



HHS Public Access

Author manuscript

IEEE J Sel Top Signal Process. Author manuscript; available in PMC 2017 June 01.

Published in final edited form as:

IEEE J Sel Top Signal Process. 2016 June ; 10(4): 672–687. doi:10.1109/JSTSP.2016.2545518.

Beyond Low Rank + Sparse: Multi-scale Low Rank Matrix Decomposition

Frank Ong and **Michael Lustig**

Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94709 USA

Abstract

We present a natural generalization of the recent low rank + sparse matrix decomposition and consider the decomposition of matrices into components of multiple scales. Such decomposition is well motivated in practice as data matrices often exhibit local correlations in multiple scales. Concretely, we propose a multi-scale low rank modeling that represents a data matrix as a sum of block-wise low rank matrices with increasing scales of block sizes. We then consider the inverse problem of decomposing the data matrix into its multi-scale low rank components and approach the problem via a convex formulation. Theoretically, we show that under various incoherence conditions, the convex program recovers the multi-scale low rank components either exactly or approximately. Practically, we provide guidance on selecting the regularization parameters and incorporate cycle spinning to reduce blocking artifacts. Experimentally, we show that the multi-scale low rank decomposition provides a more intuitive decomposition than conventional low rank methods and demonstrate its effectiveness in four applications, including illumination normalization for face images, motion separation for surveillance videos, multi-scale modeling of the dynamic contrast enhanced magnetic resonance imaging and collaborative filtering exploiting age information.

Index Terms

Multi-scale Modeling; Low Rank Modeling; Convex Relaxation; Structured Matrix; Signal Decomposition

I. INTRODUCTION

Signals and systems often exhibit different structures at different scales. Such multi-scale structure has inspired a wide variety of multi-scale signal transforms, such as wavelets [1], curvelets [2] and multi-scale pyramids [3], that can represent natural signals compactly. Moreover, their ability to compress signal information into a few significant coefficients has made multi-scale signal transforms valuable beyond compression and are now commonly used in signal reconstruction applications, including denoising [4], compressed sensing [5], [6], and signal separation [7]–[9]. By now, multi-scale modeling is associated with many success stories in engineering applications.

On the other hand, low rank methods are commonly used instead when the signal subspace needs to be estimated as well. In particular, low rank methods have seen great success in

applications, such as biomedical imaging [10], face recognition [11] and collaborative filtering [12], that require exploiting the global data correlation to recover the signal subspace and compactly represent the signal at the same time. Recent convex relaxation techniques [13] have further enabled low rank model to be adaptable to various signal processing tasks, including matrix completion [14], system identification [15] and phase retrieval [16], making low rank methods ever more attractive.

In this paper, we present a multi-scale low rank matrix decomposition method that incorporates multi-scale structures with low rank methods. The additional multi-scale structure allows us to obtain a more accurate and compact signal representation than conventional low rank methods whenever the signal exhibits multi-scale structures (see Figure 1). To capture data correlation at multiple scales, we model our data matrix as a sum of block-wise low rank matrices with increasing scales of block sizes (more detail in Section II) and consider the inverse problem of decomposing the matrix into its multi-scale components. Since we do not assume an explicit basis model, multi-scale low rank decomposition also prevents modeling errors or basis mismatch that are commonly seen with multi-scale signal transforms. In short, our proposed multi-scale low rank decomposition inherits the merits from both multi-scale modeling and low rank matrix decomposition.

Leveraging recent convex relaxation techniques, we propose a convex formulation to perform the multi-scale low rank matrix decomposition. We provide a theoretical analysis in Section V that extends the rank-sparsity incoherence results in Chandrasekaran et al. [17]. We show that the proposed convex program decomposes the data matrix into its multi-scale components exactly under a deterministic incoherence condition. In addition, in Section VI, we provide a theoretical analysis on approximate multi-scale low rank matrix decomposition in the presence of additive noise that extends the work of Agarwal et al. [18].

A major component of this paper is to introduce the proposed multi-scale low rank decomposition with emphasis on its practical performance and applications. We provide practical guidance on choosing regularization parameters for the convex method in Section IV and describe heuristics to perform cycle spinning [19] to reduce blocking artifacts in Section IX. In addition, we applied the multi-scale low rank decomposition on real datasets and considered four applications of the multi-scale low rank decomposition: illumination normalization for face images, motion separation for surveillance videos, compact modeling of the dynamic contrast enhanced magnetic resonance imaging and collaborative filtering exploiting age information. (See Section X for more detail). Our results show that the proposed multi-scale low rank decomposition provides intuitive multi-scale decomposition and compact signal representation for a wide range of applications.

Related work

Our proposed multi-scale low rank matrix decomposition draws many inspirations from recent developments in rank minimization [13], [14], [18], [20]–[24]. In particular, the multi-scale low rank matrix decomposition is a generalization of the low rank + sparse decomposition proposed by Chandrasekaran et al. [17] and Candès et al. [25]. Our multi-scale low rank convex formulation also fits into the convex demixing framework proposed

by McCoy et al. [26]–[28], who studied the problem of demixing components via convex optimization. The proposed multi-scale low rank decomposition can be viewed as a concrete and practical example of the convex demixing problem. However, their theoretical analysis assumes that each component is randomly oriented with respect to each other, and does not apply to our setting, where we observe the direct summation of the components. Bakshi et al. [29] proposed a multi-scale principal component analysis by applying principal component analysis on wavelet transformed signals, but such method implicitly constrains the signal to lie on a predefined wavelet subspace. Various multi-resolution matrix factorization techniques [30], [31] were proposed to greedily peel off components of each scale by recursively applying matrix factorization. One disadvantage of these factorization methods is that it is not straightforward to incorporate them with other reconstruction problems as models. Similar multi-scale modeling using demographic information was also used in collaborative filtering described in Vozalis and Margaritis [32].

II. MULTI-SCALE LOW RANK MATRIX MODELING

In this section, we describe the proposed multi-scale low rank matrix modeling in detail. To concretely formulate the model, we assume that we can partition the data matrix of interest Y into different scales. Specifically, we assume that we are given a multi-scale partition $\{P_i\}_{i=1}^L$ of the indices of an $M \times N$ matrix, where each block b in P_i is an order magnitude larger than the blocks in the previous scale P_{i-1} . Such multi-scale partition can be naturally obtained in many applications. For example in video processing, a multi-scale partition can be obtained by decimating both space and time dimensions. Figures 2 and 4 provide two examples of a multi-scale partition, the first one with decimation along two dimensions and the second one with decimation along one dimension. In Section X, we provide practical examples on creating these multi-scale partitions for different applications.

To easily transform between the data matrix and the block matrices, we then consider a block reshape operator $R_b(X)$ that extracts a block b from the full matrix X and reshapes the block into an $m_i \times n_i$ matrix (Figure 3), where $m_i \times n_i$ is the i th scale block matrix size determined by the user.

Given an $M \times N$ input matrix Y and its corresponding multi-scale partition and block reshape operators, we propose a multi-scale low rank modeling that models the $M \times N$ input matrix Y as a sum of matrices $\sum_{i=1}^L X_i$, in which each X_i is block-wise low rank with respect to its partition P_i that is,

$$Y = \sum_{i=1}^L X_i$$

$$X_i = \sum_{b \in P_i} R_b^T \left(U_b S_b V_b^T \right)$$

where U_b , S_b , and V_b are matrices with sizes $m_i \times r_b$, $r_b \times r_b$ and $n_i \times r_b$ respectively and form the rank- r_b reduced SVD of $R_b(X_i)$. Note that when the rank of the block matrix $R_b(X_i)$ is

zero, we have $\{U_b, S_b, V_b\}$ as empty matrices, which do not contribute to X_j . Figure 2 and 4 provide illustrations of two kinds of modeling with their associated partitions.

By constraining each block matrices to be of low rank, the multi-scale low rank modeling captures the notion that some nearby entries are more similar to each other than global entries in the data matrix. We note that the multi-scale low rank modeling is a generalization of the low rank + sparse modeling proposed by Chandrasekaren et al. [17] and Candès et al. [25]. In particular, the low rank + sparse modeling can be viewed as a 2-scale low rank modeling, in which the first scale has block size 1×1 and the second scale has block size $M \times N$. By adding additional scales between the sparse and globally low rank matrices, the multi-scale low rank modeling can capture locally low rank components that would otherwise need many coefficients to represent for low rank + sparse.

Given a data matrix Y that fits our multi-scale low rank model, our goal is to decompose the data matrix Y to its multi-scale components $\{X_i\}_{i=1}^L$. The ability to recover these multi-scale components is beneficial for many applications and allows us to, for example, extract motions at multiple scales in surveillance videos (Section X). Since there are many more parameters to be estimated than the number of observations, it is necessary to impose conditions on X_j . In particular, we will exploit the fact that each block matrix is low rank via a convex program, which will be described in detail in section III.

A. Multi-scale low rank + noise

Before moving to the convex formulation, we note that our multi-scale matrix modeling can easily account for data matrices that are corrupted by additive white Gaussian noise. Under the multi-scale low rank modeling, we can think of the additive noise matrix as the largest scale signal component and is unstructured in any local scales. Specifically if we observe instead the following

$$Y = \sum_{i=1}^L X_i + X_Z \quad (1)$$

where X_Z is an independent and identically distributed Gaussian noise matrix. Then we can define a reshape operator R_Z that reshapes the entire matrix into an $MN \times 1$ vector and the resulting matrix fits exactly to our multi-scale low rank model with $L+1$ scales. This incorporation of noise makes our model flexible in that it automatically provides a corresponding convex relaxation, a regularization parameter for the noise matrix and allows us to utilize the same iterative algorithm to solve for the noise matrix. Figure 5 provides an example of the noisy multi-scale low rank matrix decomposition.

III. PROBLEM FORMULATION AND CONVEX RELAXATION

Given a data matrix Y that fits the multi-scale low rank model, our goal is to recover the underlying multi-scale components $\{X_i\}_{i=1}^L$ using the fact that X_j is block-wise low rank.

Ideally, we would like to obtain a multi-scale decomposition with the minimal block matrix rank and solve a problem similar to the following form:

$$\begin{aligned} & \underset{X_1, \dots, X_L}{\text{minimize}} && \sum_{i=1}^L \sum_{b \in P_i} \text{rank}(R_b(X_i)) \\ & \text{Subject to} && Y = \sum_{i=1}^L X_i \end{aligned}$$

However, each rank minimization for each block is combinatorial in nature. In addition, it is not obvious whether the direct summation of ranks is a correct formulation as a 1-sparse matrix and a rank-1 matrix should intuitively not carry the same cost. Hence, the above non-convex problem is not a practical formulation to obtain the multi-scale decomposition.

Recent development in convex relaxations suggests that rank minimization problems can often be relaxed to a convex program via nuclear norm relaxation [13], [23], while still recovering the optimal solution to the original problem. In particular, Chandrasekaran et al. [17] and Candès et al., [25] showed that a low rank + sparse decomposition can be relaxed to a convex program by minimizing a nuclear norm + ℓ_1 -norm objective as long as the signal constituents are incoherent with respect to each other. In addition, Candès et al., [25] showed that the regularization parameters for sparsity and low rank should be related by the square root of the matrix size. Hence, there is hope that, along the same line, we can perform the multi-scale low rank decomposition exactly via a convex formulation.

Concretely, let us define $\|\cdot\|_{\text{nuc}}$ to be the nuclear norm, the sum of singular values, and $\|\cdot\|_{\text{msv}}$ be the maximum singular value norm. For each scale i , we consider the block-wise nuclear norm to be the convex surrogate for the block-wise ranks and define $\|\cdot\|_{(i)}$ the block-wise nuclear norm for the i th scale as

$$\|\cdot\|_{(i)} = \sum_{b \in P_i} \|R_b(\cdot)\|_{\text{nuc}}$$

Its associated dual norm $\|\cdot\|_{(i)}^*$ is then given by

$$\|\cdot\|_{(i)}^* = \max_{b \in P_i} \|R_b(\cdot)\|_{\text{msv}}$$

which is the maximum of all block-wise maximum singular values.

