Published in final edited form as:

IEEE Trans Comput Imaging. 2017 December ; 3(4): 535–550. doi:10.1109/TCI.2017.2721819.

A Fast Algorithm for Convolutional Structured Low-rank Matrix Recovery

Greg Ongie [Member, IEEE] and

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, 48109

Mathews Jacob [Senior Member, IEEE]

Department of Electrical and Computer Engineering, University of Iowa, Iowa City, IA, 52245 USA

Abstract

Fourier domain structured low-rank matrix priors are emerging as powerful alternatives to traditional image recovery methods such as total variation (TV) and wavelet regularization. These priors specify that a convolutional structured matrix, i.e., Toeplitz, Hankel, or their multi-level generalizations, built from Fourier data of the image should be low-rank. The main challenge in applying these schemes to large-scale problems is the computational complexity and memory demand resulting from a lifting the image data to a large scale matrix. We introduce a fast and memory efficient approach called the Generic Iterative Reweighted Annihilation Filter (GIRAF) algorithm that exploits the convolutional structure of the lifted matrix to work in the original un-lifted domain, thus considerably reducing the complexity. Our experiments on the recovery of images from undersampled Fourier measurements show that the resulting algorithm is considerably faster than previously proposed algorithms, and can accommodate much larger problem sizes than previously studied.

Index Terms

Structured Low-Rank Matrix Recovery; Multi-level Toeplitz Matrices; Annihilating Filter; Finite Rate of Innovation; Compressed Sensing; MRI Reconstruction

I. Introduction

Recovering image data from limited and noisy measurements is a central problem in many biomedical imaging problems, including MRI, microscopy, and tomography. The conventional approach to solve these ill-posed problems is to regularize the recovery using smoothness or sparsity priors in discrete image domain, as in compressed sensing-based MRI reconstruction [1]. In contrast to this approach, several researchers have proposed a new class of recovery algorithms that impose a low-rank penalty on a convolutional structured matrix constructed from Fourier data of the image, i.e., a matrix that is Toeplitz, Hankel, or their multi-level generalizations [2]–[8]. These convolutional structured low-rank matrix penalties are emerging as powerful alternatives to conventional discrete spatial

domain penalties because of their ability to exploit continuous domain analogs of sparsity. These methods are often called “off-the-grid”, in the sense that they do not require discrete gridding of the signal in spatial domain. This is because these methods model linear dependencies among the Fourier coefficients of the image induced by low-complexity continuous domain properties of the image (detailed below). In particular, the linear dependencies can be expressed as a nulling of the Fourier coefficients of the image by convolution with a finite impulse response filter—a so-called “annihilation relationship”—that translates directly to the low-rank property of a Toeplitz-like matrix built from the Fourier coefficients. These methods then pose image recovery either partially or entirely in Fourier domain as a structured low-rank matrix recovery problem, where the structure of the matrix encodes the particular modeling assumptions.

The linear dependencies between the Fourier coefficients exploited in structured low-rank matrix priors result from a variety of assumptions, including continuous domain analogs of sparsity [3], [6], [8], [9], correlations in the locations of the sparse coefficients in space [8], [9], multi-channel sampling [2], [10], [11], or smoothly varying complex phase [4]. For example, the LORAKS framework [4] capitalized on the sparsity and smooth phase of the continuous domain image using structured low-rank priors, which offers improved reconstructions over conventional ℓ_1 recovery. Similarly, we have shown that the Fourier coefficients of continuous domain piecewise constant images whose discontinuities are localized to zero level-set of a bandlimited function satisfy an annihilation relationship [9]; their recovery from undersampled Fourier measurements translates to a convolutional structured low-rank matrix completion problem [8]. This model, which exploits the smooth structure of the discontinuities along with sparsity, can lead to significant improvement in reconstruction quality over traditional methods such as total variation (TV) minimization; see [8], [9] for more details. Similarly, the ALOHA framework [6] reformulates the recovery of a transform sparse signal as a structured low-rank matrix recovery problem, e.g., images that are sparse under the undecimated Harr wavelet transform. Inspired by auto-calibration techniques in parallel MRI [12]–[15], the extension to recovery of parallel MRI from undersampled measurements is formulated as a structured low-rank matrix recovery problem in [2], [10]. Similar approaches have also been found very effective in auto-calibrated multishot MRI [16] and correction of echo-planar MRI data [17]. The theoretical performance of structured low-rank matrix completion methods has been studied in [3], [18], [19], showing improved statistical performance over standard discrete spatial domain recovery.

Despite improvements in empirical and statistical recovery performance over traditional methods as seen from [2]–[4], [6], [8], [16], these structured low-rank matrix recovery schemes are associated with a dramatic increase in computational complexity and memory demand; this restricts their direct extension to multi-dimensional imaging applications, such as dynamic MRI reconstruction. In particular, these algorithms involve the recovery of a large-scale Toeplitz-like matrix whose combined dimensions are several orders of magnitude larger than those of the image. For example, the dimension of the structured matrix is roughly $10^6 \times 2000$ for a realistic scale 2-D MRI reconstruction problem (see Table I), while the matrix is several orders of magnitude larger in dynamic MRI applications, making it impossible to store. Moreover, typical algorithms require a singular value decomposition

(SVD) of the dense large-scale matrix at each iteration, which is computationally prohibitive for large-scale problems. Several strategies have been introduced to minimize or avoid SVD's in these algorithms. For example, the algorithms derived in [4], [6] replace the full SVD's with more efficient truncated SVD's or matrix inversions by assuming the matrix to be recovered is low-rank, or well approximated as such. However, even with these low-rank modifications, the algorithms in [4], [6] still have considerable memory demand, since they require storing a variable having dimensions of the large-scale matrix.

In this paper we introduce a novel, fast algorithm for a class of convolutional structured low-rank matrix recovery problems arising in MRI reconstruction and other imaging contexts. What distinguishes the algorithm from other approaches is its direct exploitation of the convolutional structure of the matrix—none of the current algorithms fully exploit this structure. This enables us to evaluate the intermediate steps of the algorithm efficiently using fast Fourier transforms in the original problem domain, resulting in an algorithm with significant reductions in memory demand and computational complexity. Our approach does not require storing or performing computations on the large-scale structured matrix, nor do we need to make overly strict low-rank assumptions about the solution. The proposed approach is based on the iterative reweighted least squares (IRLS) algorithm for low-rank matrix completion [20]–[22], which we adapt to the convolutional structured matrix setting. This algorithm minimizes the nuclear norm or the non-convex Schatten- p quasi-norm of the structured matrix subject to data constraints. However, as we show, the direct extension of the IRLS algorithm to the large-scale setting does not yield a fast algorithm. We additionally propose a systematic approximation of convolutional structured matrices that radically simplifies the subproblems of the IRLS algorithm, while keeping the low-rank property of the matrix intact. The combination of these two ingredients, the IRLS algorithm and the proposed approximation of the convolutional structured matrices, we call the generalized iterative reweighted annihilating filter (GIRAF) algorithm. The name reflects the fact that the algorithm can be interpreted as alternating between (1) the robust estimation of an annihilating filter for the data, and (2) solving for the data best annihilated by the filter in a least squares sense.

Finally, we note that several existing approaches for inverse problems in imaging utilize non-convex approaches, similar to the present work. For example, an augmented Lagrangian iterative shrinkage algorithm for non-convex Schatten p -norms has been introduced for the denoising of images in [23]. However, since these schemes apply low-rank denoising on groups of non-local patches, the matrices in their setting are considerably smaller than in our context and does not possess the convolutional structure that we consider in this work. Similarly, the non-convex penalties studied in [24], [25] are discrete formulations using non-convex vector norms in the image domain; these methods cannot exploit the continuous domain properties that are modeled using structured low-rank penalties, and hence will have limited applicability in the context we consider.

II. Signal reconstruction by convolutional structured low-rank matrix recovery

A. Convolutional structured low-rank matrix models in imaging

Structured low-rank matrix approximation (SLRA) models are widely used in many branches of signal processing [26], [27]. A SLRA model assumes some property of the signal data is equivalent to the low-rank property of a matrix constructed from the data. This paper is motivated by recent convolutional SLRA models in MRI reconstruction [2], [4]–[8], [28], and related inverse problems in imaging [6], [29], [30]. In this setting, various spatial domain properties of the image (e.g., limited support, smooth phase, piecewise constant, etc.) translate into the low-rank property of a convolutional structured matrix (e.g., Toeplitz, Hankel, and their multi-level generalizations) constructed from the Fourier coefficients of the image. Recovery of the image from undersampled or corrupted measurements is then posed in Fourier domain as a convolutional structured low-rank matrix recovery problem.

As a motivating example for the SLRA approach in imaging, consider the class of signals consisting of a sparse linear combination of Dirac impulses in 1-D:

$$\rho(x) = \sum_{i=1}^r c_i \delta(x - x_i), \text{ for all } x \in [0, 1] \quad (1)$$

This signal model is “off-the-grid” in the sense that the impulse locations $\{x_i\}_{i=1}^r$ can be arbitrary points in $[0, 1]$ and are not required to lie on a discrete grid. This example and the analysis that follows is closely related classical methods in line spectral estimation, including Prony’s method [31] and its robust refinements: MUSIC [32], ESPRIT [33], matrix pencil [34], and others. See [35] for a comprehensive overview. These methods utilize the fact that the sparsity of the signal implies a Toeplitz matrix built from the Fourier coefficients of ρ is rank deficient. To see why this is, let $\mu(x)$ be a periodic bandlimited function on $[0, 1]$ having r zeros at the Dirac locations $\{x_i\}_{i=1}^r$. It is easily seen that

$$\rho(x)\mu(x) = 0 \text{ for all } x \in [0, 1] \quad (2)$$

where the equality is understood in the sense of distributions or generalized functions (see, e.g., [36]). This multiplication annihilation relationship in spatial domain translates to a convolution annihilation relationship in Fourier domain:

$$(\hat{\rho} * \hat{\mu})[k] = 0, \text{ for all } k \in \mathbb{Z} \quad (3)$$

where $\hat{\rho}$ and $\hat{\mu}$ denote the Fourier coefficients of ρ and μ . For a finite collection of low-pass Fourier coefficients $\hat{\rho}[k], |k| \leq K$ where K is a pre-determined cut-off frequency, the convolution annihilation relationship (3) can be expressed in matrix form as

$$\text{Toep}(\hat{\rho})\mathbf{h} = 0. \quad (4)$$

Here $\text{Toep}(\hat{\rho}) \in \mathbb{C}^{M \times N}$ denotes a rectangular Toeplitz matrix built from $\hat{\rho}[k]$, and $\mathbf{h} \in \mathbb{C}^N$ is a vector of the Fourier coefficients of μ , zero-padded if necessary to have length N . This shows $\mathbf{h}$ is a non-trivial nullspace vector for $\text{Toep}(\hat{\rho})$, i.e., $\text{Toep}(\hat{\rho})$ is rank deficient. Notice also that any multiple of $\mu(x)$ by a phase factor, $\gamma(x) = \mu(x)e^{j2\pi kx}$, $k \in \mathbb{Z}$, will also satisfy the above annihilation relationship, i.e., $\text{Toep}(\hat{\rho})\mathbf{h}' = 0$ where $\mathbf{h}'$ is a vector of the Fourier coefficients of γ . This implies that $\text{Toep}(\hat{\rho})$ has a large nullspace, hence is a low-rank matrix. In fact, one can show $\text{rank}[\text{Toep}(\hat{\rho})] = r$, which establishes a one-to-one correspondence between the sparsity of the signal (1) and the rank of a Toeplitz matrix built from its Fourier coefficients.

When we only have samples of $\hat{\rho}[k]$ for $k \in \Omega$ where $\Omega \subset \mathbb{Z}$ is a sampling set of arbitrary locations, we can use the low-rank property of $\text{Toep}(\hat{\rho})$ to recover $\hat{\rho}[k]$ for all $|k| \leq K$ as the solution to the following matrix completion problem:

$$\min_{\hat{\phi}} \text{rank}[\text{Toep}(\hat{\phi})] \text{ subject to } \hat{\phi}[k] = \hat{\rho}[k], \quad \forall k \in \Omega \quad (5)$$

Recovery guarantees for a convex relaxation to (5) were studied in [3] assuming the sampling locations Ω are drawn uniformly at random from the set $\{k: |k| \leq K\}$ for some fixed cut-off K . Multi-dimensional generalizations of the model (1) and the recovery program (5) were also considered in [3], and have been adapted to super-resolution imaging context as well [29].

Additionally, SLRA models are central to several MRI reconstruction tasks. An important example is in parallel MRI, where Fourier data of the image is collected simultaneously at multiple receive coils having different spatial sensitivity profiles. Several auto-calibration techniques in parallel MRI, such as GRAPPA [12], ESPIRiT [14], PRUNO [15], and related techniques [13], model the coil sensitivities as smooth functions in spatial domain, which translates to the low-rank property of a convolutional structured matrix built from data obtained at a uniformly sampled calibration region in Fourier domain; see, e.g., [13]. Estimates of the coil sensitivity maps are obtained from the nullspace of this matrix, and then used to reconstruct missing Fourier data via linear prediction. The extension to recovery of parallel MRI without a fully sampled calibration region was formulated as a structured low-rank matrix recovery problem in the SAKE framework [2]. Extensions of this work to the recovery of single coil data was proposed in [4], by modeling the image as having sparse support or smoothly varying phase. This was later incorporated into a multi-coil framework in [10]. Further generalizations of this approach were proposed in [6], which employed transform sparsity of the image, using wavelets and other derivative-like operators.

