Recovery of damped exponentials using structured low rank matrix completion

Arvind Balachandrasekaran, Student Member, IEEE, Vincent Magnotta and Mathews Jacob, Senior Member, IEEE

Abstract

We introduce a structured low rank matrix completion algorithm to recover a series of images from their undersampled measurements, where the signal along the parameter dimension at every pixel is described by a linear combination of exponentials. We exploit the exponential behavior of the signal at every pixel, along with the spatial smoothness of the exponential parameters to derive an annihilation relation in the Fourier domain. This relation translates to a low-rank property on a structured matrix constructed from the Fourier samples. We enforce the low rank property of the structured matrix as a regularization prior to recover the images. Since the direct use of current low rank matrix recovery schemes to this problem is associated with high computational complexity and memory demand, we adopt an iterative re-weighted least squares (IRLS) algorithm, which facilitates the exploitation of the convolutional structure of the matrix. Novel approximations involving two dimensional Fast Fourier Transforms (FFT) are introduced to drastically reduce the memory demand and computational complexity, which facilitates the extension of structured low rank methods to large scale three dimensional problems. We demonstrate our algorithm in the MR parameter mapping setting and show improvement over the state-of-the-art methods.

Index Terms

Hankel/Toeplitz matrix, regularized recovery, smoothness, parameter mapping.

I. INTRODUCTION

Recovering parameters of a linear combination of damped exponentials is a problem of high significance in many MR imaging applications, including MR parameter mapping [1], [2], MR spectroscopy [3], and fat/water imaging [4]. The objective is to estimate from a series of MR images the spatial maps of the exponential parameters, which are indicative of the underlying tissue microstructure or metabolism. These maps are often used as bio-markers for pathologies including neuro-degenerative and cardiovascular disorders [1]–[3]. Current approaches involve acquiring multiple images by sampling the exponential signals at different points, followed by a pixel-by-pixel fitting of the exponential model to estimate the parameters. However, the main challenge with these schemes is the long acquisition time, resulting from the need to acquire a large number of high spatial resolution images. Recently, several researchers have considered compressive sensing methods for recovering images from under-sampled Fourier measurements using priors which enforce sparsity, smoothness and low-rankness [5]–[15].

The recovery of exponential parameters from few uniform samples of a linear combination of exponentials is a classical and well researched problem in signal processing [16]. The model has been extended to include a large class of signals with finite rate of innovation or finite number of discontinuities in [17], [18]. The early work in this direction focussed on the uniform sampling setting, where the linear dependencies between the samples of the signal translated to an "annihilation relationship". This implies that the signal can be nulled by the convolution with a finite impulse response filter. Recently, several researchers have extended the framework to recover a linear combination of undamped sinusoids from a few non-uniform Fourier samples [19]–[21]. These methods compactly represent the annihilation relation as a product of a Hankel matrix formed from the signal samples and a vector of annihilation filter coefficients. With this reformulation, the Hankel matrix can be shown to be low-rank; the low-rank property can be enforced to complete the matrix from its non-uniform measurements.

Arvind Balachandrasekaran, Mathews Jacob are with the Department of Electrical and Computer Engineering and Vincent Magnotta is with the Department of Radiology, University of Iowa, Iowa City, IA, 52245 USA (e-mail:arvind-balachandrasekaran@uiowa.edu;vincent-magnotta@uiowa.edu;mathews-jacob@uiowa.edu)

This work is supported by grants NIH 1R01EB019961-01A1 and ONR N00014-13-1-0202.

In this paper, we introduce a structured matrix completion algorithm for recovering a series of MR images from their non-uniformly under-sampled Fourier measurements, where the signal along the parameter dimension at every pixel is described by a linear combination of damped exponentials. We also assume that the exponential parameters vary smoothly in space. We observe that this model is general enough to account for many applications, including MR spectroscopy, parameter mapping, and diffusion MRI. For example, in MR parameter mapping, the signal along the parameter dimension could vary as a function of echo time, repetition time, and/or spin-lock duration. In this paper, we consider the single parameter setting for simplicity. We exploit the exponential behavior at every pixel, along with the smoothness of the parameters in the spatial dimensions, to derive an annihilation relation in the k-tdomain; t denotes the parameter dimension. These 3-D convolution relations can be compactly represented using a multi-fold Toeplitz matrix formed from the k-t samples. We show that this matrix has a large null space, and hence is low rank. We enforce the low rank property of the structured matrix as a prior to recover the missing entries from the under-sampled Fourier measurements. The spatial smoothness as well as the number of exponentials in the model can be controlled by the rank of the Toeplitz matrix, which is in-turn dependent on the regularization parameter.

The straightforward implementation of the above structured Toeplitz matrix recovery scheme is associated with huge memory demand and high algorithmic complexity. Specifically, the size of the Toeplitz matrix is often several orders of magnitude greater than the size of the multidimensional signal. We introduce an algorithm based on a half-circulant approximation of the Toeplitz matrix, which eliminates the need for the explicit evaluation and storage of the structured matrix. This work is a generalization of our recent work [22], where we introduced the GIRAF (Generalized Iterative Reweighted Annihilating Filter) algorithm for recovering the missing entries of a Toeplitz/Hankel matrix, when only few of its entries are observed. The approximation of linear convolutions by circular convolutions enabled an efficient implementation of the algorithm using Fast Fourier transforms (FFTs). The circulant approximation in [22] is valid when the signal samples decay rapidly towards the boundaries. This scheme is not directly applicable in our setting, since the signal samples have significant magnitude at the first few points along the parameter dimension. We modify the GIRAF algorithm and adopt a hybrid approach to solve the problem. Specifically, we perform the 3-D linear convolution as a series of 2-D circular convolutions along the spatial dimensions and linear convolution along the parameter dimension. Such a modification allows us to apply our algorithm on large scale multi-dimensional exponential estimation problems. The preliminary version of this work is accepted as a conference paper [23]. Compared to the work [23], the theoretical and algorithmic frameworks are further developed here, in addition to the application of the problem to the recovery of single and multi-channel T_2 weighted images.

The proposed method has similarities with structured matrix priors introduced to exploit various signal properties, including finite support and smoothly varying phase [20], piece-wise smooth continuous domain images [21], and continuous domain wavelet sparsity [24]. Similar structured low rank priors have also appeared in the recovery of calibrationless multichannel data [25] and multi-shot diffusion weighted images [26]. However, none of the above 2-D methods are designed to exploit the smooth exponential structure of the 3-D dataset. In [27], a Hankel matrix is constructed by exploiting the temporal smoothness using a Fourier transform and the spatial redundancies are exploited using a wavelet transform; the exponential structure of the temporal signal is not taken into account. In addition, the recovery of each $k_y - t$ slice is performed independently, assuming Cartesian sampling. The exponential structure of the signal is exploited in [28] and [29], where a Hankel matrix is constructed at every pixel by exploiting the linear predictability of the exponential time series. Since the linear combination of pixel-wise structured low rank priors is not capable of exploiting the similarities between the pixels in the dataset, the authors additionally use low rank and joint sparsity penalties on the Casorati matrix; see section II.D for more details. The proposed formulation enables the joint exploitation of the spatial priors; this approach is computationally more efficient and requires fewer free regularization parameters.

A. Notation

We collect the different notations used through out the paper and describe them in this section for easy reference. Unless otherwise mentioned, bold upper-case letters \mathbf{X} and bold lower-case letters \mathbf{y} are used to represent matrices and vectors respectively; $[\mathbf{y}]^T$, $[\mathbf{X}]^T$ represent a transpose of the vector \mathbf{y} and matrix \mathbf{X} respectively. We denote a function that is dependent on \mathbf{r} and n by $x[\mathbf{r}, n]$. The collection of function values for all possible values of \mathbf{r} and n are denoted by the vector \mathbf{x} . The discrete Fourier transform of x is denoted by \hat{x} , while the vector corresponding to the function values is denoted by $\hat{\mathbf{x}}$. We use non-bold lower-case greek alphabets such as μ, α to represent constants. Upper-case greek alphabets Λ, Θ represent index sets containing the support of the coefficients of the filter and Fourier data; $|\Lambda|$ is used to denote the size of the set Λ . We denote the 2D and 3D convolution by * and \otimes respectively. The calligraphic letters (e.g. \mathcal{A}, \mathcal{T}) denote operators. For example, \mathcal{A} is the forward operator that models the image acquisition as in (2), while \mathcal{T} is a lifting operator that constructs a multi-fold Toeplitz structured matrix $\mathcal{T}(\hat{\rho})$ from the entries of $\hat{\rho}$.







