI theory discussed above, which will facilitate the recovery of surfaces or points living on surfaces from a few noisy samples. The main motivation is the joint recovery of an ensemble of images (e.g. images in a cardiac MRI time series) from their noisy and undersampled measurements. This approach is a non-linear generalization of the popular union of subspaces model, which represents the images as a sparse linear combination of some basis images.

We model the surface as the zero level set of a bandlimited function as in Section IV-B. For simplicity, we will explain the approach in 2D. Consider an arbitrary point \mathbf{x} on the curve $\hat{h}(\mathbf{x}) = 0$, where $h(\mathbf{x})$ is specified by (10). The space domain annihilation relation translates to $\mathbf{c}^T \phi_{\Lambda}(\mathbf{x}) = 0$, where

$$\phi_{\Lambda}(\mathbf{x}) = \left[\exp(j\mathbf{k}_1^T \mathbf{x}) \quad \dots \quad \exp(j\mathbf{k}_{|\Lambda|}^T \mathbf{x}) \right]^T \quad (28)$$

is a non-linear mapping or lifting of a point \mathbf{x} to a high-dimensional space [21] (see Fig. 4) of dimension $|\Lambda|$. Since this non-linear lifting strategy is similar to feature maps used in kernel methods [24], we term $\phi_\Lambda(\mathbf{x})$ as the feature map of the point \mathbf{x} .

Let us now consider a set of N points on the curve, denoted by $\mathbf{x}_1, \dots, \mathbf{x}_N$. The above annihilation relations can be compactly represented as $\mathbf{c}^T \Phi_\Lambda(\mathbf{X}) = 0$, where

$$\Phi_\Lambda(\mathbf{X}) = [\phi_\Lambda(\mathbf{x}_1) \quad \phi_\Lambda(\mathbf{x}_2) \quad \dots \quad \phi_\Lambda(\mathbf{x}_N)] \quad (29)$$

is the feature matrix of the points. The results in [21] show that if the number of samples exceeds a bound that depends on the complexity of the curve, it can be uniquely estimated.

We have shown in [21] that if we choose a feature map with a bandwidth of Λ_1 such that $\Lambda \subset \Lambda_1$, the feature matrix satisfies the relation similar to (17):

$$\text{rank}(\Phi_{\Lambda_1}(\mathbf{X})) \leq |\Lambda_1| - |\Lambda_1 \ominus \Lambda|. \quad (30)$$

Generalizing the SLR approach in Section IV, we recover the samples on the curve/surface from a few measurements corrupted by noise [21] by solving an optimization problem:

$$\underset{\mathbf{X}}{\text{minimize}} \quad \lambda \|\Phi_{\Lambda_1}(\mathbf{X})\|_* + \|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|^2. \quad (31)$$

We use an iterative-reweighted least-squares (IRLS) algorithm [36] that relies on the *kernel trick* [24] for recovery. Specifically, the algorithm alternates between a quadratic optimization scheme and the evaluation of the graph Laplacian matrix; the algorithm may be interpreted as a the discretization of the manifold by a graph, whose connectivity is specified by the distances on the manifold/surface. Note that unlike the approaches in Section IV, this algorithm is non-convex. Hence, the above optimization problem cannot be guaranteed to achieve the global minimum. Despite the lack of guarantees, the algorithm was able to yield good performance in applications, as seen in Fig. 4 and Fig. 8.

B. Relation to Deep Convolutional Neural Networks

One of the interesting twists of the SLR approach is its relation to deep neural networks. Specifically, in the recent theory of deep convolutional framelets [37], the authors showed that a deep neural network can be interpreted as a framelet representation whose basis can be obtained from Hankel matrix decomposition. Moreover, in a recent follow-up study [38], the authors further revealed that the rectified linear unite (ReLU) plays a key role in making the representation adaptive to input by providing combinatorial basis selection.

Specifically, for a given Hankel matrix $\mathbf{H}(f)$, let us consider two matrices Φ^\top and Ψ whose dimensions are determined such that they can multiplied to the left and the right of the Hankel matrix, respectively. Furthermore, suppose there exists matrices $\tilde{\Phi}$ and their duals $\tilde{\Psi}$ satisfying the frame condition $\tilde{\Phi}\Phi^\top = \mathbf{I}$, $\Psi\tilde{\Psi}^\top = \mathbf{I}$, where the superscript \top denotes the matrix transpose and \mathbf{I} refers to an identity matrix with appropriate size. The existence of such

matrices can be trivially shown as infinite number of orthonormal matrices satisfy the frame condition. In fact, the frame condition allows overcomplete representation. Then, it is easy to show

$$\mathbf{H}(f) = \tilde{\Phi} \mathbf{C} \tilde{\Psi}^T, \quad \text{where} \quad \mathbf{C} = \Phi^T \mathbf{H}(f) \Psi. \quad (32)$$

One of the most interesting observations in [37] is that (32) can be equivalently represented by encoder and decoder convolution structure:

$$\mathbf{C} = \Phi^T (f * \alpha(\Psi)) \quad , \quad f = (\tilde{\Phi} \mathbf{C}) * \beta(\tilde{\Psi}) \quad (33)$$

where $\alpha(\Psi)$ and $\beta(\tilde{\Psi})$ are the encoder and decoder layer multichannel convolution filters obtained by rearranging Ψ and $\tilde{\Psi}$, respectively [37]. This observation led to the findings that Φ^T and $\tilde{\Phi}$ correspond to the pooling and unpooling layers, respectively [37].

However, to satisfy the frame conditions, the number of output filter channels should increase exponentially, which is difficult to satisfy in practice [37]. Moreover, in contrast to SLR approaches, the exact decomposition of the Hankel matrix in (33) is not interesting in neural networks, since the output of the network should be different from the input due to the task-dependent processing. Moreover, the representation should be generalized well for various inputs rather than a specific input at the training phase. In a recent extension of the deep convolutional framelets [37], the authors revealed that the convolutional framelet representation is indeed combinatorial due to the combinatorial nature of ReLU. Specially, for the case of an encoder-decoder convolutional neural network (CNN) without a skipped connection it was shown that the CNN output g can be represented as follows [38]:

$$g = \sum_i \langle \mathbf{b}_i(f), f \rangle \tilde{\mathbf{b}}_i(f), \quad (34)$$

where $\mathbf{b}_i(f)$ and $\tilde{\mathbf{b}}_i(f)$ denote the i th columns of the $\mathbf{B}(f)$ and $\tilde{\mathbf{B}}(f)$ matrices given by