We then consider the following convex relaxation for our multi-scale low rank decomposition problem:

$$\begin{aligned}
& \underset{X_1, \dots, X_L}{\text{minimize}} && \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} \\
& \text{subject to} && Y = \sum_{i=1}^L X_i
\end{aligned} \tag{2}$$

where $\{\lambda_i\}_{i=1}^L$ are the regularization parameters and their selection will be described in detail in section IV.

Our convex formulation is a natural generalization of the low rank + sparse convex formulation [17], [25]. With the two sided matrix partition (Fig. 2), the nuclear norm applied to the 1×1 blocks becomes the element-wise ℓ -norm and the norm for the largest scale is the nuclear norm. With the one sided matrix partition (Fig. 4), the nuclear norm applied to $1 \times N$ blocks becomes the group-sparse norm and can be seen as a generalization of the group sparse + low rank decomposition [21]. If we incorporate additive Gaussian noise in our model as described in Section II-A, then we have a nuclear norm applied to an $MN \times 1$ vector, which is equivalent to the Frobenius norm.

One should hope that the theoretical conditions from low rank + sparse can be generalized rather seamlessly to the multi-scale counterpart. Indeed, in Section V, we show that the core theoretical guarantees in the work of Chandrasekaren et al. [17] on exact low rank + sparse decomposition can be generalized to the multi-scale setting. In section VI, we show that the core theoretical guarantees in the work of Agarwal et al. [18] on noisy matrix decomposition can be generalized to the multi-scale setting as well to provide approximate decomposition guarantees.

IV. GUIDANCE ON CHOOSING REGULARIZATION PARAMETERS

In this section, we provide practical guidance on selecting the regularization parameters $\{\lambda_i\}_{i=1}^L$. Selecting the regularization parameters $\{\lambda_i\}_{i=1}^L$ is crucial for the convex decomposition to succeed, both theoretically and practically. While theoretically we can establish criteria on selecting the regularization parameters (see Section V and VI), such parameters are not straightforward to calculate in practice as it requires properties of the signal components $\{X_i\}_{i=1}^L$ *before* the decomposition.

To select the regularization parameters $\{\lambda_i\}_{i=1}^L$ in practice, we follow the suggestions from Wright et al. [33] and Fogel et al. [34], and set each regularization parameter λ_i to be the Gaussian complexity of each norm $\|\cdot\|_{(i)}$, which is defined the expectation of the dual norm of random Gaussian matrix:

$$\lambda_i \sim E \left[\|G\|_{(i)}^* \right] \tag{3}$$

where \sim denotes equality up to some constant and G is a unit-variance independent and identically distributed random Gaussian matrix.

The resulting expression for the Gaussian complexity is the maximum singular value of a random Gaussian matrix, which has been studied extensively by Bandeira and Handel [35]. The recommended regularization parameter for scale i is given by

$$\lambda_i \sim \sqrt{m_i} + \sqrt{n_i} + \sqrt{\log\left(\frac{MN}{\max(m_i, n_i)}\right)} \quad (4)$$

For the sparse matrix scale with 1×1 block size, $\lambda_i \sim \sqrt{\log(MN)}$ and for the globally low rank scale with $M \times N$ block size, $\lambda_i \sim \sqrt{M} + \sqrt{N}$. Hence this regularization parameter selection is consistent with the ones recommended for low rank + sparse decomposition by Candès et al. [25], up to a log factor. In addition, for the noise matrix with $MN \times 1$ block size, $\lambda_i \sim \sqrt{MN}$, which has similar scaling as in square root LASSO [36]. In practice, we found that the suggested regularization parameter selection allows exact multi-scale decomposition when the signal model is matched (for example Figure 1) and provides visually intuitive decomposition for real datasets.

For approximate multi-scale low rank decomposition in the presence of additive noise, some form of theoretical guarantees for the regularization selection can be found in our analysis in Section VI. In particular, we show that if the regularization parameters λ_i is larger than the Gaussian complexity of $\|\cdot\|_{(i)}^*$ in addition to some “spikiness” parameters, then the error between recovered decomposition and the ground truth $\{X_i\}_{i=1}^L$ is bounded by the block-wise matrix rank.

V. THEORETICAL ANALYSIS FOR EXACT DECOMPOSITION

In this section, we provide a theoretical analysis of the proposed convex formulation and show that if $\{X_i\}_{i=1}^L$ satisfies a deterministic incoherence condition, then the proposed convex formulation (2) recovers $\{X_i\}_{i=1}^L$ from Y exactly.

Our analysis follows similar arguments taken by Chandrasekaran et al. [17] on low rank + sparse decomposition and generalizes them to the proposed multi-scale low rank decomposition. Before showing our main result (Theorem V.1), we first describe the subgradients of our objective function (Section V-A) and define a coherence parameter in terms of the block-wise row and column spaces (Section V-B).

A. Subdifferentials of the block-wise nuclear norms

To characterize the optimality of our convex problem, we first look at the subgradients of our objective function. We recall that for any matrix X with $\{U, S, V\}$ as its reduced SVD representation, the subdifferential of $\|\cdot\|_{\text{nuc}}$ at X is given by [23], [37],

$$\partial \|X\|_{\text{nuc}} = \left\{ UV^\top + W : W \text{ and } X \text{ have orthogonal row and column spaces and } \|W\|_{\text{msv}} \leq 1 \right\}$$

Now recall that we define the block-wise nuclear norm to be $\|\cdot\|_{(i)} = \sum_{b \in P_i} \|R_b(\cdot)\|_{\text{nuc}}$.

Then using the chain rule and the fact that $R_b(X_i) = U_b S_b V_b^\top$, we obtain an expression for the subdifferential of $\|\cdot\|_{(i)}$ at X_i as follows:

$$\partial \|X_i\|_{(i)} = \left\{ \sum_{b \in P_i} R_b^\top (U_b V_b^\top + W_b) : W_b \text{ and } R_b(X_i) \text{ have orthogonal row and column spaces and } \|W_b\|_{\text{msv}} \leq 1 \right\}$$

To simplify our notation, we define $E_i = \sum_{b \in P_i} R_b^\top (U_b V_b^\top)$ and T_i to be a vector space that contains matrices with the same block-wise row spaces or column spaces as X_i that is,

$$T_i = \left\{ \sum_{b \in P_i} R_b^\top (U_b X_b^\top + Y_b V_b^\top) : X_b \in \mathbb{C}^{n_i \times r_i}, Y_b \in \mathbb{C}^{m_i \times r_i} \right\}$$

where $m_i \times n_i$ is the size of the block matrices for scale i and r_b is the matrix rank for block b . Then, the subdifferential of each $\|\cdot\|_{(i)}$ at X_i can be compactly represented as,

$$\partial \|X_i\|_{(i)} = \left\{ E_i + W_i : W_i \in T_i^\perp \text{ and } \|W_i\|_{(i)}^* \leq 1 \right\}$$

We note that E_i can be thought of as the “sign” of the matrix X_i pointing toward the principal components, and, in the case of the sparse scale, is exactly the sign of the entries.

In the rest of the section, we will be interested in projecting a matrix X onto T_i which can be performed with the following operation:

$$\mathcal{P}_{T_i}(X) = \sum_{b \in P_i} R_b^\top \left(U_b U_b^\top R_b(X) + R_b(X) V_b V_b^\top - U_b U_b^\top R_b(X) V_b V_b^\top \right)$$

Similarly, to project a matrix X onto the orthogonal complement of T_i , we can apply the following operation:

$$\mathcal{P}_{T_i^\perp}(X) = \sum_{b \in P_i} R_b^\top \left((I - U_b U_b^\top) R_b(X) (I - V_b V_b^\top) \right)$$

where I is an appropriately sized identity matrix.

B. Incoherence

Following Chandrasekaren et al. [17], we consider a deterministic measure of incoherence through the block-wise column and row spaces of X_j . Concretely, we define the coherence parameter for the j th scale signal component X_j with respect to the i th scale to be the following:

$$\mu_{ij} = \max_{N_j \in T_j, \|N_j\|_{(j)}^* \leq 1} \|N_j\|_{(i)}^* \quad (5)$$

Using μ_{ij} as a measure of incoherence, we can quantitatively say that the j th scale signal component is incoherent with respect to the i th scale if μ_{ij} is small. In the case of low rank + sparse, Chandrasekaren et al. [17] provides excellent description of the concepts behind the coherence parameters. We refer the reader to their paper for more detail.

C. Main Result

Given the above definition of incoherence, the following theorem states our main result for exact multi-scale low rank decomposition:

Theorem V.1—If we can choose regularization parameters $\{\lambda_i\}_{i=1}^L$ such that

$$\sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i} < \frac{1}{2}, \text{ for } i=1, \dots, L$$

then $\{X_i\}_{i=1}^L$ is the unique optimizer of the proposed convex problem (2).

In particular when the number of scales $L = 2$, the condition on $\{\mu_{12}, \mu_{21}\}$ reduces to $\mu_{12}\mu_{21} < 1/4$ and the condition on $\{\lambda_1, \lambda_2\}$ reduces to $2\mu_{12} < \lambda_1/\lambda_2 < 1/(2\mu_{21})$, which is in similar form as Theorem 2 in Chandrasekaren et al. [17].

The proof for the above theorem is given in Appendix A.

VI. THEORETICAL ANALYSIS FOR APPROXIMATE DECOMPOSITION

In this section, we provide a theoretical analysis for approximate multi-scale low rank decomposition when the measurement is corrupted by additive noise as described in Section II-A. Our result follows arguments from Agarwal et al. [18] on noisy 2-scale matrix decomposition and extends it to the multi-scale setting.

Instead of using the incoherence parameter μ_{ij} defined for the exact decomposition analysis in Section V, we opt for a weaker characterization of incoherence between scales for approximate decomposition, studied in Agarwal et al. [18]. Concretely, we consider spikiness parameters α_{ij} between the j th signal component X_j and i th scale norm $\|\cdot\|_{(i)}$ such that,

$$\alpha_{ij} = \|X_j\|_{(i)}^*$$

for each $j \neq i$. Hence, if α_{ij} is small, we say X_j is not spiky with respect to the i th norm.

For analysis purpose, we also impose the constraints $\|X_j\|_{(i)}^* \leq \alpha_{ij}$ in the convex program. That is, we consider the solution from the following convex program:

$$\underset{X_1, \dots, X_L, X_{\mathcal{Z}}}{\text{minimize}} \quad \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} + \lambda_{\mathcal{Z}} \|X_{\mathcal{Z}}\|_{\text{fro}} \quad (6)$$

$$\begin{aligned} \text{subject to } Y &= \sum_{i=1}^L X_i + X_{\mathcal{Z}} \\ \|X_j\|_{(i)}^* &\leq \alpha_{ij} \text{ for } j \neq i \end{aligned} \quad (7)$$

We emphasize that the additional constraints (7) are imposed only for the purpose of theoretical analysis and are not imposed in our experimental results. In particular, for our simulation example in Figure 5, the minimizer of the convex program (2), using the recommended regularization parameters in Section IV, satisfied the constraints (7) even when the constraints were not imposed.

Let us define $\{\Delta_i\}_{i=1}^L$ and $X_{\mathcal{Z}}$ to be the errors between the ground truth components $\{X_i\}_{i=1}^L$ and $X_{\mathcal{Z}}$ and the minimizers of convex program (6). Then, equivalently, we can denote $\{X_i + \Delta_i\}_{i=1}^L$ and $X_{\mathcal{Z}+}$ as the minimizers of the convex program (6). The following theorem states our main result for approximate decomposition.

Theorem VI.1

If we choose $\{\lambda_i\}_{i=1}^L$ such that

$$\lambda_i \geq 2\lambda_{\mathcal{Z}} \frac{\|X_{\mathcal{Z}}\|_{(i)}^*}{\|X_{\mathcal{Z}}\|_{\text{fro}}} + \sum_{j \neq i} 2\alpha_{ij} \quad (8)$$

and $\lambda_{\mathcal{Z}}$ such that

$$\lambda_{\mathcal{Z}} \geq \sqrt{64 \sum_{i=1}^L \lambda_i^2 \sum_{b \in P_i} r_b} \quad (9)$$

then the error is bounded by

$$\sum_{i=1}^L \|\Delta_i\|_{\text{fro}} \lesssim \frac{\|X_{\mathcal{Z}}\|_{\text{fro}}}{\lambda_{\mathcal{Z}}} \sum_{i=1}^L \lambda_i \sqrt{\sum_{b \in P_i} r_b}$$

where \lesssim denotes inequality up to a universal constant.