In the remainder of this paper, we focus on the SLRA model that motivated this work—the continuous domain piecewise smooth image model introduced in [8], [9], [28]. We showed

that this class of signals satisfies a Fourier domain convolution annihilation relationship, provided the signal discontinuities are localized to the zero-set of a smooth bandlimited function, as in the finite-rate-of-innovation curves model [37]. This work can also be viewed as an extension of the SLRA models investigated in [38]–[41] in the context of super-resolution MRI, which was formulated for 1-D piecewise polynomial signals. Our work generalizes this model recovery of 2-D piecewise polynomial images with discontinuities supported on curves. For example, suppose $f(x, y)$ is a piecewise constant function in 2-D with discontinuities contained in the zero-set $\{(x, y) \in [0, 1]^2: \mu(x, y) = 0\}$ where μ is bandlimited. Analogous to (2), one can show the partial derivatives of f are annihilated by multiplication with μ in spatial domain:

$$(\partial_x f)\mu = 0 \text{ and } (\partial_y f)\mu = 0 \quad (6)$$

in the distributional sense; see Figure 1 for an illustration. This translates to the Fourier domain convolution annihilation relationship:

$$(\widehat{\partial_x f} * \widehat{\mu})[k_x, k_y] = 0, \quad (7)$$

$$(\widehat{\partial_y f} * \widehat{\mu})[k_x, k_y] = 0 \text{ for all } (k_x, k_y) \in \mathbb{Z}^2, \quad (8)$$

Here the expressions $\widehat{\partial_x f}$ and $\widehat{\partial_y f}$ are computed in Fourier domain by the weightings

$$\widehat{\partial_x f}[k_x, k_y] = j2\pi k_x \widehat{f}[k_x, k_y],$$

$$\widehat{\partial_y f}[k_x, k_y] = j2\pi k_y \widehat{f}[k_x, k_y] \text{ for all } (k_x, k_y) \in \mathbb{Z}^2,$$

We can represent the convolution relations (7) and (8) in matrix notation as:

$$\mathcal{G}(\widehat{f})\mathbf{h} = \begin{pmatrix} \text{Toep}_2(\widehat{\partial_x f}) \\ \text{Toep}_2(\widehat{\partial_y f}) \end{pmatrix} \mathbf{h} = 0 \quad (9)$$

where $\text{Toep}_2(\widehat{g})$ denotes the Toeplitz-like structured matrix built from any 2-D array of coefficients $\widehat{g}$ representing 2-D linear convolution with $\widehat{g}$. Figure 1 illustrates the construction of $\mathcal{G}(\widehat{f})$, which we call the *gradient weighted matrix lifting*. Similar to the 1-D setting we can show that $\mathcal{G}(\widehat{f})$ is a low-rank matrix. In [9] we proved that under certain geometric restrictions on the edge-set of f , there exists a unique bandlimited annihilating function μ_0 whose shifts in Fourier domain span the nullspace of $\mathcal{G}(\widehat{f})$. In particular, if $\mathcal{G}(\widehat{f})$

is built assuming $\mathbf{h} \in \mathbb{C}^N$ is an array of size $\Lambda \subset \mathbb{Z}^2$, $|\Lambda| = N$, then for sufficiently large M one has

$$\text{rank}(\mathcal{G}(\hat{f})) = N - S, \quad (10)$$

where S is the number of integer shifts of $\hat{\mu}_0$ inside the index set Λ (see Prop. 6 in [9]). This establishes a correspondence between the rank of $\mathcal{G}(\hat{f})$ and the complexity of the edge-set of f , as measured by the bandwidth of the annihilating function μ_0 .

In [8] we proposed to recover the Fourier coefficients of a piecewise constant image using a structured low-rank matrix completion formulation similar to (5). This approach has several applications to under-sampled MRI reconstruction, since MRI measurements are accurately modeled as the multi-dimensional Fourier coefficients of the underlying image. See also [4]–[7] for similar Fourier domain SLRA models for MRI reconstruction based on alternative image models.

B. Problem Formulation

A SLRA model supposes the data $\mathbf{x}_0$ to be recovered is such that a structured matrix $\mathcal{T}(\mathbf{x}_0)$ constructed from $\mathbf{x}_0$ is low-rank. In many SLRA models, including those for robust spectral estimation [3] and the super-resolution piecewise constant image model [9] introduced above, it is also known that $\mathcal{T}(\mathbf{x}_0)$ is the unique rank minimizer subject to certain data constraints. This suggests we can attempt to recover $\mathbf{x}_0$ from its linear measurements $\mathbf{A}\mathbf{x}_0 = \mathbf{b}$ by solving the following rank minimization problem:

$$\min_{\mathbf{x}} \text{rank}[\mathcal{T}(\mathbf{x})] \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{b}. \quad (11)$$

However, it is well-known that (11) is NP-hard in general [42]. Many authors have investigated tractable methods to obtain exact or approximate minimizers to the structured low-rank matrix recovery problem (11), including non-convex methods based on alternating projections (also known as Cadzow's method) [2], [4], [43], convex relaxation methods [7], [44], or local optimization techniques [45]. We discuss several of these methods in the Supplementary Materials.

We choose to formulate the structured low-rank matrix recovery problem as a family of relaxations to (11):

$$\min_{\mathbf{x}} \|\mathcal{T}(\mathbf{x})\|_p^p \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{b}, \quad (12)$$

where $\|\cdot\|_p$ denotes the family of Schatten- p quasi-norms with $0 < p \leq 1$, defined for an arbitrary matrix $\mathbf{X}$ by

$$\|X\|_p = \left(\sum_i \sigma_i(X)^p \right)^{\frac{1}{p}}, \quad (13)$$

where $\sigma_i(X)$ are the singular values of X . In the case $p = 1$, this penalty is in fact a norm, coinciding nuclear norm of a matrix. Equivalently, the Schatten- p quasi-norms can be defined in terms of the following trace formula:¹

$$\|X\|_p = \text{Tr}[(X * X)^{\frac{p}{2}}]^{\frac{1}{p}}. \quad (14)$$

We also define the penalty

$$\|X\|_0^0 = \sum_i \log(\sigma_i(X)) = \frac{1}{2} \log \det(X * X) \quad (15)$$

where $\log$ denotes the natural logarithm, which can be viewed as the limiting case² of (13)

as $p \rightarrow 0$. Note that the penalty $\|\cdot\|_1$ is convex, but the penalty $\|\cdot\|_p$ is non-convex for $0 < p < 1$.

In the case where the measurements are corrupted by noise, i.e. $\mathbf{b} = \mathbf{A}\mathbf{x}_0 + \mathbf{n}$, where $\mathbf{n}$ is assumed to be a vector of i.i.d. complex white Gaussian noise, we relax the equality constraint in (12) by incorporating a data fidelity term into the objective:

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \lambda \|\mathcal{T}(\mathbf{x})\|_p^p. \quad (16)$$

Here $\lambda > 0$ is a tunable regularization parameter balancing data fidelity and the degree to which $\mathcal{T}(\mathbf{x})$ is assumed to be low-rank.

¹The fractional q -th power of a positive semi-definite matrix $\mathbf{Y} \in \mathbb{C}^{n \times n}$ is defined as $\mathbf{Y}^q = \sum_i \lambda_i^q \mathbf{v}_i \mathbf{v}_i^*$, where $\{\mathbf{v}_i\}_{i=1}^n$ is a orthonormal set of eigenvectors for $\mathbf{Y}$ with associated non-negative eigenvalues $\{\lambda_i\}_{i=1}^n$.

²More precisely, for any scalar $t > 0$ we have $\lim_{p \rightarrow 0} \frac{t^p - 1}{p} = \log(t)$. Therefore if $X \in \mathbb{C}^{M \times N}$ has no zero singular values,

$$\lim_{p \rightarrow 0} \frac{1}{p} \left(\|X\|_p^p - \frac{N}{p} \right) = \lim_{p \rightarrow 0} \sum_{i=1}^N \frac{\sigma_i(X)^p - 1}{p} = \sum_{i=1}^N \log(\sigma_i(X))$$

C. Convolutional structured matrix lifting model

By a *matrix lifting* we mean any linear operator $\mathcal{T}$ mapping a vector $\mathbf{x} \in \mathbb{C}^m$ to a “lifted matrix” $\mathcal{T}(\mathbf{x}) \in \mathbb{C}^{M \times N}$. In this work we focus on a general class of matrix liftings having a similar convolution structure to (4) and (9). For subsets Λ and Γ of $\mathbb{Z}^d$, to be defined in the sequel, define $\text{Toep}_d(\mathbf{y})$ to be the matrix representing linear convolution with the d -dimensional array $\mathbf{y} = \{y[\mathbf{k}]: \mathbf{k} \in \mathbb{C} \subset \mathbb{Z}^d\}$, such that

$$\text{Toep}_d(\mathbf{y})\mathbf{h} = \mathbf{y} * \mathbf{h}, \quad (17)$$

where $\mathbf{h} = \{h[\mathbf{k}]: \mathbf{k} \in \Lambda \subset \mathbb{Z}^d\}$ and $\mathbf{y} * \mathbf{h}$ is the array defined by

$$(\mathbf{y} * \mathbf{h})[\mathbf{k}] = \sum_{\ell \in \Lambda} y[\mathbf{k} - \ell]h[\ell], \quad \mathbf{k} \in \Gamma \subset \mathbb{Z}^d. \quad (18)$$

Here the convolution is restricted to the set of “valid” indices $\Gamma \subset \mathbb{Z}^d$ satisfying $\mathbf{k} \in \Gamma$ only if $\mathbf{k} - \ell \in \Lambda$ for all $\ell \in \Lambda$, i.e., the set of indices for which the sum in (18) is well-defined. We call $\mathbf{h}$ a filter, and Λ the *filter support*, the *data support*, and Γ the set of *valid indices*; see Figure 2 for an illustration in dimension $d = 2$. When the index sets Λ and Γ are rectangular, the type of matrix structure exhibited by $\text{Toep}_d(\mathbf{y})$ has been called *multi-level Toeplitz* [46], and we adopt this term here as well. For example, $\text{Toep}_2(\mathbf{y})$ is a block Toeplitz matrix with Toeplitz blocks, meaning the blocks are arranged in a Toeplitz pattern and each block is itself Toeplitz.

In this work we consider structured matrix liftings that have a vertical block structure where each block is multi-level Toeplitz:

$$\mathcal{T}(\mathbf{x}) = \begin{pmatrix} \mathcal{T}_1(\mathbf{x}) \\ \vdots \\ \mathcal{T}_K(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \text{Toep}_d(\mathbf{M}_1\mathbf{x}) \\ \vdots \\ \text{Toep}_d(\mathbf{M}_K\mathbf{x}) \end{pmatrix} \quad (19)$$

Here each $\mathbf{M}_j$, $1 \leq j \leq K$ denotes some linear transformation. Typical choices include the identity, a reshaping operator, or element-wise multiplication by a set of weights. For example, the gradient weighted lifting (9) has two blocks where $\mathbf{M}_1$ and $\mathbf{M}_2$ are diagonal matrices representing weightings by Fourier derivatives $j2\pi k_x$ and $j2\pi k_y$, respectively. Different weighting schemes have also been considered in related formulations [5], [8], [30].

If a lifted matrix $\mathcal{T}(\mathbf{x})$ in the form (19) is rank deficient, then every non-trivial vector $\mathbf{h} \in \text{null } \mathcal{T}(\mathbf{x})$ can be interpreted as an *annihilating filter* for each $\mathbf{y}_j = \mathbf{M}_j\mathbf{x}$, $1 \leq j \leq K$, in the sense that

$$(\mathbf{y}_j * \mathbf{h})[\mathbf{k}] = 0 \text{ for all } \mathbf{k} \in \Gamma. \quad (20)$$

In other words, if the structured matrix (19) is low-rank, it means there exists a large collection of linearly independent annihilating filters for the data defining the matrix. This generalizes the annihilating filter formulation that is central to finite-rate-of-innovation modeling [47].

Finally, we note that any (multi-level) Hankel matrix can be rewritten as a (multi-level) Toeplitz matrix through a permutation of its rows and columns, which has no effect on the rank of the matrix. In this way we can also incorporate (block) multi-level Hankel matrix liftings into the model (19), such as those proposed in [3], [6].