(a) Construction of the Toeplitz matrix $\mathcal{T}(\hat{\rho})$ of dimension $|\Delta| \times |\Lambda|$.

(b) Illustration of the index sets for the construction of Toeplitz matrix.

(c) Illustration of the index sets for the minimal and assumed fil-

Fig. 1. (a) Illustration of the construction of the matrix $\mathcal{T}(\hat{\rho})$ from the dataset $\hat{\rho}$: The rows of the Toeplitz matrix correspond to the cuboid shaped neighborhoods of the Fourier samples. The number of columns is equal to the size of the filter support $(|\Lambda|)$. Similarly, the number of rows is equal to the number of valid linear convolutions between $\hat{\rho}$ and the filter, denoted by Δ . b) Illustration of the relation between the filter and signal supports and the matrix dimensions: The Fourier samples $\hat{\rho}[\mathbf{k}, n]$ and the filter coefficients $d[\mathbf{k}, n]$ are assumed to be supported on the rectangular sets Γ and Λ , respectively. The 3-D convolution between them is valid in the dotted rectangular region (in green) and the set of valid indices is represented by $\Delta = \Gamma : \Lambda$. c) The rectangular sets containing the coefficients of the minimal and assumed filters are represented by Θ and Λ respectively. The number of linearly independent null space vectors of $\mathcal{T}(\hat{\rho})$ is given by all possible valid shifts of Θ in Λ , denoted by $\Lambda : \Theta$; this implies that $\mathcal{T}(\hat{\rho})$ is low rank and we enforce this property to estimate the missing entries of the matrix.

II. RECOVERY USING ANNIHILATION RELATIONS

A. Measurement model

We consider the recovery of a series of images ρ from its multichannel Fourier measurements b. The multichannel Fourier data $\hat{\rho}_i$ corresponding to the *i*th image frame can be modeled as

$$\mathbf{b}_{ij} = \underbrace{\mathbf{S}_i \mathbf{F} \mathbf{C}_j \mathbf{F}^*}_{\mathbf{A}_{ij}} \widehat{\boldsymbol{\rho}}_i + \boldsymbol{\eta}_{ij}, \ j = 1 \dots N_{coils}$$
(1)

where \mathbf{b}_{ij} and η_{ij} are the under-sampled Fourier measurements and zero mean white gaussian noise corresponding to the i^{th} frame and j^{th} coil respectively, \mathbf{C}_j is the receiver coil sensitivity of the j^{th} coil, \mathbf{S}_i is the sampling mask for the i^{th} frame and \mathbf{F} is the 2D discrete Fourier transform (DFT) matrix. (1) can be compactly written as

$$\mathbf{b} = \mathcal{A}(\widehat{\boldsymbol{\rho}}) + \boldsymbol{\eta} \tag{2}$$

where $\hat{\rho} = [\hat{\rho}_1, \hat{\rho}_2, \dots, \hat{\rho}_T] \in \mathbb{C}^{B \times T}$ is the Fourier data in Casorati matrix form [30] with the *i*th column representing the vector of Fourier data at time instant t_i . \mathcal{A} is a linear operator representing Fourier under-sampling and multiplication of coil sensitivities with $\hat{\rho}$.

B. Annihilation property of smooth exponentials

We model the signal at the spatial location $\mathbf{r} = (x, y)$ as a linear combination of L exponentials:

$$\rho[\mathbf{r}, n] = \sum_{i=1}^{L} \alpha_i(\mathbf{r}) \ \beta_i(\mathbf{r})^n, \tag{3}$$

where $\alpha_i(\mathbf{r}) \in \mathbb{C}$ are the amplitudes, $\beta_i(\mathbf{r}) \in \mathbb{C}$ is the exponential parameter that is dependent on the underlying physiology, *n* refers to the signal index along the parameter dimension and *L* is the number of exponentials at the voxel. For instance, in T_2 mapping applications the exponential parameters at the pixel location indexed by \mathbf{r} are given by $\beta_i(\mathbf{r}) = \exp\left(\frac{-\Delta T}{T_{2,i}(\mathbf{r})}\right)$. Here ΔT is the difference in echo times between two frames and $T_{2,i}$ is the relaxation parameter of the *i*th tissue component (e.g. gray matter, CSF or white matter).

The exponential signal, described in (3), at each pixel location can be annihilated by a 1-D FIR filter $g[\mathbf{r}, n]$ [16]:

$$\sum_{m=0}^{L} \rho[\mathbf{r},m] \ g[\mathbf{r},n-m] = 0, \quad \forall \mathbf{r}.$$
(4)

where (4) represents a 1-D convolution between the signal $\rho[\mathbf{r}, n]$ and the L+1 tap filter $g[\mathbf{r}, n]$. Since the exponential parameters vary from pixel to pixel, the filter $g[\mathbf{r}, n]$ also varies with the spatial location \mathbf{r} .

In practice, the exponential parameters vary smoothly as a function of space. This implies that the coefficients of the filter $g[\mathbf{r}, n]$ can be assumed to be smooth functions of the spatial variable \mathbf{r} . Taking the 2-D Fourier transform of (4) along the spatial dimensions, we obtain the following annihilation relation in the Fourier domain:

$$\hat{\rho}[\mathbf{k},n] \otimes d[\mathbf{k},n] = 0.$$
⁽⁵⁾

where $\hat{\rho}[\mathbf{k}, n] \stackrel{\mathcal{F}_{2D}}{\leftrightarrow} \rho[\mathbf{r}, n]$ and $d[\mathbf{k}, n] \stackrel{\mathcal{F}_{2D}}{\leftrightarrow} g[\mathbf{r}, n]$ are the spatial Fourier coefficients of $\rho[\mathbf{r}, n]$ and $g[\mathbf{r}, n]$, respectively. Here, \otimes denotes 3-D convolution.

Since the filter coefficients of $g[\mathbf{r}, n]$ vary smoothly as a function of space, we assume $d[\mathbf{k}, n]$ to be a 3-D FIR filter, whose coefficients are support limited in the rectangular set $\Lambda \subset \mathbb{Z}^3$; the size of Λ (spatial bandwidth of $d[\mathbf{k}, n]$) controls the spatial smoothness of the parameters, while the bandwidth along the parameter dimension is dependent on the number of exponentials in the signal model.

We focus on the recovery of the Fourier coefficients of the signal specified by $\hat{\rho}$ within the rectangular set $\Gamma \subset \mathbb{Z}^3$. For simplicity, we assume $\hat{\rho}$ has T frames, each of dimension $P \times Q$. The set Γ is illustrated by the red cuboid in Fig. 1.(b). The 3-D convolution (5) can be compactly written as

$$\mathcal{T}(\widehat{\boldsymbol{\rho}}) \mathbf{d} = 0 \tag{6}$$

where \mathcal{T} is a linear operator that maps a 3-D dataset $\hat{\rho}$ into a lifted matrix $\mathcal{T}(\hat{\rho}) \in \mathbb{C}^{m \times s}$. The construction of the matrix is illustrated in Fig. 1.(a). Similarly d represents the vectorized 3-D filter $d[\mathbf{k}, n]$. Here, $s = |\Lambda|$ is the number of the columns of the matrix, where Λ denotes the support of the filter d indicated by the blue cuboid in Fig. 1.(b). The number of rows in the matrix is denoted by $m = |\Gamma : \Lambda|$, which corresponds to the number of valid linear convolutions between $\hat{\rho}$ and the filter; the convolutions are valid in the green dotted cuboid in Fig. 1.(b). The set $\Gamma : \Lambda$ is specified by

$$\Gamma : \Lambda = \{ \mathbf{v} \mid \Lambda + \mathbf{v} \subseteq \Gamma; \mathbf{v} \in \mathbb{Z}^3 \}.$$
(7)

The blue dotted cuboid in Fig. 1.(b) indicates a valid shift of the blue solid cuboid, whose support is given by Λ . We observe that $\mathcal{T}(\hat{\rho})$ has a multi-fold Toeplitz structure since the matrix-vector product in (6) corresponds to a 3-D convolution.

C. Dimensions of the fundamental subspaces of $\mathcal{T}(\hat{\rho})$

We denote the index set of the filter with the smallest support (termed as the minimal filter), which satisfies the annihilation relation, to be Θ . In practice, the support Θ (denoted by the purple cuboid in Fig. 1.(c) is unknown. In such cases, the support set of the filter is overestimated to Λ , such that $\Theta \subset \Lambda$. Let the dimensions of Λ be $N_1 \times N_2 \times M$.