Hence, when the spikiness parameters are negligible and $X_{\mathcal{Z}} = \sigma G$, where G is an independent, identically distributed Gaussian noise matrix with unit variance and σ is the noise standard deviation, choosing $\lambda_{\mathcal{Z}} \sim E[\|G\|_{\text{fro}}] \sim \sqrt{MN}$ and

$\lambda_i \sim E[\|G\|_{(i)}^*] \sim \sqrt{m_i} + \sqrt{n_i} + \sqrt{\log(MN/\max(m_i, n_i))}$ ensures the condition is satisfied with high probability. This motivates the recommended regularization selection in Section IV.

The proof for the above theorem is given in Appendix B and follows arguments from Agarwal et al. [18] on noisy matrix decomposition and Belloni et al. [36] on square root LASSO.

VII. AN ITERATIVE ALGORITHM FOR SOLVING THE MULTI-SCALE LOW RANK DECOMPOSITION

In the following, we will derive an iterative algorithm that solves for the multi-scale low rank decomposition via the Alternating Direction of Multiple Multipliers (ADMM) [38]. While the proposed convex formulation (2) can be formulated into a semi-definite program, first-order iterative methods are commonly used when solving for large datasets for their computational efficiency and scalability. A conceptual illustration of the algorithm is shown in Figure 6.

To formally obtain update steps using ADMM, we first formulate the problem into the standard ADMM form with two separable objectives connected by an equality constraint,

$$\begin{aligned} & \underset{X_i, Z_i}{\text{minimize}} && \mathbf{I}\left\{Y = \sum_{i=1}^L X_i\right\} + \sum_{i=1}^L \lambda_i \|Z_i\|_{(i)} \\ & \text{subject to} && X_i = Z_i \end{aligned} \quad (10)$$

where $\mathbf{I}\{\cdot\}$ is the indicator function.

To proceed, we then need to obtain the proximal operators [39] for the two objective functions $\mathbf{I}\left\{Y = \sum_{i=1}^L X_i\right\}$ and $\sum_{i=1}^L \lambda_i \|Z_i\|_{(i)}$. For the data consistency objective $\mathbf{I}\left\{Y = \sum_{i=1}^L X_i\right\}$, the proximal operator is simply the projection operator to the set. To obtain the proximal operator for the multi-scale nuclear norm objective $\sum_{i=1}^L \lambda_i \|X_i\|_{(i)}$, we

first recall that the proximal operator for the nuclear norm $\|X\|_{\text{nuc}}$ with parameter λ is given by the singular value soft-threshold operator [23],

$$\text{SVT}_{\lambda}(X) = U \max\left(\sum -\lambda, 0\right) V^{\top} \quad (11)$$

Since we defined the block-wise nuclear norm for each scale i as $\sum_{b \in P_i} \|R_b(\cdot)\|_{\text{nuc}}$, the norm is separable with respect to each block and its proximal function with parameter λ_i is given by the block-wise singular value soft-threshold operator,

$$\text{BLOCKSVT}_{\lambda_i}(X) = \sum_{b \in P_i} R_b^{\top}(\text{SVT}_{\lambda_i}(R_b(X))) \quad (12)$$

which simply extracts every blocks in the matrix, performs singular value thresholding and puts the blocks back to the matrix. We note that for 1×1 blocks, the block-wise singular value soft-threshold operator reduces to the element-wise soft-threshold operator and for $1 \times N$ blocks, the block-wise singular soft-threshold operator reduces to the joint soft-threshold operator.

Putting everything together and invoking the ADMM recipe [38], we have the following algorithm to solve our convex multi-scale low rank decomposition (2):

$$\boxed{\begin{aligned} X_i &\leftarrow (Z_i - U_i) + \frac{1}{L} \left(Y - \sum_{i=1}^L (Z_i - U_i) \right) \\ Z_i &\leftarrow \text{BLOCKSVT}_{\lambda_i/\rho}(X_i + U_i) \\ U_i &\leftarrow U_i - (Z_i - X_i) \end{aligned}} \quad (13)$$

where ρ is the ADMM parameter that only affects the convergence rate of the algorithm.

The resulting ADMM update steps are similar in essence to the intuitive update steps in Figure 6, and alternates between data consistency and enforcing multi-scale low rank. The major difference of ADMM is that it adds a dual update step with U_i which bridges the two objectives and ensures the convergence to the optimal solution. Under the guarantees of ADMM, in the limit of iterations, X_i and Z_i converge to the optimal solution of the convex program (2) and U_i converges to a scaled version of the dual variable. In practice, we found that ~ 1000 iterations are sufficient without any visible change for imaging applications. Finally, we note that because the proximal operator for the multi-scale nuclear norm is computationally simple, other proximal operator based algorithms [39] can also be used.

VIII. COMPUTATIONAL COMPLEXITY

Given the iterative algorithm (13), one concern about the multi-scale low rank decomposition might be that it is significantly more computationally intensive than other

low rank methods as we have many more SVD's and variables to compute for. In this section, we show that because we decimate the matrices at each scale geometrically, the theoretical computational complexity of the multi-scale low rank decomposition is similar to other low rank decomposition methods, such as the low rank + sparse decomposition.

For concreteness, let us consider the multi-scale partition with two-sided decimation shown in Figure 2 and have block sizes $m_j = 2^{i-1}$ and $n_j = 2^{i-1}$. Similar to other low rank methods, the SVD's dominate the per iteration complexity for the multi-scale low rank decomposition. For an $M \times N$ matrix, each SVD costs $\#flops(M \times N SVD) = O(MN^2)$. The per iteration complexity for the multi-scale low rank decomposition is dominated by the summation of all the SVD's performed for each scale, which is given by,

$$\begin{aligned} & \#flops(M \times N SVD) + 4\#flops(M/2 \times N/2 SVD) + \dots \\ & = O(MN^2) + O(MN^2)/2 + O(MN^2)/4 + \dots \\ & \leq 2O(MN^2) \approx \#flops(M \times N SVD) \end{aligned} \quad (14)$$

Hence, the per-iteration computational complexity of the multi-scale low rank with two-sided decimated partition is on the order of a $M \times N$ matrix SVD. In general, one can show that the per-iteration complexity for arbitrary multi-scale partition is at most $\log(N)$ times the full matrix SVD.

While theoretically, the computation cost for small block sizes should be less than bigger block sizes, we found that in practice the computation cost for computing the small SVD's can dominate the per-iteration computation. This is due to the overhead of copying small block matrices and calling library functions repeatedly to compute the SVD's.

Since we are interested in thresholding the singular values and in practice many of the small block matrices are zero as shown in Section X, one trick of reducing the computation time is to quickly compute an upper bound on the maximum singular value for block matrices before the SVD's. Then if the upper bound for the maximum singular value is less than the threshold, we know the thresholded matrix will be zero and can avoid computing the SVD. Since for any matrix X , its maximum singular value is bounded by the square root of any matrix norm on $X^T X$ [40], there are many different upper bounds that we can use. In particular, we choose the maximum row norm and consider the following upper bound,

$$\sigma_{\max}(X) \leq \sqrt{\max_i \sum_j |X_{ik} X_{jk}|} \quad (15)$$

Using this upper bound, we can identify many below-the-threshold matrices before computing the SVD's at all. In practice, we found that the above trick provides a modest speedup of $3 \sim 5\times$.

IX. HEURISTICS FOR TRANSLATION INVARIANT DECOMPOSITION

Similar to wavelet transforms, one drawback of the multi-scale low rank decomposition is that it is not translation invariant, that is, shifting the input changes the resulting decomposition. In practice, this translation variant nature often creates blocking artifacts near the block boundaries, which can be visually jarring for image or video applications. One solution to remove these artifacts is to introduce overlapping partitions of the matrix so that the overall algorithm is translation invariant. However, this vastly increases both memory and computation especially for large block sizes. In the following, we will describe a cycle spinning approach that we used in practice to reduce the blocking artifacts with only slight increase in per-iteration computation.

Cycle spinning [19] has been commonly used in wavelet denoising to reduce the blocking artifacts due to the translation variant nature of the wavelet transform. To minimize artifacts, cycle spinning averages the denoised results from all possible shifted copies of the input, thereby making the entire process translation invariant. Concretely, let \mathcal{S} be the set of all shifts possible in the target application, SHIFT_s denote the shifting operator by s , and DENOISE be the denoising operator of interest. Then the cycle spinned denoising of the input X is given by:

$$\frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \text{SHIFT}_{-s} (\text{DENOISE} (\text{SHIFT}_s (X))) \quad (16)$$

In the context of multi-scale low rank decomposition, we can make the iterative algorithm translation invariant by replacing the block-wise singular value thresholding operation in each iteration with its cycle spinning counterpart. In particular, for our ADMM update steps, we can replace the Z_i step to:

$$Z_i \leftarrow \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \text{SHIFT}_{-s} (\text{BLOCKSVT}_{\lambda_i/\rho} (\text{SHIFT}_s (X_i + U_i))) \quad (17)$$

To further reduce computation, we perform random cycle spinning in each iteration as described in Figueiredo et al. [41], in which we randomly shifts the input, performs block-wise singular value thresholding and then unshifts back:

$$Z_i \leftarrow \text{SHIFT}_{-s} (\text{BLOCKSVT}_{\lambda_i/\rho} (\text{SHIFT}_s (X_i + U_i))) \quad (18)$$

where s is randomly chosen from the set \mathcal{S} .

Using random cycle spinning, blocking artifacts caused by thresholding are averaged over iterations and in practice, reduces distortion significantly. Figure 7 shows an example of the multi-scale low rank decomposition with and without random cycle spinning applied on a simulated data that does not fall on the partition grid. The decomposition with random cycle

spinning vastly reduces blocking artifacts that appeared in the one without random cycle spinning.

X. APPLICATIONS

To test for practical performance, we applied the multi-scale low rank decomposition on four different real datasets that are conventionally used in low rank modeling: illumination normalization for face images (Section X-A), motion separation for surveillance videos (Section X-B), multi-scale modeling of dynamic contrast enhanced magnetic resonance imaging (Section X-C) and collaborative filtering exploiting age information (Section X-D). We compared our proposed multi-scale low rank decomposition with low rank + sparse decomposition for the first three applications and with low rank matrix completion for the last application. Randomly cycle spinning was used for multi-scale low rank decomposition for all of our experiments. Regularization parameters λ_j were chosen exactly as

$\sqrt{m_i} + \sqrt{n_i} + \sqrt{\log(MN/\max(m_i, n_i))}$ for multi-scale low rank and $\max(m_i, n_i)$ for low rank + sparse decomposition. Our simulations were written in the C programming language and ran on a 20-core Intel Xeon workstation. Some results are best viewed in video format, which are available as supplementary materials.

In the spirit of reproducible research, we provide a software package (in C and partially in MATLAB) to reproduce most of the results described in this paper. The software package can be downloaded from: https://github.com/frankong/multi_scale_low_rank.git

A. Multi-scale Illumination Normalization for Face Recognition Pre-processing

Face recognition algorithms are sensitive to shadows or occlusions on faces. In order to obtain the best possible performance for these algorithms, it is desired to remove illumination variations and shadows on the face images. Low rank modeling are often used to model faces and is justified by approximating faces as convex Lambertian surfaces [11].

Low rank + sparse decomposition [25] was recently proposed to capture uneven illumination as sparse errors and was shown to remove small shadows while capturing the underlying faces as the low rank component. However, most shadows are not sparse and contain structure over different lighting conditions. Here, we propose modeling shadows and illumination changes in different face images as block-low rank as illumination variations are spatially correlated in multiple scales.