III. Iteratively reweighted least squares algorithms for structured low-rank matrix recovery

We adapt an iteratively reweighted least squares (IRLS) approach to minimizing (12) or (16), originally proposed in [20], [21] in the low-rank matrix completion setting; see also [48] for an alternative matrix factorization approach to Schatten- p norm minimization. The IRLS approach is motivated by the observation that the Schatten- p quasi-norm of any matrix $\mathbf{X}$ may be re-expressed as a weighted Frobenius norm:

$$\|\mathbf{X}\|_p^p = \text{Tr}[(\mathbf{X} * \mathbf{X}) \underbrace{(\mathbf{X} * \mathbf{X})^{\frac{p}{2}-1}}_{\mathbf{H}}] = \|\mathbf{X}\sqrt{\mathbf{H}}\|_F^2, \quad (21)$$

for all $0 < p \leq 1$, provided $\mathbf{X}$ has no zero singular values so that $\mathbf{H}$ is well-defined. Here $\sqrt{\mathbf{H}}$ denotes any matrix³ satisfying $\mathbf{H} = (\sqrt{\mathbf{H}})(\sqrt{\mathbf{H}})^*$. This suggests an iterative algorithm for minimizing the Schatten- p quasi-norm that alternates between updating a weight matrix $\mathbf{H}$ and solving a weighted least squares problem with $\mathbf{H}$ fixed. In particular, substituting $\mathbf{X} = \mathcal{T}(\mathbf{x})$ in the IRLS- p algorithm presented in [21] gives the following iterative scheme for solving (16):

$$\mathbf{H}_n = [\mathcal{T}(\mathbf{x}^{(n-1)})^* \mathcal{T}(\mathbf{x}^{(n-1)})]^{\frac{p}{2}-1}$$

$$\mathbf{x}^{(n)} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathcal{T}(\mathbf{x})\sqrt{\mathbf{H}_n}\|_F^2.$$

³One choice of $\sqrt{\mathbf{H}}$ is the standard matrix square root $\mathbf{H}^{\frac{1}{2}} = \sum_i \lambda_i^{\frac{1}{2}} \mathbf{v}_i \mathbf{v}_i^*$ where $(\lambda_i, \mathbf{v}_i)$ denotes an eigen-pair, but we will show later it is computationally advantageous to use a different choice of square root in our algorithm.

To ensure the matrix defining $\mathbf{H}_n$ is invertible, and to improve numerical stability, following the approach in [21] we introduce a smoothing parameter $\epsilon_n > 0$ to the $\mathbf{H}$ -update as

$$\mathbf{H}_n = \left(\mathcal{T}(\mathbf{x}^{(n-1)}) * \mathcal{T}(\mathbf{x}^{(n-1)}) + \epsilon_n \mathbf{I} \right)^{-\frac{1}{2}}, \text{ satisfying } \epsilon_n \rightarrow 0 \text{ as } n \rightarrow \infty. \text{ Pseudo-code}$$

summarizing this approach is given in Algorithm 1. In the Appendix we show that the IRLS- p algorithm can be derived as a majorization-minimization (MM) algorithm [49] for the objective (16). Due to properties of MM algorithms, this ensures that the IRLS- p iterates monotonically decrease the cost function (12). Moreover, in the convex case $p = 1$ (i.e. nuclear norm regularization), the iterates are guaranteed to converge to the global minimum of (16); in the non-convex case $p \in [0, 1)$, convergence to the global optimality cannot be ensured, but the iterates are still guaranteed to converge to a stationary point. We refer the reader to [21], [22] for a more detailed convergence analysis of the IRLS- p algorithm in the context of low-rank matrix completion.

Algorithm 1

IRLS- p algorithm for structured low-rank matrix recovery

```

Initialize  $\mathbf{x}^{(0)}$  and choose  $\epsilon_0 > 0$ ;
for  $n = 1$  to  $I_{max}$  do
     $\mathbf{H}_n = [\mathcal{T}(\mathbf{x}^{(n-1)}) * \mathcal{T}(\mathbf{x}^{(n-1)}) + \epsilon_{n-1} \mathbf{I}]^{\frac{p}{2}-1}$ ;
     $\mathbf{x}^{(n)} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathcal{T}(\mathbf{x}) \sqrt{\mathbf{H}_n}\|_F^2$ ;
    Choose  $\epsilon_n$  such that  $0 < \epsilon_n \leq \epsilon_{n-1}$ ;
end

```

A. Challenges with the direct IRLS approach

While the IRLS- p algorithm has shown several advantages in the low-rank matrix completion setting [20]–[22], the direct adaptation of IRLS- p to the structured matrix liftings considered in this work is computationally prohibitive for large-scale problems. To see why, first note that the weight matrix update requires computing an inverse power of the Gram matrix $\mathcal{T}(\mathbf{x}) * \mathcal{T}(\mathbf{x}) \in \mathbb{C}^{N \times N}$. Computing this Gram matrix directly by matrix multiplication will be costly when the inner dimension $M \gg N$, and inverting the resulting matrix can be unstable due to round-off errors. Instead, a more stable and efficient method to compute the inverse is by an SVD of the $M \times N$ matrix $\mathcal{T}(\mathbf{x})$, which requires $\mathcal{O}(MN^2)$ flops to compute. In the case that $\mathcal{T}(\mathbf{x})$ is approximately low-rank, one can instead compute only the top r singular values and singular vectors either using deterministic methods, such as Lanczos bidiagonalization [50], or by randomized methods [51], at a reduced cost of $\mathcal{O}(MNr)$ flops. However, these approaches are still dominated by a computational cost that is linear in M , and will still be prohibitively slow when M is large or when the matrix $\mathcal{T}(\mathbf{x})$ is not sufficiently low-rank. Note these are the same fundamental costs involved in other SVD-based algorithms for structured low-rank matrix recovery; see the Supplementary Materials for more details.

The IRLS- p algorithm additionally requires solving the following least-squares problem at each iteration:

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \lambda \|\mathcal{T}(\mathbf{x})\sqrt{\mathbf{H}}\|_F^2. \quad (22)$$

To give an idea of the costs involved in solving (22) via an iterative solver, consider the case where the matrix lifting has the form $\mathcal{T}(\mathbf{x}) = \text{Toep}_d(\mathbf{x})$, i.e., $\mathcal{T}(\mathbf{x})\mathbf{h} = \mathbf{x} * \mathbf{h}$. The challenge is that standard iterative methods for solving (22), such as the CG or LSQR algorithm [52], would require $\mathcal{O}(N)$ multi-dimensional FFTs per iteration, where $N = |\Lambda|$ is the total number of filter coefficients. However, N can be on the order of hundreds or thousands for the problems considered in this work. Therefore, standard methods for solving the least-squares problem (22) will also be prohibitively costly for large-scale problems.

IV. Proposed GIRAF Algorithm

As observed in the previous section, the direct IRLS algorithm will not scale well to large problem instances because of two challenges: 1) computing the weight matrix update requires a large-scale SVD, and 2) computing a solution to the least squares problem requires a prohibitive number of FFTs. In this section we propose a novel approximation of the problem formulation (16) to overcome these difficulties. The main idea is to *approximate the structured matrix lifting* $\mathcal{T}$ in a systematic way such that the complexity of the resulting IRLS subproblems simplify, while preserving the rank structure of the lifting as best as possible.

A. Half-circulant approximation of a Toeplitz matrix

Our approximation is based on the observation that every multi-level Toeplitz matrix can be embedded in a larger multi-level circulant matrix. Specifically, the d -level Toeplitz matrix $\text{Toep}_d(\mathbf{y})$ built with filter support $\Lambda \subset \mathbb{Z}^d$ and valid index set $\Gamma \subset \mathbb{Z}^d$ can always be expressed as:

$$\text{Toep}_d(\mathbf{y}) = \mathbf{P}_\Gamma \mathbf{C}(\mathbf{y}) \mathbf{P}_\Lambda^* \quad (23)$$

Here $\mathbf{C}(\mathbf{y}) \in \mathbb{C}^{L \times L}$ is a matrix representing convolution with the array $\mathbf{y}$, $\mathbf{P}_\Gamma \in \mathbb{C}^{M \times L}$ is a matrix representing restriction to valid set Γ , and $\mathbf{P}_\Lambda^* \in \mathbb{C}^{L \times N}$ represents a zero-padding outside the filter support Λ . The matrices $\mathbf{P}_\Gamma$ and $\mathbf{P}_\Lambda^*$ act as row-restriction and column-restriction operators, respectively; see Figure (3) for an illustration. Because of the restriction matrix $\mathbf{P}_\Gamma$, we can assume $\mathbf{C}(\mathbf{y})$ represents a multi-dimensional *circular convolution*, i.e., $\mathbf{C}(\mathbf{y})$ is multi-level circulant, provided the convolution takes place on a sufficiently large rectangular grid to avoid wrap-around boundary effects. In particular, the rectangular circular convolution grid should be at least as large as the data support $\Gamma \subset \mathbb{Z}^d$. Without loss of generality, from now on we assume the data support is rectangular so that it coincides with the circular convolution grid.

To simplify subsequent derivations, in this section we assume the structured matrix lifting $\mathcal{T}(\mathbf{x})$ consists of a single multi-level Toeplitz block: $\mathcal{T}(\mathbf{x}) = \text{Toep}_d(\mathbf{M}\mathbf{x})$. According to (23), $\mathcal{T}(\mathbf{x})$ can be written as

$$\mathcal{T}(\mathbf{x}) = \mathbf{P}_\Gamma \mathbf{C}(\mathbf{M}\mathbf{x}) \mathbf{P}_\Lambda^* \quad (24)$$

We propose approximating the structured matrix lifting $\mathcal{T}(\mathbf{x})$ with the surrogate lifting $\widetilde{\mathcal{T}}(\mathbf{x})$ defined by

$$\widetilde{\mathcal{T}}(\mathbf{x}) = \mathbf{C}(\mathbf{M}\mathbf{x}) \mathbf{P}_\Lambda^* \quad (25)$$

i.e., we omit the left-most row restriction operator $\mathbf{P}_\Gamma$ from $\mathcal{T}(\mathbf{x})$ so that $\widetilde{\mathcal{T}}(\mathbf{x})$ is the full vertical section of the multi-level circulant matrix $\mathbf{C}(\mathbf{M}\mathbf{x})$. In general, we say a matrix $\mathbf{X} \in \mathbb{C}^{M \times N}$, $N \geq M$, is (multi-level) *half-circulant* if it can be written as $\mathbf{X} = \mathbf{C}\mathbf{P}^*$, where $\mathbf{C} \in \mathbb{C}^{M \times M}$ is (multi-level) circulant and $\mathbf{P} \in \mathbb{C}^{N \times M}$ is a restriction matrix, i.e., $\mathbf{X}$ is obtained by selecting N full columns from $\mathbf{C}$.

Inserting the approximation (25) into (16), we propose solving

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \lambda \|\widetilde{\mathcal{T}}(\mathbf{x})\|_p^p. \quad (26)$$

using the IRLS- p algorithm. In other words, rather than penalizing the Schatten- p norm of the exact multilevel Toeplitz lifted matrix, we penalize its multi-level half-circulant approximation instead. We can justify this approach by considering the effect the half-circulant approximation (25) has on the singular values of the lifting. Recall that $\mathbf{C}(\mathbf{y})$ denotes circular convolution with the array $\mathbf{y}$ on a rectangular grid $\mathbb{Z}^d$ containing the index set Γ . Hence, after a rearrangement of rows, we can always write $\mathcal{T}(\mathbf{x})$ submatrix of $\widetilde{\mathcal{T}}(\mathbf{x})$:

$$\widetilde{\mathcal{T}}(\mathbf{x}) = \begin{bmatrix} \mathcal{T}(\mathbf{x}) \\ \mathcal{B}(\mathbf{x}) \end{bmatrix} = \begin{bmatrix} \mathbf{P}_\Gamma \mathbf{C}(\mathbf{M}\mathbf{x}) \mathbf{P}_\Lambda^* \\ \mathbf{P}_{\Gamma^C} \mathbf{C}(\mathbf{M}\mathbf{x}) \mathbf{P}_\Lambda^* \end{bmatrix}, \quad (27)$$

where Γ^C is the set complement of Γ inside the circular convolution grid. In other words, $\widetilde{\mathcal{T}}(\mathbf{x})$ is the unique matrix obtained by augmenting the rows of $\mathcal{T}(\mathbf{x})$ to make it (multi-level) half-circulant.

As a consequence of the embedding (27), the singular values of $\mathcal{T}(\mathbf{x})$ are bounded by those of $\widetilde{\mathcal{T}}(\mathbf{x})$:

$$\sigma_i(\mathcal{T}(\mathbf{x})) \leq \sigma_i(\widetilde{\mathcal{T}}(\mathbf{x})) \text{ for } i = 1, \dots, N. \quad (28)$$

(see Corollary 3.1.3 in [53]). Hence, the Schatten- p quasi-norm of $\mathcal{T}(\mathbf{x})$ is also bounded by that of $\widetilde{\mathcal{T}}(\mathbf{x})$:

$$\|\mathcal{T}(\mathbf{x})\|_p^p \leq \|\widetilde{\mathcal{T}}(\mathbf{x})\|_p^p \text{ for all } 0 \leq p \leq 1. \quad (29)$$

Moreover, if in minimizing $\|\widetilde{\mathcal{T}}(\mathbf{x})\|_p^p$ we obtain a low-rank matrix $\widetilde{\mathcal{T}}(\mathbf{x}')$ with rank r , the singular value bounds (28) show that $\mathcal{T}(\mathbf{x}')$ is also low-rank with rank r . This shows that it is reasonable to use $\|\widetilde{\mathcal{T}}(\mathbf{x})\|_p^p$ as a surrogate penalty for $\|\mathcal{T}(\mathbf{x})\|_p^p$.