When the size of the filter is overestimated to Λ , it will result in multiple linearly independent vectors **d** in the null space of $\mathcal{T}(\hat{\rho})$. Specifically, if $c[\mathbf{k}, n]$ is the minimal filter, then any FIR filter of the form

$$d[\mathbf{k}, n] = c[\mathbf{k}, n] \otimes e[\mathbf{k}, n]$$
(8)

will also satisfy the following annihilation relation $\hat{\rho}[\mathbf{k}, n] \otimes d[\mathbf{k}, n] = 0$, or equivalently $\mathcal{T}(\hat{\rho})\mathbf{d} = 0$. Here, $e[\mathbf{k}, n]$ is any FIR filter such that $d[\mathbf{k}, n]$ is support limited to Λ . The number of such filters $(d[\mathbf{k}, n])$ is specified by the

set of all valid shifts of Θ in Λ , denoted by $\Lambda : \Theta$ [21]; this set is indicated by the orange cuboid in Fig. 1.(c). The corresponding shifted filters are linearly independent [21].

The above discussion shows that the dimension of the kernel of $\mathcal{T}(\hat{\rho})$ is at least $|\Lambda:\Theta|$. i.e.,

$$\dim\left(\ker\left(\mathcal{T}\left(\widehat{\boldsymbol{\rho}}\right)\right)\right) \ge |\Lambda:\Theta| \tag{9}$$

where Λ is the assumed filter size and $|\Lambda : \Theta|$ denotes the cardinality of the set $\Lambda : \Theta$. By the rank-nullity theorem, the rank of $\mathcal{T}(\hat{\rho})$ or the dimension of the image space of $\mathcal{T}(\hat{\rho})$ is specified by

$$\operatorname{rank}\left(\mathcal{T}\left(\widehat{\boldsymbol{\rho}}\right)\right) = \dim\left(\operatorname{im}\left(\mathcal{T}\left(\widehat{\boldsymbol{\rho}}\right)\right)\right) \leq |\Lambda| - |\Lambda:\Theta|.$$
(10)

Consider two datasets $\hat{\rho}_1$ and $\hat{\rho}_2$, where the size of Θ (purple cuboid in Fig. 1.(c)) is smaller for $\hat{\rho}_2$. Note that a smaller minimal filter corresponds to a dataset with fewer exponentials and smoother parameters. Since the number of valid shifts indicated by the set $|\Lambda : \Theta|$ is higher for the dataset $\hat{\rho}_2$, we have $\operatorname{rank}(\mathcal{T}(\hat{\rho}_2)) < \operatorname{rank}(\mathcal{T}(\hat{\rho}_1))$. Hence, the rank of the Toeplitz matrix can be used as a measure of complexity of the dataset.

Since we expect the matrix $\mathcal{T}(\hat{\rho})$ to be low rank, we use the rank prior as well as the Toeplitz structure of the matrix to recover the exponential image series from under-sampled Fourier measurements. The support of the different index sets along with the construction of the Toeplitz matrix is illustrated in Fig. 1.

D. Structured low-rank recovery from few measurements

Since (2) is an ill-posed problem, we employ the structured low rank matrix prior proposed in the previous sub-sections for the recovery of missing Fourier samples. We formulate the recovery of the Fourier data $\hat{\rho}$ from the under-sampled measurements **b** as the following structured matrix completion problem:

$$\min_{\widehat{\rho}} \operatorname{rank} \left[\mathcal{T}(\widehat{\rho}) \right] \text{ such that } \mathbf{b} = \mathcal{A}(\widehat{\rho}) + \eta$$
(11)

Since the problem specified in (11) is NP hard, we relax the rank function with a Schatten p ($0 \le p \le 1$) norm. The relaxed objective function is then given by

$$\widehat{\boldsymbol{\rho}}^{\star} = \arg\min_{\widehat{\boldsymbol{\rho}}} \|\mathcal{T}(\widehat{\boldsymbol{\rho}})\|_{p} + \frac{\mu}{2} \|\mathcal{A}(\widehat{\boldsymbol{\rho}}) - \mathbf{b}\|_{2}^{2}$$
(12)

where μ is a regularization parameter that balances the weight given to the Schatten norm of the matrix and the data consistency term in (12). $\mathcal{T}(\hat{\rho}) \in \mathbb{C}^{m \times s}$ is a multifold Toeplitz matrix formed from the Fourier samples $\hat{\rho}$. $\|\mathbf{X}\|_p$ is the Schatten p norm, defined as $\|\mathbf{X}\|_p := \frac{1}{p} \text{Tr}[(\mathbf{X}^H \mathbf{X})^{\frac{p}{2}}] = \frac{1}{p} \text{Tr}[(\mathbf{X}\mathbf{X}^H)^{\frac{p}{2}}] = \frac{1}{p} \sum_i \sigma_i^p$; σ_i are the singular values of \mathbf{X} . When p = 1, the Schatten norm reduces to the convex nuclear norm and for $(0 \le p < 1)$, the Schatten norm is a non-convex penalty; When $p \to 0$, $\|\mathbf{X}\|_p := \sum_i \log \sigma_i$. In section III, we will focus on the algorithm to solve the optimization problem (12).

E. Relation to pixel-wise structured low-rank priors in [28]:

The recovery of MR parameter weighted images considered in [28] is closely related to the proposed work. In [28], the sum of structured low rank priors, formed at every pixel, is considered. Specifically, for every pixel, a Hankel matrix is constructed using the temporal signal at that pixel. The images are reconstructed by exploiting the low rank structure of all these matrices. The objective function is specified by

$$\{\bar{\boldsymbol{\rho}}_m\} = \arg\min_{\bar{\boldsymbol{\rho}}} \sum_{i=1}^{l} \operatorname{rank}[\mathcal{T}(\bar{\boldsymbol{\rho}}(\mathbf{r}_i))], \text{ such that } \mathcal{A}(\bar{\boldsymbol{\rho}}) = \mathbf{b}$$
(13)

where $\bar{\rho}_m$ is the set of images to be recovered, l is the total number of pixels, A is a linear operator and **b** is a vector of measurements.

We now consider a special case of our setting, where the spatial dimensions of Λ (assumed filter size) are the same as that of the dataset (i.e. $P = N_1$ and $Q = N_2$), which is related to the above model. Since the spatial dimensions of the filter are the same as the dataset, no spatial smoothness is assumed on the annihilation filter coefficients. In this case, the dimension of the Toeplitz matrix $\mathcal{T}(\rho)$, where ρ is a signal of interest, is specified

by $(T - M + 1) \times (N_1 \cdot N_2 \cdot M)$. The Toeplitz matrix after a re-arrangement of the columns has the following structure:

$$\mathcal{T}(\boldsymbol{\rho}) = \left(\left| \mathcal{T}(\boldsymbol{\rho}(\mathbf{r}_1)) \right| \left| \mathcal{T}(\boldsymbol{\rho}(\mathbf{r}_2)) \right| \dots \left| \left| \mathcal{T}(\boldsymbol{\rho}(\mathbf{r}_l)) \right| \right)$$
(14)

In (14), each of the Toeplitz matrices $\mathcal{T}(\rho(\mathbf{r}_i)) \in \mathbb{C}^{(T-M+1)\times M}$, whose entries correspond to the temporal signal at location \mathbf{r}_i . We observe that enforcing a low rank prior on the special case considered in (14) results in a more constrained approach than the pixel-wise low rank penalty in (13). In particular, the global low-rank prior considered in (14) enables the exploitation of correlations between the columns of $\mathcal{T}(\rho)$, in addition to the annihilation relations. In contrast, the pixel-wise approach in (13) is not capable of exploiting these correlations. An additional low-rank prior on the Casorati matrix of the images or a wavelet prior has to be used as in [28] to exploit these correlations. Since the special case of our formulation is already capable of exploiting these correlations, we do not require additional priors. More importantly, we consider spatially bandlimited filters, which account for the smoothness of the exponential parameters. We observe that this property offers a 3 dB improvement (see Table I) over the special case considered in (14). In summary, the structured low-rank prior considered in this work qualitatively consolidates and unifies the multiple diverse priors used in [28].

Another key benefit of the above special case is the reduced computational complexity. When an iterative reweighted least squares (IRLS) based approach [31] is employed, each step of the algorithm (13) requires the eigen decomposition of as many Gram matrices as pixels in the dataset. In contrast, the use of the prior (14) requires the eigen decomposition of only one Gram matrix at each iteration. As the dimension of the Gram matrices in both the cases is the same and equal to $(T - M + 1) \times (T - M + 1)$, the computational complexity of the special case is orders of magnitude lower than that of the pixel-wise structured low rank strategy.