We considered face images from the Yale B face database [42]. Each face image was of size 192×168 with 64 different lighting conditions. The images were then reshaped into a $32, 256 \times 64$ matrix and both multi-scale low rank and low rank + sparse decomposition were applied on the data matrix. For low rank + sparse decomposition, we found that the best separation result was obtained when each face image was normalized to the maximum value. For multi-scale low rank decomposition, the original unscaled image was used. Only the space dimension was decimated as we assumed there was no ordering in different illumination conditions. The multi-scale matrix partition can be visualized as in Figure 4.

Figure 8 shows one of the comparison results. Multi-scale low rank decomposition recovered almost shadow-free faces. In particular, the sparkles in the eyes were represented in the 1×1 block size and the larger illumination changes were represented in bigger blocks, thus capturing most of the uneven illumination changes. In contrast, low rank + sparse decomposition could only recover from small illumination changes and still contained the larger shadows in the globally low rank component.

B. Multi-scale Motion Separation for Surveillance Videos

In surveillance video processing, it is desired to extract foreground objects from the video. To be able to extract foreground objects, both the background and the foreground dynamics have to be modeled. Low rank modeling have been shown to be suitable for slowly varying videos, such as background illumination changes. In particular, if the video background only changes its brightness over time, then it can be represented as a rank-1 matrix.

Low rank + sparse decomposition [25] was proposed to foreground objects as sparse components and was shown to separate dynamics from background components. However, sparsity alone cannot capture motion compactly and often results in ghosting artifacts occurring around the foreground objects as shown in Figure 9. Since video dynamics are correlated locally at multiple scales in space and time, we propose using the multi-scale low rank modeling with two sided decimation to capture different scales of video dynamics over space and time.

We considered a surveillance video from Li et al. [43]. Each video frame was of size 144×176 and the first 200 frames were used. The video frames were then reshaped into a $25,344 \times 200$ matrix and both multi-scale low rank and low rank + sparse decomposition were applied on the data matrix.

Figure 9 shows one of the results. Multi-scale low rank decomposition recovered a mostly artifact free background video in the globally low rank component whereas low rank + sparse decomposition exhibits ghosting artifact in certain segments of the video. For the multi-scale low rank decomposition, body motion was mostly captured in the $16 \times 16 \times 16$ scale while fine-scale motion was captured in $4 \times 4 \times 4$ scale.

C. Multi-scale Low Rank Modeling for Dynamic Contrast Enhanced Magnetic Resonance Imaging

In dynamic contrast enhanced magnetic resonance imaging (DCE-MRI), a series of images over time is acquired after a T_1 contrast agent was injected into the patient. Different tissues then exhibit different contrast dynamics over time, thereby allowing radiologists to characterize and examine lesions. Compressed sensing Magnetic Resonance Imaging [44] is now a popular research approach used in three dimensional DCE-MRI to speed up acquisition. Since the more compact we can represent the image series, the better our compressed reconstruction result becomes, an accurate modeling of the dynamic image series is desired to improve the compressed sensing reconstruction results for DCE-MRI.

When a region contains only one type of tissue, then the block matrix constructed by stacking each frame as columns will have rank 1. Hence, low rank modeling [10], and

locally low rank modeling [45] have been popular models for DCE-MRI. Recently, low rank + sparse modeling [46] have also been proposed to model the static background and dynamics as low rank and sparse matrices respectively. However, dynamics in DCE-MRI are almost never sparse and often exhibit correlation across different scales. Hence, we propose using a multi-scale low rank modeling to capture contrast dynamics over multiple scales.

We considered a fully sampled dynamic contrast enhanced image data. The data was acquired in a pediatric patient with 20 contrast phases, $1 \times 1.4 \times 2 \text{ mm}^3$ resolution, and 8s temporal resolution. The acquisition was performed on a 3T GE MR750 scanner with a 32-channel cardiac array using an RF-spoiled gradient-echo sequence. We considered a 2D slice of size 154×112 were then reshaped into a $17, 248 \times 20$ matrix. Both multi-scale low rank and low rank + sparse decomposition were applied on the data matrix.

Figure 10 shows one of the results. In the multi-scale low rank decomposition result, small contrast dynamics in vessels were captured in 4×4 blocks while contrast dynamics in the liver were captured in 16×16 blocks. The biggest block size captured the static tissues and interestingly the respiratory motion. Hence, different types of contrast dynamics were captured compactly in their suitable scales. In contrast, the low rank + sparse modeling could only provide a coarse separation of dynamics and static tissue, which resulted in neither truly sparse nor truly low rank components.

D. Multi-scale Age Grouping for Collaborative Filtering

Collaborative filtering is the task of making predictions about the interests of a user using available information from all users. Since users often have similar taste for the same item, low rank modeling is commonly used to exploit the data similarity to complete the rating matrix [14], [22], [23]. On the other hand, low rank matrix completion does not exploit the fact that users with similar demographic backgrounds have similar taste for similar items. In particular, users of similar age should have similar taste. Hence, we incorporated the proposed multi-scale low rank modeling with matrix completion by partitioning users according to their age and compared it with the conventional low rank matrix completion. Our method belongs to the general class of collaborative filtering methods that utilize demographic information [32].

To incorporate multi-scale low rank modeling into matrix completion, we change the data

consistency constraint in problem (2) to $[Y]_{jk} = \left[\sum_{i=1}^L X_i \right]_{jk}$ for observed jk entries, and correspondingly, the update step for $\{X_i\}_{i=1}^L$ in equation (13) is changed to

$[X_i]_{jk} \leftarrow \left[(Z_i - U_i) + \frac{1}{L} \left(Y - \sum_{i=1}^L (Z_i - U_i) \right) \right]_{jk}$ for observed jk entries and $[X_i]_{jk} \leftarrow [Z_i - U_i]_{jk}$ for unobserved jk entries. We emphasize that our theoretical analysis does not cover matrix completion and the presented collaborative filtering application is mainly of empirical interest.

To compare the methods, we considered the 100K Movie-Lens dataset, in which 943 users rated 1682 movies. The resulting matrix was of size 1682×943 , where the first dimension

represented movies and the second dimension represented users. The entire matrix had 93.7% missing entries. Test data was further generated by randomly undersampling the rating matrix by 5. The algorithms were then run on the test data and root mean squared errors were calculated over all available entries. To obtain a multi-scale partition of the matrix, we sorted the users according to their age along the second dimension and partitioned them evenly into age groups.

Figure 11 shows a multi-scale low rank reconstructed user rating matrix. Using multiple scales of block-wise low rank matrices, correlations in different age groups were captured. For example, one of the scales shown in Figure 11 captures the tendency that younger users rated Star Wars higher whereas the more senior users rated Gone with the Wind higher. The multi-scale low rank reconstructed matrix achieved a root mean-squared-error of 0.9385 compared to a root mean-squared-error of 0.9552 for the low rank reconstructed matrix.

XI. DISCUSSION

We have presented a multi-scale low rank matrix decomposition method that combines both multi-scale modeling and low rank matrix decomposition. Using a convex formulation, we can solve for the decomposition efficiently and exactly, provided that the multi-scale signal components are incoherent. We provided a theoretical analysis of the convex relaxation for exact decomposition, which extends the analysis in Chandrasekaran et al. [17], and an analysis for approximate decomposition in the presence of additive noise, which extends the analysis in Agarwal et al. [18]. We also provided empirical results that the multi-scale low rank decomposition performs well on real datasets.

We would also like to emphasize that our recommended regularization parameters empirically perform well even with the addition of noise, and hence in practice does not require manual tuning. While some form of theoretical guarantees for the regularization parameters are provided in the approximate decomposition analysis, complete theoretical guarantees are not provided, especially for noiseless situations, and would be valuable for future work.

Our experiments show that the multi-scale low rank decomposition improves upon the low rank + sparse decomposition in a variety of applications. We believe that more improvement can be achieved if domain knowledge for each applications is incorporated with the multi-scale low rank decomposition. For example, for face shadow removal, prior knowledge of the illumination angle might be able to provide a better multi-scale partition. For movie rating collaborative filtering, general demographic information and movie types can be used to construct multi-scale partitions in addition to age information.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work is supported by NIH grants R01EB019241, R01EB009690, and P41RR09784, Sloan research fellowship, Okawa research grant and NSF GRFP.

APPENDIX A PROOF OF THEOREM V.1

In this section, we provide a proof of Theorem V.1 and show that if $\{X_i\}_{i=1}^L$ satisfies a deterministic incoherence condition, then the proposed convex formulation (2) recovers $\{X_i\}_{i=1}^L$ from Y exactly. Our proof makes use of the dual certificate common in such proofs. We will begin by proving a technical lemma collecting three inequalities.

Lemma A.1

For $i = 1, \dots, L$, the following three inequalities hold,

$$\|\mathcal{P}_{T_i}(X)\|_{(i)}^* \leq \|X\|_{(i)}^* \quad \text{for any matrix } X \quad (19)$$

$$\|\mathcal{P}_{T_i^\perp}(X)\|_{(i)}^* \leq \|X\|_{(i)}^* \quad \text{for any matrix } X \quad (20)$$

$$\|N_j\|_{(i)}^* \leq \mu_{ij} \|N_j\|_{(j)}^* \quad \text{for } j \neq i \text{ and } N_j \in T_j \quad (21)$$

Proof

To show the first inequality (19), we recall that $\|X\|_{(i)}^* = \max_{b \in P_i} \|R_b(X)\|_{\text{msv}}$. Then, using the variational representation of the maximum singular value norm, we obtain,

$$\begin{aligned} \|\mathcal{P}_{T_i}(X)\|_{(i)}^* &= \max_{b \in P_i} \max_{u, v} u^\top R_b(\mathcal{P}_{T_i}(X))v \\ &= \max_{b \in P_i} \max_{\substack{u \in \text{col}(R_b(X_i)) \text{ or} \\ v \in \text{row}(R_b(X_i))}} u^\top R_b(X)v \\ &\leq \max_{b \in P_i} \max_{u, v} u^\top R_b(X)v = \|X\|_{(i)}^* \end{aligned}$$

where col and row denote the column and row spaces respectively.

Similarly, we obtain the second inequality (20):

$$\begin{aligned} \|\mathcal{P}_{T_i^\perp}(X)\|_{(i)}^* &= \max_{b \in P_i} \max_{\substack{u \in \text{col}^\perp(R_b(X_i)) \text{ and} \\ v \in \text{row}^\perp(R_b(X_i))}} u^\top R_b(X)v \\ &\leq \max_{b \in P_i} \max_{u, v} u^\top R_b(X)v = \|X\|_{(i)}^* \end{aligned}$$

The third inequality (21) follows from the incoherence definition that $\mu_{ij} \geq \|N_j\|_{(i)}^* / \|N_j\|_{(j)}^*$ for any non-zero N_j . ■

Next, we will show that if we can choose some parameters to “balance” the coherence between the scales, then the block-wise row/column spaces $\{T_i\}_{i=1}^L$ are independent, that is $\sum_{i=1}^L T_i$ is a direct sum. Consequently, each matrix N in the span of $\{T_i\}_{i=1}^L$ has a unique decomposition $N = \sum_{i=1}^L N_i$, where $N_i \in T_i$.

Proposition A.2

If we can choose some positive parameters $\{\lambda_i\}_{i=1}^L$ such that

$$\sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i} < 1, \quad \text{for } i=1, \dots, L \quad (22)$$

then we have

$$T_i \cap \sum_{j \neq i} T_j = \{0\}, \quad \text{for } i=1, \dots, L \quad (23)$$

In particular when $L = 2$, the condition on $\{\mu_{12}, \mu_{21}\}$ reduces to $\mu_{12}\mu_{21} < 1$, which coincides with Proposition 1 in Chandrasekaran et al. [17]. We also note that given μ_{ij} , we can obtain $\{\lambda_i\}_{i=1}^L$ that satisfies the condition $\sum_{j \neq i} \mu_{ij} \lambda_j < \lambda_i$ by solving a linear program.

