

An Algorithm for Parallel Reconstruction of Jointly Sparse Tensors with Applications to Hyperspectral Imaging*

Qun Li

Microsoft Corporation
11025 NE 8th St., Bellevue, WA 98004
qul@microsoft.com

Edgar A. Bernal

United Technologies Research Center
411 Silver Ln., East Hartford, CT 06118
bernalea@utrc.utc.com

Abstract

A wide range of Compressive Sensing (CS) frameworks have been proposed to address the task of color and hyperspectral image sampling and reconstruction. Methods for reconstruction of jointly sparse vectors that leverage joint sparsity constraints such as the Multiple Measurement Vector (MMV) approach have been shown to outperform Single Measurement Vector (SMV) frameworks which don't exploit such constraints. Recent work has shown that exploiting joint sparsity while simultaneously preserving the high-dimensional structure of the data results in further performance improvements. In this paper, we introduce a parallelizable extension of a previously proposed serial tensorial MMV approach which, like its predecessor, exploits joint sparsity constraints multiple data dimensions simultaneously, but that, in contrast, is parallelizable in nature. We demonstrate empirically that the proposed method provides better reconstruction fidelity of hyperspectral imagery given a fixed number of measurements, and that it is also more computationally efficient than the current state of the art.

1. Introduction

The acquisition and storage of hyperspectral imagery impose stringent hardware constraints which result in increased cost and complexity of hyperspectral imagers. Hyperspectral imaging is the process of using specialized sensors to collect image information across the electromagnetic spectrum, often beyond the visible electromagnetic wavelength range. A hyperspectral image can be represented as a three-dimensional data cube where the first and second dimensions correspond to spatial data and the third dimension corresponds to the spectral bands. Objects have their own respective fingerprints known as spectral signatures; consequently, there is a wide range of applications, e.g., remote

sensing, astronomy, mineralogy, agriculture, healthcare and surveillance, that rely on decomposing a two-dimensional image of a scene into its spectral bands in order to enable object identification within the scene. The main disadvantages of hyperspectral imaging are related to the intrinsic high dimensionality of the data, which imposes storage, computational, and sensitivity constraints. These factors drive up the complexity and cost of traditional hyperspectral imagers. Fortunately, there is significant redundancy in hyperspectral images along both the spatial and the spectral dimensions which can be exploited by judicious sampling and reconstruction techniques, which in turn gives rise to sensing equipment with reduced complexity and cost.

Compressive Sensing (CS) techniques have been proposed which exploit redundancies in hyperspectral data in order to improve the efficiency of the sampling and reconstruction processes. CS techniques aim at reconstructing sparse signals from a small number of linear measurements, relative to the original signal dimensionality. CS is often posed as an optimization problem whose objective consists in finding the solution of an underdetermined system of equations under a set of sparsity constraints. A vector $\mathbf{x} \in \mathbb{R}^N$ is called s -sparse if it has at most s non-zero entries. In mathematical terms, the linear sampling process of \mathbf{x} results in a set of samples $\mathbf{y} \in \mathbb{R}^m$ obtained by computing $\mathbf{y} = \mathbf{A}\mathbf{x}$. In this expression, $\mathbf{A} \in \mathbb{R}^{m \times N}$, where $m < N$, denotes the sampling matrix. The reconstruction process attempts to recover \mathbf{x} by finding a $\hat{\mathbf{z}} \in \mathbb{R}^N$ that satisfies

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \|\mathbf{z}\|_0 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{A}\mathbf{z}. \quad (1)$$

These sampling and recovery processes constitute what is traditionally termed the Single Measurement Vector (SMV) framework. The ℓ_0 minimization problem in Eq. 1 is nonpolynomial-deterministic- (NP-) hard, and thus, computationally intractable, so it is often approximated by optimizing its ℓ_1 equivalent or by implementing greedy algorithms.

When the signal to be reconstructed consists of a set of vectors with common support, simultaneous sampling of

*This work was done while the authors were with PARC, A Xerox Company

the vectors can lead to improved reconstruction. This is the premise of the Multiple Measurement Vector (MMV) approach [5]. Specifically, assume the signal to be reconstructed consists of a set of q vectors $\mathbf{x}_j \in \mathbb{R}^N$ with common support, that is, $\mathbf{X} = [\mathbf{x}_1 | \mathbf{x}_2 | \dots | \mathbf{x}_q] \in \mathbb{R}^{N \times q}$. Simultaneous sampling results in a set of q vectorial measurements that can be represented in the form of a matrix $\mathbf{Y} \in \mathbb{R}^{m \times q}$, where $m < N$, obtained by computing $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{Y} = [\mathbf{y}_1 | \mathbf{y}_2 | \dots | \mathbf{y}_q]$. The optimization task in the MMV recovery problem is often formulated as the following optimization task:

$$\hat{\mathbf{Z}} = \arg \min_{\mathbf{Z}} \|\mathbf{Z}\|_{2,0} \quad \text{s.t.} \quad \mathbf{Y} = \mathbf{A}\mathbf{Z}, \quad (2)$$

where $\|\mathbf{Z}\|_{2,0}$ denotes the ℓ_0 norm of the vector whose elements are ℓ_2 norms of the rows of matrix \mathbf{Z} [5].

1.1. Related Work

Approximations to the NP-hard problem from Eq. 2 have been formulated. For example, convex-relaxation-based [5, 6, 17, 16] and approximation [15] alternatives have been proposed. The optimization-based approach from [15] addresses two problems simultaneously, namely providing an approximate solution to Eq. 2 and applying the MMV recovery framework to the task of color image reconstruction. This task, along with that of hyperspectral image sampling and reconstruction (HSI), had previously been formulated as a group-sparse recovery problem [14, 11]. The direct extension of [15] to HSI is trivial. Alternative solutions to Eq. 2 have also been proposed which enforce stricter constraints on the MMV problem. Specifically, the work in [10, 4] imposes a low-rank constraint on the measurement matrix in addition to the traditional joint sparsity constraints.