III. OPTIMIZATION ALGORITHM

The minimization of (12) using classical low-rank matrix recovery schemes is challenging due to the large size of the structured matrix, which often requires several orders of magnitude more memory, when compared to the original 3-D dataset. Current structured low-rank methods [20], [24] employ approaches originally designed for low-rank matrix recovery and do not exploit the structure of the matrix; the direct use of these 2-D algorithms to our 3-D setting is difficult due to the large memory demand and computational complexity.

We modify the GIRAF algorithm [22] to the 3-D setting to minimize the computational complexity. Specifically, we employ an IRLS based algorithm [31] to solve the optimization problem in (12). This approach allows us to use efficient approximations for operations involving the Toeplitz matrix using fast Fourier transforms; these modifications quite significantly reduce the computational complexity and memory demand. To derive the basic idea of the algorithm, we use the following identity to express the Schatten p norm of a matrix as a weighted Frobenius norm:

$$\|\mathbf{Y}\|_{p} = \frac{1}{p} \operatorname{Tr}[\underbrace{(\mathbf{Y}\mathbf{Y}^{*})^{\frac{p}{2}-1}}_{\mathbf{H}} \mathbf{Y}\mathbf{Y}^{*}]$$
(15)

$$=\frac{1}{p}\|\mathbf{H}^{\frac{1}{2}}\mathbf{Y}\|_{F}^{2} \tag{16}$$

Hence the solution to the minimum Schatten norm can be obtained by alternating between the update of a weight matrix **H** and the solution to a weighted least squares problem. In our case, we set $\mathbf{Y} = \mathcal{T}(\hat{\boldsymbol{\rho}})$ in (15), which decouples (12) into two sub-problems. At the n^{th} iteration, the sub-problems are given by

$$\widehat{\boldsymbol{\rho}}^{(n)} = \arg\min_{\widehat{\boldsymbol{\rho}}} \| (\sqrt{\mathbf{H}})^{(n-1)} \mathcal{T}(\widehat{\boldsymbol{\rho}}) \|_F^2 + \frac{\mu p}{2} \| \mathcal{A}(\widehat{\boldsymbol{\rho}}) - \mathbf{b} \|_2^2$$
(17)

$$\mathbf{H}^{(n)} = [\underbrace{\mathcal{T}(\widehat{\boldsymbol{\rho}}^{(n)}) \, \mathcal{T}(\widehat{\boldsymbol{\rho}}^{(n)})^{*}}_{\mathbf{P}} + \epsilon^{(n)} \, \mathbf{I}]^{\frac{p}{2}-1}$$
(18)

where $\epsilon^{(n)} \to 0$ is added to stabilize the inverse. Hence to solve (12), we employ an alternating minimization scheme that cycles between the sub-problems (17) and (18) till the cost of (12) between successive iterates is below a tolerance threshold. In the next two sub-sections, we describe an efficient implementation of the two sub-problems.

A. Least squares solution of (17)

Let the rows of $\sqrt{\mathbf{H}}$ be denoted by $[(\mathbf{h}^{(1)})^T, \dots, (\mathbf{h}^{(M)})^T]^T$. Substituting for $\sqrt{\mathbf{H}}$ in (17), we obtain

$$\widehat{\boldsymbol{\rho}}^* = \arg\min_{\widehat{\boldsymbol{\rho}}} \sum_{i=1}^M \|\mathbf{h}^{(i)} \mathcal{T}(\widehat{\boldsymbol{\rho}})\|_2^2 + \frac{\mu p}{2} \|\mathcal{A}(\widehat{\boldsymbol{\rho}}) - \mathbf{b}\|_2^2$$
(19)

The term $\mathbf{h}^{(i)} \mathcal{T}(\hat{\rho})$ in (19) represents a 3-D linear convolution between the 3-D sequences $\mathbf{h}^{(i)}$ and $\hat{\rho}$.

In the GIRAF algorithm [22], the linear convolutions were approximated by circular convolutions, so that they could be efficiently implemented using Fast Fourier transforms (FFTs). The approximations were valid due to the rapid decay of the Fourier coefficients towards the boundaries. However, in our case as the signal at every voxel follows an exponential curve, the magnitude of the Fourier coefficients are high at the first few points along the parameter dimension. Hence, the direct application of the GIRAF scheme to our setting gives poor results.

We introduce a hybrid strategy to improve the approximations, while keeping the memory demand and computational complexity low. In particular, we approximate the 3-D linear convolution as a series of 2-D circular convolutions along the spatial dimensions and a linear convolution along the parameter dimension. Denoting the l^{th} frame of $\hat{\rho}$ and h as $\hat{\rho}_l[\mathbf{k}]$ and $h_l[\mathbf{k}]$ respectively, we rewrite the convolution relation as

$$h[\mathbf{k}, n] \otimes \hat{\rho}[\mathbf{k}, n] = \sum_{m} \sum_{\mathbf{p}} \hat{\rho}[\mathbf{k} - \mathbf{p}, n - m] h[\mathbf{p}, m]$$
$$= \sum_{m} \underbrace{\sum_{m} \sum_{\mathbf{p}} \hat{\rho}_{n-m} [\mathbf{k} - \mathbf{p}] h_{m}[\mathbf{p}],}_{g_{n-m,m} = \hat{\rho}_{n-m} * h_{m}}$$
(20)

where \otimes denotes 3-D convolution and * denotes 2-D convolution. Since the spatial Fourier coefficients of $\hat{\rho}[\mathbf{k}]$ decay rapidly towards the boundaries, the 2-D linear convolutions $g_{j,l} = \hat{\rho}_j * h_l$ can be approximated as 2-D circular convolutions and efficiently computed using fast Fourier transforms as shown in [22]. After the 2-D convolutions are evaluated for all feasible combinations $g_{j,l}$, we compute the outer sum.

Now, we express the aforementioned idea in compact matrix notations. Let **h** and $\hat{\rho}$ consist of M and T frames respectively. We consider an arbitrary filter h of spatial dimensions $N_1 \times N_2$ and denote its i^{th} frame by \mathbf{h}_i . Now $(\mathbf{h} \mathcal{T}(\hat{\rho}))$ can be expanded as,

$$\mathbf{h} \, \mathcal{T}(\widehat{\boldsymbol{\rho}}) = \begin{pmatrix} \mathbf{h}_M & \dots & \mathbf{h}_1 \end{pmatrix} \begin{pmatrix} \mathbf{T}(\widehat{\boldsymbol{\rho}}_1) & \dots & \mathbf{T}(\widehat{\boldsymbol{\rho}}_{T-M+1}) \\ \vdots & \vdots & \vdots \\ \mathbf{T}(\widehat{\boldsymbol{\rho}}_M) & \dots & \mathbf{T}(\widehat{\boldsymbol{\rho}}_T) \end{pmatrix}$$
(21)

In the above equation, $\mathbf{T}(\hat{\rho}_j)$ represents a Toeplitz matrix formed from the samples of $\hat{\rho}_j$. This matrix can be expressed in terms of a larger circulant matrix [22] in the following way:

$$\mathbf{T}(\widehat{\boldsymbol{\rho}}_j) \approx \mathbf{P}^*_{\Lambda_s} \, \mathbf{C}(\widehat{\boldsymbol{\rho}}_j) \tag{22}$$

Here $\mathbf{C}(\widehat{\rho}_j) \in \mathbb{C}^{L \times L}$ is a circulant matrix formed from the Fourier samples $\widehat{\rho}_j$, Λ_s is the support of a frame of the filter and $\mathbf{P}^*_{\Lambda_s} \in \mathbb{C}^{N_1 N_2 \times L}$ corresponds to zero padding operation outside the filter support Λ_s . Note that the support of \mathbf{h}_l is often much smaller than that of $\widehat{\rho}_j$.

Using the approximation in (22), we can efficiently evaluate $\mathbf{h}_l \mathbf{T}(\hat{\rho}_j)$, which is the 2-D linear convolution between the l^{th} frame of \mathbf{h} and j^{th} frame of $\hat{\rho}$, as

$$\mathbf{h}_{l} \mathbf{T}(\widehat{\boldsymbol{\rho}}_{j}) \approx \mathbf{h}_{l} \mathbf{P}_{\Lambda_{s}}^{*} \mathbf{C}(\widehat{\boldsymbol{\rho}}_{j}) \\ = [\widehat{\boldsymbol{\rho}}_{j}]^{T} \underbrace{\mathbf{C}(\mathbf{h}_{l} \mathbf{P}_{\Lambda_{s}}^{*})}_{\mathbf{C}_{l}}.$$
(23)

where we have used the commutative property of convolution to arrive at the expression in (23). Here, $\mathbf{C}_{l} = \mathbf{C}(\mathbf{h}_{l}\mathbf{P}^{*}_{\Lambda_{s}})$ is a circulant matrix formed from the zero padded filter coefficients \mathbf{h}_{l} . Hence, the product $[\widehat{\boldsymbol{\rho}}_{j}]^{T} \mathbf{C}_{l}$ denotes the 2-D circular convolution between $[\widehat{\boldsymbol{\rho}}_{j}]^{T}$ and the zero-padded filter coefficients.