Empirically, we find that when the data support is sufficiently large relative to the filter size, i.e., when $M \gg N$ where $\mathcal{T}(\mathbf{x}) \in \mathbb{C}^{M \times N}$, then using the surrogate lifting $\widetilde{\mathcal{T}}(\mathbf{x})$ results in negligible approximation errors (see Figure 6). When the data support is small, we recommend solving for the data array $\mathbf{x}$ on a slightly larger “oversampled” grid $\mathbf{x}'$ to keep the approximation error low. Typically we take $\mathbf{x}' = \mathbf{x} + 2\Lambda$, i.e. we pad the data support with an extra margin having the size of the filter support. After solving the problem on the oversampled grid $\mathbf{x}'$, we finally restrict the solution to the desired data support $\mathbf{x}$. We study the effect of using an oversampled grid in Section V (see Figure 7).

We now show the IRLS algorithm applied to (26) results in subproblems with significantly reduced complexity due to the half-circulant structure of the approximated lifting. In particular, we will show the algorithm can be interpreted as alternating between: (1) solving for an annihilating filter for the data, and (2) solving for the data best annihilated by the filter in a least-squares sense. Due to this interpretation, we call this approach the Generic Iterative Reweighted Annihilating Filter (GIRAF) algorithm.

B. GIRAF least squares subproblem

1) Reformulation as least squares annihilation—For a fixed weight matrix $\mathbf{H}$, the GIRAF least squares subproblem has the form

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\widetilde{\mathcal{T}}(\mathbf{x})\sqrt{\mathbf{H}}\|_F^2 \quad (30)$$

Under the half-circulant assumption (25), we can re-express the Frobenius norm above as

$$\|\widetilde{\mathcal{T}}(\mathbf{x})\mathbf{H}^{\frac{1}{2}}\|_F^2 = \sum_{i=1}^N \|\widetilde{\mathcal{T}}(\mathbf{x})\mathbf{h}_i\|_2^2 = \sum_{i=1}^N \|\mathbf{C}(\mathbf{M}\mathbf{x})\mathbf{P}_{\Lambda}^*\mathbf{h}_i\|_2^2 \quad (31)$$

where $\mathbf{h}_i$ is the i th column of $\sqrt{\mathbf{H}}$. Because circular convolution is commutative, we can write $\mathbf{C}(\mathbf{M}\mathbf{x})\mathbf{P}_\Lambda^*\mathbf{h}_i = \mathbf{C}_i\mathbf{M}\mathbf{x}$, where $\mathbf{C}_i = \mathbf{C}(\mathbf{P}_\Lambda^*\mathbf{h}_i)$ is a multi-level circulant matrix representing circular convolution with the zero-padded filter $\mathbf{h}_i$ on the rectangular grid $\mathbb{C} \subset \mathbb{Z}^d$. Therefore, $\mathbf{C}_i = \mathbf{F}\mathbf{D}_i\mathbf{F}^*$ where $\mathbf{F}$ is the DFT on $\mathbb{C}$, and $\mathbf{D}_i$ the diagonal matrix having entries given by the array $\mathbf{d}_i = \mathbf{F} * \mathbf{P}_\Lambda^*\mathbf{h}_i$, the inverse DFT of the zero-padded filter $\mathbf{h}_i$. This allows us to simplify the right-hand side of (31) as

$$\sum_{i=1}^N \|\mathbf{C}_i\mathbf{M}\mathbf{x}\|_2^2 = \mathbf{x}^* \mathbf{M}^* \left(\sum_{i=1}^N \mathbf{C}_i^* \mathbf{C}_i \right) \mathbf{M}\mathbf{x} = \mathbf{x}^* \mathbf{M}^* \mathbf{F} \left(\sum_{i=1}^N \mathbf{D}_i^* \mathbf{D}_i \right) \mathbf{F}^* \mathbf{M}\mathbf{x} = \|\mathbf{D}^{\frac{1}{2}} \mathbf{F}^* \mathbf{M}\mathbf{x}\|_2^2,$$

where we have set $\mathbf{D} = \sum_{i=1}^N \mathbf{D}_i^* \mathbf{D}_i$, which is again diagonal with entries $\mathbf{d} \in \mathbb{C}^{|\mathbb{C}|}$ given by

$$\mathbf{d} = \sum_{i=1}^N |\mathbf{d}_i|^2, \quad \mathbf{d}_i = \mathbf{F}^* \mathbf{P}_\Lambda^* \mathbf{h}_i \text{ for all } i = 1, \dots, N, \quad (32)$$

where $|\cdot|$ is applied element-wise. Therefore, (30) transforms into the weighted least squares problem:

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{D}^{\frac{1}{2}} \mathbf{F}^* \mathbf{M}\mathbf{x}\|_2^2. \quad (33)$$

Observe that by making use of the half-circulant assumption we have effectively reduced the working dimension of the least-squares problem back down to the dimension of the original decision variable $\mathbf{x}$.

Computing the weights $\mathbf{d}$ according to the formula (32) can be costly for large-scale problems since it requires N large-scale FFTs. A more efficient approach is to pre-compute the filter $\mathbf{h}$ defined by

$$\mathbf{h} = \sum_{i=1}^N (\tilde{\mathbf{h}}_i * \mathbf{h}_i) \quad (34)$$

where $\tilde{\mathbf{h}}_i$ is the reversed, conjugated filter defined by $\tilde{\mathbf{h}}_i[\mathbf{k}] = \overline{\mathbf{h}_i[-\mathbf{k}]}$ for all $\mathbf{k} \in \Lambda$. Note that $\mathbf{h}$ has coefficients supported within $2\Lambda := \{\mathbf{k} + \mathbf{\ell}, \mathbf{\ell} \in \Lambda\}$ since each $\mathbf{h}_i$ is supported within Λ . Therefore, applying the DFT convolution theorem to (32), $\mathbf{d}$ can be obtained by $\mathbf{d} = \mathbf{F}^* \mathbf{P}_{2\Lambda}^* \mathbf{h}$, which after computing $\mathbf{h}$, requires only one large-scale FFT. We call filter $\mathbf{h}$ defined in (34) the *re-weighted annihilating filter* and $\mathbf{d}$ the *annihilation weights*.

2) ADMM solution of least-squares annihilation—The linear least-squares problem (33) can be readily solved with an off-the-shelf solver, such as the CG or LSQR algorithm [52]. However, the problem is poorly conditioned when the annihilation weights $\mathbf{d}$ are close to zero, which is expected to happen as the iterations progress. Therefore, solving (33) efficiently during the course of the GIRAF algorithm will require a robust preconditioning strategy. However, because the transformation $\mathbf{M}$ is not always invertible, designing an all-purpose preconditioner for (33) is challenging. Instead, we adopt an approach which allows us to solve a series of subproblems with predictably good conditioning. The approach is based on the following variable splitting:

$$\min_{\mathbf{x}} \left\| \mathbf{A}\mathbf{x} - \mathbf{b} \right\|_2^2 + \lambda \left\| \mathbf{D}^{\frac{1}{2}} \mathbf{y} \right\|_2^2 \quad \text{subject to} \quad \mathbf{F}\mathbf{y} = \mathbf{M}\mathbf{x}. \quad (35)$$

The equality constrained problem (35) can be efficiently solved with the alternating directions method of multipliers (ADMM) algorithm [54], which results in the following iterative scheme:

$$\mathbf{y}^{(n)} = \arg \min_{\mathbf{y}} \left\| \mathbf{D}^{\frac{1}{2}} \mathbf{y} \right\|_2^2 + \gamma \left\| \mathbf{y} - \mathbf{F}^* \mathbf{M} \mathbf{x}^{(n-1)} + \mathbf{q}^{(n-1)} \right\|_2^2, \quad (36)$$

$$\mathbf{x}^{(n)} = \arg \min_{\mathbf{x}} \left\| \mathbf{A}\mathbf{x} - \mathbf{b} \right\|_2^2 + \gamma \lambda \left\| \mathbf{y}^{(n)} - \mathbf{F}^* \mathbf{M} \mathbf{x} + \mathbf{q}^{(n-1)} \right\|_2^2, \quad (37)$$

$$\mathbf{q}^{(n)} = \mathbf{q}^{(n-1)} + \mathbf{F} \mathbf{y}^{(n)} - \mathbf{M} \mathbf{x}^{(n)}, \quad (38)$$

where $\mathbf{q}$ represents a vector of Lagrange multipliers, and $\gamma > 0$ is a fixed parameter that can be tuned to improve the conditioning of the subproblems. Subproblems (36) and (37) are both quadratic and can be solved efficiently: (36) has the exact solution

$$\mathbf{y}^{(n)} = (\mathbf{D} + \gamma \mathbf{I})^{-1} \left(\gamma \mathbf{F}^* [\mathbf{M} \mathbf{x}^{(n-1)} - \mathbf{q}] \right). \quad (39)$$

Since $\mathbf{D} + \gamma \mathbf{I}$ is diagonal, its inverse acts as an element-wise division. Likewise, the solution of (37) can be obtained as

$$\mathbf{x}^{(n)} = (\mathbf{A}^* \mathbf{A} + \lambda \gamma \mathbf{M}^* \mathbf{M})^{-1} [\mathbf{A}^* \mathbf{b} + \lambda \gamma \mathbf{M}^* \mathbf{F} (\mathbf{y}^{(n)} + \mathbf{q}^{(n-1)})]. \quad (40)$$

In many problems of interest, both $\mathbf{A}^*\mathbf{A}$ and $\mathbf{M}^*\mathbf{M}$ are diagonal, which means $\mathbf{x}^{(n)}$ is also obtained by an efficient element-wise division. In this case, aside from the element-wise operations, the computational cost of one pass of ADMM iterations (36)–(38) is three FFTs. When either $\mathbf{A}^*\mathbf{A}$ or $\mathbf{M}^*\mathbf{M}$ are not diagonal, an approximate solution to (37) can be found by a few passes of an iterative solver instead. Even with inexact updates of the $\mathbf{x}$ -subproblem, we are still guaranteed convergence of the overall ADMM algorithm under fairly broad conditions; see, e.g., [54], [55].

Update (39) suggests choosing γ adaptively before solving each ADMM subproblem according to the magnitude of the annihilation weights $\mathbf{d} = \{d_j\}_{j=1}^L$ of $\mathbf{D}$ to ensure good conditioning and fast convergence of the ADMM scheme. We recommend choosing $\gamma = (\max_j d_j)/\delta$ for $\delta \geq 1$; we study the effect of δ in Section V.

C. GIRAF re-weighted annihilating filter subproblem

1) Construction of Gram matrix—At each iteration, the GIRAF algorithm requires updating a weight matrix $\mathbf{H}$ according to:

$$\mathbf{H} = \left(\widetilde{\mathcal{T}}(\mathbf{x})^* \widetilde{\mathcal{T}}(\mathbf{x}) + \epsilon \mathbf{I} \right)^{\frac{p}{2} - 1} \quad (41)$$

Rather than computing $\mathbf{H}$ via a full or partial SVD of the tall $M \times N$ matrix $\widetilde{\mathcal{T}}(\mathbf{x})$, as is recommended in [21], we propose computing $\mathbf{H}$ via the eigendecomposition of the smaller $N \times N$ Gram matrix $\mathbf{G} = \widetilde{\mathcal{T}}(\mathbf{x})^* \widetilde{\mathcal{T}}(\mathbf{x})$. This is because the half-circulant approximation (25) allows us to compute $\mathbf{G}$ efficiently in a matrix-free manner, as we now show. From (25), we have

$$\mathbf{G} = \mathbf{P}_\Lambda \mathbf{C}(\mathbf{M}\mathbf{x})^* \mathbf{C}(\mathbf{M}\mathbf{x}) \mathbf{P}_\Lambda^*. \quad (42)$$

The restriction matrices $\mathbf{P}_\Lambda$ and $\mathbf{P}_\Lambda^*$ in (42) extract an $N \times N$ block of $\mathbf{C}(\mathbf{M}\mathbf{x})^* \mathbf{C}(\mathbf{M}\mathbf{x})$ corresponding to the intersection of the rows and columns indexed by the filter support set Λ . Note that the product $\mathbf{C}(\mathbf{M}\mathbf{x})^* \mathbf{C}(\mathbf{M}\mathbf{x}) = \mathbf{C}(\mathbf{g})$ is again a multi-level circulant matrix whose entries are generated by the array $\mathbf{g} = \mathbf{F} |\mathbf{F}^* \mathbf{M}\mathbf{x}|^2$ with the operation $|\cdot|^2$ is understood element-wise. Therefore, we can build $\mathbf{G}$ by performing a sliding-window operation that extracts every patch of size Λ from the array $\mathbf{g}$. In particular, because of the restriction matrices $\mathbf{P}_\Lambda$ and $\mathbf{P}_\Lambda^*$ in (42), we only need to consider patches coming from $\mathbf{g}$ restricted to the index set $2\Lambda = \{\mathbf{k} + \ell \mathbf{k}, \ell \in \Lambda\}$. Aside from this sliding-window operation, the main cost in computing $\mathbf{G}$ is two FFTs.