$$\mathbf{B}(f) = \mathbf{E}^1 \Sigma^1(f) \mathbf{E}^2 \dots \Sigma^{\kappa-1}(f) \mathbf{E}^{\kappa}, \quad (35)$$

$$\tilde{\mathbf{B}}(f) = \mathbf{D}^1 \tilde{\Sigma}^1(f) \mathbf{D}^2 \dots \tilde{\Sigma}^{\kappa-1}(f) \mathbf{D}^{\kappa}, \quad (36)$$

where $\Sigma^l(f)$ and $\tilde{\Sigma}^l(f)$ denote diagonal matrices with 0 and 1 values that are determined by the ReLU output in the previous convolution steps; \mathbf{E}^l (resp. \mathbf{D}^l) denotes the l -th layer encoder (resp. decoder) matrix that is composed of pooling matrix Φ^l (unpooling matrix $\tilde{\Phi}^l$) and encoder filter matrix Ψ^l (resp. decoder filter matrix $\tilde{\Psi}^l$). Similar basis representation holds for the encoder-decoder CNNs with skipped connection [38].

When there are no ReLU nonlinearities and pooling and filter matrices satisfy the frame condition for each l , the representation (34) is indeed a convolutional frame representation of f as in (33), ensuring perfect signal reconstruction. Even with the ReLU, there exists conjugate filter sets that can allow perfect reconstruction condition [37]. However, the

explicit dependence of (35) and (36) on the input f due to the ReLU nonlinearity makes the neural network representation much richer and adaptive to different input signals, since the resulting frame representation in (34) depends on specific input. In this regard, the role of ReLU may be similar to the sparsity patterns in compressed sensing MRI, since the sparsity pattern determines the difference basis representation of CS for given inputs. Furthermore, the number of distinct linear representations increases exponentially with the network depth, width, and the skipped connection, thanks to the combinatorial nature of ReLU nonlinearities [38]. This exponentially large *expressivity* of the neural network is another important advantage, which may, with the combination of the aforementioned *adaptivity*, explain the origin of the success of deep neural networks for image reconstruction.

VIII. APPLICATION OF SLR FOR IMPROVED MRI

The main benefit of the SLR framework is its ability to account for a wide variety of signal models, facilitating the exploitation of extensive redundancies between their Fourier samples. This feature makes SLR applicable to a wide variety of MRI applications, listed below.

A. Highly accelerated MRI

During the past few years, different SLR priors were introduced for highly accelerated MRI, each designed to exploit specific signal properties.

1) Support/sparsity priors [7]: The signal is assumed to be sparse or support limited to a region in low-rank modeling of local k -space neighborhoods (LORAKS), which results in annihilation conditions in k -space. The LORAKS scheme formulated the CS-MRI problem by using the block Hankel matrix of the images as the prior. The LORAKS framework also exploits phase constraints, which are detailed in [7].

2) Transform domain sparsity [3], [9], [10]: The annihilating filter-based low-rank Hankel matrix (ALOHA) approach considers general signals that can also be sparse in a transform domain. Specifically, the signal f is modeled as $Lf = w$, where L denotes a constant coefficient linear differential equation, often called the whitening operator [14]:

$$L := b_K \partial^K + b_{K-1} \partial^{K-1} + \dots + b_1 \partial + b_0, \quad (37)$$

and w is an innovation signal composed of a stream of Diracs or differentiated Diracs. Evaluating the Fourier transform, we have $\hat{w} := \hat{l} \hat{f}$ where $\hat{l}(f) = b_K (i2\pi f)^K + \dots + b_1 (i2\pi f) + b_0$ and the associated Hankel matrix $\mathbf{H}(\hat{w})$ from the innovation spectrum \hat{w} becomes rank deficient. The Fourier coefficients are interpolated similar to (15) by minimizing the nuclear norm of a block Hankel matrix whose entries are the Fourier coefficients of \hat{f} weighted by \hat{l} .

3) Piecewise smooth signal model [3], [16]: The GIRAF algorithm generalized the piecewise polynomial 1-D model [3] to recover piecewise constant multi-dimensional signals from their sparse Fourier coefficients by minimizing the nuclear norm of the

vertically stacked block Hankel matrices in (12) (see Figure 3.(a)), as described in Section IV-B. Each of the block Hankel matrices correspond to weighted Fourier coefficients of the signal.

This model was recently extended to represent the image as a linear combination of piecewise constant (f_{pwc}) and piecewise linear components (f_{pwl}) [20]:

$f(\mathbf{x}) = f_{\text{pwc}}(\mathbf{x}) + f_{\text{pwl}}(\mathbf{x})$. The recovery is posed as:

$$\min_{g_1, g_2 \in \mathbb{C}^n} \left\| \mathcal{F}(\widehat{\nabla} g_1) \right\|_* + \left\| \mathcal{S}(\widehat{\nabla}_2 g_2) \right\|_* + \lambda \left\| P_{\Omega}(\widehat{g}_1 + \widehat{g}_2) - P_{\Omega}(\widehat{f}) \right\|^2,$$

where \mathcal{S} is the matrix obtained by the vertical stacking of the block Hankel matrices of three second degree partial derivatives, denoted by the vector $\widehat{\nabla}_2 f$. We demonstrate the benefit of some of the flavors of the proposed scheme in Fig. 5

B. Calibration-free correction of trajectory and phase errors

The multichannel framework introduced in Section V-C provides a versatile tool to solve several calibration problems in MRI, where different k-space regions are acquired from different excitations; these datasets will differ in phase errors, which often manifest as artifacts. As discussed in Section II-A2, (2) can also be used to model the acquisition of multi-shot Fourier data. Specifically, the data is split into multiple groups or virtual channels, each corresponding to different distortions. The joint recovery of these groups or virtual channels are performed as in (24).

1) Correction of phase errors in multi-shot diffusion MRI [18]: Diffusion MRI (DMRI) is a valuable tool for assessing brain connectivity and tissue micro-structure. High-resolution DMRI is often acquired using multi-shot EPI schemes, where different Fourier regions are acquired from different radio-frequency excitation pulses. For example, the even lines in a two-shot EPI sequence are collected from one shot, while the odd lines are collected in the second shot. Subtle physiological motion between the shots in the presence of diffusion encoding gradients manifest as motion induced phase errors between the shots. Previous works [18] have shown that these phase errors can be corrected using the SLR scheme, even when the data is acquired using 4–8 shots. The results in Fig. 6.(a) shows the improved reconstructions offered by this scheme.