Fig. 2. Effect of approximations on the run time of the proposed algorithm and comparison of SNR: The approximations introduced in the proposed method enable efficient computation of the sub-problems (17) and (18) using fast Fourier transforms (FFTs). This results in a faster convergence (7.5 fold speed up) to the same solution as the one obtained using the IRLS (direct) method.

We propose to implement the circular convolutions using fast Fourier transforms to minimize the computational complexity. Specifically, we compute (23) efficiently as

$$[\widehat{\boldsymbol{\rho}}_j]^T \mathbf{C}_l = [\widehat{\boldsymbol{\rho}}_j]^T \underbrace{\mathbf{F}^* \mathbf{D}_l \mathbf{F}}_{\mathbf{C}_l}$$
(24)

where \mathbf{D}_l is a diagonal matrix with diagonal entries $\mu_l \stackrel{\mathcal{F}_{2D}}{\leftrightarrow} \mathbf{h}_l \mathbf{P}^*_{\Lambda_s}$ and \mathbf{F} denotes the 2-D discrete Fourier transform matrix.

Using (24), we simplify (21) as,

$$\mathbf{hT}(\widehat{\boldsymbol{\rho}}) \approx [\widehat{\boldsymbol{\rho}}]^T \mathbf{Q}^* \underbrace{\begin{pmatrix} \mathbf{D}_M & \cdots & \mathbf{0} \\ \mathbf{D}_{M-1} & \mathbf{D}_M & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{D}_1 & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_1 & \mathbf{D}_M \\ \vdots & \vdots & \cdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{D}_1 \end{pmatrix}}_{[\mathbf{D}(\mathbf{h})]^T} \mathbf{Q}^*$$

where each of the matrices \mathbf{D}_l are diagonal matrices and the dimension of $\mathbf{D}(\mathbf{h})^T$ is $PQT \times PQk$ with k = T - M + 1. Here, $[\hat{\boldsymbol{\rho}}]^T = [\hat{\boldsymbol{\rho}}_1, \dots, \hat{\boldsymbol{\rho}}_T]^T$ and $\mathbf{Q}^* = \mathbf{I} \bigotimes \mathbf{F}^*$ is a block diagonal matrix with the diagonal blocks being the inverse 2-D Fourier transform matrix. Thus, we can express $\|\mathbf{h} \mathbf{T}(\hat{\boldsymbol{\rho}})\|^2$ as $\|\mathbf{D}(\mathbf{h})\mathbf{Q}^* \hat{\boldsymbol{\rho}}\|^2$.

Using the above relation and substituting for the first term in (19) we obtain,

$$\widehat{\boldsymbol{\rho}}^* = \arg\min_{\widehat{\boldsymbol{\rho}}} \sum_{i=1}^M \|\mathbf{D}(\mathbf{h}^{(i)})\mathbf{Q}^*\widehat{\boldsymbol{\rho}}\|_2^2 + \frac{\mu p}{2} \|\mathcal{A}(\widehat{\boldsymbol{\rho}}) - \mathbf{b}\|_2^2$$
(25)

The above equation can be solved by taking its gradient and setting it to zero. The gradient is given by

$$2\mathbf{Q}\left(\underbrace{\sum_{i=1}^{M} (\mathbf{D}(\mathbf{h}^{(i)}))^* \mathbf{D}(\mathbf{h}^{(i)})}_{\mathbf{G}}\right) \mathbf{Q}^* \widehat{\boldsymbol{\rho}} + \mu p \ \mathcal{A}^* \mathcal{A} \widehat{\boldsymbol{\rho}} = \mu p \ \mathcal{A}^* \mathbf{b}$$
(26)

Note that prior to solving (26), **G** can be precomputed efficiently. Denote $\mathbf{E}^* = [\mathbf{D}_M, \mathbf{D}_{M-1}, \dots, \mathbf{D}_1]$. In order to populate the entries of $(\mathbf{D}(\mathbf{h}^{(i)}))^*\mathbf{D}(\mathbf{h}^{(i)})$, we need to compute one product $(\mathbf{E}\mathbf{E}^*)^{(i)}$, and the sum of k = (T - M + 1) sparse matrices. Each sparse matrix contains a shifted version of $(\mathbf{E}\mathbf{E}^*)^{(i)}$ as the only non-zero block. After precomputing **G**, we only need a few iterations of conjugate gradient (cg) algorithm to solve (26).



Fig. 3. Effect of approximations introduced in the proposed method on the recovery of coil combined T_2 weighted images from 30 percent uniform random Fourier measurements: One frame (4th Echo, TE = 40ms) of the image sequence corresponding to the ground truth is shown in (a) along with the frame of the sampling mask in (e). The 4th echo of the reconstructed image sequence using proposed method (b) is compared with those obtained using IRLS (direct) in (c) and GIRAF in (d). The corresponding error images are shown in (f)-(h). The estimated T_2 maps, which were multiplied by a mask to remove the background and the CSF region, are shown in (i)-(iv) and the corresponding error maps are shown in (v)-(vii).

Algorithm 1: Proposed algorithm for the recovery of exponential image time series

Initialize $\hat{\rho}^{(0)}$ and choose $\epsilon^{(0)} > 0$; for n = 1 to N_{max} or tolerance threshold reached do Step1: Weight matrix update: Compute each block $\mathbf{R}_{p,q}$ of the Gram matrix \mathbf{R} using (28) and (29); Compute the eigen values and eigen vectors $\{\lambda^{(i)}, \mathbf{u}^{(i)}\}_{i=1}^{M}$ from the eigen decomposition of \mathbf{R} ; Evaluate the null space vectors: $\mathbf{h}^{(i)} = \sqrt{\alpha_i}(\mathbf{u}^{(i)})^*$, where $\alpha^{(i)} = (\lambda^{(i)} + \epsilon^{(n-1)})^{\frac{p}{2}-1}$; Step2: Least Squares update: Solve the least squares problem: $\hat{\rho}^{(n)} = \arg\min_{\hat{\rho}} \sum_{i=1}^{M} \|\mathbf{D}(\mathbf{h}^{(i)})\mathbf{Q}^*\hat{\rho}\|_2^2 + \frac{\mu p}{2} \|\mathcal{A}(\hat{\rho}) - \mathbf{b}\|_2^2$ using Conjugate gradient (CG) method; Choose $\epsilon^{(n)}$ such that $0 < \epsilon^{(n)} \le \epsilon^{(n-1)}$; end

B. Weight-Update

The first step in computing the weight matrix $\mathbf{H}^{\frac{1}{2}}$ involves forming the Gram matrix $\mathbf{R} = \mathcal{T}(\hat{\rho}) \mathcal{T}(\hat{\rho})^*$. The direct computation of \mathbf{R} requires the evaluation and storage of the lifted matrix $\mathcal{T}(\hat{\rho})$, which will be a computationally

expensive and memory intensive operation. Instead we propose an efficient way to compute the Gram matrix. Specifically, we assume \mathbf{R} to be partitioned in the following way:

$$\begin{pmatrix} \mathbf{R}_{1,1} & \mathbf{R}_{1,2} & \dots & \mathbf{R}_{1,M} \\ \mathbf{R}_{2,1} & \mathbf{R}_{2,2} & \dots & \mathbf{R}_{2,M} \\ \vdots & \vdots & \cdots & \vdots \\ \vdots & \vdots & \cdots & \vdots \\ \mathbf{R}_{M,1} & \mathbf{R}_{M,2} & \cdots & \mathbf{R}_{M,M} \end{pmatrix}$$
(27)

where the above matrix has M column and row partitions and $\mathbf{R}_{i,j}$ is a matrix block of dimension $N_1N_2 \times N_1N_2$. We obtain a general expression for the matrix block corresponding to the p^{th} row and q^{th} column partition of \mathbf{R} as

$$\mathbf{R}_{p,q} = \sum_{i=1}^{k} \mathbf{T}(\widehat{\boldsymbol{\rho}}_{p+i-1}) \mathbf{T}(\widehat{\boldsymbol{\rho}}_{q+i-1})^{*}$$
(28)