Traditional CS theory is well-suited for sampling and reconstruction of one-dimensional signals, and naïve extensions of the CS framework to multidimensional problems typically rely on vectorial representations of the data, which result in increased computational and memory requirements. Additionally, methodologies that vectorize high-dimensional data fail to preserve the inherent tensorial nature of the data, which results in failure to exploit correlations across the different dimensionalities. Examples of high-dimensional data that are not appropriately represented by vectors, and thus not adequately recovered by vectorial methods, include color and hyperspectral images, as well as video streams. Extensions of the CS theory to multi-dimensional data with applications in HSI have recently been addressed [2, 9, 7]. Most existing tensorial CS approaches, however, still fail to leverage joint sparsity constraints.

To our knowledge, until the method in [3] was introduced, no existing algorithm was capable of exploiting joint sparsity across different signal elements, while at the same time preserving the dimensionality of the data. In this paper, we introduce a parallelizable variant of the algorithm

from [3] that surpasses it both in terms of reconstruction performance and computational efficiency.

1.2. Contributions

We propose a parallelizable version of the approach proposed in [3] which is, to our knowledge, the state-of-the-art with regards to MMV recovery of a multidimensional signal with jointly sparse components. The method in [3] is tensorial in nature and, as such, achieves better performance than competing methods because it is able to preserve the multidimensional nature of the data being reconstructed. One of its limitations lies in its sequential multi-stage implementation, given that it requires one reconstruction stage for each of the tensorial modes. This *modus operandi* results in decreased fidelity since the reconstruction error from early stages propagates into subsequent stages. In contrast, the algorithm proposed in this paper decouples the reconstruction stages for the multiple tensor modes, which results in improved reconstruction quality and reduced computational complexity: since the reconstruction stages are effectively independent, they can be implemented in a parallel manner. The proposed method, like its predecessor, is capable of leveraging the tensorial nature of the data in the sampling and reconstruction stages; its parallel nature, on the other hand, enables it to provide improved reconstruction performance and achieve increased computational efficiency.

2. Notation and Terminology

As per the convention used in [13][2], we use *Italic* characters to represent scalar values (*e.g.*, a, B), **bold-face** characters to represent vectors (*e.g.*, \mathbf{a}, \mathbf{b}), capital *italic bold-face* characters to represent matrices (*e.g.*, \mathbf{A}, \mathbf{B}) and capital *calligraphic* characters to represent tensors (*e.g.*, \mathcal{A}, \mathcal{B}). A tensor is a multidimensional array which may be used to represent high-dimensional signals such as color and hyperspectral images and videos. Tensor $\mathcal{X} \in \mathbb{R}^{N_1 \times \dots \times N_d}$ has order d (or is d -dimensional), and the dimensionality of its i^{th} mode (also denoted mode i) is N_i . In short, the order or number of dimensions of a tensor corresponds to its number of modes, and each mode has its own associated dimensionality. The remainder of this section provides a brief overview of relevant multilinear algebra concepts.

Kronecker product The Kronecker product of matrices $\mathbf{A} \in \mathbb{R}^{I \times J}$ and $\mathbf{B} \in \mathbb{R}^{K \times L}$, denoted by $\mathbf{A} \otimes \mathbf{B}$, is a matrix of size $(I \cdot K) \times (J \cdot L)$ which is computed as $\mathbf{A} \otimes \mathbf{B} = \begin{pmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \dots & a_{1J}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \dots & a_{2J}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{I1}\mathbf{B} & a_{I2}\mathbf{B} & \dots & a_{IJ}\mathbf{B} \end{pmatrix}$.

Mode- i fiber and mode- i unfolding The mode- i fibers of tensor $\mathcal{X} = [x_{\alpha_1, \dots, \alpha_i, \dots, \alpha_d}] \in \mathbb{R}^{N_1 \times \dots \times N_i \times \dots \times N_d}$ are obtained

by fixing every index but α_i . The mode- i unfolding of \mathcal{X} , denoted $\mathbf{X}_{(i)}$, is a matrix of size $N_i \times (N_1 \cdots N_{i-1} \cdot N_{i+1} \cdots N_d)$ whose columns are the mode- i fibers of \mathcal{X} . We use $R_i(\mathcal{X})$ to denote the space formed by the mode- i fibers of \mathcal{X} .

Mode- i product The mode- i product between tensor $\mathcal{X} = [x_{\alpha_1, \dots, \alpha_i, \dots, \alpha_d}] \in \mathbb{R}^{N_1 \times \dots \times N_i \times \dots \times N_d}$ and matrix $\mathbf{U} = [u_{j, \alpha_i}] \in \mathbb{R}^{J \times N_i}$ is a tensor $\mathcal{X} \times_i \mathbf{U} \in \mathbb{R}^{N_1 \times \dots \times N_{i-1} \times J \times N_{i+1} \times \dots \times N_d}$, and is defined element-wise as

$$(\mathcal{X} \times_i \mathbf{U})_{\alpha_1, \dots, \alpha_{i-1}, j, \alpha_{i+1}, \dots, \alpha_d} = \sum_{\alpha_i=1}^{N_i} x_{\alpha_1, \dots, \alpha_i, \dots, \alpha_d} u_{j, \alpha_i}.$$

3. Tensorial compressive sensing of hyperspectral images

In the context of this paper, a hyperspectral image is represented as a third-order tensorial data cube, the first two tensorial modes of which correspond to the spatial dimensions, and the third mode of which is associated with the spectral dimension of the image. More specifically, the sampled hyperspectral cube representing the image is a tensor of order three $\mathcal{X} \in \mathbb{R}^{N_r \times N_c \times N_b}$, where N_r and N_c are the number of rows and columns of the image, respectively, and N_b is the number of spectral bands. As in most CS frameworks, the proposed algorithm comprises a sampling and a reconstruction stage. The *sampling stage* is identical to that of [3], and involves the acquisition of a set of linear measurements from the hyperspectral image of a scene of interest. In the *reconstruction stage*, the hyperspectral representation of the scene is estimated by *independently* recovering each tensorial mode from the measurements acquired in the sampling stage. As in [3], we pose the image reconstruction problem in the context of a tensorial MMV framework; in contrast, the reconstruction stage of the proposed algorithm is parallelizable, and differs from that in [3] which is sequential in nature since it requires the reconstruction of certain modes to be completed before the reconstruction of other modes can take place.