2) Update of re-weighted annihilating filter—According to section IV-B, the GIRAF least squares subproblem can be interpreted as an annihilation of the data subject to a filter $\mathbf{h}$ determined by the columns $\mathbf{h}_1, \dots, \mathbf{h}_N$ of $\sqrt{\mathbf{H}}$. We now show how to update $\mathbf{h}$ directly from

an eigendecomposition of the Gram matrix $\mathbf{G} = \widetilde{\mathcal{T}}(\mathbf{x})^* \widetilde{\mathcal{T}}(\mathbf{x})$, rather than forming $\sqrt{\mathbf{H}}$ explicitly. Let $\mathbf{G} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^*$ where the columns $\{\mathbf{v}_i\}_{i=1}^N$ of $\mathbf{V}$ are an orthonormal basis of eigenvectors and $\mathbf{\Lambda}$ is diagonal matrix of the associated eigenvalues $\{\lambda_i\}_{i=1}^N$. Then the weight matrix update (41) reduces to

$$\mathbf{H} = [\mathbf{V}(\mathbf{\Lambda} + \varepsilon \mathbf{I})\mathbf{V}^*]^{\frac{p}{2}-1} = \mathbf{V}(\mathbf{\Lambda} + \varepsilon \mathbf{I})^{\frac{p}{2}-1} \mathbf{V}^*. \quad (43)$$

One choice of the matrix square root $\sqrt{\mathbf{H}}$ is

$$\sqrt{\mathbf{H}} = \mathbf{V}(\mathbf{\Lambda} + \varepsilon \mathbf{I})^{-\frac{q}{2}} = [(\lambda_1 + \varepsilon)^{-\frac{q}{2}} \mathbf{v}_1, \dots, (\lambda_M + \varepsilon)^{-\frac{q}{2}} \mathbf{v}_M]. \quad (44)$$

where $q = 1 - \frac{p}{2}$. However, the least squares subproblem of the algorithm only needs as input the filter $\mathbf{h}$ defined in (34), which is determined by the columns $\mathbf{h}_i$ of $\sqrt{\mathbf{H}}$. Making the substitutions $\mathbf{h}_i = (\lambda_i + \varepsilon)^{-\frac{q}{2}} \mathbf{v}_i$ in (34) gives the following update:

$$\mathbf{h} = \sum_{i=1}^N (\lambda_i + \varepsilon)^{-q} (\widetilde{\mathbf{v}}_i * \mathbf{v}_i); \quad q = 1 - \frac{p}{2}. \quad (45)$$

Note that the weight $(\lambda_i + \varepsilon)^{-q}$, $q > 0$, is large only when the filter $\mathbf{v}_i$ is close to the null space of $\widetilde{\mathcal{T}}(\mathbf{x})$, i.e., when $\mathbf{v}_i$ is an annihilating filter. Therefore, $\mathbf{h}$ can be thought of as a weighted average of all the annihilating filters for $\widetilde{\mathcal{T}}(\mathbf{x})$. Notice that the only effect of changing the Schatten- p penalty parameter $0 < p < 1$ is to change the exponent q in the computation of (45). Smaller values of p (larger values of q) are more likely to promote low-rank solutions at the iterations progress, since the weights on the filters close to the null space will be higher; see Figure 4 for an illustration of this effect.

Algorithm 2GIRAF- p algorithm for structured low-rank matrix recovery

Initialize $\mathbf{x}^{(0)}$ and choose $\epsilon_0 > 0$;
for $n = 1$ **to** $I_{\max}$ **do**
 Step 1: Annihilating Filter Update
 Build $\mathbf{G} = \check{\mathcal{T}}(\mathbf{x}^{(n-1)})^* \check{\mathcal{T}}(\mathbf{x}^{(n-1)})$ using (42);
 Find eigendecomposition $(\lambda_i, \mathbf{v}_i)_{i=1}^N$ of $\mathbf{G}$;
 Compute re-weighted annihilating filter
 $\mathbf{h} = \sum_{i=1}^N (\lambda_i + \epsilon_{n-1})^{\frac{p}{2}-1} (\tilde{\mathbf{v}}_i * \mathbf{v}_i)$;
 Convert filter to weights $\mathbf{D} = \text{diag}(\mathbf{F}^* \mathbf{P}_{2\Lambda}^* \mathbf{h})$
 Step 2: Least Squares Annihilation
 Solve least-squares problem:
 $\mathbf{x}^{(n)} = \arg \min_{\mathbf{x}} \|\mathbf{D}^{\frac{1}{2}} \mathbf{F}^* \mathbf{M} \mathbf{x}\|^2 + \lambda \|\mathbf{A} \mathbf{x} - \mathbf{b}\|^2$
 by ADMM iterations (36)–(38) or CG;
 Choose ϵ_n such that $0 < \epsilon_n \leq \epsilon_{n-1}$;
end

D. Extension to liftings with multiple blocks

To simplify the derivation of the GIRAF algorithm we assumed that the original matrix lifting $\mathcal{T}(\mathbf{x})$ consists of a single multi-level Toeplitz block. However, the GIRAF algorithm can easily be modified to accommodate liftings with a vertical block structure, as in (19), by applying the half-circulant approximation to each block. In this case, one can show the GIRAF least-squares problem simplifies to

$$\min_{\mathbf{x}} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2^2 + \lambda \sum_{i=1}^K \|\mathbf{D}^{\frac{1}{2}} \mathbf{F}^* \mathbf{M}_i \mathbf{x}\|_2^2. \quad (46)$$

where $\mathbf{D} = \text{diag}(\mathbf{d})$, and the annihilation weights $\mathbf{d}$ are computed the same as in (32) and (45) from an eigendecomposition of the Gram matrix $\mathbf{G} = \sum_{i=1}^K \check{\mathcal{T}}_i(\mathbf{x})^* \check{\mathcal{T}}_i(\mathbf{x})$.

E. Implementation Details

An overview of the GIRAF algorithm is given in Algorithm 2. Empirically we find that several of the heuristics in [44] for setting the smoothing parameter ϵ_n work well for the GIRAF algorithm, as well. In particular, we set $\epsilon_0 = \lambda_{\max}/100$ where $\lambda_{\max}$ is the maximum eigenvalue of the Gram matrix $\check{\mathcal{T}}(\mathbf{x}^{(0)})^* \check{\mathcal{T}}(\mathbf{x}^{(0)})$, which is obtained as a by-product of the first iterate of the algorithm. We recommend decreasing the smoothing parameter ϵ_n exponentially as $\epsilon_n = \epsilon_0 (\eta)^{-n}$, where $\eta > 1$ is a fixed parameter. We find that $\eta = 1.3$

suitable for a wide range of problem instances. The ADMM approach for solving the GIRAF least-squares subproblem also requires a conditioning parameter δ . We typically set $\delta = 100$, which was found to represent a desirable trade-off between speed and accuracy for a wide range of problem instances; see Figure 5.

V. Numerical experiments

In this section we perform several experiments to investigate properties of the GIRAF algorithm, and compare with other algorithms for structured low-rank matrix recovery. All experiments were conducted in MATLAB 8.5.0 (R2015a) on a Linux desktop computer with a Intel Xeon 3.20GHz CPU and 24 GB RAM.

A. Behavior of GIRAF algorithm

1) Half-circulant approximation, oversampled grid, and choice of p —The GIRAF algorithm relies on a half-circulant approximation (25) of Toeplitz matrix liftings. Here we investigate the error induced by this approximation in a 1-D signal recovery setting, where the matrix lifting is single-level Toeplitz: $\mathcal{T}(x) = \text{Toep}(x)$. We randomly undersample the exact Fourier coefficients $x_0 = (\hat{\rho}[k]: |k| < 64)$ of a 1-D stream of r Diracs $\rho(x)$ as in (1), and attempt to recover the missing samples by applying the IRLS- p to the Toeplitz matrix lifting $\mathcal{T}(x)$ having dimensions 113×15 , and by using the GIRAF- p algorithm, which is equivalent to applying IRLS- p to the half-circulant approximated lifting $\tilde{\mathcal{T}}(x)$ having dimensions 157×15 . We observe the IRLS-0 algorithm generally outperforms the IRLS-1, consistent with the results in [21]. Furthermore, the quality of the GIRAF- p reconstructions mimics those obtained with IRLS- p , but with small approximation errors. However, the approximation error is noticeably less severe in the non-convex $p = 0$ case.

We investigate this phenomenon in another experiment shown in Figure 7. Using the same experimental setup, we varied the working grid size of the problem, i.e., we varied the size of the oversampled reconstruction grid $\tilde{\mathcal{T}}$ defined in Section IV-A; this only changes the number of rows M in the lifting $\tilde{\mathcal{T}}(x)$. We vary $M = 127, \dots, 255$, corresponding to a maximum oversampling factor of 2. For the convex GIRAF-1 algorithm, we find that the reconstruction error stagnates with respect to the oversampling factor. However, for the non-convex GIRAF- p algorithms, $p < 1$, the reconstruction diminishes significantly with the oversampling factor. We additionally perform a similar experiment using the gradient weighted lifting scheme (9) on synthetic data that is piecewise constant and known to be low-rank in the lifted domain (the **SL** dataset shown in Figure 9) with data size 201×201 , filter size 25×25 , and 50% random samples. In this case we also obtain similar results—the nonconvex GIRAF algorithms achieve substantially smaller approximation errors that decrease with the oversampling factor. While this suggests one should use a very large oversampled grid to eliminate approximation errors, there is a trade-off with computational cost, since the oversampled grid represents the working dimension of the problem. The extent of oversampling will depend on the problem setting and required precision.

These experiments suggest it is better to use a non-convex GIRAF- p , $0 < p < 1$, algorithm over the convex GIRAF-1 algorithm in order to mitigate errors due to the half-circulant

approximation. In the remainder of our experiments we focus on the GIRAF-0 algorithm, since it yielded the lowest reconstruction error in most cases.

2) ADMM approach to GIRAF least-squares solution—In Figure 5 we compare the computation time of the proposed ADMM approach (36)–(38) for solving the GIRAF least-square subproblem (33) against standard CG and LSQR solvers. Here we investigate the GIRAF-0 algorithm applied to a Fourier domain recovery experiment using the gradient weighted lifting scheme (9). We take the measurement operator $\mathbf{A}$ to be a Fourier domain sampling, and sample 25% of the data uniformly at random, using the Shepp-Logan phantom (data size 256×256 and filter size of 35×35). The metric used to evaluate each approach is the normalized mean-square difference $\text{NMSD} = \frac{\|\mathbf{x} - \mathbf{x}_*\|^2}{\|\mathbf{x}_*\|^2}$ where $\mathbf{x}$ is the current iterate, and $\mathbf{x}_*$ is the true solution, which was obtained by running CG algorithm to high precision. The proposed ADMM approach shows nearly a ten-fold increase in over CG and LSQR. The rate at which the ADMM approach reduces the NMSD shows some sensitivity to the parameter δ , but only in the high-accuracy regime ($\text{NMSD} < 10^{-4}$).

B. Recovery Experiments

To demonstrate the benefits of the GIRAF algorithm for structured low-rank matrix recovery, we focus on the problem of undersampled MRI reconstruction in 2-D. In this setting, the goal is to recover an array $\mathbf{x}_0$ of the Fourier coefficients of an image from missing or corrupted Fourier samples. In addition to the direct IRLS- p algorithm (see Algorithm 1), we compare GIRAF against the following algorithms proposed for structured low-rank matrix recovery:

- **Alternating projections (AP).** Also known as Cadzow's method [56], this approach was adopted for noisy finite-rate-of-innovation (FRI) signal recovery in [37], [57], [58], and for auto-calibrated parallel MRI reconstruction in [2]. Reconstruction in an AP algorithm is posed as

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \quad \text{subject to} \quad \begin{cases} \mathbf{X} = \mathcal{T}(\mathbf{x}) \\ \text{rank } \mathbf{X} \leq r, \end{cases} \quad (47)$$

which is solved by alternately projecting $\mathbf{X}$ onto: (1) the set of matrices with rank less than or equal to r , (2) the space of linear structured matrices specified by the range of $\mathcal{T}$, and (3) the data fidelity constraint set. Note that the AP algorithm requires a rank parameter r . A novel extension of the AP algorithm incorporating proximal-smoothing is used in the LORAKS framework [4], [10], [59] for structured low-rank based MRI reconstruction. We use **AP-PROX** (AP with proximal smoothing) to refer to the algorithm introduced in [4], to distinguish it from the structured low-rank matrix models also introduced in [4]; see the Supplementary Materials for more details.

- **Singular value thresholding (SVT).** Proposed in [60] for general low-rank matrix recovery by nuclear norm minimization, and adapted to Hankel structured matrix case in [44], and for 2-D spectral estimation in [3]. The SVT algorithm

minimizes the objective (12) or (16) (with $p = 1$) via iterative soft thresholding of the singular values of the lifted matrix.