where k := T - M + 1. To compute $\mathbf{R}_{p,q}$, we use the relation in (22) and simplify $\mathbf{T}(\widehat{\rho}_i)\mathbf{T}(\widehat{\rho}_j)^*$ as

$$\mathbf{P}_{i,j} = \mathbf{T}(\widehat{\boldsymbol{\rho}}_i)\mathbf{T}(\widehat{\boldsymbol{\rho}}_j)^* = \mathbf{P}^*_{\Lambda_s} \underbrace{\mathbf{C}(\widehat{\boldsymbol{\rho}}_i)\mathbf{C}(\widehat{\boldsymbol{\rho}}_j)^*}_{\mathbf{C}(\mathbf{g})} \mathbf{P}_{\Lambda_s}$$
(29)

where the entries of $\mathbf{C}(\mathbf{g})$ are obtained from the array \mathbf{g} . The entries of \mathbf{g} are given by $\mathbf{F}(\rho_i \circ \operatorname{conj}(\rho_j))$, where ρ_i and ρ_j are the images corresponding to the Fourier samples $\hat{\rho}_i$ and $\hat{\rho}_j$ respectively, \mathbf{F} denotes a 2-D DFT matrix, conj denotes the conjugate operation and \circ denotes point-wise multiplication. Hence the entries of every row of $\mathbf{P}_{i,j}$ can be populated by performing a sliding window operation that extracts and vectorizes a $N_1 \times N_2$ patch from a $(2N_1 - 1) \times (2N_2 - 1)$ neighborhood.

Next the weight matrix $\mathbf{H}^{\frac{1}{2}}$ is efficiently computed from the eigen decomposition of **R**. Let **U** represent the orthogonal basis of eigen vectors $\mathbf{u}^{(i)}$ and Λ be a diagonal matrix containing the eigen values $\lambda^{(i)}$. Then the eigen decomposition of **R** is given by $\mathbf{U}\Lambda\mathbf{U}^*$. Substituting for **R** in (18) and simplifying further we obtain,

$$\mathbf{H} = [\mathbf{U}(\mathbf{\Lambda} + \epsilon \mathbf{I})\mathbf{U}^*]^{\frac{p}{2}-1} = \mathbf{U}(\mathbf{\Lambda} + \epsilon \mathbf{I})^{\frac{p}{2}-1}\mathbf{U}^*.$$

Hence, one choice of the matrix square root $\mathbf{H}^{\frac{1}{2}}$ is

$$\mathbf{H}^{\frac{1}{2}} = (\mathbf{\Lambda} + \epsilon \mathbf{I})^{\frac{p}{4} - \frac{1}{2}} \mathbf{U}^* = \left[(\mathbf{h}^{(1)})^T, \dots, (\mathbf{h}^{(M)})^T \right]^T$$

where $\mathbf{h}^{(i)} = \sqrt{\alpha^{(i)}} (\mathbf{u}^{(i)})^*$ and $\alpha^{(i)} = (\lambda^{(i)} + \epsilon)^{\frac{p}{2}-1}$.

C. Implementation details:

The details of the alternating minimization algorithm to solve (17) and (18) are described in Algorithm (1). We initialize and vary the value of ϵ as described in [31]. Specifically, we initialize ϵ as $\epsilon^{(0)} = \lambda_{max}/100$, where λ_{max} is the largest eigen value of the gram matrix $\mathcal{T}(\hat{\rho}_{init})^{*}(\hat{\rho}_{init})^{*}$, which is formed from the initial guess $\hat{\rho}_{init} = \hat{\rho}^{(0)}$. Every iteration, we decrease the value of ϵ as $\epsilon^{(n)} = \epsilon^{(n-1)}/\gamma$, $\gamma > 1$. For all our experiments we set $\gamma = 1.4$. We run the optimization algorithm for different values of the regularization parameter μ and choose the value which results in the best signal to noise ratio (SNR), where SNR := $20 \log \frac{\|X_g\|_2}{\|X_g - X_r\|_2}$. Here X_g and X_r are the ground truth and the reconstructed images respectively. To observe the behavior of μ across different acceleration factors, we reconstructed images at different acceleration factors ranging from six to twelve. We observed that the optimal μ estimated was fairly constant across different acceleration factors. In the absence of ground truth data it is not possible to compute the SNR. Hence choosing the regularization parameter using the above-mentioned approach is not feasible. In such cases, we could employ heuristic techniques such as the L-curve method [32] or more sophisticated methods [33], [34], where risk functions approximating the SNR are used to estimate the optimal regularization parameter. We implemented the algorithm in MATLAB, which runs on a Linux workstation with a 3.6Ghz Intel Xeon CPU and 32GB RAM.



Fig. 4. Comparison of the proposed method with different reconstruction schemes on the recovery of coil combined data from 30 percent uniform random Fourier measurements: One frame (4th Echo, TE = 40ms) of the image sequence is shown in (a)-(e) along with the frame of sampling mask in (f). The corresponding error images are shown in (g)-(j). The estimated T_2 maps, specifically the enclosed rectangular region is zoomed and is shown in (i)-(v) with the corresponding error maps shown in (vi)-(ix). Here the maps were multiplied by a mask to remove the background and the CSF region. The improvements offered by the proposed scheme can be easily appreciated from the T_2 error images and the estimated T_2 maps.

IV. EXPERIMENTS AND RESULTS

We demonstrate our algorithm on a fully sampled axial 2-D dataset, which was acquired on a Siemens 3T Trio scanner with 12 coils using a turbo spin echo sequence. The scan parameters were: TR = 2500 ms, slice thickness = 5 mm, Matrix size = 128×128 and FOV = 22×22 cm². By varying the echo times (TE) from 10 to 120 ms, we acquired T_2 weighted images at twelve equispaced TE.

A. State-of-the-art methods used for comparison:

We compare the proposed method with three state-of-the-art methods: k - t low rank [35], blind compressed sensing (BCS) [14], and the pixel-wise structured low-rank prior in (13), which we refer to as HLR-Voxel. An IRLS algorithm [31] was employed to solve the nuclear norm minimization in (13). We compared the methods for both single and multi-channel recovery experiments.

We also demonstrate the improved speed up offered by the proposed scheme by comparing it with traditional IRLS as well as the multidimensional GIRAF algorithm [22]. For the traditional IRLS (direct) method, we solve (17) and (18) directly without introducing FFT based approximations. After the images are reconstructed, we estimate the T_2 maps by fitting a mono exponential model to every pixel.

B. Single channel recovery:

We demonstrate the proposed method on the recovery of single channel T_2 -weighted data from 30% uniform random measurements and compare it with the three aforementioned methods. To create the single channel data, we performed a principal component analysis (PCA) on the original multi-channel data and selected the most significant component to obtain a coil compressed data. For the proposed method, a filter of size $122 \times 122 \times 2$



Fig. 5. Comparison of the proposed method with different reconstruction schemes on the recovery of multi channel data at an acceleration factor of 12: In the first row of A), B) and C) three frames corresponding to 2nd Echo (TE=20ms), 9th Echo (TE=90ms) and 11th Echo (TE=110ms) are shown in (a)-(e) along with the frame of the sampling mask in (f). The corresponding error images are shown in (g)-(j) in the second row of A), B) and C). In D), the mean signal decay is plotted over a gray matter (Region A) and white matter (Region B) region for the ground truth and all the competing methods. The signal within the red ROI was corrupted due to some non-idealities in the acquisition. In E), the estimated T_2 maps, specifically the enclosed rectangular region is zoomed and is shown in (i)-(v) with the corresponding error maps shown in (vi)-(ix). Here the maps were multiplied by a mask to remove the background and the CSF region. We observe that the reconstructions from the proposed method have fewer errors, which can be appreciated from the error maps of the T_2 weighted images as well with the noise-like artifacts in the T_2 maps.



Fig. 6. Plot showing the SNR of the reconstructions at all the echo times for the proposed and the competing methods.

was used to recover the images. The value of p = 0.6 for the Schatten norm was chosen for both the k - t low rank and the proposed algorithm. The results from the different methods are shown in Fig. 4. We observe that the T_2 weighted images and the estimated T_2 maps from the proposed method have fewer errors and artifacts than the competing methods.

TABLE IEFFECT OF FILTER SIZE ON SNR OF T_2 WEIGHTED IMAGES.