3.1. Sampling stage

As in [3], the sampling of \mathcal{X} involves the computation of a set of mode- i products between \mathcal{X} and sampling matrices \mathbf{U}_i , for $i = 1, 2$, one for each spatial tensor mode. The entries of the sampling matrices are randomly generated. The outcome of the sampling stage is a tensor \mathcal{Y} that can be computed as

$$\mathcal{Y} = \mathcal{X} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2, \quad (3)$$

where $\mathbf{U}_1 \in \mathbb{R}^{m_r \times N_r}$ and $\mathbf{U}_2 \in \mathbb{R}^{m_c \times N_c}$, where m_r and m_c are the number of measurements along the mode 1 and mode 2 (e.g., columns and rows) respectively, and $\mathcal{Y} \in \mathbb{R}^{m_r \times m_c \times N_b}$.

It should be appreciated that the outcome of the sampling stage is still a tensor of order 3, in contrast with vectorial methods (see, e.g., [15]) where the dimensions of the measurement matrix are effectively $(m_r \cdot m_c) \times N_b$.

3.2. Reconstruction stage

The reconstruction stage obtains $\hat{\mathcal{X}}$, an estimate of the hyperspectral image cube \mathcal{X} , by solving a set of ℓ_1 -minimization tasks that rely on knowledge of the set of samples \mathcal{Y} and the set of sampling matrices. Recovering the original data $\mathcal{X} \in \mathbb{R}^{N_r \times N_c \times N_b}$ from samples $\mathcal{Y} \in \mathbb{R}^{m_r \times m_c \times N_b}$ involves reconstruction of the modes that are compressed in the sampling stage. We introduce two alternative approaches to the reconstruction stage, which can be seen as parallelizable MMV extensions of the sequential tensorial CS recovery stages proposed in [3]. The following discussion applies equally to both approaches and is equivalent to the parallelizable tensorial reconstruction framework first proposed in [9]. Fig. 1 visually illustrates the steps involved in the reconstruction stage.

Let matrix $\mathbf{A}_k \in \mathbb{R}^{m_r \times m_c}$ denote the k -th slice of \mathcal{Y} corresponding to the k -th spectral band, for $k = 1, 2, \dots, N_b$ (see ‘Slicing’ stage in Fig. 1.) Decompose \mathbf{A}_k such that

$$\mathbf{A}_k = \sum_{i=1}^{r_k} \mathbf{b}_{ik}^{(r)} \mathbf{b}_{ik}^{(c)\top}, \quad (4)$$

where r_k is the rank of \mathbf{A}_k , $\mathbf{b}_{ik}^{(r)} \in R_1(\mathbf{A}_k) \subseteq \mathbf{U}_1 R_1(\mathcal{X}) \subseteq \mathbb{R}^{m_r}$, and $\mathbf{b}_{ik}^{(c)} \in R_2(\mathbf{A}_k) \subseteq \mathbf{U}_2 R_2(\mathcal{X}) \subseteq \mathbb{R}^{m_c}$ (see ‘Decomposition’ stage in Fig. 1.) Here $R_1(\mathbf{A}_k)$ and $R_2(\mathbf{A}_k)$ denote the row and column spaces of \mathbf{A}_k respectively, and $R_1(\mathcal{X})$ and $R_2(\mathcal{X})$ denote the spaces formed by mode-1 and mode-2 fibers of \mathcal{X} respectively. The decomposition from Eq. 4 can be achieved via techniques such as matrix singular value decomposition (SVD) or QR decomposition [1]. Theoretically, all *exact* matrix decomposition methods will yield identical reconstruction results. This rank decomposition of \mathbf{A}_k is expressed as the multiplication of two matrices, as illustrated in Fig. 1. These matrices are denoted as left and right matrices, and their columns consist of $\mathbf{b}_{ik}^{(r)}$ and $\mathbf{b}_{ik}^{(c)}$ for $i = 1, \dots, r_k$, respectively. In Fig. 1, single columns with index i , namely $\mathbf{b}_{ik}^{(r)}$ in the left matrix and $\mathbf{b}_{ik}^{(c)}$ in the right matrix, are highlighted in red.

Let $\mathbf{f}_{ik}^{(r)\star} \in R_1(\mathcal{X}) \subset \mathbb{R}^{N_r}$ be a solution to

$$\mathbf{f}_{ik}^{(r)\star} = \arg \min \|\mathbf{f}_{ik}^{(r)}\|_1, \quad \text{s.t.} \quad \mathbf{U}_1 \mathbf{f}_{ik}^{(r)} = \mathbf{b}_{ik}^{(r)}, \quad \text{for } i = 1, \dots, r_k, \quad (5)$$

and let $\mathbf{f}_{ik}^{(c)\star} \in R_2(\mathcal{X}) \subset \mathbb{R}^{N_c}$ be a solution to

$$\mathbf{f}_{ik}^{(c)\star} = \arg \min \|\mathbf{f}_{ik}^{(c)}\|_1, \quad \text{s.t.} \quad \mathbf{U}_2 \mathbf{f}_{ik}^{(c)} = \mathbf{b}_{ik}^{(c)}, \quad \text{for } i = 1, \dots, r_k. \quad (6)$$




Figure 1. Visualization of the reconstruction stage. See Sec. 3.2 for a detailed explanation of the figure.