2) Compensation of trajectory errors [39]: MRI images can also suffer from artifacts, resulting from k-space data of different excitations experiencing different distortions. A typical example is radial imaging, where the shifted radial spokes cause streaking artifacts. A matrix completion using the low-rank prior [39] is used to jointly recover the artifact-free images and the uncorrupted calibration data. Several other groups have shown the utility of the low-rank based approach for the correction of trajectory errors in EPI and radial setting (not referenced here due to space constraints).

C. Correction of k-space outliers [17], [40]

Many MRI artifacts from the instability of an MR system, patient motion, inhomogeneities of gradient fields, etc. are manifested as outliers in k-space data. Because MR artifacts usually appear as sparse k-space components, the artifact-corrupted MRI measurements $\hat{z}(f)$ can be modeled as $\hat{z}(f) = \hat{x}(f) + \hat{s}(f)$, where $\hat{x}(f)$ is a k-space data of an artifact-free image and $\hat{s}(f)$ is the sparse k-space outlier [17]. If the unknown signal can be sparsified by applying a whitening operator (37) by performing a lifting to a Hankel structured matrix, we can see that

$$\mathbf{H}(\hat{y}) = \underbrace{\mathbf{H}(\hat{l} \odot \hat{x})}_{\text{low-rank}} + \underbrace{\mathbf{H}(\hat{l} \odot \hat{s})}_{\text{sparse}}, \quad (38)$$

where \hat{l} denotes the spectrum of the whitening operator.

Note that the second term in Eq. (38) is sparse, because the lifted Hankel matrix from sparse components is still sparse, as illustrated in Fig. 7. Thus, Eq. (38) becomes a structure for sparse + low-rank decomposition, and a modified version of robust principal component analysis (RPCA) was used to decompose the low-rank and sparse component of the Hankel matrix $\mathbf{H}(\hat{y})$. After the sparse + low-rank decomposition, the weighted k-space for the low-rank component is returned to original k-space by performing an un-weighting [17].

D. Recovery of exponentials

1) Parameter mapping [12], [19]: MR parameter mapping methods, which estimate the T_1 and T_2 relaxation constants, can enable tissue characterization and hence are clinically very valuable. However, the the long scan time resulting from the need for multiple temporal frames makes these schemes challenging for routine clinical use.

Under-sampling each temporal frame, followed by sparse recovery, has been a popular approach to reduce scan time. The spatio-temporal signal can be modeled as an exponential with smoothly varying parameters as described in Section IV-C, which facilitates its recovery from under-sampled measurements.

2) Field inhomogeneity compensation [28], [29]: MRI schemes such EPI are associated with long readouts. The EPI signal at the time point t after the excitation is the Fourier sample of the image modulated by a time-evolving exponential:

$\rho[\mathbf{r}, n] = f(\mathbf{r})\beta(\mathbf{r})^n$, where the exponential parameter $\beta(\mathbf{r}) = e^{-\left(\mathbf{R}_2^*(\mathbf{r}) + j\omega(\mathbf{r})\right)T}$ results from the field inhomogeneity and T is the sampling rate in time; $f(\mathbf{r})$ is the true image. The standard inverse Fourier transform reconstruction ignoring the exponential term thus recovers the field inhomogeneity distorted image. We used the approach described in Section IV-C to recover the entire time series $\rho[\mathbf{r}, n]$ from the Fourier measurements. Post-recovery, the image $\rho[\mathbf{r}, 0]$ is chosen as the un-distorted signal. Since there are two complex unknowns $f(\mathbf{r})$ and $\beta(\mathbf{r})$ at each pixel, we use two shifted EPI readouts to recover the time series [29]. The results of this approach are shown in Fig. 6.(b).

E. Free-breathing and ungated cardiac MRI [21]—In cardiac MRI clinical practice, breath-held cardiac cine MRI is the standard protocol to evaluate cardiac function. Many subject groups cannot tolerate the breath-holds, which disqualifies such patients for cardiac MRI exams. While free-breathing and ungated cardiac MRI is the ideal protocol, a main challenge is the quite significant acceleration needed to facilitate this scheme. The work in [21] used the surface recovery strategy in Section VII-A to recover free-breathing and ungated data from highly undersampled measurements with great success as shown in Fig. 6.(b).

IX. CONCLUSION & FUTURE WORK

The SLR formulation provides a flexible framework to exploit different continuous domain signal priors, which are difficult for current discrete compressive sensing frameworks to capture. The framework comes with theoretical guarantees and fast algorithms and software, making it readily applicable in multidimensional imaging problems including MRI. The recent results showed that there are important links between SLR frameworks and machine learning approaches. The proposed framework is demonstrated in several challenging MRI applications.

In spite of its flexibility and many applications, there are still remaining open challenges, some of which will be discussed below.

Sampling patterns:

The single-channel theoretical results (16) and (18) guarantee the recovery of the Fourier coefficients of the signals on a grid from a randomly chosen subset of coefficients. The results exhibit a \log^4 dependence on the final grid size, suggesting increasing sample demand with larger grid sizes. However, we note from Sections III-A and III-B that once the null space is identified, the signal extrapolation to any grid size is possible from the low-frequency Fourier samples [10, Theorem 1]. This strongly suggests a variable-density sampling approach, which may significantly reduce the sampling demand, and which was also confirmed by the empirical results in [16]. More theoretical work in this area is needed to determine the best sampling patterns. While empirical results demonstrate the great benefit of SLR schemes in multichannel signal recovery from non-uniform samples, this problem is not well studied from a theoretical perspective. However, the SLR framework can still be adapted to work well through the synergistic combination of multiple signal priors (e.g., sparsity and multichannel annihilation priors), achieved by composite lifting [18]. More theoretical work is needed on this front to improve the understanding of this problem.

Theoretical guarantees and noise performance:

Most of the above analysis assumes that the measurements are noise-free and the underlying SLR matrix is low-rank. For example, the multichannel annihilation relationship in (20) is valid only for the noiseless k -space measurements. The model mismatch introduces a slow decay of the singular value spectrum of the Hankel matrix. The above theoretical results are extended to scenarios involving measurement noise in specific SLR models in [14], [16].

However, note that SLR formulations can be obtained from many different redundancies within the signal [5], [7], [18], [29]; the generalization of the theoretical robustness results in [14], [16] to general SLR schemes deserves further investigation.

Links between machine learning and SLR:

Although the aforementioned link between the SLR and deep learning is interesting, there are still open questions, and more theoretical work is necessary to understand this connection. The generalizability, optimization landscape, and expressivity (see the extensive list of references in [38]) in the context of machine-learning based reconstruction are not still well-understood, which may be an important research direction.

