



# Sparse non-negative tensor factorization using columnwise coordinate descent

Ji Liu <sup>\*</sup>, Jun Liu, Peter Wonka, Jieping Ye

Department of Computer Science and Engineering, Arizona State University, Tempe, AZ 85287, United States

## ARTICLE INFO

### Article history:

Received 4 September 2009

Received in revised form

5 February 2011

Accepted 28 May 2011

Available online 15 June 2011

### Keywords:

Sparse

Non-negative

Tensor factorization

Columnwise coordinate descent

## ABSTRACT

Many applications in computer vision, biomedical informatics, and graphics deal with data in the matrix or tensor form. Non-negative matrix and tensor factorization, which extract data-dependent non-negative basis functions, have been commonly applied for the analysis of such data for data compression, visualization, and detection of hidden information (factors). In this paper, we present a fast and flexible algorithm for sparse non-negative tensor factorization (SNTF) based on columnwise coordinate descent (CCD). Different from the traditional coordinate descent which updates one element at a time, CCD updates one column vector simultaneously. Our empirical results on higher-mode images, such as brain MRI images, gene expression images, and hyperspectral images show that the proposed algorithm is 1–2 orders of magnitude faster than several state-of-the-art algorithms.

© 2011 Elsevier Ltd. All rights reserved.

## 1. Introduction

Non-negative matrix and tensor factorization (NMF/NTF) aim to extract data-dependent non-negative basis functions [13,6,20,27], so that the target data can be expressed by the linear or multi-linear combination of these non-negative components. They have been commonly applied for the analysis of such data for data compression, visualization, and detection of hidden information (factors), e.g., in face recognition [24], psychometric [19] and image analysis [25]. Additionally, the basis can be constrained to be sparse which typically leads to an even more meaningful decomposition of the data. As a result, many researchers focused on sparse non-negative matrix factorization (SNMF) [13,14,4,9] in the past few years.

A tensor, as a more general “matrix”, can be used to express more complicated intrinsic structures of higher-mode data. Thus, sparse non-negative tensor factorization (SNTF) is a natural extension of the SNMF problem. Recently, SNTF began to receive more attention. It is used in high-mode medical data analysis [18], psychometric [19], etc. The SNTF problem is not as well studied as the matrix case. In comparison with SNMF, SNTF has two additional challenges. First, a tensor usually contains a much larger data set than a matrix, thus SNMF needs to pay more attention to computing efficiency than other factors.

The other challenge lies in how to deal with the so-called “core tensor” in SNMF. Because of the special structure, tensor factorization always contains implicitly or explicitly a core tensor, which does not exist in matrix factorization. How to efficiently and effectively deal with it is one key problem in SNTF. We can either fix it as an identity [18], or incorporate it into the optimization procedure [17]. The former approach is not flexible in handling the unbalanced target tensor data, while the latter one is computationally very expensive, which makes it unsuitable for large high-mode tensor data.

In this paper, we propose a fast and flexible SNTF algorithm, which iteratively updates one basis at a time. We employ the idea of coordinate descent (CD) for the updating. CD has recently been shown to be very efficient in solving the sparse learning problem [22]. It updates one element of the basis at a time and the algorithm cycles through all elements until convergence. The key to its high efficiency lies in the closed-form solution for each update. In the proposed algorithm, we identify the independent groups of elements among different bases, and update at one time one column vector which consists of elements from all bases, rather than one element from one group. We call it “columnwise coordinate descent” (CCD). In addition, we design a flexible way to deal with the core tensor problem mentioned above. We apply the proposed algorithms to three types of higher-mode images such as brain MRI images, gene expression pattern images, and hyperspectral images. Our experiments show that the proposed algorithm is 1–2 orders of magnitude faster than several recent SNMF and SNTF algorithms while achieving the same objective value.

<sup>\*</sup> Corresponding author. Tel.: +1 480 965 9932.  
E-mail address: [ji.liu@asu.edu](mailto:ji.liu@asu.edu) (J. Liu).

The rest of the paper is structured as follows: the related work is presented in Section 2; Section 3 introduces the proposed CCD algorithm; we present the experimental results in Section 4; finally, we conclude this paper in Section 5.

## 2. Related work

Since Lee and Seung [13] started a flurry of research on non-negative matrix factorization (NMF), this field received broad attention. In addition to the use of different objective functions such as the least squares [4] and Kullback–Leibler [13], the main difference among various algorithms lies in the update rule. The update rule directly influences the convergence speed and the quality of the factorization. The multiplicative update rule proposed by Lee and Seung [14] was considered to be the classical one, although its convergence speed was quite slow. Gonzalez and Zhang [7] used an interior-point gradient method to accelerate the multiplicative update. Recently, a quasi-Newton optimization approach was employed as the update rule by Zduniedk and Cichocki [28]. Lin [15] employed a projected gradient bound-constrained optimization method which has better convergence properties than the multiplicative update. Recently, Kim and Park [10] proposed a novel formulation of sparse non-negative matrix factorization (SNMF) and used alternating non-negativity-constrained least squares for the computation. Their results showed that it achieved better clustering performance with a lower computational cost than other existing NMF algorithms. See [1] for a more detailed review on NMF.