Class size	CNID		CND (JD)
inter size	SINK	Inter size	SNR (dB)
128x128x10	28.05	102x102x11	30.80
122x122x10	30.30	102x102x10	31.21
114x114x10	31.00	102x102x7	31.13
108x108x10	31.12	102x102x4	30.96
102x102x10	31.21	102x102x2	30.78
100x100x10	31.20	102x102x1	29.88

(a) Varying spatial dimension (b) Varying temporal dimensions

C. Multi-channel recovery:

In Fig. 5, we compare the different methods on the recovery of multi-channel T_2 -weighted data from twelve fold under-sampled Fourier measurements. The data was retrospectively under-sampled using a combination of uniform Cartesian and a pseudo-random variable density sampling patterns. Specifically, we uniformly undersampled the xand y directions by a factor of 2 and refer this sampling mask as a 2×2 uniform Cartesian mask. To increase the incoherence between the frames, we also shifted every frame of the mask by zero or one unit (done randomly) along the x and y directions. We achieved an acceleration factor of twelve by combining the four fold uniform Cartesian mask with a three fold pseudo-random variable density undersampling pattern. Three frames of the sampling mask corresponding to the three echoes are shown in (f) in Fig. 5 A), B) and C). We used a Schatten p = 0.7 for both the proposed and the low rank methods. Also for the proposed method, a filter of size $102 \times 102 \times 10$ was used in the recovery of images. Reconstructions corresponding to three echos with TE = 20ms, TE = 90ms and TE =110ms respectively and the T_2 maps for all the methods are shown in A) through C) and E) respectively in Fig. 5. We observe that the T_2 weighted images from the proposed method have fewer errors than the competing methods. Also, the T_2 maps corresponding to the proposed method are a lot smoother and have fewer artifacts, especially in the enclosed rectangular region, than those obtained from the competing methods. In Fig. 5 D), we plot the mean signal decay over two regions of interests for all the methods. We observe that the signal corresponding to the proposed method matches closely to the ground truth signal in both cases. We also observe that the mean signal decay from the ROI (in red) in the gray matter is not exactly an exponential function. Nevertheless, the proposed signal model (3) approximates the signal as a linear combination of exponentials and captures the signal decay quite well. This suggests that the performance of the algorithm may degrade gradually, when the signal deviates from an exponential model. In Fig. (6), we plot the Signal to Noise ratio (SNR) of the reconstructions at each echo time for the proposed and the competing methods. We observe that for all the methods, the SNR increases for the first few echoes and then decreases for the remaining echoes, which could be due to non-ideal acquisition conditions. Nevertheless, the SNR of the images obtained from the proposed method is higher at all echo times, when compared to the state-of-the-art methods.

D. Effect of filter size on image recovery:

We study the effect of filter size or equivalently the dimensions of the Toeplitz matrix on the SNR of the T_2 weighted images recovered from twelve fold under-sampled multi-channel Fourier data in Table I. We study the effect of varying the spatial dimensions of the filter on the SNR in Table Ia. We observe that the filters with a smaller spatial support $(102 \times 102 \times 10)$ provide improved results than larger filters, thus demonstrating the benefit of exploiting spatial smoothness. A filter with large spatial dimensions $(128 \times 128 \times 10)$ fails to exploit any spatial smoothness. We demonstrate the benefit of exploiting the annihilation relations, which takes into account the exponential structure of the signal along the parameter dimension at every pixel, in Ib. We observe that a filter having multiple taps along the temporal dimension $(102 \times 102 \times 10)$ results in reconstructions with a better SNR than those obtained using a filter with size $(102 \times 102 \times 1)$, which just exploits joint sparsity. From these experiments, we also note that varying the spatial support of the filter has a higher impact on the SNR than the temporal support.

E. Effect of approximations on image recovery:

We study the effect of the approximations, introduced in the proposed method, on the recovery of coil compressed T_2 -weighted data from 30% uniform random Fourier measurements. We note that the approximations introduced in the proposed method enable efficient computation of the sub-problems (17) and (18) using Fast Fourier transforms (FFT). This results in a faster convergence to the solution when compared to the IRLS-direct method, as shown in Fig. 2. Specifically, we observe a 7.5 fold speed up due to the proposed algorithm. In Fig. 3, we observe that the T_2 weighted images and the T_2 maps corresponding to the proposed method have similar SNR and image quality compared to those obtained using the IRLS-direct method. These results demonstrate the effectiveness of the approximations introduced in the proposed method. From the figure, we also observe that the images and the maps obtained from the GIRAF algorithm have a lot of errors and artifacts. This is because the approximations in the GIRAF algorithm break down in our setting thus resulting in poor T_2 estimates.

V. DISCUSSION & CONCLUSION

We introduced a novel structured matrix recovery algorithm to recover an image series with smoothly varying exponential parameters from under-sampled Fourier measurements. As the proposed method exploits the spatial smoothness of the parameters and the exponential structure of the signal along the parameter dimension at every pixel, it results in improved reconstructions over the other state-of-the-art methods. The comparisons on T_2 estimation problems in the context of MR parameter mapping demonstrate the potential of this scheme, with reduced errors in both the reconstructed images and T_2 maps compared to state of the art methods.

As the size of the filter is not known apriori, we treated it as an optimization parameter and chose the dimension that resulted in the best SNR. We observed that the spatial dimensions of the filter had a greater effect on the SNR than the temporal dimensions. Specifically, a filter with smaller spatial support $(102 \times 102 \times 10)$ provided improved reconstructions with higher SNR than a filter of size $128 \times 128 \times 10$, which failed to incorporate any spatial smoothness. Hence as the proposed matrix prior incorporates spatial smoothness, it eliminates the need for additional spatial regularizers or priors to further constrain the image recovery. Similarly, the reconstructions using a filter with multiple taps along the temporal dimension $(102 \times 102 \times 10)$ had higher SNR than the filter with one tap, which exploits the joint sparsity of the Casorati matrix formed from the Fourier samples. We also observed that for both the multi-channel and coil combined data, the filter sizes yielding reconstructions with highest SNR were different. The size of the filter reflects the complexity of the model which is usually dependent on the number of measured samples, number of coils etc. Since the model complexity is different for both the datasets, it resulted in different sizes of the filter.

To solve the optimization problem, we employ an iterative least squares (IRLS) based strategy, which decouples the original problem into two sub-problems. We adopt a hybrid approach to keep the memory demand and computational complexity low. Specifically, we introduce novel approximations, which allow us to solve the subproblems using FFTs. This resulted in a faster convergence to the same solution as the one obtained using the IRLS (direct) method. The proposed algorithm was approximately 7.5 times faster than the IRLS (direct) method. We also observed that the GIRAF algorithm broke down in our setting resulting in poor reconstructed images and T_2 maps.

The proposed framework may be extended to the multi-dimensional parameter setting. For instance, in MR parameter mapping, if the signal along the parameter dimension at every pixel varies as a function of both TR and TE, then the filter coefficients will be dependent on both exponential parameters T_1 and T_2 . Hence, annihilation relations similar to (4) and (5) can be derived, which can be compactly represented using a low rank Toeplitz matrix. The low rank property of the Toeplitz matrix can then be enforced to recover the images from under-sampled Fourier measurements. We plan to investigate this problem in the future.

The results in Fig.(5).D indicate that the performance of the algorithm degrades gradually, when the voxel time profiles deviate from the exponential signal model. Specifically, such signals may be reasonably approximated as a linear combination of few exponentials. In addition, the use of the low-rank penalty that only requires the matrix to be approximately low-rank, rather than a low-rank constraint, also allows the signal to deviate from the exponential model. We did not account for the presence of artifacts due to stimulated echoes [36]–[40] which are formed when a turbo spin echo sequence is used. Further investigation is required to study the impact of such errors.