Acknowledgement

JCY is supported by National Research Foundation of Korea under Grant NRF-2016R1A2B3008104. MJ is supported by grants NIH 1R01EB019961-01A1 and R01 EB019961-02S1.

Biographies

Mathews Jacob is a professor in the Department of Electrical and Computer Engineering and is heading the Computational Biomedical Imaging Group (CBIG) at the University of Iowa. His research interests include image reconstruction, image analysis and quantification in the context of magnetic resonance imaging. He received his Ph.D degree from the Biomedical Imaging Group at the Swiss Federal Institute of Technology in 2003. He was a Beckman postdoctoral fellow at the University of Illinois at Urbana Champaign between 2003 and 2006. He is the recipient of the CAREER award from the National Science Foundation and the Research Scholar Award from American Cancer Society. He is currently the associate editor of the IEEE Transactions on Medical Imaging and IEEE Transactions on Computational Imaging, and is the general chair of IEEE ISBI 2020. (Email: mathews-jacob@uiowa.edu)

Merry Mani is an Assistant Professor in the Department of Radiology at the University of Iowa. Her research focuses on the development of fast and efficient MRI techniques to study neurological disorders, with a focus on diffusion based microstructure imaging. She received her PhD from the University of Rochester in 2014. She has been a postdoctoral fellow at the Magnetic Resonance Research Facility at the University of Iowa from 2014–2018, developing diffusion MRI technique for ultra-high field imaging. She is a recipient of the Young Investigator Award from the Brain and Behavior Research Foundation. (Email: merry-mani@uiowa.edu)

Jong Chul Ye is currently a Professor of the Dept. of Bio/Brain Engineering and Adjunct Professor at Dept. of Mathematical Sciences of Korea Advanced Institute of Science and Technology (KAIST), Korea. He is currently an associate editor for IEEE Trans. on Medical Imaging, a Senior Editor of IEEE Signal Processing Magazine, and a section editor of BMC Biomedical Engineering. He is a vice-chair of IEEE SPS Technical Committee for Computational Imaging, and a general co-chair (with Mathews Jacob) for 2020 IEEE Symp. on Biomedical Imaging (ISBI), Iowa City. His group was the first place winner of the 2009

Recon Challenge at the ISMRM workshop, and the second place winner at 2016 Low Dose CT Grand Challenge organized by the American Association of Physicists in Medicine (AAPM). His current research interests include machine learning and sparse recovery for various imaging reconstruction problems in x-ray CT, MRI, optics, ultrasound, etc. (Email: jong.ye@kaist.ac.kr)

REFERENCES

- [1]. Doneva M, "An overview of mathematical models for computational MRI," *IEEE Signal Processing Magazine*, 2020.
- [2]. H B, de Beer R, WM B, JH C, and van Ormondt D, "Application of linear prediction and singular value decomposition (LPSVD) to determine NMR frequencies and intensities from the FID." *Magnetic Resonance in Medicine*, vol. 2, no. 1, pp. 86–89, 1985. [PubMed: 3831681]
- [3]. Liang Z-P, Haacke EM, and Thomas CW, "High-resolution inversion of finite Fourier transform data through a localised polynomial approximation," *Inverse Problems*, vol. 5, no. 5, 1989.
- [4]. Zhang J, Liu C, and Moseley ME, "Parallel reconstruction using null operations (pruno)," *Magnetic Resonance in Medicine*, vol. 66, no. 5, 2011.
- [5]. Shin PJ, Larson PE, Ohliger MA, Elad M, Pauly JM, Vigneron DB, and Lustig M, "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]
- [6]. Uecker M, Lai P, Murphy MJ, Virtue P, Elad M, Pauly JM, Vasanawala S, and Lustig M, "ESPIRIT - an eigenvalue approach to autocalibrating parallel MRI: where sense meets grappa," *Magn. Reson. Med*, vol. 71, no. 3, pp. 990–1001, 2014. [PubMed: 23649942]
- [7]. Haldar JP, "Low-Rank Modeling of Local-Space Neighborhoods (LORAKS) for Constrained MRI," *IEEE Trans. on Medical Imaging*, vol. 33, no. 3, pp. 668–681, 2014. [PubMed: 24595341]
- [8]. Qu X, Mayzel M, Cai J-F, Chen Z, and Orekhov V, "Accelerated NMR spectroscopy with low-rank reconstruction," *Angewandte Chemie International Edition*, vol. 54, no. 3, pp. 852–854, 2015. [PubMed: 25389060]
- [9]. Jin KH, Lee D, and Ye JC, "A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix," *IEEE Transactions on Computational Imaging*, vol. 2, no. 4, pp. 480–495, 2016.
- [10]. Ongie G. and Jacob M, "Off-the-grid recovery of piecewise constant images from few Fourier samples," *SIAM Journal on Imaging Sciences*, vol. 9, no. 3, pp. 1004–1041, 2016. [PubMed: 29973971]
- [11]. Cao P, Shin PJ, Park I, Najac C, Marco-Rius I, Vigneron DB, Nelson SJ, Ronen Sabrina M., and Larson PEZ, "Accelerated high bandwidth MR spectroscopic imaging using compressed sensing," *Magnetic Resonance in Medicine*, vol. 76, no. 2, pp. 369–379, 2016. [PubMed: 27228088]
- [12]. Peng X, Ying L, Liu Y, Yuan J, Liu X, and Liang D, "Accelerated exponential parameterization of T2 relaxation with model-driven low rank and sparsity priors," *Magnetic Resonance in Medicine*, vol. 76, no. 6, pp. 1865–78, 2016. [PubMed: 26762702]
- [13]. Haldar JP and Zhuo J, "P-LORAKS: Low-rank modeling of local k-space neighborhoods with parallel imaging data," *Magnetic resonance in medicine*, vol. 75, no. 4, pp. 1499–1514, 2016. [PubMed: 25952136]
- [14]. Ye JC, Kim JM, Jin KH, and Lee K, "Compressive sampling using annihilating filter-based low-rank interpolation," *IEEE Transactions on Information Theory*, vol. 63, no. 2, pp. 777–801, Feb. 2017.
- [15]. 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]
- [16]. 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, 2018. [PubMed: 30140146]