Proof

Suppose by contradiction that there exists $\{\lambda_i\}_{i=1}^L$ such that $\sum_{j \neq i} \mu_{ij} \lambda_j / \lambda_i < 1$, but $T_i \cap \sum_{j \neq i} T_j \neq \{0\}$. Then there exists $\{N_i \in T_i\}_{i=1}^L$ such that $\sum_{i=1}^L \lambda_i N_i = 0$ and not all N_i zero. But this leads to a contradiction because for $i = 1, \dots, L$,

$$\begin{aligned} \|N_i\|_{(i)}^* &= \left\| - \sum_{j \neq i} \frac{\lambda_j}{\lambda_i} N_j \right\|_{(i)}^* \\ &\leq \sum_{j \neq i} \frac{\lambda_j}{\lambda_i} \mu_{ij} \|N_j\|_{(j)}^* \end{aligned} \quad (24)$$

$$\leq \left(\sum_{j \neq i} \frac{\lambda_j}{\lambda_i} \mu_{ij} \right) \max_{j \neq i} \|N_j\|_{(j)}^* \quad (25)$$

$$< \max_{j \neq i} \|N_j\|_{(j)}^* \quad (26)$$

where we have used equation (21) for the first inequality (24), Holder's inequality for second inequality (25) and $\sum_{j \neq i} \mu_{ij} \lambda_j / \lambda_i < 1$ for the last inequality. Hence, none of $\{\|N_i\|_{(j)}^*\}_{i=1}^L$ is the largest of the set, which is a contradiction. ■

Our next theorem shows an optimality condition of the convex program (2) in terms of its dual solution.

Theorem A.3

(Lemma 4.2 [33]). $\{X_i\}_{i=1}^L$ is the unique minimizer of the convex program (2) if there exists a matrix Q such that for $i = 1, \dots, L$,

- 1) $\mathcal{P}_{T_i}(Q) = \lambda_i E_i$
- 2) $\|\mathcal{P}_{T_i^\perp}(Q)\|_{(i)}^* < \lambda_i$

Proof

Consider any non-zero perturbation $\{\Delta_i\}_{i=1}^L$ to $\{X_i\}_{i=1}^L$ such that $\{X_i + \Delta_i\}_{i=1}^L$ stays in the feasible set, that is $\sum_{i=1}^L \Delta_i = 0$. We will show that $\sum_{i=1}^L \lambda_i \|X_i + \Delta_i\|_{(i)} > \sum_{i=1}^L \lambda_i \|X_i\|_{(i)}$.

We first decompose Δ_i into orthogonal parts with respect to T_i that is,

$\Delta_i = \mathcal{P}_{T_i}(\Delta_i) + \mathcal{P}_{T_i^\perp}(\Delta_i)$. We also consider a specific subgradient $G = [G_1 \dots G_L]^\top$ of

$\sum_{i=1}^L \lambda_i \|\cdot\|_{(i)}$ at $\{X_i\}_{i=1}^L$ such that $\|\mathcal{P}_{T_i^\perp}(G_i)\|_{(i)}^* \leq \lambda_i$, and

$\langle \mathcal{P}_{T_i^\perp}(\Delta_i), \mathcal{P}_{T_i^\perp}(G_i) \rangle = \lambda_i \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)}$. Then, from the definition of subgradient and the

fact that $\sum_{i=1}^L \Delta_i = 0$, we have,

$$\begin{aligned} \sum_{i=1}^L \lambda_i \|X_i + \Delta_i\|_{(i)} &\geq \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} + \langle \Delta_i, G_i \rangle \\ &= \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} + \langle \Delta_i, G_i \rangle - \langle \Delta_i, Q \rangle \end{aligned}$$

Applying the orthogonal decomposition with respect to T_i and using

$\mathcal{P}_{T_i}(G_i) = \mathcal{P}_{T_i}(Q) = \lambda_i E_i$, we have,

$$\begin{aligned} \sum_{i=1}^L \lambda_i \|X_i + \Delta_i\|_{(i)} &\geq \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} + \langle \mathcal{P}_{T_i^\perp}(\Delta_i), \mathcal{P}_{T_i^\perp}(G_i) \rangle \\ &\quad - \langle \mathcal{P}_{T_i^\perp}(\Delta_i), \mathcal{P}_{T_i^\perp}(Q) \rangle \end{aligned}$$

Using Holder's inequality and the assumption for the subgradient G_i , we obtain,

$$\begin{aligned}
\sum_{i=1}^L \lambda_i \|X_i + \Delta_i\|_{(i)} &\geq \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} + \lambda_i \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \\
&\quad - \|\mathcal{P}_{T_i^\perp}(Q)\|_{(i)}^* \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \\
&> \sum_{i=1}^L \lambda_i \|X_i\|_{(i)}
\end{aligned}$$

■

With Proposition A.2 and Theorem A.3, we are ready to prove Theorem V.1.

Proof of Theorem V.1

Since $\sum_{j \neq i} \mu_{ij} \lambda_j / \lambda_i < 1/2$, by Proposition A.2, $T_i \cap \sum_{j \neq i} T_j = \{0\}$ for all i . Thus, there is a unique matrix Q in $\sum_{i=1}^L T_i$ such that $\mathcal{P}_{T_i}(Q) = \lambda_i E_i$. In addition, Q can be uniquely expressed as a sum of elements in T_i . That is, $Q = \sum_{i=1}^L Q_i$ with $Q_i \in T_i$. We now have a matrix Q that satisfies the first optimality condition. In the following, we will show that it also satisfies the second optimality condition $\|\mathcal{P}_{T_i^\perp} Q\|_{(i)}^* < \lambda_i$.

If the vector spaces $\{T_i\}_{i=1}^L$ are orthogonal, then Q_i is exactly $\lambda_i E_i$. Because they are not necessarily orthogonal, we express Q_i as $\lambda_i E_i$ plus a correction term $\lambda_i \epsilon_i$. That is, we express $Q_i = \lambda_i (E_i + \epsilon_i)$. Putting Q_i 's back to Q , we have

$$Q = \sum_{i=1}^L \lambda_i (E_i + \epsilon_i) \quad (27)$$

Combining the above equation (27) with the first optimality condition (A.3), $\mathcal{P}_{T_i}(Q) = \lambda_i E_i$, we have $\sum_{j=1}^L \lambda_j \mathcal{P}_{T_i}(E_j + \epsilon_j) = \lambda_i E_i$. Since $\mathcal{P}_{T_i}(E_i + \epsilon_i) = E_i + \epsilon_i$, rearranging the equation, we obtain the following recursive expression for ϵ_i :

$$\epsilon_i = - \mathcal{P}_{T_i} \left(\sum_{j \neq i} \frac{\lambda_j}{\lambda_i} (E_j + \epsilon_j) \right) \quad (28)$$

We now obtain a bound on $\|\mathcal{P}_{T_i^\perp}(Q)\|_{(i)}^*$ in terms of ϵ_i .

$$\begin{aligned}
\|\mathcal{P}_{T_i^\perp}(Q)\|_{(i)}^* &= \|\mathcal{P}_{T_i^\perp}(\sum_{j \neq i} \lambda_j (E_j + \epsilon_j))\|_{(i)}^* \\
&\leq \|\sum_{j \neq i} \lambda_j (E_j + \epsilon_j)\|_{(i)}^*
\end{aligned} \quad (29)$$

$$\leq \sum_{j \neq i} \mu_{ij} \lambda_j (1 + \|\varepsilon_j\|_{(j)}^*) \quad (30)$$

$$\leq \left(\sum_{j \neq i} \mu_{ij} \lambda_j \right) \max_{j \neq i} (1 + \|\varepsilon_j\|_{(j)}^*) \quad (31)$$

where we obtain equation (29) from equation (20), equation (30) from equation (21) and the last inequality (31) from Holder's inequality.

Similarly, we obtain a recursive expression for $1 + \|\varepsilon_i\|_{(i)}^*$ using equation (28)

$$\begin{aligned} 1 + \|\varepsilon_i\|_{(i)}^* &= 1 + \|\mathcal{P}_{T_i} \left(\sum_{j \neq i} \frac{\lambda_j}{\lambda_i} (E_j + \varepsilon_j) \right)\|_{(i)}^* \\ &\leq 1 + \left\| \sum_{j \neq i} \frac{\lambda_j}{\lambda_i} (E_j + \varepsilon_j) \right\|_{(i)}^* \end{aligned} \quad (32)$$

$$\leq 1 + \sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i} (1 + \|\varepsilon_j\|_{(j)}^*) \quad (33)$$

$$\leq 1 + \left(\sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i} \max_{j \neq i} (1 + \|\varepsilon_j\|_{(j)}^*) \right) \quad (34)$$

where we obtain equation (32) from equation (19), equation (33) from equation (21) and the last inequality (34) from Holder's inequality.

Taking the maximum over i on both sides and rearranging, we have

$$\max_i (1 + \|\varepsilon_i\|_{(i)}^*) \leq \frac{1}{1 - \max_i \sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i}}$$

Putting the bound back to equation (31), we obtain

$$\begin{aligned} \|\mathcal{P}_{T_i^\perp}(Q)\|_{(i)}^* &\leq \lambda_i \frac{\sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i}}{1 - \max_i \sum_{j \neq i} \mu_{ij} \frac{\lambda_j}{\lambda_i}} \\ &< \lambda_i \end{aligned} \quad (35)$$

where we used $\sum_{j \neq i} \mu_{ij} \lambda_j / \lambda_i < 1/2$ in the last inequality.

Thus, we have constructed a dual certificate Q that satisfies the optimality conditions (A.3) and $\{X_i\}_{i=1}^L$ is the unique optimizer of the convex problem (2).

APPENDIX B PROOF OF THEOREM VI

In this section, we provide a proof of Theorem VI, showing that as long as we can choose our regularization parameters accordingly, we obtain a solution from the convex program (6) that is close to the ground truth $\{X_i\}_{i=1}^L$.

We will begin by proving a technical lemma collecting three inequalities. Throughout the section, we will assume $X_{\mathcal{X}}$ is non-zero for simplicity, so that the subgradient of $\|X_{\mathcal{X}}\|_{\text{fro}}$ is exactly $X_{\mathcal{X}} / \|X_{\mathcal{X}}\|_{\text{fro}}$.

Lemma B.1

For $i=1, \dots, L$, the following three inequalities hold,

$$\|X_i\|_{(i)} - \|X_i + \Delta_i\|_{(i)} \leq \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} - \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \quad (36)$$

$$\sum_{i=1}^L \lambda_i \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \leq 3 \sum_{i=1}^L \lambda_i \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} \quad (37)$$

$$\|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} \leq \sqrt{2 \sum_{b \in \mathcal{P}_i} r_b} \|\Delta_i\|_{\text{fro}} \quad (38)$$

Proof

We will prove the inequalities in order.

Let us choose a subgradient $G_i = E_i + W_i$ of $\|X_i\|_{(i)}$ at X_i such that

$\langle W_i, \mathcal{P}_{T_i^\perp}(\Delta_i) \rangle = \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)}$. Then, from the definition of the subgradient, we have,

$$\begin{aligned} \|X_i + \Delta_i\|_{(i)} &\geq \|X_i\|_{(i)} + \langle G_i, \Delta_i \rangle \\ &= \|X_i\|_{(i)} + \langle E_i, \mathcal{P}_{T_i}(\Delta_i) \rangle + \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \\ &= \|X_i\|_{(i)} - \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} + \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \end{aligned} \quad (39)$$

where we used Holder's inequality for the last inequality (39). Re-arranging, we obtain the first result (36).

For the second inequality, we note that since $\sum_{i=1}^L X_i + \Delta_i + X_{\mathcal{X}} + \Delta_{\mathcal{X}} = Y$, we have

$\Delta_{\mathcal{X}} = -\sum_{i=1}^L \Delta_i$. From the definition of subgradient, we obtain,

$$\begin{aligned} \lambda_{\mathcal{X}} \|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} &\geq \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} + \lambda_{\mathcal{X}} \left\langle \frac{X_{\mathcal{X}}}{\|X_{\mathcal{X}}\|_{\text{fro}}}, \Delta_{\mathcal{X}} \right\rangle \\ &= \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} - \sum_{i=1}^L \lambda_{\mathcal{X}} \left\langle \frac{X_{\mathcal{X}}}{\|X_{\mathcal{X}}\|_{\text{fro}}}, \Delta_i \right\rangle \\ &\geq \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} - \sum_{i=1}^L \lambda_{\mathcal{X}} \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\|X_{\mathcal{X}}\|_{\text{fro}}} \|\Delta_i\|_{(i)} \end{aligned} \quad (40)$$

$$\geq \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} - \sum_{i=1}^L \frac{\lambda_i}{2} \|\Delta_i\|_{(i)} \quad (41)$$

$$\begin{aligned} &\geq \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} - \sum_{i=1}^L \frac{\lambda_i}{2} \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} \\ &\quad - \frac{\lambda_i}{2} \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \end{aligned} \quad (42)$$

where we obtain equation (40) from Holder's inequality, equation (41) from the condition of λ_i (8) and equation (42) from the triangle inequality.