A matrix $\mathbf{G}_k \in \mathbb{R}^{N_r \times N_c}$ is then formed by

$$\mathbf{G}_k = \sum_{i=1}^{r_k} \mathbf{f}_{ik}^{(r)*} \mathbf{f}_{ik}^{(c)*\top}, \quad k = 1, \dots, N_b. \quad (7)$$

Then $\hat{\mathcal{X}} \in \mathbb{R}^{N_r \times N_c \times N_b}$ is the tensor whose k -th slice corresponding to the k -th spectral band is \mathbf{G}_k , for $k = 1, 2, \dots, N_b$. The arrow labeled ‘Reconstruction’ in Fig. 1 involves solving Eq. 5 and Eq. 6 and recovering the previously sampled/compressed dimensions from the left and right matrices of \mathbf{A}_k . Single columns corresponding to index i , namely $\mathbf{f}_{ik}^{(r)*}$ in the matrix recovered from the left matrix and $\mathbf{f}_{ik}^{(c)*}$ in the matrix recovered from the right matrix, are highlighted in green. Multiplication of the two reconstructed matrices results in matrix \mathbf{G}_k (see ‘Composition’ stage in Fig. 1,) which corresponds to the k -th spectral band of the final reconstructed tensor $\hat{\mathcal{X}}$, obtained by stacking \mathbf{G}_k (see ‘Assembly’ stage in Fig. 1.) Theorem 3.3 in [8] guarantees the uniqueness of the recovered $\hat{\mathcal{X}}$. The fundamental difference between the reconstruction methods proposed in this paper and those first introduced in [9] lies in the way the optimization tasks from Eqs. 5 and 6 are solved: the algorithms introduced herein exploit the joint sparsity between the different components in the hyperspectral data, while the algorithm from [9] reconstructs each vector independently. The main difference between the reconstruction methods introduced herein and those proposed in [3], which also exploit the joint sparsity across components in the data, lies in the parallelizability of the newly introduced algorithms, which is not present in the algorithm from [3].

3.2.1 Joint recovery across the spectral dimension

The first alternative to performing the reconstruction step from Eqs. 5 and 6 exploits the correlation of the data across the spectral dimension. To that end, the solution to the set of ℓ_1 -minimization tasks can be framed as an MMV joint

recovery problem. For simplicity, consider first the case where $r_1 = r_2 = \dots = r_k$ for all k , namely, each matrix \mathbf{A}_k in Eq. 4 has the same rank; the more general scenario will be discussed later. Instead of solving the $2 \sum_{k=1}^{N_b} r_k$ minimization tasks from Eq. 5 and Eq. 6 individually (as is done in [9]), which would fail to leverage the joint sparsity property across the spectral dimension, r_k jointly sparse matrices of dimensions $N_r \times N_b$ can be recovered from r_k measurement matrices of dimensions $m_r \times N_b$, and r_k jointly sparse matrices of dimensions $N_c \times N_b$ can be recovered from r_k measurement matrices of dimensions $m_c \times N_b$. Any recovery method for jointly sparse matrices can be employed to achieve this goal. In this paper, we pose this problem as an optimization task which is effectively the tensorial extension of the method in [5], or, alternatively, the parallelizable version of the method in [3], namely:

$$\begin{aligned} \mathbf{F}_i^{(r)*} &= [\mathbf{f}_{i1}^{(r)}, \mathbf{f}_{i2}^{(r)}, \dots, \mathbf{f}_{iN_b}^{(r)}] = \arg \min \|\mathbf{F}_i^{(r)}\|_{2,0} \quad \text{s.t.} \\ \mathbf{B}_i^{(r)} &= [\mathbf{b}_{i1}^{(r)}, \mathbf{b}_{i2}^{(r)}, \dots, \mathbf{b}_{iN_b}^{(r)}] = \mathbf{U}_1 \mathbf{F}_i^{(r)}, \text{ for } i = 1, 2, \dots, r_k. \end{aligned} \quad (8)$$

$$\begin{aligned} \mathbf{F}_i^{(c)*} &= [\mathbf{f}_{i1}^{(c)}, \mathbf{f}_{i2}^{(c)}, \dots, \mathbf{f}_{iN_b}^{(c)}] = \arg \min \|\mathbf{F}_i^{(c)}\|_{2,0} \quad \text{s.t.} \\ \mathbf{B}_i^{(c)} &= [\mathbf{b}_{i1}^{(c)}, \mathbf{b}_{i2}^{(c)}, \dots, \mathbf{b}_{iN_b}^{(c)}] = \mathbf{U}_2 \mathbf{F}_i^{(c)}, \text{ for } i = 1, 2, \dots, r_k. \end{aligned} \quad (9)$$

where each column $\mathbf{f}_{ik}^{(r)}$ of $\mathbf{F}_i^{(r)*}$ and $\mathbf{f}_{ik}^{(c)}$ of $\mathbf{F}_i^{(c)*}$ correspond to the i -th largest singular value of \mathbf{A}_k for $k = 1, 2, \dots, N_b$. Fig. 2 visually illustrates the formation of the MMV matrices in the reconstruction stage just described.