- [17]. Jin KH, Um J-Y, Lee D, Lee J, Park S-H, and Ye JC, "MRI artifact correction using sparse+low-rank decomposition of annihilating filter-based Hankel matrix," *Magnetic Resonance in Medicine*, vol. 78, no. 1, pp. 327–340, 2017. [PubMed: 27464787]
- [18]. 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]
- [19]. Bilgic B, Kim TH, Liao C, Manhard MK, Wald LL, Haldar JP, and Setsompop K, "Improving parallel imaging by jointly reconstructing multi-contrast data," *Magnetic Resonance in Medicine*, vol. 80, pp. 619–632, 2018. [PubMed: 29322551]
- [20]. Hu Y, Liu X, and Jacob M, "A generalized structured low-rank matrix completion algorithm for MR image recovery," *IEEE Transactions On Medical Imaging*, 2018.
- [21]. Poddar S, Mohsin Y, Ansah D, Thattaliyath B, Ashwath R, and Jacob M, "Free-breathing cardiac MRI using bandlimited manifold modelling," *IEEE Trans. on Computational Imaging*, vol. 35, no. 4, pp. 1106–1115, 2018.
- [22]. Haldar JP and Setsompop K, "Linear predictability in MRI reconstruction: Leveraging shift-invariant Fourier structure for faster and better imaging," *IEEE Signal Processing Magazine, Special Issue on Computational MRI (this issue)*, 2020.
- [23]. Candes E. and Granda C, "Towards a mathematical theory of super-resolution," *Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, 2014.
- [24]. Scholkopf B. and Smola AJ, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [25]. Vetterli M, Marziliano P, and Blu T, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [26]. Dologlou I, van Ormondt D, and Carayannis G, "MRI scan time reduction through non-uniform sampling and SVD-based estimation," *Signal Processing*, vol. 55, no. 2, pp. 207–219, 1996.
- [27]. Guo D, Lu H, and Qu X, "A fast low rank hankel matrix factorization reconstruction method for non-uniformly sampled magnetic resonance spectroscopy," *IEEE Access*, vol. 5, pp. 16033–16039, 2017.
- [28]. Nguyen H, Sutton B, Morrison R, and Do M, "Joint estimation and correction of geometric distortions for EPI functional MRI using harmonic retrieval," *IEEE Trans Med Imaging*, vol. 28, pp. 423–34, 2009. [PubMed: 19244014]
- [29]. Balachandrasekaran A, Mani M, and Jacob M, "Calibration-free b0 correction of epi data using structured low rank matrix recovery," *IEEE Transactions on Medical Imaging*, in press.
- [30]. Candes E, Romberg J, and Tao T, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [31]. Pan H, Blu T, and Dragotti PL, "Sampling curves with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 62, no. 2, pp. 458–471, 2014.
- [32]. Gurelli MI and Nikias CL, "EVAM: An eigenvector-based algorithm for multichannel blind deconvolution of input colored signals," *IEEE Transactions on Signal Processing*, vol. 43, no. 1, pp. 134–149, 1995.
- [33]. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, and Haase A, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn. Reson. Med*, vol. 47, no. 6, pp. 1202–1210, 2002. [PubMed: 12111967]
- [34]. Morrisson R, Jacob M, and Do M, "Multichannel estimation of coil sensitivities in parallel MRI," *IEEE International Symposium on Biomedical Imaging*, 2007.
- [35]. Srebro N, "Learning with matrix factorizations," Ph.D. dissertation, Dept. of Elect. Eng., Comput. Sci., Massachusetts Inst. of Technol., Cambridge, MA, USA, 2004.
- [36]. Mohan K. and Fazel M, "Iterative reweighted algorithms for matrix rank minimization," *Journal of Machine Learning Research*, vol. 13, no. Nov, pp. 3441–3473, 2012.
- [37]. Ye J, 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.
- [38]. Ye JC and Sung WK, "Understanding geometry of encoder-decoder CNNs," in *Proceedings of the 36th International Conference on Machine Learning*, ser. *Proceedings of Machine Learning*

Research, Chaudhuri K. and Salakhutdinov R, Eds., vol. 97. Long Beach, California, USA: PMLR, 09–15 Jun 2019, pp. 7064–7073.

- [39]. Jiang W, Larson PEZ, and Lustig M, “Simultaneous auto-calibration and gradient delays estimation (SAGE) in non-Cartesian parallel MRI using low-rank constraints,” *Magnetic Resonance in Medicine*, vol. 80, pp. 2006–2016, 2019.
- [40]. Bydder M, Rapacchi S, G. O O., Guye M, and Ranjeva JP, “Trimmed autocalibrating k-space estimation based on structured matrix completion.” *Magnetic Resonance Imaging*, vol. 43, no. 88, pp. 88–94, 2017. [PubMed: 28716683]

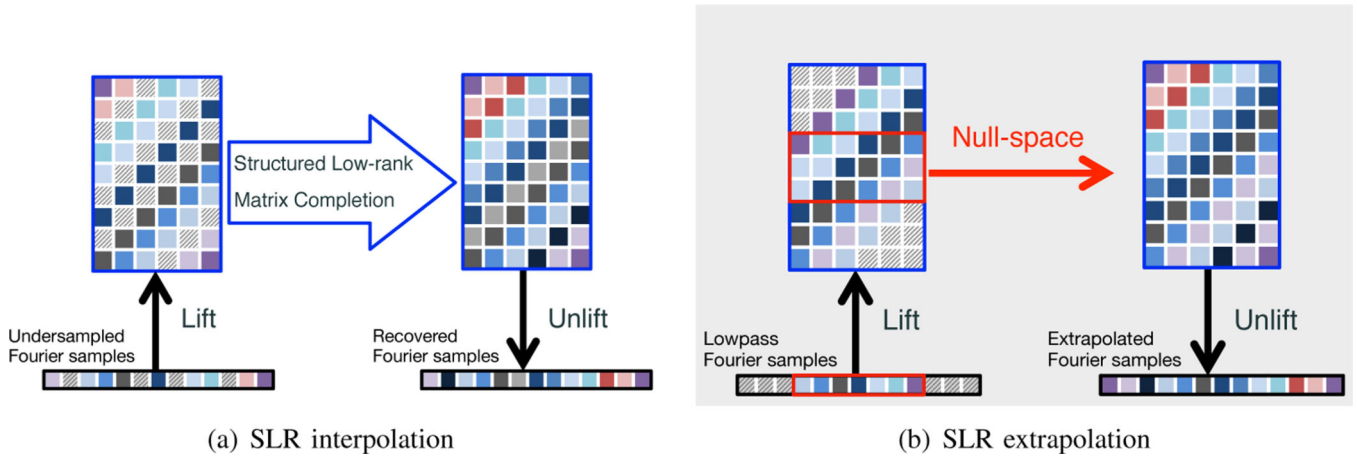
**Fig. 1.**

Illustration of SLR-based interpolation and extrapolation methods in the context of 1-D FRI.