- **Singular value thresholding with factorization heuristic (SVT+UV).** This approach was proposed in [61] for general structured low-rank matrix recovery, and was adopted by the present authors in [8] for structured low-rank based super-resolution MRI reconstruction. It is also adopted in the ALOHA framework [30] for a variety of imaging applications, including structured low-rank based MRI reconstruction. The SVT+UV approach also minimizes the objective (12) or (16) (with $p = 1$), but uses the matrix factorization characterization of the nuclear norm:

$$\begin{aligned} \|X\|_* = \min_{\substack{U \in \mathbb{C}^{M \times R} \\ V \in \mathbb{C}^{N \times R} \\ X = UV^*}} \frac{1}{2} (\|U\|_F^2 + \|V\|_F^2), \quad (48) \end{aligned}$$

which holds true provided the inner dimension R of U and V^* satisfies $R \geq \text{rank } X$. This allows efficient solutions of the SVT subproblems by matrix inverses; see the Supplementary Materials for more details.

To aid in reproducibility of our results, we give the implementation details and pseudo-code for all these algorithms in the Supplementary Materials. We note only the AP, AP-PROX, and SVT+UV algorithms were previously used for the type of large-scale SLRA problems that we consider. However, we also include comparisons against SVT and IRLS for benchmark purposes.

1) GIRAF for LORAKS recovery—First, to demonstrate the benefit of the GIRAF algorithm for an existing SLRA approach in MRI, we apply the GIRAF to the LORAKS constrained MRI reconstruction framework [4]. In Figure 8 we compare the performance of the GIRAF to the AP-PROX algorithm proposed in [4], [62], using code and data distributed by the authors⁴. For simplicity we restrict our comparisons to the “C-LORAKS” matrix lifting, which is based on a spatial sparsity assumption; this is equivalent to a single block two-level Toeplitz lifting (23) with $M = I$ where I is the identity matrix. The LORAKS framework assumes filters with approximately circular support and we implement filters with the same size and support in the GIRAF algorithm to ensure a fair comparison. For the LORAKS algorithm we used the default parameters distributed with the code. We tested both algorithms on the dataset provided with the LORAKS code (180×180 Fourier coefficients of the Shepp-Logan phantom with phase) and the provided Fourier domain sampling masks (a uniform random sampling pattern with and a structured sampling pattern, accounting for $\approx 63\%$ of the data). We find the GIRAF algorithm converges an order of magnitude faster than the LORAKS algorithm (2–5 s versus 50–200+ s) and, in these experiments, to a solution with lower relative error as measured by the normalized mean

⁴<http://mr.usc.edu/download/loraks/>

square error $\text{NMSE} = \|\mathbf{x} - \mathbf{x}_0\|^2 / \|\mathbf{x}_0\|^2$, where $\mathbf{x}$ is the solution obtained from the algorithm, and $\mathbf{x}_0$ is the ground truth data. In addition to testing the default LORAKS rank cutoff parameter $r = 35$ suggested for this problem setting, we also varied r to study the effect on the reconstruction. We observe that the LORAKS approach shows significant variation in final NMSE depending on the rank estimate r . Similar behavior was observed in the experiments in [5]. A benefit of the GIRAF algorithm is that it does not require this explicit rank estimate. However, we do not claim the GIRAF algorithm will always outperform the LORAKS AP-PROX algorithm for appropriately tuned r in terms of NMSE, and here we did not optimize over all possible r . The comparison made here is only to demonstrate the GIRAF algorithm is significantly faster than the LORAKS AP-PROX algorithm when considering the same matrix lifting, and yields similar results in terms of NMSE and visual quality.

2) GIRAF for piecewise constant image recovery—Here we test the GIRAF-0 algorithm for the exact recovery of images belonging to the piecewise constant SLRA model proposed in [9], which uses the gradient weighted matrix lifting described in (9). We adapt the AP, AP-PROX, SVT, SVT+UV, and IRLS-0 algorithms to this setting. To compare computation time among algorithms, we first experiment with recovering simulated piecewise constant images from their undersampled Fourier coefficients. The datasets used in these experiments are shown in Figure 9. In each experiment, we sample uniformly at random in Fourier domain at the specified undersampling factor (USF). Since the simulated Fourier coefficients are noise-free, each algorithm incorporates equality data constraints, as in the formulation (12). To ensure a fair comparison, the algorithms that use a rank estimate (AP, SVT+UV) were passed either the exact rank of the simulated data, which was calculated either using the known bandwidth of the level-set polynomial describing the edge-set of the image (see [9]), or the index for which the normalized singular values were less than 10^{-2} . The results of these experiments are shown in Table I. We report the CPU time and number of iterations for each algorithm to reach $\text{NMSE} = \|\mathbf{x} - \mathbf{x}_0\|^2 / \|\mathbf{x}_0\|^2 \leq 10^{-4}$, where $\mathbf{x}$ is the reconstruction at the current iterate, and $\mathbf{x}_0$ is the ground truth data. We use the NMSE as a error metric since the competing algorithms solve different cost functions, and cannot be compared on the basis of how well they minimize a common cost function. Observe that for small to medium problem sizes (**PWC1**, and **PWC2**), the GIRAF algorithm is competitive or significantly faster than state-of-the-art methods. For the large-scale problems (**SL**, and **BRAIN**) the GIRAF algorithm converges orders of magnitude faster than competing algorithms, demonstrating its superior scalability. GIRAF is also successful on all the “hard” problem instances where SVT fails to converge below the set NMSE tolerance.

In Figure 10 we show the results of a similar recovery experiment, but where the Fourier samples are corrupted with noise. We test on the **SL** dataset with $\text{USF} = 0.65$, adding complex white Gaussian noise to the Fourier samples such that the signal-to-noise ratio (SNR) is approximately 22 dB. Here we test GIRAF against AP-PROX and SVT-UV in their regularized formulations, (16), tuning regularization parameters to obtain the optimal NMSE in each case. In this experiment and subsequent ones we report the time each algorithm took to reach 1% of the final NMSE, where the final NMSE is obtained by running each

algorithm until the relative change in NMSE between iterates is less than 10^{-6} . We observe the GIRAF algorithm shows similar run-time as in the noise-free setting, and converges to a solution with smaller NMSE.

3) Application to compressed sensing MRI reconstruction—In Figure 11 we demonstrate the GIRAF algorithm for the recovery of real MRI data from undersampled Fourier data using the gradient weighted lifting scheme. The datasets we experiment on were obtained from a fully sampled four-coil parallel MRI acquisition, compressed into a single virtual coil using an SVD-based technique [63]. We then retrospectively undersample the single virtual coil. To compensate for complex phase in the data, we estimate the phase of the image from very few of its low-pass Fourier samples, and incorporate this into the measurement operator $\mathbf{A}$, as recommended in [1]. To demonstrate the potential benefit of a Fourier domain SLRA approach over standard discrete spatial domain penalties, we compare against a standard total variation (TV) regularized reconstruction, implemented with an efficient ADMM/Split-Bregman approach [64]. We use the GIRAF-0 algorithm with a filter size of 55×55 . Observe that the GIRAF reconstruction shows nearly a 1dB improvement in $\text{SNR} = 10 \log_{10}(\text{NMSE})$ over the TV reconstruction. We use a relatively high undersampling factor to illustrate the difference in reconstruction quality over TV; the reconstruction quality may not be appropriate for radiological evaluations. The runtime of GIRAF algorithm on these datasets is roughly 1–2 minutes. While this is slower than a TV reconstruction, which takes roughly 5–10 seconds, it is still substantially faster than previous reported single-coil SLRA-based MRI reconstruction schemes [4], [30]. Moreover, for these experiments we could not compare against the AP-PROX or SVT+UV algorithms because of memory limitations.

One advantage of the GIRAF algorithm is that it scales well with filter size over competing algorithms. We show in Figure 12 that using larger filter sizes directly translates to improved image quality in the case of the gradient weighted lifting, with modest increases in run-time.

VI. Discussion and conclusion

We introduced the GIRAF algorithm for multi-level Toeplitz/Hankel low-rank matrix recovery problems arising in MRI reconstruction and other imaging applications. The algorithm is based on the IRLS approach for low-rank matrix completion, combined with a novel half-circulant approximation of multi-level Toeplitz matrices. This approximation dramatically reduces the computational complexity and memory demands of the direct IRLS algorithm for large problem sizes. Due to the monotonicity of the Schatten- p regularizer with respect to the half-circulant approximation, we show it suffices to regularize the approximated matrix, rather than the original multi-level Toeplitz matrix. Compared to previous approaches the GIRAF algorithm is an order of magnitude faster on several realistic problem instances. This is because in the GIRAF algorithm performs most of its computations in the original “un-lifted” problem domain with FFTs, rather than in the “lifted” matrix domain. An important feature of the GIRAF algorithm is that it can accommodate larger filters required by more sophisticated image priors, such as the off-the-grid piecewise constant image model proposed in [9]. This enables SLRA/annihilating filter approaches for a much wider class of imaging problems, including realistic large-scale

multidimensional datasets, such as those encountered in multi-dimensional MRI reconstruction tasks.

Unlike some previous approaches, the GIRAF algorithm does not require strict low-rank approximations, nor does it require an estimate of the underlying rank or model-order. Arguably the optimal selection of the ε_0 smoothing parameter used in the GIRAF algorithm implicitly performs model-order selection. However, we propose a method to automatically set this parameter based on spectral properties of the initialization to the algorithm. Similar strategies could be used for other algorithms that require a rank estimate to reduce their dependence on the rank.

The multi-level Toeplitz matrix lifting model considered in this work may seem narrow in application, but in fact a variety of existing SLRA models belong to this class, giving the GIRAF algorithm fairly wide applicability. Single-level Toeplitz/Hankel matrix liftings form the foundation for several modern approaches in spectral estimation [58], direction-of-arrival estimation [65], and system identification [44]. Multi-level Toeplitz/Hankel matrix liftings have been used for multi-dimensional spectral estimation [3], [66], and video inpainting [30], [67]. Of particular interest is the application of the framework to super-resolution microscopy [68], which is of high significance in biological imaging. SLRA methods have been seen to be very promising in this context [29], offering more accurate localizations than state-of-the-art methods; the GIRAF algorithm could potentially enable the extension of these methods to recover the whole image rather than small patches, due to its low memory demand. Similarly, most of the SLRA models proposed for MRI reconstruction have a block multi-level Toeplitz/Hankel matrix structure. For example, the SAKE framework for parallel MRI reconstruction [2] considers a matrix lifting with a horizontal block structure, where each block is a multi-level Hankel matrix. Similarly, the matrix liftings considered in the LORAKS [4], [10] and ALOHA [6] frameworks have block multi-level Toeplitz/Hankel structure. In this work we restricted our attention to matrix liftings having a vertical block structure. While our current algorithm is not readily applicable to these settings, the GIRAF framework could also potentially be generalized to more general block-rectangular liftings, such as those used parallel MRI context [2], [10]; we save this as a topic for future work.

Finally, the approximation intrinsic to the GIRAF approach means that the algorithm might not be appropriate for all SLRA problems. For instance, if the rank of the lifted matrix is expected to be very small and is known in advance, singular value thresholding or alternating projection algorithms may be more efficient or more accurate. However, for the large-scale SLRA problems encountered in imaging, where the rank of the lifted matrix is typically unknown and possibly large, our numerical experiments show the GIRAF approach represents a desirable trade-off between accuracy and runtime versus other state-of-the-art algorithms.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work is supported by grants NIH 1R01EB019961-01A1, NSF CCF-1116067, ONR N00014-13-1-0202, and ACS RSG-11-267-01-CCE.

Appendix Majorization-minimization formulation of IRLS-p algorithm

Majorization-minimization (MM) algorithms iteratively minimize a given cost function by minimizing a surrogate cost function that upper bounds the original at each iteration [49]. More precisely, if $f(\mathbf{x})$ is any real-valued cost function on a domain Ω , a function $g(\mathbf{x}; \mathbf{x}_0)$ defined on $\Omega \times \Omega$ is said to majorize $f(\mathbf{x})$ at the fixed value $\mathbf{x}_0 \in \Omega$ if

$$g(\mathbf{x}; \mathbf{x}_0) \geq f(\mathbf{x}), \text{ for all } \mathbf{x} \in \Omega \quad (49)$$

$$g(\mathbf{x}_0; \mathbf{x}_0) = f(\mathbf{x}_0) \quad (50)$$

If (49) and (50) hold for all $\mathbf{x}_0 \in \Omega$ then g is called a *majorizer* for f . An MM algorithm finds a minimizer for f by sequentially solving for the iterates

$$\mathbf{x}^{(n+1)} = \arg \min_{\mathbf{x} \in \Omega} g(\mathbf{x}; \mathbf{x}^{(n)}). \quad (51)$$

Using properties (49) and (50) it is easy to show the cost function f must monotonically decrease at each iteration. Furthermore, if f is strongly convex, the iterates $\mathbf{x}^{(n)}$ are guaranteed to converge to the unique global minimizer of f ; when convexity is violated, the iterates are still guaranteed to converge to a stationary point [69].