Each matrix $\mathbf{G}_k \in \mathbb{R}^{N_r \times N_c}$ can then be obtained by first multiplying the k -th column of $\mathbf{F}_i^{(r)*}$ and the transpose of the k -th column of $\mathbf{F}_i^{(c)*}$ and then summing up the result over all r_k such multiplications. More specifically, \mathbf{G}_k can be computed as follows:




Figure 2. Formation of the MMV measurement matrices in the process of joint recovery across the spectral dimension. The i -th column of each left matrix (highlighted in red) forms the columns of matrix $\mathbf{B}_i^{(r)}$, and the i -th column of each right matrix (highlighted in purple) forms the columns of matrix $\mathbf{B}_i^{(c)}$, for $i = 1, 2, \dots, r_k$. These matrices correspond to the left and right matrices obtained by decomposing \mathbf{A}_k as illustrated in Fig. 1 and as described in Eq. 4.

$$\mathbf{G}_k = \sum_{i=1}^{r_k} \mathbf{f}_{ik}^{(r)} \mathbf{f}_{ik}^{(c)} \quad \text{for } k = 1, 2, \dots, N_b. \quad (10)$$

Then $\hat{\mathcal{X}} = \mathcal{G} \in \mathbb{R}^{N_r \times N_c \times N_b}$ can be obtained by stacking all \mathbf{G}_k for $k = 1, \dots, N_b$ along the spectral dimension. The parallelizable nature of this approach stems from the fact that Eqs. 8 and 9 can be solved independently.

In the more general case where r_k is not the same for all k , at most $r_{\tilde{k}}$, where $\tilde{k} = \arg \max_k r_k$ for $k = 1, 2, \dots, N_b$ jointly sparse matrices of dimensions $N_r \times N_i$, where $N_i \leq N_b$, can be recovered from $r_{\tilde{k}}$ measurement matrices of dimensions $m_r \times N_i$, and $r_{\tilde{k}}$ jointly sparse matrices of dimensions $N_c \times N_i$ can be recovered from $r_{\tilde{k}}$ measurement matrices of dimensions $m_c \times N_i$. In other words, the number of MMV problems is determined by the \mathbf{A}_k with the largest rank, and the number of columns in $\mathbf{F}_i^{(r)*}$ and $\mathbf{F}_i^{(c)*}$ can range from 1 to N_b .

3.2.2 Joint recovery across the spectral and spatial dimensions

The second procedure to perform the reconstruction step from Eq. 5 and Eq. 6 exploits the correlation across both the spectral and the spatial dimensions. To that end, the solution to Eq. 5 is reduced to a single recovery step of a jointly sparse matrix of size $N_r \times (\sum_{k=1}^{N_b} r_k)$ from a single measurement matrix of size $m_r \times (\sum_{k=1}^{N_b} r_k)$; similarly, the solution to Eq. 6 is reduced to a single recovery step of a

jointly sparse matrix of size $N_c \times (\sum_{k=1}^{N_b} r_k)$ from a single measurement matrix of size $m_c \times (\sum_{k=1}^{N_b} r_k)$. For simplicity, we again assume $r_1 = r_2 = \dots = r_k$ for all k . The discussion that follows is easily extendable to the more general case.

$$\begin{aligned} \mathbf{F}^{(r)*} &= [\mathbf{f}_{11}^{(r)}, \mathbf{f}_{12}^{(r)}, \dots, \mathbf{f}_{1N_b}^{(r)}, \mathbf{f}_{21}^{(r)}, \mathbf{f}_{22}^{(r)}, \dots, \mathbf{f}_{2N_b}^{(r)}, \dots, \mathbf{f}_{r_k1}^{(r)}, \mathbf{f}_{r_k2}^{(r)}, \dots, \mathbf{f}_{r_kN_b}^{(r)}] \\ &= \arg \min \|\mathbf{F}^{(r)}\|_{2,0} \quad \text{s.t.} \\ \mathbf{B}^{(r)} &= [\mathbf{b}_{11}^{(r)}, \mathbf{b}_{12}^{(r)}, \dots, \mathbf{b}_{1N_b}^{(r)}, \mathbf{b}_{21}^{(r)}, \mathbf{b}_{22}^{(r)}, \dots, \mathbf{b}_{2N_b}^{(r)}, \dots, \mathbf{b}_{r_k1}^{(r)}, \mathbf{b}_{r_k2}^{(r)}, \dots, \mathbf{b}_{r_kN_b}^{(r)}] \\ &= \mathbf{U}_1 \mathbf{F}^{(r)}. \end{aligned} \quad (11)$$

$$\begin{aligned} \mathbf{F}^{(c)*} &= [\mathbf{f}_{11}^{(c)}, \mathbf{f}_{12}^{(c)}, \dots, \mathbf{f}_{1N_b}^{(c)}, \mathbf{f}_{21}^{(c)}, \mathbf{f}_{22}^{(c)}, \dots, \mathbf{f}_{2N_b}^{(c)}, \dots, \mathbf{f}_{r_k1}^{(c)}, \mathbf{f}_{r_k2}^{(c)}, \dots, \mathbf{f}_{r_kN_b}^{(c)}] \\ &= \arg \min \|\mathbf{F}^{(c)}\|_{2,0} \quad \text{s.t.} \\ \mathbf{B}^{(c)} &= [\mathbf{b}_{11}^{(c)}, \mathbf{b}_{12}^{(c)}, \dots, \mathbf{b}_{1N_b}^{(c)}, \mathbf{b}_{21}^{(c)}, \mathbf{b}_{22}^{(c)}, \dots, \mathbf{b}_{2N_b}^{(c)}, \dots, \mathbf{b}_{r_k1}^{(c)}, \mathbf{b}_{r_k2}^{(c)}, \dots, \mathbf{b}_{r_kN_b}^{(c)}] \\ &= \mathbf{U}_1 \mathbf{F}^{(c)}. \end{aligned} \quad (12)$$

$\hat{\mathcal{X}} = \mathcal{G} \in \mathbb{R}^{N_r \times N_c \times N_b}$ can be recovered as discussed in Sec. 3.2.1, that is, by obtaining \mathbf{G}_k first according to Eq. 10 and then stacking all \mathbf{G}_k for $k = 1, \dots, N_b$ along the spectral dimension. The parallelizable nature of this approach stems from the fact that Eqs. 11 and 12 can be solved independently. Fig. 3 visually illustrates the formation of the MMV matrices in the reconstruction stage just described.