(a) In SLR interpolation, the data is acquired on a non-uniformly sub-sampled Fourier grid. The SLR interpolation scheme relies on a lifting of the signal samples to a Hankel matrix, which has missing entries indicated by the hashed boxes. The one-to-one relation between the rank of a matrix and the continuous domain sparsity of the space domain signal is used to pose the recovery of missing samples as a structured low-rank matrix completion (SLRMC) problem in the lifted matrix domain. Specifically, the algorithm determines the matrix with the lowest rank that satisfies the Hankel structure and is consistent with the known matrix entries. Post-recovery, the matrix is unlifted to obtain the Fourier samples of the signal. (b) In SLR extrapolation problems, the low-frequency Fourier coefficients of the signal are uniformly sampled. The central fully known matrix region is used to estimate the null space of the matrix, which is used to linear-predict/extrapolate the missing high-frequency samples. The SLR algorithms that exploit the different signal structures differ only in the structure of the lifted matrix; the algorithms are essentially the same.

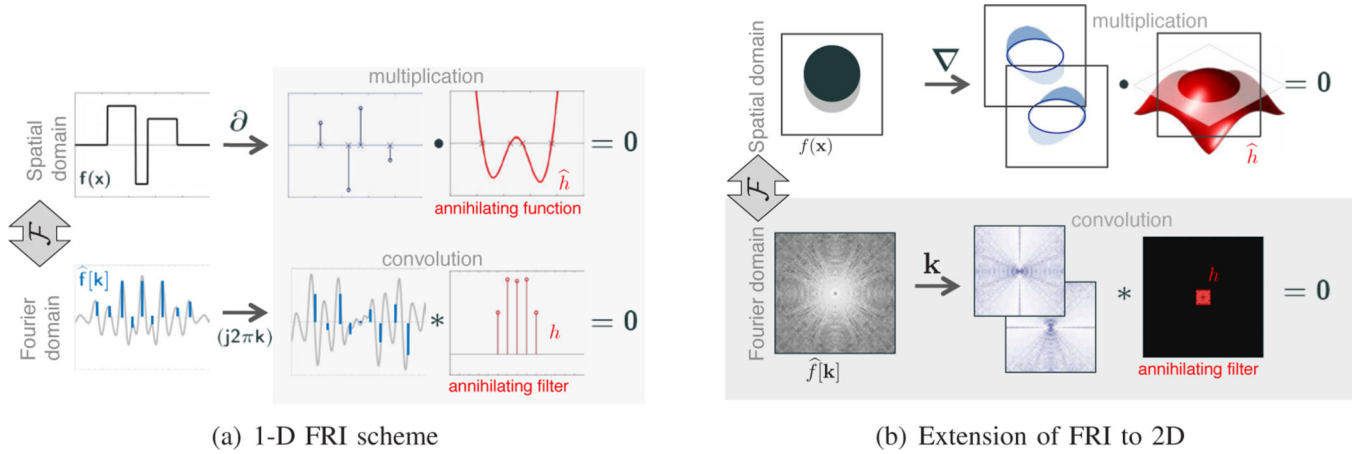


Fig. 2. Illustration of 1-D and 2-D FRI relations. (a) We illustrate Prony’s results in the light blue box, which demonstrates the recovery of Diracs at arbitrary locations. The Fourier coefficients of this signal consist of complex exponentials denoted by y_k in the bottom row. Prony’s results show that complex exponentials can be annihilated by the convolution with a filter h . This Fourier domain convolution relation can also be viewed as a multiplication-based annihilation relation in the space domain. These results can be extended to the recovery of piecewise constant 1-D signals by considering the derivative. Specifically, the derivative of piecewise constant signals consists of a linear combination of Diracs, which brings the problem to Prony’s setting. (b). Extension of 1-D FRI to the 2-D setting. The gradient of the piecewise constant signal consists of a series of Diracs supported on curves and the location of image edges. We observe that the gradient can be annihilated by the multiplication by a function \hat{h} , whose zeros overlap with the image edges. If the \hat{h} is bandlimited, the convolution with the Fourier coefficients of the image with the filter coefficients will also be zero. We generalize the notion of 1-D sparsity in FRI with the bandwidth of \hat{h} ; a more bandlimited \hat{h} will correspond to smoother curves, as shown in [10]. The above framework can be used to recover continuous domain piecewise constant images, whose edges are localized to bandlimited zero-measure curves of infinite support. This simple piecewise constant model is extended to more general piecewise smooth models in [20].