We now construct a majorizer for the Schatten p -norm $\left\| \cdot \right\|_p^p$, including values of $p \in [0, 1)$

where the penalty is non-convex. Let H_+^N denote the space of Hermitian positive semidefinite $N \times N$ matrices and H_{++}^N the space of Hermitian positive definite $N \times N$ matrices. If $f: [0, \infty) \rightarrow \mathbb{R}$ is any scalar function, and $\mathbf{Y} \in H_+^N$, we can define $f(\mathbf{Y})$ by $f(\mathbf{Y}) = \sum_i f(\lambda_i) \mathbf{P}_i$ where $\mathbf{Y}$ has the spectral decomposition $\mathbf{Y} = \sum_i \lambda_i \mathbf{P}_i$, i.e., $\lambda_i \geq 0$ are the eigenvalues of $\mathbf{Y}$ and $\mathbf{P}_i$ are the orthogonal projectors onto the corresponding eigenspaces. In particular, the q th power $\mathbf{Y}^q$, $q \in (0, 1]$, of any $\mathbf{Y} \in H_+^N$ is defined as

$$\mathbf{Y}^q = \sum_i \lambda_i^q \mathbf{P}_i. \quad (52)$$

Likewise, for any $\mathbf{Y} \in H_{++}^N$, $\log(\mathbf{Y})$ is defined as

$$\log(\mathbf{Y}) = \sum_i \log(\lambda_i) \mathbf{P}_i \quad (53)$$

Our majorizer is derived from the following trace inequalities:

Proposition 1

If $\mathbf{Y}_0 \in H_{++}^N$ then for all $\mathbf{Y} \in H_{++}^N$ we have

$$\text{Tr}[\mathbf{Y}^q] \leq \text{Tr}[\mathbf{Y}_0^q + q\mathbf{Y}_0^{q-1}(\mathbf{Y} - \mathbf{Y}_0)] \quad (54)$$

for all $q \in (0, 1]$. If additionally $\mathbf{Y} \in H_{++}^N$, then

$$\text{Tr}[\log(\mathbf{Y})] \leq \text{Tr}[\log(\mathbf{Y}_0) + (\mathbf{Y}_0)^{-1}(\mathbf{Y} - \mathbf{Y}_0)] \quad (55)$$

Proof

These are special cases of Klein's inequality (see, e.g., Proposition 2.5.2 in [70]), which states that for any concave function $f: [0, \infty) \rightarrow \mathbb{R}$ differentiable on $(0, \infty)$, and any $\mathbf{Y} \in H_{++}^N, \mathbf{Y}_0 \in H_{++}^N$ we have

$$\text{Tr}[f(\mathbf{Y}) - f(\mathbf{Y}_0) - f'(\mathbf{Y}_0)(\mathbf{Y} - \mathbf{Y}_0)] \leq 0 \quad (56)$$

where $f(\mathbf{X}) := \sum_i f(\lambda_i) \mathbf{P}_i$ when $\mathbf{X}$ has the spectral decomposition $\mathbf{X} = \sum_i \lambda_i \mathbf{P}_i$. Choosing $f(t) = t^q$ and $f(t) = \log(t)$ gives the desired inequalities. $\square$

Let $\mathbf{X}, \mathbf{X}_0 \in \mathbb{C}^{M \times N}$ where $\mathbf{X}$ is arbitrary and $\mathbf{X}_0$ has no zero singular values. Substituting $\mathbf{Y} = \mathbf{X}^* \mathbf{X} \in H_{++}^N$, $\mathbf{Y}_0 = \mathbf{X}_0^* \mathbf{X}_0 \in H_{++}^N$, and $q = p/2$ into (54), it follows that the function

$$\begin{aligned}
g_p(\mathbf{X}; \mathbf{X}_0) &= \frac{1}{2} \text{Tr}[(\mathbf{X}_0^H \mathbf{X}_0)^{\frac{p}{2}} + \frac{p}{2} (\mathbf{X}_0^H \mathbf{X}_0)^{\frac{p}{2}-1} (\mathbf{X}^H \mathbf{X} - \mathbf{X}_0^H \mathbf{X}_0)] = \frac{1}{2} \text{Tr}[(\mathbf{X}_0^H \mathbf{X}_0)^{\frac{p}{2}-1} \mathbf{X}^H \mathbf{X}] \\
&+ C(\mathbf{X}_0)
\end{aligned}
\tag{57}$$

$$= \frac{1}{2} \left\| \mathbf{X} \sqrt{(\mathbf{X}_0^H \mathbf{X}_0)^{\frac{p}{2}-1}} \right\|_F^2 + C(\mathbf{X}_0), \tag{58}$$

where $C(\mathbf{X}_0)$ is a term depending only on $\mathbf{X}_0$, satisfies the majorization relations

$$g_p(\mathbf{X}; \mathbf{X}_0) \geq \|\mathbf{X}\|_p^p, \text{ for all } \mathbf{X} \in \mathbb{C}^{M \times N} \tag{59}$$

$$g_p(\mathbf{X}; \mathbf{X}_0) = \|\mathbf{X}_0\|_p^p \tag{60}$$

for all $p \in (0, 1]$. A similar argument using inequality (55) shows that (59) and (60) can be extended to hold for the $p = 0$ case, with the majorizer $g_0(\mathbf{X}; \mathbf{X}_0)$ defined by (58) with $p = 0$. Finally, substituting $\mathbf{X} = \mathcal{T}(\mathbf{x})$ gives the following MM scheme for minimizing the Schatten- p penalty:

$$\mathbf{x}^{(n)} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\lambda}{2} \|\mathcal{T}(\mathbf{x}) \sqrt{\mathbf{H}_n}\|_F^2 \tag{61}$$

where $\mathbf{H}_n$ is determined by the previous iterate $\mathbf{x}^{(n-1)}$ according to

$$\mathbf{H}_n = [\mathcal{T}(\mathbf{x}^{(n-1)})^* \mathcal{T}(\mathbf{x}^{(n-1)})]^{\frac{p}{2}-1}, \tag{62}$$

To ensure the matrix in (62) is invertible, we can introduce a smoothing parameter ε_n and update instead

$$\mathbf{H}_n = [\mathcal{T}(\mathbf{x}^{(n-1)})^* \mathcal{T}(\mathbf{x}^{(n-1)}) + \varepsilon_n \mathbf{I}]^{\frac{p}{2}-1}, \tag{63}$$

where we take $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$. This gives the IRLS- p algorithm as shown in Algorithm 1.

References

1. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*. 2007; 58(6):1182–1195. [PubMed: 17969013]
2. Shin PJ, Larson PE, Ohliger MA, Elad M, Pauly JM, Vigneron DB, Lustig M. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. *Magnetic Resonance in Medicine*. 2014; 72(4):959–970. [PubMed: 24248734]
3. Chen Y, Chi Y. Robust spectral compressed sensing via structured matrix completion. *Information Theory, IEEE Transactions on*. 2014; 60(10):6576–6601.
4. Halder JP. Low-rank modeling of local k-space neighborhoods (LORAKS) for constrained MRI. *IEEE Transactions on Medical Imaging*. 2014; 33(3):668–681. [PubMed: 24595341]
5. Halder JP. Low-rank modeling of local k-space neighborhoods: from phase and support constraints to structured sparsity. *SPIE Optical Engineering+ Applications International Society for Optics and Photonics*. 2015:959710–959710.
6. Jin KH, Lee D, Ye JC. A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix. *IEEE Transactions on Computational Imaging*. 2016; 2(4):480–495.
7. Jin KH. A novel k-space annihilating filter method for unification between compressed sensing and parallel MRI. *International Symposium on Biomedical Imaging (ISBI)*. Apr.2015
8. Ongie, G., Jacob, M. *Sampling Theory and Applications (SampTA)*. Washington DC: 2015. Recovery of piecewise smooth images from few Fourier samples.
9. Ongie, G., Jacob, M. Off-the-grid recovery of piecewise constant images from few Fourier samples. *SIAM Journal on Imaging Science*. 2016. vol To appear pre-print available online: <http://arxiv.org/abs/1510.00384>
10. Halder JP, Zhuo J. P-LORAKS: Low-rank modeling of local k-space neighborhoods with parallel imaging data. *Magnetic resonance in medicine*. 2015
11. Kim TH, Setsompop K, Halder JP. LORAKS makes better SENSE: Phase-constrained partial fourier SENSE reconstruction without phase calibration. *Magnetic Resonance in Medicine*. 2016
12. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic resonance in medicine*. 2002; 47(6):1202–1210. [PubMed: 12111967]
13. Morrison, RL., Jacob, M., Do, MN. *Biomedical Imaging: From Nano to Macro, 2007 ISBI 2007 4th IEEE International Symposium on*. IEEE; 2007. Multichannel estimation of coil sensitivities in parallel MRI; p. 117–120.
14. Lustig M, Pauly JM. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic resonance in medicine*. 2010; 64(2):457–471. [PubMed: 20665790]
15. Zhang J, Liu C, Moseley ME. Parallel reconstruction using null operations. *Magnetic Resonance in Medicine*. 2011; 66(5):1241–1253. [PubMed: 21604290]
16. Mani M, Jacob M, Kelley D, Magnotta V. Multi-shot multi-channel diffusion data recovery using structured low-rank matrix completion. 2016 arXiv preprint arXiv:1602.07274.
17. Lee J, Jin KH, Ye JC. Reference-free single-pass EPI Nyquist ghost correction using annihilating filter-based low rank Hankel matrix (ALOHA). *Magnetic Resonance in Medicine*. 2016
18. Ye JC, Kim JM, Jin KH, Lee K. Compressive sampling using annihilating filter-based low-rank interpolation. *IEEE Transactions on Information Theory*. 2016
19. Ongie, G., Biswas, S., Jacob, M. *Image Processing (ICIP), 2016 IEEE International Conference on*. IEEE; 2016. Structured low-rank recovery of piecewise constant signals with performance guarantees; p. 963–967.
20. Fornasier M, Rauhut H, Ward R. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM Journal on Optimization*. 2011; 21(4):1614–1640.

21. Mohan K, Fazel M. Iterative reweighted algorithms for matrix rank minimization. *The Journal of Machine Learning Research*. 2012; 13(1):3441–3473.
22. Lai MJ, Xu Y, Yin W. Improved iteratively reweighted least squares for unconstrained smoothed ℓ_q minimization. *SIAM Journal on Numerical Analysis*. 2013; 51(2):927–957.
23. Xie Y, Gu S, Liu Y, Zuo W, Zhang W, Zhang L. Weighted Schatten p -norm minimization for image denoising and background subtraction. *IEEE transactions on image processing*. 2016; 25(10): 4842–4857.
24. Hintermuller M, Valkonen T, Wu T. Limiting aspects of non-convex TV^p models. *SIAM J Imaging Sciences*. 2015:2581–2621.
25. Erway, JB., Plemmons, RJ., Adhikari, L., Marcia, RF. Information Theory and Its Applications (ISITA), 2016 International Symposium on. IEEE; 2016. Trust-region methods for nonconvex sparse recovery optimization; p. 275-279.
26. Chu MT, Funderlic RE, Plemmons RJ. Structured low rank approximation. *Linear algebra and its applications*. 2003; 366:157–172.
27. Markovsky I. Structured low-rank approximation and its applications. *Automatica*. 2008; 44(4): 891–909.
28. Ongie G, Jacob M. Super-resolution MRI using finite rate of innovation curves. *IEEE International Symposium on Biomedical Imaging: ISBI 2015*. 2015
29. Min J, Carlini L, Unser M, Manley S, Ye JC. Fast live cell imaging at nanometer scale using annihilating filter based low rank Hankel matrix approach. *SPIE Optical Engineering+ Applications International Society for Optics and Photonics*. 2015 95 970V–95 970V.
30. Jin KH, Ye JC. Annihilating filter-based low-rank hankel matrix approach for image inpainting. *Image Processing, IEEE Transactions on*. 2015; 24(11):3498–3511.
31. De Prony BGR. Essai experimental et analytique: sur les lois de la dilatabilité de fluides élastique et sur celles de la force expansive de la vapeur de l'alkool, a différentes températures. *Journal de l'école polytechnique*. 1795; 1(22):24–76.
32. Schmidt RO. Multiple emitter location and signal parameter estimation. *Antennas and Propagation, IEEE Transactions on*. 1986; 34(3):276–280.
33. Paulraj, A., Roy, R., Kailath, T. Circuits, Systems and Computers, 1985 Nineteeth Asilomar Conference on. IEEE; 1985. Estimation of signal parameters via rotational invariance techniques-esprit; p. 83-89.
34. Hua Y, Sarkar TK. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *IEEE Transactions on Acoustics, Speech, and Signal Processing*. 1990; 38(5):814–824.
35. Stoica P, Moses RL. Introduction to spectral analysis Prentice hall Upper Saddle River, NJ. 1997; 1
36. Strichartz, RS. A guide to distribution theory and Fourier transforms. World Scientific; 2003.
37. Pan H, Blu T, Dragotti PL. Sampling curves with finite rate of innovation. *Signal Processing, IEEE Transactions on*. 2014; 62(2)
38. Liang ZP, Haacke EM, Thomas CW. High-resolution inversion of finite fourier transform data through a localised polynomial approximation. *Inverse Problems*. 1989; 5(5):831.
39. Haacke EM, Liang ZP, Izen SH. Constrained reconstruction: A superresolution, optimal signal-to-noise alternative to the fourier transform in magnetic resonance imaging. *Medical Physics*. 1989; 16(3):388–397. [PubMed: 2739620]
40. Haacke E, Liang ZP, Izen S. Superresolution reconstruction through object modeling and parameter estimation. *Acoustics, Speech and Signal Processing, IEEE Transactions on*. Apr; 1989 37(4):592–595.
41. Haacke EM, Liang ZP, Boada FE. Image reconstruction using POCS, model constraints, and linear prediction theory for the removal of phase, motion, and gibbs artifacts in magnetic resonance and ultrasound imaging. *Optical Engineering*. 1990; 29(5):555–566.
42. Candès EJ, Recht B. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*. 2009; 9(6):717–772.
43. Cai JF, Liu S, Xu W. A fast algorithm for reconstruction of spectrally sparse signals in super-resolution. *Proc SPIE*. 2015; 9597:95 970A–95 970A-7.