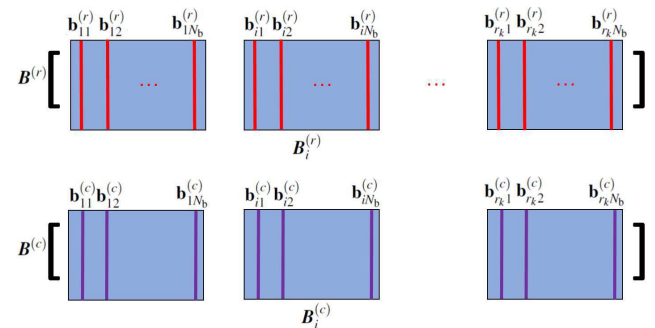


Figure 3. Formation of the MMV measurement matrices in the process of joint recovery across the spectral and spatial dimensions. All $\mathbf{B}_i^{(r)}$ and all $\mathbf{B}_i^{(c)}$, for $i = 1, 2, \dots, r_k$, in Fig. 2 are aligned to form $\mathbf{B}^{(r)}$ and $\mathbf{B}^{(c)}$ respectively.

4. Experimental results

We tested the performance of the proposed algorithms both in terms of reconstruction accuracy and execution time. We compare our results to those achieved by the method introduced in [15], which we denote SL2,0, and believe to be state-of-the-art in terms of vectorial only MMV reconstruction when only imposing the sparsity constraint on the data to be sensed. In addition, we evaluate the proposed parallelizable methods against the previously proposed serial versions from [3], which we believe to be state-of-the-art with regards to tensorial MMV reconstruction. The legends in the figures in this section refer to the proposed first (Sec. 3.2.1) and second (Sec. 3.2.2) variants of the proposed algorithm as TMMV1P and TMMV2P respectively, which stand for parallelizable Tensorial MMV exploiting correlation across the spectral dimension, and across the spectral and spatial dimensions, respectively. We refer to the serial counterparts from [3] as TMMV1S and TMMV2S respectively, which stand for serial Tensorial MMV exploiting correlation across the spectral dimension, and across the spectral and spatial dimensions, respectively.

In order to test the performance of the algorithms, acquisition and reconstruction of hyperspectral images from the AVIRIS Yellowstone dataset introduced in [12] was simulated. The AVIRIS data consists of five calibrated and corresponding 16-bit raw images acquired over Yellowstone, WY in 2006, as well as two additional 12-bit uncalibrated images. Each image is a 512-line scene containing 224 spectral bands. Reconstruction accuracy was measured in terms of Normalized Mean Squared Error (NMSE) between the reconstructed and original image, and execution time was measured in seconds on a Windows 10 machine with 32GB of RAM and an Intel i7 2.60GHz processor. The implementation of both algorithms was done in Matlab R2016b. Each data point in every figure corresponds to an average of ten runs.

In order to keep the computational complexity of the problem manageable for the SL2,0 algorithm, the hyperspectral images were cropped to 32×32 pixels and 32 bands. Let $n = m_r = m_c$ denote the number of measurements along both spatial dimensions of the hyperspectral image; the entries of the sampling matrices are drawn from a Gaussian distribution with zero-mean and standard deviation $\sqrt{\frac{1}{n}}$. For simplicity, we set the number of measurements for both spatial dimensions to be equal and use the same sampling matrix for both dimensions ($U_1 = U_2 = U$); that is, the randomly constructed Gaussian matrix U is of size $n \times 32$ for each mode. Therefore, the sampling matrix for the method in [15] is equivalent to $U \otimes U$ of size $n^2 \times 32^2$, which results in a total number of measurements of $32n^2$ across the spectral bands. We refer to $\frac{n^2}{32^2}$ as the normalized number of measurements. We vary the normal-

ized number of measurements from 0.025 to 0.5 in steps of 0.025.

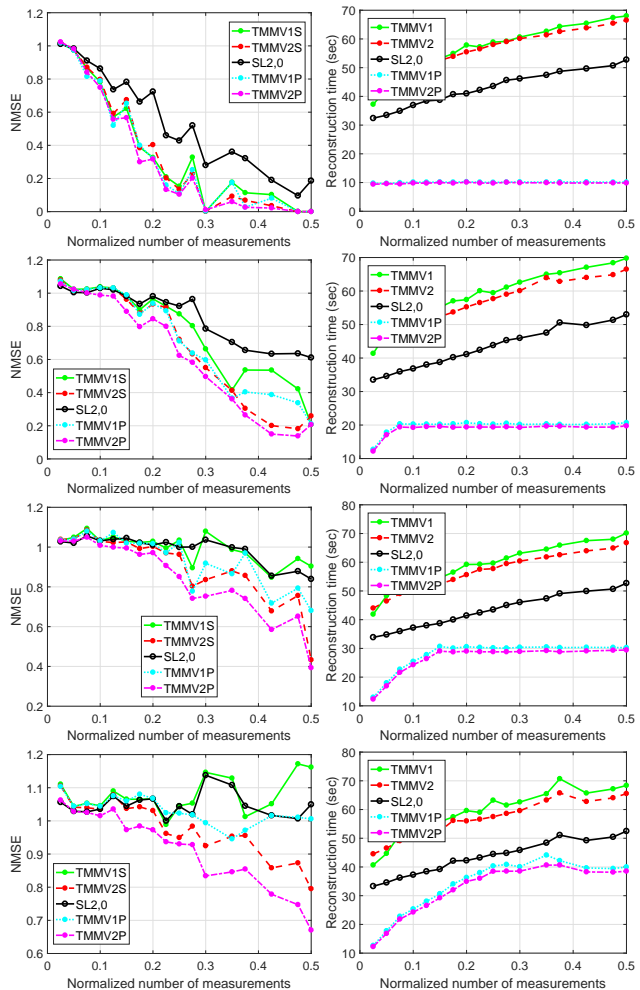


Figure 4. Performance comparison (in terms of reconstruction accuracy - left, and execution times - right) between proposed reconstruction approach, the approach in [15], and the approaches in [3] with row and column sparsity of 4 (row 1), 8 (row 2), 12 (row 3) and 16 (row 4).