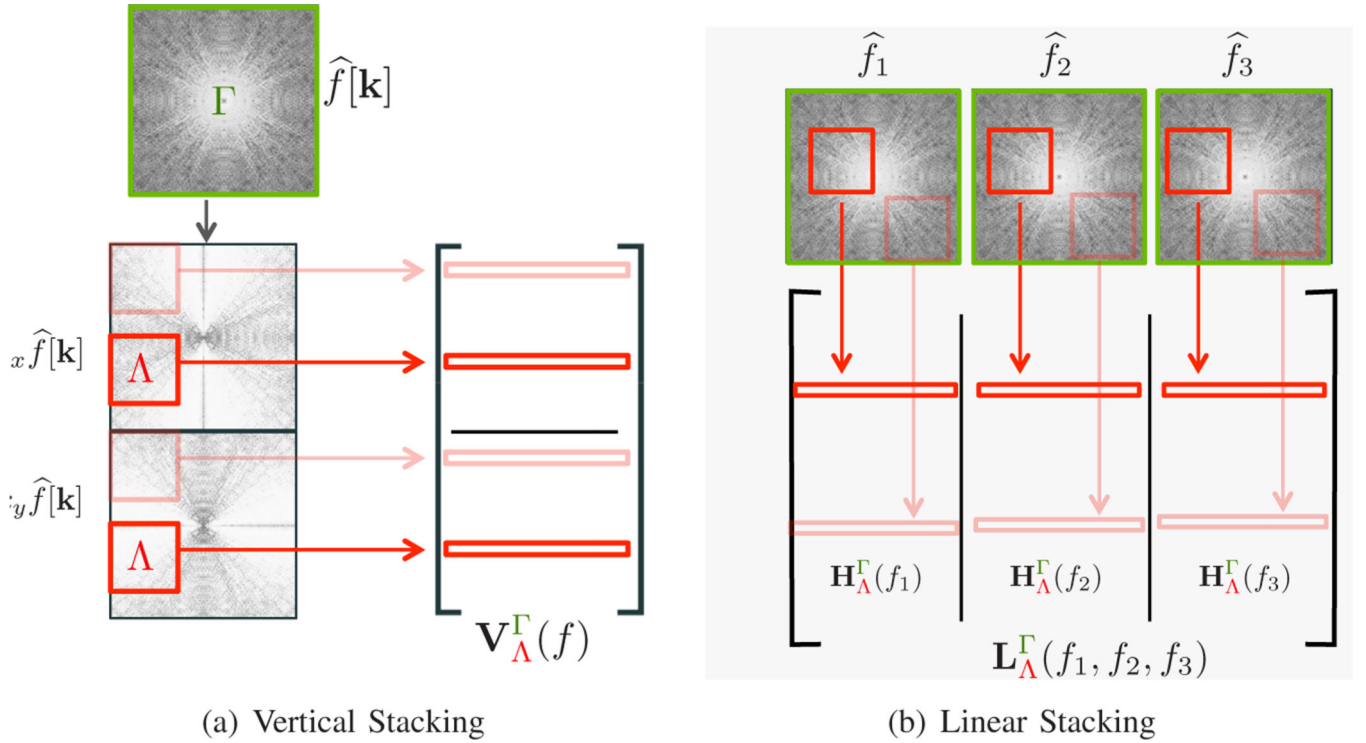


Fig. 3. Illustration of the matrix lifting operations in 2-D. (a). Vertical stacking of block Hankel matrices in (12), which is the matrix form of the convolutional relations in (11). The block Hankel matrices $\mathbf{H}_\Lambda^\Gamma(i\omega_x \hat{f})$ are constructed from the Fourier coefficients of $i\omega_x \hat{f}$. Each row of $\mathbf{H}_\Lambda^\Gamma(i\omega_x \hat{f})$ corresponds to a patch of size Λ , drawn from the the Fourier coefficients supported in Γ . The number of valid shifts of Λ within Γ is denoted by $\Gamma \ominus \Lambda$ in (11). (b) Linear stacking of block Hankel matrices in (22). The block Hankel matrices of the different channels f_i are concatenated in a linear fashion to construct the composite lifted matrix. When the low-frequency Fourier samples are fully acquired, the corresponding rows of the matrices are fully available. This facilitates the estimation of the null-space from these regions, which can be used to recover the remaining k-space regions as described in Section III. When the Fourier grid is sub-sampled, the matrices can be completed using nuclear norm minimization as discussed in Section IV.




Fig. 4. Illustration of the non-linear lifting operation in 2-D, which will map closed union of bandlimited curves to union of sub-spaces. (a) Each closed curve is mapped to a subspace. The SLR scheme relies on the low-rank structure of the subspace where the lifted points live in to recover the curves or denoise the points living on union of bandlimited curves. The utility of this scheme in the denoising of shapes is illustrated in (b). The recovery is posed as a nuclear norm minimization of the feature maps, which was solved using an iterative reweighted least-squares algorithm exploiting the kernel trick [21]. This approach is extended to higher dimensions to recover free breathing and ungated cardiac MRI data, illustrated in Fig. 8.


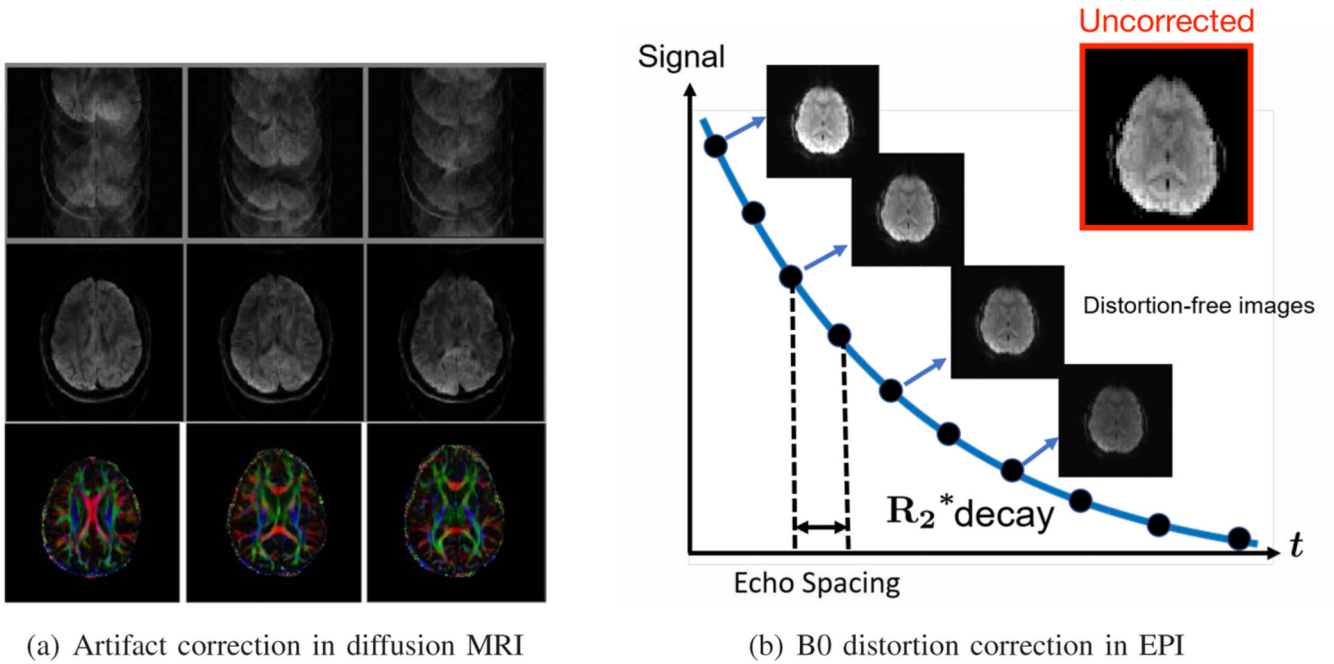
**Fig. 5.**

Illustration of different flavors of SLR recovery in a single coil setting. The Fourier transform of the image was sampled on the radial grid, which corresponds to an acceleration of 4.85. The columns correspond to (a) GIRAF, which assumes a piecewise constant image model, where the nuclear norm of the block Hankel matrix in (12) is minimized, (b) GSLR, where the image is modeled as the sum of a piecewise constant and piecewise linear functions as in (38), (c) S-LORAKS and (d) G-LORAKS [7], (e) total generalized variation (TGV), which is a discrete model that represents the images as the sum of piecewise linear and piecewise constant factors, and (f) total variation (TV). We note that the Fourier domain models (a)&(b) offer better reconstructions compared to their discrete counterparts (f) and (e), respectively, which can be appreciated from the zoomed images as well as the error images. We note that both (a) and (b) consider the low-rank structure of Hankel matrices constructed from weighted Fourier samples; this allows these methods to exploit the continuous sparsity of the edges, similar to discrete approaches such as TV or wavelet methods. The performance of the algorithms can be further improved by combining multiple SLR priors to exploit different signal priors (e.g., piecewise smoothness, multichannel sampling, smoothness of phase) [18]. These combinations can be either accounted for by different regularization terms or combined into a single nuclear norm penalty of a more complex structured matrix obtained by vertical and horizontal stackings to exploit the redundancies.