Since $\{X_i + \Delta_i\}_{i=1}^L$ and $X_{\mathcal{X}} + \Delta_{\mathcal{X}}$ achieves the minimum objective function, we have,

$$\begin{aligned} &\lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} + \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} \\ &\geq \lambda_{\mathcal{X}} \|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} + \sum_{i=1}^L \lambda_i \|X_i + \Delta_i\|_{(i)} \end{aligned}$$

Substituting equation (39) and (42), we obtain,

$$\begin{aligned}
& \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} + \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} \\
\geq & \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} - \sum_{i=1}^L \frac{\lambda_i}{2} \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} - \frac{\lambda_i}{2} \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \\
& + \sum_{i=1}^L \lambda_i \|X_i\|_{(i)} - \lambda_i \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} + \lambda_i \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \tag{43}
\end{aligned}$$

Cancelling and re-arranging, we obtain the desired inequality (37),

$$\sum_{i=1}^L \lambda_i \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)} \leq 3 \sum_{i=1}^L \lambda_i \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)}$$

For the third inequality, recall that for any rank- r matrix X , its nuclear norm $\|X\|_{\text{nuc}}$ is upper bounded by $\sqrt{r}\|X\|_{\text{fro}}$. Moreover, the projection of any matrix Y to the column and row space T of a rank r matrix is at most rank- $2r$, that is $\text{rank}(PT(Y)) \leq 2r$. Hence, we obtain,

$$\begin{aligned}
\|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)} &= \sum_{b \in \mathcal{P}_i} \|R_b(\mathcal{P}_{T_i}(\Delta_i))\|_{\text{nuc}} \\
&\leq \sum_{b \in \mathcal{P}_i} \sqrt{2r_b} \|R_b(\Delta_i)\|_{\text{fro}} \\
&\leq \sqrt{\sum_{b \in \mathcal{P}_i} 2r_b} \|\Delta_i\|_{\text{fro}}
\end{aligned}$$

where the last inequality follows from Cauchy-Schwartz inequality and the fact that

$$\sum_{b \in \mathcal{P}_i} \|R_b(\Delta_i)\|_{\text{fro}}^2 = \|\Delta_i\|_{\text{fro}}^2.$$

With these three inequalities, we now proceed to prove Theorem VI.

Proof of Theorem VI

From the optimality of $\{X_i + \Delta_i\}_{i=1}^L$, we have the following inequality,

$$\begin{aligned}
& \lambda_{\mathcal{X}} \|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} + \sum_{i=1}^L \lambda_i \|X_i + \Delta_i\|_{(i)} \\
& \leq \lambda_{\mathcal{X}} \|X_{\mathcal{X}}\|_{\text{fro}} + \sum_{i=1}^L \lambda_i \|X_i\|_{(i)}
\end{aligned}$$

Re-arranging, we obtain,

$$\begin{aligned}
& \|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} - \|X_{\mathcal{X}}\|_{\text{fro}} \\
& \leq \frac{1}{\lambda_{\mathcal{X}}} \sum_{i=1}^L \lambda_i (\|X_i\|_{(i)} - \|X_i + \Delta_i\|_{(i)})
\end{aligned}$$

For convenience, let us define $\Lambda_T = \sum_{i=1}^L \lambda_i \|\mathcal{P}_{T_i}(\Delta_i)\|_{(i)}$ and $\Lambda_{T^\perp} = \sum_{i=1}^L \lambda_i \|\mathcal{P}_{T_i^\perp}(\Delta_i)\|_{(i)}$. Then, from Lemma B.1 equation (36), we obtain,

$$\|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} - \|X_{\mathcal{X}}\|_{\text{fro}} \leq \frac{1}{\lambda_{\mathcal{X}}} (\Lambda_T - \Lambda_{T^\perp}) \quad (44)$$

We would like to keep only $\Delta_{\mathcal{X}}$ on the left hand side and cancel $X_{\mathcal{X}}$. To do this, we multiply both sides of equation (44) with $\|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} + \|X_{\mathcal{X}}\|_{\text{fro}}$. Then, using $(a+b)(a-b) = a^2 - b^2$, we expand the left hand side as:

$$\begin{aligned} \text{L.H.S.} &= \|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}}^2 - \|X_{\mathcal{X}}\|_{\text{fro}}^2 \\ &= \|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 + 2\langle X_{\mathcal{X}}, \Delta_{\mathcal{X}} \rangle \end{aligned}$$

Recall that $\Delta_{\mathcal{X}} = -\sum_{i=1}^L \Delta_i$, we obtain the following lower bound for the left hand side:

$$\begin{aligned} \text{L.H.S.} &= \|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 - 2\langle X_{\mathcal{X}}, \sum_{i=1}^L \Delta_i \rangle \\ &\geq \|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 - 2 \sum_{i=1}^L \|X_{\mathcal{X}}\|_{(i)}^* \|\Delta_i\|_{(i)} \end{aligned} \quad (45)$$

$$\geq \|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 - \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \sum_{i=1}^L \lambda_i \|\Delta_i\|_{(i)} \quad (46)$$

$$\geq \|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 - \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} (\Lambda_T + \Lambda_{T^\perp}) \quad (47)$$

where we used Holder's inequality for equation (45), the condition for λ_j for equation (46), and the triangle inequality for (47).

We now turn to upper bound the right hand side. We know

$\|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} \leq \|X_{\mathcal{X}}\|_{\text{fro}} + (\Lambda_T - \Lambda_{T^\perp})/\lambda_{\mathcal{X}}$ from equation (44). Hence, we obtain,

$$\begin{aligned} \text{R.H.S.} &= (\|X_{\mathcal{X}} + \Delta_{\mathcal{X}}\|_{\text{fro}} + \|X_{\mathcal{X}}\|_{\text{fro}}) \frac{1}{\lambda_{\mathcal{X}}} (\Lambda_T - \Lambda_{T^\perp}) \\ &\leq (2\|X_{\mathcal{X}}\|_{\text{fro}} + \frac{1}{\lambda_{\mathcal{X}}} (\Lambda_T - \Lambda_{T^\perp})) + \frac{1}{\lambda_{\mathcal{X}}} (\Lambda_T - \Lambda_{T^\perp}) \end{aligned}$$

Using Lemma B.1 equation (37), we have,

$$\begin{aligned}
\text{R.H.S.} &\leq 2 \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} (\Lambda_T - \Lambda_{T\perp}) + \frac{1}{\lambda_{\mathcal{X}}^2} (\Lambda_T - \Lambda_{T\perp})^2 \\
&\leq 2 \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} (\Lambda_T - \Lambda_{T\perp}) + 16 \frac{1}{\lambda_{\mathcal{X}}^2} \Lambda_T^2
\end{aligned} \tag{48}$$

Combining and simplifying the lower bound (47) and the upper bound (48), we obtain,

$$\|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 \leq 3 \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \Lambda_T + 16 \frac{1}{\lambda_{\mathcal{X}}^2} \Lambda_T^2 \tag{49}$$

We will now lower bound $\|\Delta_{\mathcal{X}}\|_{\text{fro}}^2 = \|\sum_{i=1}^L \Delta_i\|_{\text{fro}}^2$ by individual terms:

$$\begin{aligned}
\|\sum_{i=1}^L \Delta_i\|_{\text{fro}}^2 &= \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 + \langle \Delta_i, \sum_{j \neq i} \Delta_j \rangle \\
&\geq \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 - \|\Delta_i\|_{(i)} \sum_{j \neq i} \|\Delta_j\|_{(i)}^*
\end{aligned}$$

where we used Holder's inequality for the last inequality.

Now, using the assumption that both $\|X_j\|_{(i)}^*$ and $\|X_j + \Delta_j\|_{(i)}^*$ are bounded by α_{ij} . We have,

$$\begin{aligned}
\|\sum_{i=1}^L \Delta_i\|_{\text{fro}}^2 &\geq \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 - \|\Delta_i\|_{(i)} \sum_{j \neq i} 2\alpha_{ij} \\
&\geq \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 - \lambda_i \|\Delta_i\|_{(i)} \\
&\geq \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 - \Lambda_T - \Lambda_{T\perp}
\end{aligned} \tag{50}$$

$$\geq \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 - 4\Lambda_T \tag{51}$$

where we used the triangle inequality for equation (50) and Lemma B.1 equation (37) for equation (51).

Substituting the lower bound back to equation (49), we have

$$\sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 \leq \left(3 \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} + 4\right) \Lambda_T + 16 \frac{1}{\lambda_{\mathcal{X}}^2} \Lambda_T^2 \quad (52)$$

We now turn to upper bound the equation. From Lemma B.1 equation (38), we know that

$$\Lambda_T \leq \sum_{i=1}^L \lambda_i \sqrt{2 \sum_{b \in \mathcal{P}_i} r_b} \|\Delta_i\|_{\text{fro}}. \text{ Hence, we have,}$$

$$\begin{aligned} 16\Lambda_T^2 &\leq 16 \left(\sum_{i=1}^L \lambda_i \sqrt{2 \sum_{b \in \mathcal{P}_i} r_b} \|\Delta_i\|_{\text{fro}} \right)^2 \\ &\leq 16 \left(\sum_{i=1}^L \lambda_i^2 2 \sum_{b \in \mathcal{P}_i} r_b \right) \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 \end{aligned} \quad (53)$$

$$\leq \frac{1}{2} \lambda_{\mathcal{X}}^2 \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 \quad (54)$$

where we used Cauchy-Schwartz's inequality for equation (53) and the condition for $\lambda_{\mathcal{X}}$ for equation (54). Hence, substituting back to equation (52), rearranging and ignoring constants,

$$\text{we have, } \sum_{i=1}^L \|\Delta_i\|_{\text{fro}}^2 \lesssim \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \sum_{i=1}^L \lambda_i \sqrt{\sum_{b \in \mathcal{P}_i} r_b} \|\Delta_i\|_{\text{fro}}$$

Completing the squares with respect to $\|\Delta_i\|_{\text{fro}}$ gives us,

$$\begin{aligned} &\sum_{i=1}^L \left(\|\Delta_i\|_{\text{fro}} - \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \lambda_i \sqrt{\sum_{b \in \mathcal{P}_i} r_b} \right)^2 \\ &\lesssim \sum_{i=1}^L \left(\frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \lambda_i \sqrt{\sum_{b \in \mathcal{P}_i} r_b} \right)^2 \end{aligned}$$

Using the triangle inequality to lower bound the left hand side, we obtain


$$\begin{aligned} &\sum_{i=1}^L \left| \|\Delta_i\|_{\text{fro}} - \frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \lambda_i \sqrt{\sum_{b \in \mathcal{P}_i} r_b} \right| \\ &\lesssim \sqrt{\sum_{i=1}^L \left(\frac{\|X_{\mathcal{X}}\|_{\text{fro}}}{\lambda_{\mathcal{X}}} \lambda_i \sqrt{\sum_{b \in \mathcal{P}_i} r_b} \right)^2} \end{aligned}$$

Using the fact that ℓ_1 -norm is larger than the ℓ_2 -norm, and re-arranging give us the desired result,

$$\sum_{i=1}^L \|\Delta_i\|_{\text{fro}} \lesssim \frac{\|X_{\mathcal{Z}}\|_{\text{fro}}}{\lambda_{\mathcal{Z}}} \sum_{i=1}^L \lambda_i \sqrt{\sum_{b \in \mathcal{P}_i} r_b}$$

■

Biographies



Frank Ong is currently pursuing the Ph.D. degree in the department of Electrical Engineering and Computer Sciences at the University of California, Berkeley. He received his B.Sc. at the same department in 2013. His research interests lie in signal processing, medical image reconstruction, compressed sensing and low rank methods.