Fig. 4 includes reconstruction errors (left column) and execution times (right column) as a function of the normalized number of measurements for all five methods and for different values of column- and row-sparsity (4, 8, 12 and 16 for first, second, third and fourth rows, respectively). The plots on the left column show that the tensorial methods, TMMV1S, TMMV2S, TMMV1P, and TMMV2P achieve improved reconstruction error since they more effectively exploit correlations in the image while at the same time preserving the spatial structure of the image. In addition, the proposed parallelizable methods, namely TMMV1P and TMMV2P, both outperform their sequential counterparts TMMV1S and TMMV2S respectively. We

believe the reasons for this performance improvement are twofold. In terms of reconstruction error, the serial version reconstructs two tensorial modes sequentially, which transfers the reconstruction error in the first tensor mode to the next stage, causing the errors to accumulate and possibly magnify between stages. In contrast, the parallelizable TMMV methods decouple the reconstruction of the two tensor modes, so there is no error accumulation across the reconstructed modes. In terms of computational complexity, the proposed approach benefits from the reduced number of reconstruction problems that are solved, which is determined by the rank of the measurement matrices. We note that, as its name indicates, the parallelizable approach can benefit from multi-threading; however, for comparison purposes, we implemented and tested a non-parallelized version of the algorithm. The results show that despite its superiority, there is significant room for improvement for the parallelizable version of the algorithm when multiple processing cores are available, regardless of the rank of the original signal.

The fact that TMMV2P is generally superior to TMMV1P shows that exploiting spatial and spectral correlations together leads to a slight gain in both computational efficiency and accuracy. Both TMMV1P and TMMV2P are faster than SL2,0 and their sequential counterparts since the measurement matrix for the tensorial approaches has lower dimensionality than its vectorial counterpart. As expected, as the number of measurements being acquired increases, reconstruction error improves at a slight cost of computational complexity. We note that, in contrast with the results reported in [3], SL2,0 is more computationally efficient than both TMMV1S and TMMV2S. We believe this is because the experimental validation in [3] evaluated the performance of the algorithms on the reconstruction of 3-channel color images, which seems to indicate that SL2,0 scales better as the number of channels grows.

As stated, the NMSE reported in Fig. 4 for each data point is an average across 10 trials; however, some tasks impose constraints on the largest NMSE that can be allowed. Consequently, additional insight into the reconstruction performance of the algorithms can be gained from evaluating and comparing reconstruction success rates. A reconstruction is deemed successful when the NMSE is smaller than some predetermined threshold; in this case, we set that threshold to 1×10^{-3} , which seems to be standard in previous work in the CS literature. Success rate is the ratio between the number of successful reconstructions and the total number of trials for each data point. Fig. 5 shows the success rate comparison. The plots show that the proposed methods are more likely to achieve a successful reconstruction that satisfies the NSP constraint according to CS methodology with the same number of measurements.

Fig. 6 includes reconstruction errors (left column) and

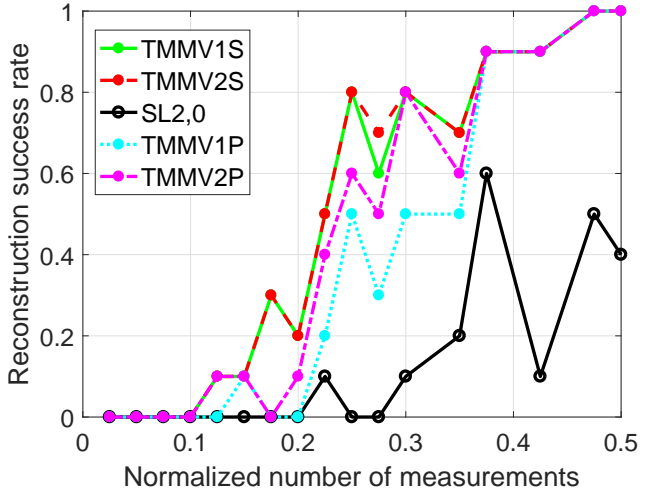


Figure 5. Success rate comparison between proposed reconstruction approach, the approach in [15], and the approaches in [3] with row and column sparsity of 4.

execution times (right column) as a function of different values of column- and row-sparsity for different numbers of normalized measurements (0.125, 0.25, 0.375 and 0.5 for first, second, third and fourth rows, respectively). The results are indicative of similar trends as those observed in Fig. 4. It can also be seen that the reconstruction time of the parallelizable methods, while always being the smallest among all methods being compared, increases with the degree of sparsity; this is because the matrix rank of each spectral slice grows as the degree of sparsity increases. In contrast, the degree of sparsity doesn't have a significant effect on the execution times for the method in [15] and the serial methods [3]. As expected, the reconstruction error grows as the sparsity level increases for all methods.

5. Conclusion

We proposed two different variants of a parallelizable reconstruction algorithm that exploits joint sparsity across different modes of a hyperspectral image while at the same time preserving its high-dimensional nature throughout the sampling and reconstruction stages of a novel CS framework. We demonstrated empirically that the newly introduced methods outperform previously proposed approaches that vectorize the data [15], as well as methods that, while still respecting the tensorial nature of the data, reconstruct the different modes sequentially [3]. The advantages of the proposed methods in terms of both computational efficiency and reconstruction accuracy were observed to be consistent across a wide range of experimental settings.