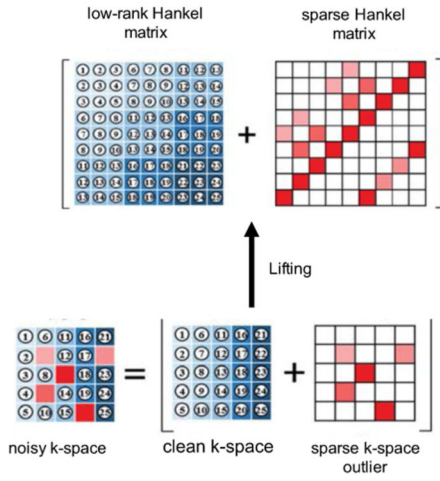


(a) Artifact correction in diffusion MRI


(b) B0 distortion correction in EPI

Fig. 6.

Example applications enabled by the SLR framework. (a) High-resolution four-shot diffusion MRI enabled by the SLR framework. The top row corresponds to uncorrected four-shot diffusion weighted MRI data from a single diffusion direction. Subtle subject motion between the shots results in phase errors between the measured Fourier datasets. In addition, eddy current artifacts also manifest as shifts in Fourier space between odd/even lines. These errors manifest as Nyquist ghosting artifacts in the uncorrected diffusion weighted images (DWI) in the top row. The SLR algorithm called MUSSELS [18] recovers the eight images, corresponding to each shot and odd/even lines separately, exploiting the phase relations between the images. The sum-of-squares combined DWI data from the above eight images for each slice are shown in the second row. The information from sixty such directions provides the fractional anisotropy diffusion MRI maps shown in the bottom row. (b) Correction of B0 distortions in EPI data [29]: the long readouts in EPI often result in spatial distortions, resulting from the inhomogeneity of the main (B0) magnetic field. Ignoring the B0-induced magnetization evolution during the long EPI readouts will be associated with spatial distortions. We reformulated the EPI distortion correction as a time-series recovery problem, where multiple images corresponding to different segments of the readouts are recovered exploiting the exponential structure of the signal. Note that each of the segments are highly undersampled. The missing Fourier samples are filled in using a structured low-rank matrix completion.



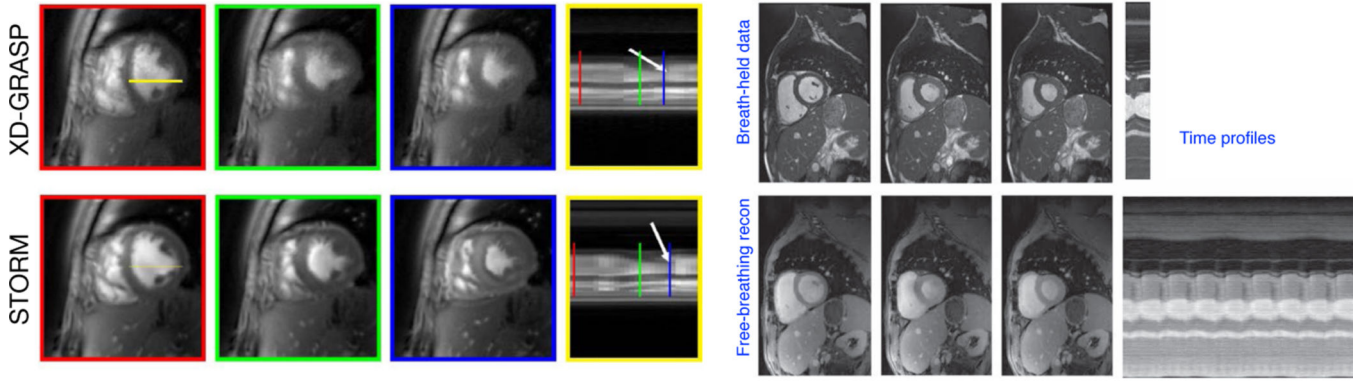
(a) Sparse + low rank decomposition.



(b) MR k-space artifact removal example.

Fig. 7.

(a) Lifting of image with sparse outlier. Noiseless image becomes low-rank Hankel matrix, whereas the noisy image becomes a sparse Hankel matrix so that robust principal component analysis can decompose the sparse and low-rank Hankel matrices. Then, by unlifting the Hankel matrix, we can obtain the noiseless k-space data and the clean images. (b) Example of MR artifact removals: (left) noisy images, (center) noisy k-space, and (right) decomposed sparse k-space outliers.



(a) Comparison with XD-GRASP

(b) Comparison with breath-held cine

Fig. 8.

Illustration of STORM algorithm [21] described in Section VIII-E. The STORM algorithm exploits the manifold structure of free-breathing and ungated images in high-dimensional space. Specifically, each image can be viewed as a smooth function of two parameters: the cardiac phase and the respiratory phase. A non-linear transformation specified by (29) transforms the data to a subspace. We recover the images from their undersampled measurements by exploiting this structure, which is compactly captured by nuclear norm of the feature matrix. We solve the optimization problem using an IRLS algorithm that uses the kernel trick, which eliminates the need to explicitly evaluate the complex features. (a) shows the comparisons of STORM reconstructed images with XD-GRASP, which is an explicit binning strategy; XD-GRASP bin the images to distinct cardiac/respiratory phases, followed by the total variation recovery of the images. We observe that the implicit binning of the data offered by STORM results in reduced blurring and improved fidelity (b) shows the comparisons of STORM reconstructions (bottom row) with classical breath-held acquisitions (top-row) that bins the data from different cardiac cycles. We observe that the image quality is comparable, whole STORM offers real-time imaging capabilities, allowing us to visualize cardiac and respiratory functions simultaneously. (a) Comparison