Michael (Miki) Lustig is an Associate Professor in EECS. He joined the faculty of the EECS Department at UC Berkeley in Spring 2010. He received his B.Sc. in Electrical Engineering from the Technion, Israel Institute of Technology in 2002. He received his Msc and Ph.D. in Electrical Engineering from Stanford University in 2004 and 2008, respectively. His research focuses on medical imaging, particularly Magnetic Resonance Imaging (MRI), and very specifically, the application of compressed sensing to rapid and high-resolution MRI, MRI pulse sequence design, medical image reconstruction, inverse problems in medical imaging and sparse signal representation.

References

1. Mallat, SG. A theory for multiresolution signal decomposition: the wavelet representation. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1989. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=192463
2. Candes, E., Demanet, L., Donoho, D., Ying, L. Fast discrete curvelet transforms. Multiscale Modeling & Simulation. 2006. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/05064182X>

3. Simoncelli EP, Freeman WT, Adelson EH, Heeger DJ. Shiftable multiscale transforms. *IEEE Transactions on Information Theory*. Mar; 1992 38(2):587–607. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=119725>.
4. Donoho DL, Johnstone IM. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*. Sep; 2006 81(3):1–31. [Online]. Available: <http://www.jstor.org/stable/2337118?origin=crossref>.
5. Donoho DL. Compressed sensing. *IEEE Transactions on Information Theory*. Apr; 2006 52(4):1289–1306. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1614066>.
6. Candes EJ, Romberg J, Tao T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*. 2006; 52(2):489–509. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1580791>.
7. Donoho DL, Huo X. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*. Nov; 2001 47(7):2845–2862. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=959265>.
8. Starck JL, Moudden Y, Bobin J, Elad M, Donoho DL. Morphological component analysis. *Optics & Photonics 2005*. Aug.2005 5914:59 140Q-59–140Q-15. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=870615>.
9. Chen SS, Donoho DL, Saunders MA. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*. 1998; 20(1):33–61. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/S1064827596304010>.
10. Liang, ZP. Spatiotemporal imaging with partially separable functions; The Proceedings of the International Symposium of Biomedical Imaging: From Nano to Macro IEEE. 2007. p. 988-991. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4193454>
11. Basri R, Jacobs DW. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Feb; 2003 25(2):218–233. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1177153>.
12. Goldberg K, Roeder T, Gupta D, Perkins C. Eigentaste: A Constant Time Collaborative Filtering Algorithm. *Information Retrieval*. 2001; 4(2):133–151. [Online]. Available: <http://link.springer.com/10.1023/A:1011419012209>.
13. Fazel, M. Ph D dissertation. Stanford University; Mar. 2002 Matrix rank minimization with applications. [Online]. Available: <https://faculty.washington.edu/mfazel/thesis-final.pdf>
14. Candès EJ, Recht B. Exact Matrix Completion via Convex Optimization. *Foundations of Computational Mathematics*. Dec; 2009 9(6):717–772. [Online]. Available: <http://link.springer.com/10.1007/s10208-009-9045-5>.
15. 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*. Jul; 2013 34(3):946–977. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/110853996>.
16. Candès EJ, Eldar YC, Strohmer T, Voroninski V. Phase Retrieval via Matrix Completion. *SIAM Journal on Imaging Sciences*. Feb; 2013 6(1):199–225. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/110848074>.
17. Chandrasekaran V, Sanghavi S, Parrilo PA. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*. 2011; 21(2):572–596. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/090761793>.
18. Agarwal A, Negahban S, Wainwright MJ. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics*. Apr; 2012 40(2):1171–1197. [Online]. Available: <http://projecteuclid.org/euclid.aos/1342625465>.
19. Coifman, RR., Donoho, DL. Wavelets and Statistics. New York, NY: Springer New York; 1995. Translation-Invariant De-Noising; p. 125-150.[Online]. Available: http://link.springer.com/10.1007/978-1-4612-2544-7_9
20. Recht B, Xu W, Hassibi B. Null space conditions and thresholds for rank minimization. *Mathematical Programming*. Oct; 2010 127(1):175–202. [Online]. Available: <http://link.springer.com/10.1007/s10107-010-0422-2>.

21. Xu H, Caramanis C, Sanghavi S. Robust PCA via outlier pursuit. *Advances in Neural Information Processing Systems*. 2010; 23:2496–2504. [Online]. Available: <http://papers.nips.cc/paper/4005-robust-pca-via-outlier-pursuit>.
22. Candes EJ, Tao T. The Power of Convex Relaxation: Near-Optimal Matrix Completion. *IEEE Transactions on Information Theory*. May; 2010 56(5):2053–2080. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5452187>.
23. Recht, B., Fazel, M., Parrilo, PA. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*. 2010. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/070697835>
24. Hsu D, Kakade SM, Zhang T. Robust Matrix Decomposition With Sparse Corruptions. *IEEE Transactions on Information Theory*. 2011; 57(11):7221–7234. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5934412>.
25. Candès EJ, Li X, Ma Y, Wright J. Robust principal component analysis? *Journal of the ACM*. May; 2011 58(3):1–37. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1970392.1970395>.
26. McCoy, MB., Tropp, JA. The achievable performance of convex demixing. arXiv preprint arXiv: 1309.7478v1. Sep. 2013 [Online]. Available: <http://arxiv.org/abs/1309.7478v1>
27. McCoy MB, Tropp JA. Sharp recovery bounds for convex demixing, with applications. *Foundations of Computational Mathematics*. Jan; 2014 14(3):1–51. [Online]. Available: <http://link.springer.com/10.1007/s10208-014-9191-2>.
28. McCoy MB, Cevher V, Dinh QT, Asaei A. Convexity in source separation: Models, geometry, and algorithms. *IEEE Signal Processing Magazine*. 2014; 31(3):87–95. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6784106.
29. Bakshi BR. Multiscale PCA with application to multivariate statistical process monitoring. *AICChE Journal*. Jul; 1998 44(7):1596–1610. [Online]. Available: <http://doi.wiley.com/10.1002/aic.690440712>.
30. Kondor, R., Teneva, N., Garg, V. Multiresolution Matrix Factorization; in *The 31st International Conference on Machine Learning*. 2014. p. 1620-1628.[Online]. Available: <http://jmlr.org/proceedings/papers/v32/kondor14.html>
31. Kakarala, R., Ogunbona, PO. Signal analysis using a multiresolution form of the singular value decomposition. *Image Processing*. 2001. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=918566
32. Vozalis M, Margaritis K. Using SVD and demographic data for the enhancement of generalized Collaborative Filtering. *Information Sciences*. Aug; 2007 177(15):3017–3037. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0020025507001223>.
33. Wright, J., Ganesh, A., Min, K., Ma, Y. Compressive principal component pursuit. *Information and Inference: A Journal of the IMA*. 2013. [Online]. Available: <http://imaia.oxfordjournals.org/content/2/1/32.short>
34. Foygel R, Mackey L. Corrupted sensing: Novel guarantees for separating structured signals. *IEEE Transaction on Information Theory*. 2014; 60(2):1223–1247. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6712045.
35. Bandeira, AS., van Handel, R. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. arXiv preprint arXiv:1408.6185v3. Aug. 2014 [Online]. Available: <http://arxiv.org/abs/1408.6185v3>
36. Belloni, A., Chernozhukov, V., Wang, L. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*. 2011. [Online]. Available: <http://biomet.oxfordjournals.org/content/98/4/791.short>
37. Watson, GA. Characterization of the subdifferential of some matrix norms. *Linear Algebra and Its Applications*. 1992. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0024379592904072>
38. Boyd S. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends® in Machine Learning*. 2010; 3(1):1–122. [Online]. Available: https://web.stanford.edu/~boyd/papers/pdf/admm_distr_stats.pdf.
39. Parikh, N., Boyd, S. Proximal algorithms. *Foundations and Trends in optimization*. 2013. [Online]. Available: http://web.stanford.edu/~boyd/papers/pdf/prox_algs.pdf

40. Horn, RA., Johnson, CR. Matrix Analysis. Cambridge University Press; Oct. 2012 [Online]. Available: <https://books.google.com/books?id=5I5AYeeh0JUC>
41. Figueiredo MAT, Nowak RD. An EM algorithm for wavelet-based image restoration. IEEE Transactions on Image Processing. Aug; 2003 12(8):906–916. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1217267>. [PubMed: 18237964]
42. Georghiades AS, Belhumeur PN, Kriegman DJ. From few to many: illumination cone models for face recognition under variable lighting and pose. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2001; 23(6):643–660. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=927464>.
43. Li L, Huang W, Gu IYH, Tian Q. Statistical Modeling of Complex Backgrounds for Foreground Object Detection. IEEE Transactions on Image Processing. Nov; 2004 13(11):1459–1472. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1344037>. [PubMed: 15540455]
44. 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. [Online]. Available: <http://doi.wiley.com/10.1002/mrm.21391>. [PubMed: 17969013]
45. Zhang T, Cheng JY, Potnick AG, Barth RA, Alley MT, Uecker M, Lustig M, Pauly JM, Vasanawala SS. Fast pediatric 3D free-breathing abdominal dynamic contrast enhanced MRI with high spatiotemporal resolution. Journal of Magnetic Resonance Imaging. Feb; 2015 41(2):460–473. [Online]. Available: <http://doi.wiley.com/10.1002/jmri.24551>. [PubMed: 24375859]
46. Otazo R, Candès E, Sodickson DK. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. Magnetic Resonance in Medicine. Mar; 2015 73(3):1125–1136. [Online]. Available: <http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=24760724&retmode=ref&cmd=prlinks>. [PubMed: 24760724]

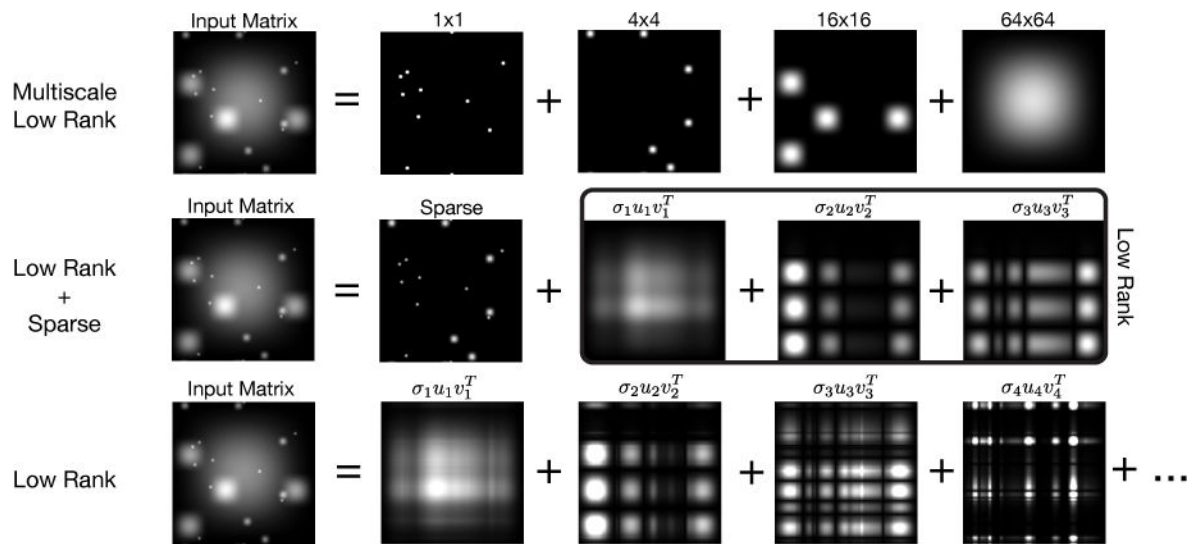


Fig. 1.

An example of our proposed multi-scale low rank decomposition compared with other low rank methods. Each blob in the input matrix is a rank-1 matrix constructed from an outer product of hanning windows. Only the multi-scale low rank decomposition exactly separates the blobs to their corresponding scales and represents each blob as compactly as possible.