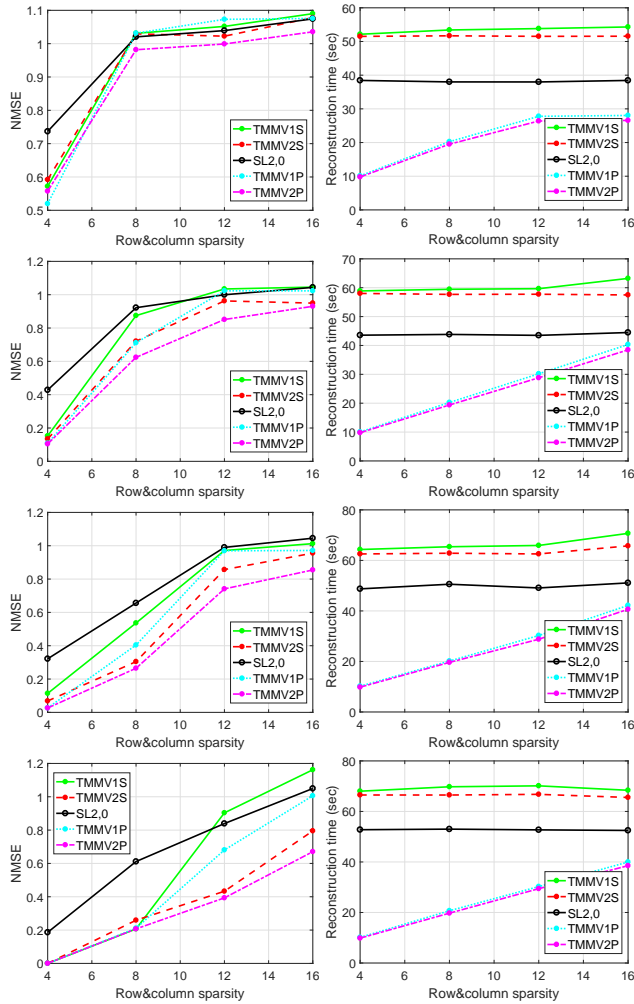


Figure 6. Performance comparison between proposed reconstruction approach and the approach in [15] with the numbers of normalized measurements being fixed to 0.125 (row 1), 0.25 (row 2), 0.375 (row 3) and 0.5 (row 4) respectively.

References

- [1] S. J. Axler. *Linear Algebra Done Right*. Springer, New York, 1997. 3
- [2] E. A. Bernal and Q. Li. Hybrid vectorial and tensorial compressive sensing for hyperspectral imaging. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2454–2458, April 2015. 2
- [3] E. A. Bernal and Q. Li. Tensorial compressive sensing of jointly sparse matrices with applications to color imaging. In *2017 IEEE International Conference on Image Processing*, Sept 2017. 2, 3, 4, 6, 7
- [4] K. Chang, Y. Liang, C. Chen, Z. Tang, and T. Qin. Color image compressive sensing reconstruction by using inter-channel correlation. In *2016 Visual Communications and Image Processing (VCIP)*, pages 1–4, Nov 2016. 2
- [5] J. Chen and X. Huo. Theoretical results on sparse representations of multiple-measurement vectors. *IEEE Transactions on Signal Processing*, 54(12):4634–4643, Dec 2006. 2, 4
- [6] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Transactions on Signal Processing*, 53(7):2477–2488, July 2005. 2
- [7] X. Ding, W. Chen, and I. Wassell. Nonconvex compressive sensing reconstruction for tensor using structures in modes. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4658–4662, March 2016. 2
- [8] S. Friedland, Q. Li, and D. Schonfeld. Compressive sensing of sparse tensors. *IEEE Transactions on Image Processing*, 23(10):4438–4447, Oct 2014. 4
- [9] S. Friedland, Q. Li, D. Schonfeld, and E. A. Bernal. Two algorithms for compressed sensing of sparse tensors. *CoRR*, abs/1404.1506, 2014. 2, 3, 4
- [10] A. Gogna, A. Shukla, H. K. Agarwal, and A. Majumdar. Split bregman algorithms for sparse / joint-sparse and low-rank signal recovery: Application in compressive hyperspectral imaging. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 1302–1306, Oct 2014. 2
- [11] G. Hennenfent, E. van den Berg, M. P. Friedlander, and F. J. Herrmann. New insights into one-norm solvers from the pareto curve. *GEOPHYSICS*, 73(4):A23–A26, 2008. 2
- [12] A. Kiely and M. Klimesh. Exploiting calibration-induced artifacts in lossless compression of hyperspectral imagery. *Geoscience and Remote Sensing, IEEE Transactions on*, 47(8):2672–2678, Aug 2009. 6
- [13] Q. Li and E. A. Bernal. Hybrid tenso-vectorial compressive sensing for hyperspectral imaging. *Journal of Electronic Imaging*, 25(3):033001, 2016. 2
- [14] A. Majumdar and R. K. Ward. Compressive color imaging with group-sparsity on analysis prior. In *2010 IEEE International Conference on Image Processing*, pages 1337–1340, Sept 2010. 2
- [15] A. Majumdar, R. K. Ward, and T. Aboulnasr. Algorithms to approximately solve np hard row-sparse mmv recovery problem: Application to compressive color imaging. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(3):362–369, Sept 2012. 2, 3, 6, 7, 8
- [16] M. Mishali and Y. C. Eldar. Reduce and boost: Recovering arbitrary sets of jointly sparse vectors. *IEEE Transactions on Signal Processing*, 56(10):4692–4702, Oct 2008. 2
- [17] E. van den Berg and M. P. Friedlander. Theoretical and empirical results for recovery from multiple measurements. *IEEE Transactions on Information Theory*, 56(5):2516–2527, May 2010. 2