44. Fazel M, Pong TK, Sun D, Tseng P. Hankel matrix rank minimization with applications to system identification and realization. *SIAM Journal on Matrix Analysis and Applications*. 2013; 34(3): 946–977.
45. Markovsky I, Usevich K. Structured low-rank approximation with missing data. *SIAM Journal on Matrix Analysis and Applications*. 2013; 34(2):814–830.
46. Voevodin VV, Tyrtshnikov EE. Numerical methods for solving problems with toeplitz matrices. *USSR Computational Mathematics and Mathematical Physics*. 1981; 21(3):1–14.
47. Vetterli M, Marziliano P, Blu T. Sampling signals with finite rate of innovation. *Signal Processing, IEEE Transactions on*. 2002; 50(6):1417–1428.
48. Shang F, Liu Y, Cheng J. Tractable and scalable Schatten quasi-norm approximations for rank minimization. *Proc 19th Int Conf Artif Intell Statist*. 2016:620–629.
49. Hunter DR, Lange K. A tutorial on MM algorithms. *The American Statistician*. 2004; 58(1):30–37.
50. Larsen RM. Lanczos bidiagonalization with partial reorthogonalization. *DAIMI Report Series*. 1998; 27(537)
51. Halko N, Martinsson PG, Tropp JA. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*. 2011; 53(2):217–288.
52. Paige CC, Saunders MA. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Transactions on Mathematical Software (TOMS)*. 1982; 8(1):43–71.
53. Horn, RA., Johnson, CR. *Topics in matrix analysis*. Cambridge Univ Press; Cambridge etc: 1991.
54. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*. 2011; 3(1):1–122.
55. Eckstein J, Bertsekas DP. On the Douglas—Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*. 1992; 55(1):293–318.
56. Cadzow JA. Signal enhancement—a composite property mapping algorithm. *Acoustics, Speech and Signal Processing, IEEE Transactions on*. 1988; 36(1):49–62.
57. Blu T, Dragotti PL, Vetterli M, Marziliano P, Coulot L. Sparse sampling of signal innovations. *Signal Processing Magazine, IEEE*. 2008; 25(2):31–40.
58. Condat L, Hirabayashi A. Cadzow denoising upgraded: A new projection method for the recovery of Dirac pulses from noisy linear measurements. *Sampling Theory in Signal and Image Processing*. 2015; 14(1):17.
59. Kim, TH., Haldar, JP. International Symposium on Biomedical Imaging (ISBI). IEEE; 2015. SMS-LORAKS: Calibrationless simultaneous multislice MRI using low-rank matrix modeling; p. 323–326.
60. Cai J, Osher S. Fast singular value thresholding without singular value decomposition. *UCLA CAM Report*. 2010; 5
61. Signoretto, M., Cevher, V., Suykens, JA. IEEE Conference on Decision and Control. 2013. An SVD-free approach to a class of structured low rank matrix optimization problems with application to system identification. no EPFL-CONF-184990
62. Haldar JP. Low-rank modeling of local k-space neighborhoods (LORAKS): Implementation and examples for reproducible research. *Technical Report USC-SIPI-414*. 2014
63. Zhang T, Pauly JM, Vasanawala SS, Lustig M. Coil compression for accelerated imaging with cartesian sampling. *Magnetic Resonance in Medicine*. 2013; 69(2):571–582. [PubMed: 22488589]
64. Goldstein T, Osher S. The split Bregman method for L1-regularized problems. *SIAM journal on Imaging Sciences*. 2009; 2(2):323–343.
65. Pal P, Vaidyanathan P. A grid-less approach to underdetermined direction of arrival estimation via low rank matrix denoising. *Signal Processing Letters, IEEE*. 2014; 21(6):737–741.
66. Condat, L., Boulanger, J., Pustelnik, N., Sahnoun, S., Sengmanivong, L. Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on. IEEE; 2014. A 2-D spectral analysis method to estimate the modulation parameters in structured illumination microscopy; p. 604–607.
67. Ding, T., Sznajder, M., Camps, O., et al. Computer Vision, 2007 ICCV 2007 IEEE 11th International Conference on. IEEE; 2007. A rank minimization approach to video inpainting; p. 1–8.

68. Sage D, Kirshner H, Pengo T, Stuurman N, Min J, Manley S, Unser M. Quantitative evaluation of software packages for single-molecule localization microscopy. *Nature methods*. 2015; 12(8):717–724. [PubMed: 26076424]
69. Jacobson MW, Fessler JA. An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms. *IEEE Transactions on Image Processing*. 2007; 16(10):2411–2422. [PubMed: 17926925]
70. Ruelle, D. *Statistical Mechanics: Rigorous Results*. 2nd. World Scientific; 1999.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 1.

2-D annihilation relationship for piecewise constant images and construction of gradient weighted lifting. If the piecewise constant image f (top) has edges supported in the zero-set of a bandlimited function μ , then the Fourier coefficients of the partial derivatives of f are annihilated by convolution by a filter h determined by the Fourier coefficient of μ . (bottom) From a 2-D array of Fourier coefficients $\hat{f}$, the matrix $\mathcal{G}(\hat{f})$ is constructed by weighting $\hat{f}$ to create two arrays of the simulated Fourier derivatives $\widehat{\partial_x f}$ and $\widehat{\partial_y f}$, then extracting each patch of size Λ by a sliding window operation.

Fig. 2.

Illustration of index sets used in construction of multi-level Toeplitz matrix liftings (17) in two-dimensions ($d=2$). Here Δ is the support of the data y , Λ (in red) is the support of the filter h with index $(0, 0)$ in black, and Γ (interior of dashed line) represents the valid set where the linear convolution $y * h$ is well-defined.

Fig. 3. Example of Half-circulant approximation in 1-D

We approximate the Toeplitz matrix lifting $\mathcal{T}(x)$ with the half-circulant matrix $\tilde{\mathcal{T}}(x)$ obtained by adding rows to $\mathcal{T}(x)$ to make it the full vertical section of a circulant matrix $C(Mx)$

Fig. 4.

Effect of Schatten- p penalty on the GIRAF re-weighted annihilating filter. We show the eigenvalues of $G = \widetilde{\mathcal{T}}(x)^* \widetilde{\mathcal{T}}(x)$ during one iteration of the GIRAF algorithm applied to the gradient weighted lifting (9) (top left), and the different re-weightings of the eigenvalues specified by different Schatten- p penalties (top right). The set of annihilation weights determined by the resulting re-weighted annihilating filter given by (45) are shown in the bottom row, normalized to $[0, 1]$. Note that for smaller values of p the annihilation weights are closer to zero near the edges because the re-weighted annihilating filter gives higher weight to filters in the nullspace of $\widetilde{\mathcal{T}}(x)$

Fig. 5.

Computation time comparison for solvers of the GIRAF least-squares problem. We study the least-squares problem arising in iterations 1 and 8 of the GIRAF algorithm applied to a Fourier recovery experiment. The proposed ADMM approach shows nearly ten-fold increase in the rate at which NMSD is decreased over CG and LSQR, provided the pre-conditioning parameter δ is chosen correctly.

Fig. 6.

Comparison of solutions obtained by the IRLS- p and the GIRAF- p algorithm, $p = 1, 0$, for recovery of Fourier coefficients of r Diracs, as in (1), from non-uniform random Fourier samples. The top row shows an experiment ($r = 4$, 50% random samples) where both IRLS-1 and IRLS-0 succeed in recovering the signal. In this case the GIRAF-1 algorithm shows non-negligible approximation errors, but the GIRAF-0 result is close to exact. The bottom row shows an experiment ($r = 6$, 33% random samples) where IRLS-1 fails, but IRLS-0 is successful; similarly, the GIRAF-1 recovery fails, while the GIRAF-0 result is again close to exact. Shown in the right column are the singular values of the exact matrix lifting $\mathcal{T}(\mathbf{x}^*)$ where $\mathbf{x}^*$ is the solution obtained from each algorithm. Note GIRAF-0 solutions are still approximately low-rank under the exact lifting.

Fig. 7.

Reconstruction error versus oversampling factor for GIRAF- p algorithm on the recovery of synthetic low-rank data. The quality of the half-circulant approximation (25) used in the GIRAF- p algorithms improves by working on an oversampled grid, i.e., increasing the number of rows in the half-circulant matrix lifting. The left plot corresponds to the recovery Fourier coefficients of Diracs in 1-D, using the same settings as in Figure 6; we report the average NMSE over 50 random trials. The right plot shows recovery of Fourier coefficients of a synthetic 2-D piecewise constant image using the gradient weighted lifting (9). Observe that the NMSE stagnates for the $p = 1$ case (convex penalty). However, the NMSE diminishes towards zero with the oversampling factor when $p = 0.5$ and $p = 0$ (non-convex penalties). This suggests it is important to use the non-convex versions of GIRAF on a sufficiently large oversampled grid to mitigate errors due to the half-circulant approximation.

Fig. 8.

Comparison of GIRAF with LORAKS AP-PROX algorithm. Plotted is the per iteration NMSE against elapsed CPU time in the recovery of synthetic data from the indicated from random Fourier sampling (left) and a structured grid-like Fourier sampling (right) using the C-LORAKS matrix lifting [62]. The AP-PROX algorithm requires a rank estimate r , which needs fine-tuning to obtain the best NMSE. By contrast, the GIRAF algorithm does not require a rank estimate, and in these experiments, converged to a solution with lower NMSE.

Fig. 9. Synthetic piecewise constant images used in structured low-rank recovery experiments. Shown below the images are the normalized singular value spectrum of the gradient weighted lifted matrix constructed from the fully sampled Fourier coefficients, where the dotted line indicates the rank estimate r .

Fig. 10.

Comparison of GIRAF with competing algorithms for structured low-rank matrix recovery with noisy data using a piecewise constant SLRA model. Plotted is the per iteration NMSE against elapsed CPU time in the recovery of synthetic data from noisy random uniform Fourier samples (USF=0.65, sample SNR=22dB).

Fig. 11.

Undersampled MRI reconstruction using GIRAF algorithm with piecewise constant SLRA model. We retrospectively undersample in Fourier domain from variable density random Cartesian lines (left) uniform random sampling mask with a calibration region (right), both with USF=0.50. In each case we show the total variation (TV) regularized recovery and the SLRA recovery using GIRAF. Error images are shown in the bottom row.

Fig. 12.

Effect of filter size on reconstruction; zoomed for detail. We reconstruct Fourier data from 20% random variable density samples using the GIRAF-0 algorithm and the gradient weighted lifting scheme. The reconstruction shows a +1dB change in SNR by increasing the filter size from 15×15 to 45×45 , with only a four-fold increase in run-time.

Convergence time of algorithms for structured low-rank recovery of piecewise constant images from noise-free random fourier samples.

TABLE I

Problem		AP		SVT		SVT+UV		IRLS-0		GIRAF-0	
dataset	data size	filter	rank	USF	# iter	time (s)	# iter	time (s)	# iter	time (s)	# iter
PWC1	65×65	9×9	32	0.50	17	5.4	11	1.1	11	0.8	3
PWC1	65×65	9×9	32	0.33	39	12.8	–	Inf	19	1.4	5
PWC2	129×129	17×17	118	0.50	19	48.0	9	25.2	15	21.6	3
PWC2	129×129	17×17	118	0.33	90	228.0	–	Inf	24	36.2	5
SL	201×201	25×25	336	0.65	31	762.0	13	338.7	11	150.9	3
SL	201×201	25×25	336	0.50	88	2151.7	–	Inf	36	494.4	4
BRAIN	255×255	45×45	1250	0.65	–	Mem	–	Mem	23	4270.3	3
BRAIN	255×255	45×45	1250	0.50	–	Mem	–	Mem	–	Inf	5

Stopping criteria: NMSE 10^{-4} .

Key: Inf - Algorithm converged above NMSE threshold. Mem - Not enough memory to run algorithm.

Computer specifications: Intel Xeon 3.20GHz CPU and 24 GB RAM.