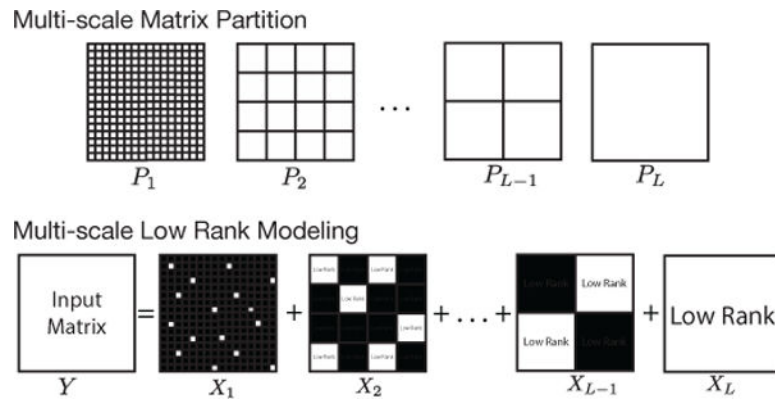


Fig. 2. Illustration of a multi-scale matrix partition and its associated multi-scale low rank modeling. Since the zero matrix is a matrix with the least rank, our multi-scale modeling naturally extends to sparse matrices as 1×1 low rank matrices.




Fig. 3. Illustration of the block reshape operator R_b . R_b extracts block b from the full matrix and reshapes it into an $m_i \times n_i$ matrix. Its adjoint operator R_b^\top takes an $m_i \times n_i$ matrix and embeds it into a full-size zero matrix.

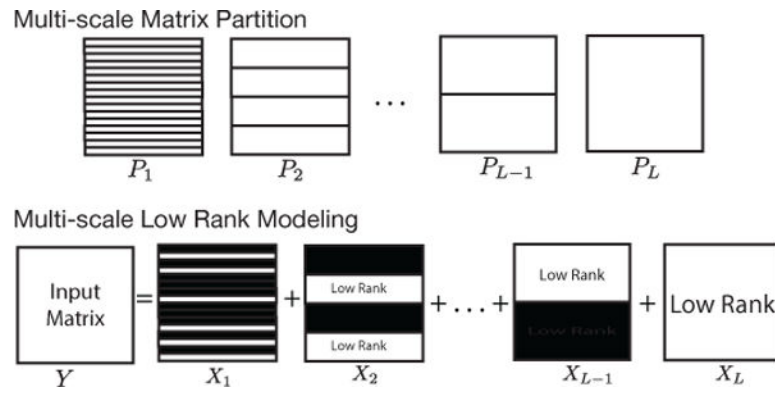


Fig. 4. Illustration of another multi-scale matrix partition and its associated multi-scale low rank modeling. Here, only the vertical dimension of the matrix is decimated. Since a $1 \times N$ matrix is low rank if and only if it is zero, our multi-scale modeling naturally extends to group sparse matrices.




Fig. 5. An example of the multi-scale low rank decomposition in the presence of additive Gaussian noise by solving the convex program (2).




Fig. 6. A conceptual illustration of how to obtain a multi-scale low rank decomposition. First, we extract each block from the input matrix and perform a thresholding operation on its singular value to recover the significant components. Then, we subtract these significant components from our input matrix, thereby enabling the recovery of weaker, previously submerged components.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Fig. 7.

An example of the multi-scale low rank decomposition with and without random cycle spinning. Each blob in the input matrix Y is a rank-1 matrix constructed from an outer product of hanning windows and is placed at random positions. Blocking artifacts can be seen in the decomposition without random cycle spinning while vastly diminished in the random cycle spinned decomposition.

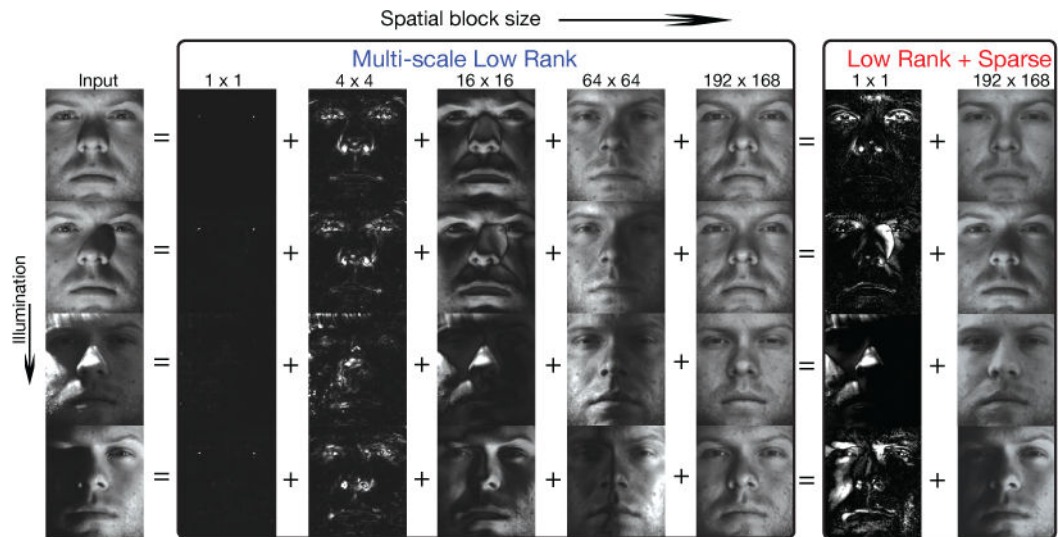


Fig. 8. Multi-scale low rank versus low rank + sparse on faces with uneven illumination. Multi-scale low rank decomposition recovers almost shadow-free faces, whereas low rank + sparse decomposition can only remove some shadows.

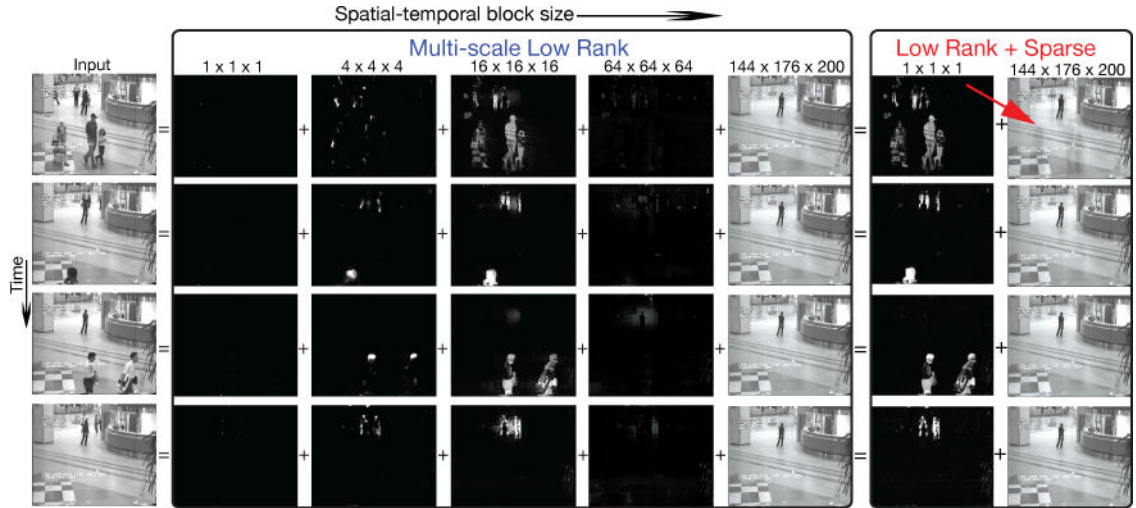


Fig. 9. Multi-scale low rank versus low rank + sparse decomposition on a surveillance video. For the multi-scale low rank, body motion is mostly captured in the $16 \times 16 \times 16$ scale while fine-scale motion is captured in $4 \times 4 \times 4$ scale. Background video component is captured in the globally low rank component and is almost artifact-free. Low rank + sparse decomposition exhibits ghosting artifacts as pointed by the red arrow because they are neither globally low rank or sparse.

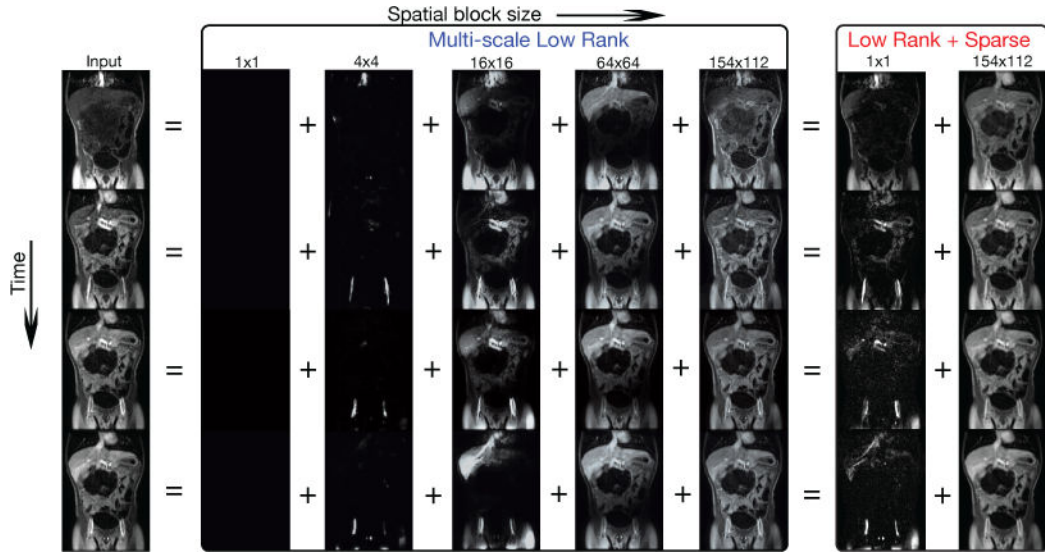


Fig. 10. Multi-scale low rank versus low rank + sparse decomposition on a dynamic contrast enhanced magnetic resonance image series. For the multi-scale result, small contrast dynamics in vessels are captured in 4×4 blocks while contrast dynamics in the liver are captured in 16×16 blocks. The biggest block size captures the static tissues and interestingly the respiratory motion. In contrast, the low rank + sparse modeling could only provide a coarse separation of dynamics and static tissue, which result in neither truly sparse nor truly low rank components.

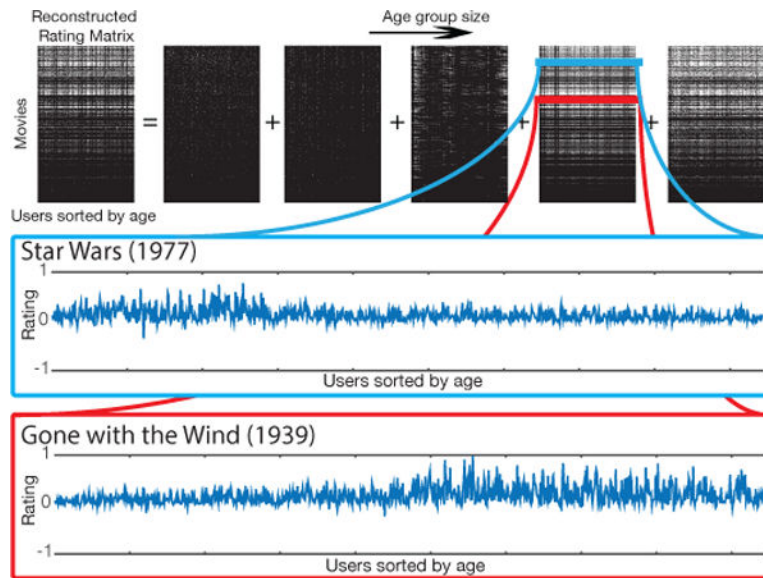


Fig. 11. Multi-scale low rank reconstructed matrix of the 100K MovieLens dataset. The extracted signal scale component captures the tendency that younger users rated Star Wars higher whereas the more senior users rated Gone with the Wind higher.