REFERENCES

- [1] M. Ugander, P. S. Bagi, A. J. Oki, B. Chen, L.-Y. Hsu, A. H. Aletras, S. Shah, A. Greiser, P. Kellman, and A. E. Arai, "Myocardial edema as detected by pre-contrast T1 and T2 CMR delineates area at risk associated with acute myocardial infarction," *JACC: Cardiovascular Imaging*, vol. 5, no. 6, pp. 596–603, 2012.
- [2] A. Borthakur, M. Sochor, C. Davatzikos, J. Q. Trojanowski, and C. M. Clark, "T1ρ MRI of Alzheimer's disease," *Neuroimage*, vol. 41, no. 4, pp. 1199–1205, 2008.
- [3] A. Di Costanzo, F. Trojsi, M. Tosetti, T. Schirmer, S. M. Lechner, T. Popolizio, and T. Scarabino, "Proton MR spectroscopy of the brain at 3T: an update," *European radiology*, vol. 17, no. 7, pp. 1651–1662, 2007.
- [4] P. A. Hardy, R. S. Hinks, and J. A. Tkach, "Separation of fat and water in fast spin-echo MR imaging with the three-point Dixon technique," *Journal of Magnetic Resonance Imaging*, vol. 5, no. 2, pp. 181–185, 1995.
- [5] L. Feng, R. Otazo, H. Jung, J. H. Jensen, J. C. Ye, D. K. Sodickson, and D. Kim, "Accelerated cardiac T2 mapping using breath-hold multiecho fast spin-echo pulse sequence with k-t FOCUSS," *Magnetic Resonance in Medicine*, vol. 65, no. 6, pp. 1661–1669, 2011.
- [6] J. V. Velikina, A. L. Alexander, and A. Samsonov, "Accelerating MR parameter mapping using sparsity-promoting regularization in parametric dimension," *Magnetic Resonance in Medicine*, vol. 70, no. 5, pp. 1263–1273, 2013.
- [7] F. H. Petzschner, I. P. Ponce, M. Blaimer, P. M. Jakob, and F. A. Breuer, "Fast MR parameter mapping using k-t principal component analysis," *Magnetic Resonance in Medicine*, vol. 66, no. 3, pp. 706–716, 2011.
- [8] C. Huang, C. G. Graff, E. W. Clarkson, A. Bilgin, and M. I. Altbach, "T2 mapping from highly undersampled data by reconstruction of principal component coefficient maps using compressed sensing," *Magnetic Resonance in Medicine*, vol. 67, no. 5, pp. 1355–1366, 2012.
- [9] B. Zhao, W. Lu, T. K. Hitchens, F. Lam, C. Ho, and Z.-P. Liang, "Accelerated MR parameter mapping with low-rank and sparsity constraints," *Magnetic Resonance in Medicine*, vol. 74, no. 2, pp. 489–498, 2015.
- [10] B. Zhao, F. Lam, and Z.-P. Liang, "Model-based MR parameter mapping with sparsity constraints: parameter estimation and performance bounds," *IEEE Transactions on Medical Imaging*, vol. 33, no. 9, pp. 1832–1844, 2014.
- [11] M. Doneva, P. Börnert, H. Eggers, C. Stehning, J. Sénégas, and A. Mertins, "Compressed sensing reconstruction for magnetic resonance parameter mapping," *Magnetic Resonance in Medicine*, vol. 64, no. 4, pp. 1114–1120, 2010.
- [12] T. Zhang, J. M. Pauly, and I. R. Levesque, "Accelerating parameter mapping with a locally low rank constraint," *Magnetic Resonance in Medicine*, vol. 73, no. 2, pp. 655–661, 2015.
- [13] Y. Zhu, Q. Zhang, Q. Liu, Y.-X. J. Wang, X. Liu, H. Zheng, D. Liang, and J. Yuan, "PANDA-T1ρ: Integrating principal component analysis and dictionary learning for fast T1ρ mapping," *Magnetic Resonance in Medicine*, vol. 73, no. 1, pp. 263–272, 2015.
- [14] S. Bhave, S. G. Lingala, C. P. Johnson, V. A. Magnotta, and M. Jacob, "Accelerated whole-brain multi-parameter mapping using blind compressed sensing," *Magnetic Resonance in Medicine*, 2015.
- [15] W. Li, M. Griswold, and X. Yu, "Fast cardiac T1 mapping in mice using a model-based compressed sensing method," *Magnetic Resonance in Medicine*, vol. 68, no. 4, pp. 1127–1134, 2012.
- [16] P. Stoica and R. L. Moses, Introduction to spectral analysis. Prentice hall Upper Saddle River, 1997, vol. 1.
- [17] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, June 2002.
- [18] Z.-P. Liang, E. M. Haacke, and C. W. Thomas, "High-resolution inversion of finite Fourier transform data through a localised polynomial approximation," *Inverse Problems*, vol. 5, no. 5, p. 831, 1989.
- [19] Y. Chen and Y. Chi, "Robust spectral compressed sensing via structured matrix completion," *IEEE Transactions on Information Theory*, vol. 60, no. 10, pp. 6576–6601, 2014.
- [20] J. Haldar, "Low-rank modeling of local k-space neighborhoods LORAKS for constrained MRI," *IEEE Transactions on Medical Imaging*, no. 3, pp. 668–681, March 2014.
- [21] G. Ongie and M. Jacob, "Off-the-Grid Recovery of Piecewise Constant Images from Few Fourier Samples," SIAM J Imaging Sciences, in press, 2015.
- [22] G. Ongie and M. Jacob, "GIRAF: A fast algorithm for structured low-rank matrix recovery." [Online]. Available: http://arxiv.org/abs/1609.07429
- [23] A. Balachandrasekaran and M. Jacob, "Novel Structured Low-rank algorithm to recover spatially smooth exponential image time series," *ArXiv e-prints*, Mar. 2017.
- [24] K. H. Jin, D. Lee, and J. C. Ye, "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.
- [25] P. J. Shin, P. E. Larson, M. A. Ohliger, M. Elad, J. M. Pauly, D. B. Vigneron, and M. Lustig, "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," *Magnetic Resonance in Medicine*, vol. 72, no. 4, pp. 959–970, 2014.
- [26] M. Mani, M. Jacob, D. Kelley, and V. Magnotta, "Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (mussels)," *Magnetic Resonance in Medicine*, 2016.
- [27] D. Lee, K. H. Jin, E. Y. Kim, S.-H. Park, and J. C. Ye, "Acceleration of MR parameter mapping using annihilating filter-based low rank Hankel matrix," *Magnetic Resonance in Medicine*, 2016.
- [28] X. Peng, L. Ying, Y. Liu, J. Yuan, X. Liu, and D. Liang, "Accelerated exponential parameterization of T2 relaxation with model-driven low rank and sparsity priors MORASA," *Magnetic Resonance in Medicine*, 2016.
- [29] H. M. Nguyen, X. Peng, M. N. Do, and Z.-P. Liang, "Denoising MR spectroscopic imaging data with low-rank approximations," *IEEE TBE*, pp. 78–89, 2013.
- [30] Z.-P. Liang, "Spatiotemporal imaging with partially separable functions," in 2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro. IEEE, 2007, pp. 988–991.
- [31] K. Mohan and M. Fazel, "Iterative reweighted algorithms for matrix rank minimization," *Journal of Machine Learning Research*, vol. 13, no. Nov, pp. 3441–3473, 2012.

- [32] P. C. Hansen and D. P. OLeary, "The use of the l-curve in the regularization of discrete ill-posed problems," SIAM Journal on Scientific Computing, vol. 14, no. 6, pp. 1487–1503, 1993.
- [33] S. Ramani, T. Blu, and M. Unser, "Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Transactions on Image Processing*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [34] S. Ramani, Z. Liu, J. Rosen, J.-F. Nielsen, and J. A. Fessler, "Regularization parameter selection for nonlinear iterative image restoration and mri reconstruction using gcv and sure-based methods," *IEEE Transactions on Image Processing*, vol. 21, no. 8, pp. 3659–3672, 2012.
- [35] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob, "Accelerated dynamic MRI exploiting sparsity and low-rank structure:kt SLR," *IEEE Transactions on Medical imaging*, vol. 30, no. 5, pp. 1042–1054, 2011.
- [36] N. Ben-Eliezer, D. K. Sodickson, and K. T. Block, "Rapid and accurate t2 mapping from multi-spin-echo data using bloch-simulationbased reconstruction," *Magnetic Resonance in Medicine*, vol. 73, no. 2, pp. 809–817, 2015.
- [37] S. Majumdar, S. Orphanoudakis, A. Gmitro, M. O'donnell, and J. Gore, "Errors in the measurements of t2 using multiple-echo mri techniques. i. effects of radiofrequency pulse imperfections," *Magnetic Resonance in Medicine*, vol. 3, no. 3, pp. 397–417, 1986.
- [38] S. Majumdar, S. Orphanoudakis, A. Gmitro, M. O'Donnell, and J. Gore, "Errors in the measurements of t2 using multiple-echo mri techniques. ii. effects of static field inhomogeneity," *Magnetic Resonance in Medicine*, vol. 3, no. 4, pp. 562–574, 1986.
- [39] A. Crawley and R. Henkelman, "Errors in t2 estimation using multislice multiple-echo imaging," Magnetic Resonance in Medicine, vol. 4, no. 1, pp. 34–47, 1987.
- [40] J. Hennig, "Multiecho imaging sequences with low refocusing flip angles," *Journal of Magnetic Resonance (1969)*, vol. 78, no. 3, pp. 397–407, 1988.