Tensor factorization is a natural extension of matrix factorization. It has been applied successfully in face recognition [24], psychometric [19], and image analysis [25]. Two popular models have been studied for tensor factorization including the Parafac model [8,3] and the Tucker model [23]. Most work focus on the Parafac model, since it is simpler and easier to understand from the matrix perspective. Welling, M. and Weber, M. [26] proposed a non-negative tensor factorization algorithm based on the multiplicative updating rule [14], a natural extension of matrix factorization; Kim et al. [11] proposed a non-negative tensor factorization algorithm based on the alternating large-scale non-negativity constrained least squares; A non-negative sparse factorization algorithm using Kullback–Leibler divergence [13] was introduced by FitzGerald et al. [20]; Mørup et al. [16] proposed a SNTF algorithm based on the Parafac model, which employs the  $L_1$  norm penalty as the sparseness constraint like ours, and applies the multiplicative updating rule like most other algorithms [6,20]. They further employ over relaxed bound optimization strategy to accelerate the computing speed; recently, Cichocki et al. [5] presented an algorithm using alpha and beta divergences.

It is interesting to note that the Parafac model is just a specific example of the Tucker model when the core tensor is fixed to be an identity. A key issue in applying the Tucker model is how to construct the core tensor. Lathauwer et al. [12] presented the well-known HOSVD (higher-order singular value decomposition) factorization based on the SVD algorithm for matrices. Without a non-negative requirement, it forced all factors to be orthogonal so that the core tensor could be computed through a unique and explicit expression. Bro and Andersson [2] implemented a non-negative Tucker model factorization, but the core tensor was not guaranteed to be non-negative. Recently, Mørup et al. [17] proposed a non-negative sparse tensor factorization algorithm, which incorporated the core tensor into the optimization. However, the optimization process for the core tensor dominates the computational cost per iteration, resulting in overloaded computing time for convergence.

## 3. Algorithm

We detail the proposed algorithm in this section. We first introduce the SNMF problem in Section 3.1. Then, the more general SNTF problem is presented in Section 3.2. We handle the core tensor issue in Section 3.3.

### 3.1. SNMF

SNMF approximates a matrix  $A$  by a matrix  $\hat{A}$  so that  $\hat{A}$  is the product of two matrices, i.e.,

$$A \sim \hat{A} = MN^T.$$

We formulate the SNMF problem as the following optimization:

$$\begin{aligned} \min_{M,N} : \quad & \frac{1}{2} \|A - MN^T\|^2 + \lambda_1 |M| + \lambda_2 |N| \\ \text{s.t.} \quad & M \geq 0, \quad N \geq 0, \end{aligned} \quad (1)$$

where  $A \in \mathbb{R}^{p \times q}$ ,  $M \in \mathbb{R}^{p \times r}$  and  $N \in \mathbb{R}^{q \times r}$ . Here, the  $L_1$  norm penalty  $\lambda_1 |M| + \lambda_2 |N|$  is used as the sparseness constraint. Without loss of generality, we can express the target matrix  $A$  as  $A \approx MIN^T$ , where  $I$  is an identity matrix with the rank  $r$ . It is worthwhile to point out that  $I$  can be considered as a core tensor for the matrix  $\hat{A}$ , when  $\hat{A}$  is regarded as a 2-mode tensor. We will discuss the more general tensor case in the next subsection.

### 3.2. SNTF

Before formulating the SNTF problem, we first introduce some tensor fundamentals and notations. Tensors are multi-linear mappings over a set of vector spaces. An  $N$ -mode tensor is defined as  $\mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ . Its elements are denoted as  $a_{i_1 i_2 \dots i_N}$ , where  $1 \leq i_n \leq I_n, 1 \leq n \leq N$ . For instance, a vector is a 1-mode tensor and a matrix is a 2-mode tensor. It is sometimes convenient to unfold a tensor into a matrix. The unfolding of a tensor  $\mathcal{A}$  along the  $n$ th mode is defined as

$$A_{(n)} \in \mathbb{R}^{I_n \times (I_1 \dots I_{n-1} I_{n+1} \dots I_N)}.$$

The mode- $n$  product of a tensor  $\mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  and a matrix  $U \in \mathbb{R}^{J \times I_n}$  is denoted by  $\mathcal{A} \times_n U$ . Its result is still a  $N$ -mode tensor  $\mathcal{B} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N}$ . Its elements are defined as

$$b_{i_1 \dots i_{n-1} j i_{n+1} \dots i_N} = \sum_{i_n} a_{i_1 \dots i_{n-1} i_n i_{n+1} \dots i_N} u_{j i_n}. \quad (2)$$

The mode- $n$  product  $\mathcal{B} = \mathcal{A} \times_n U$  can be computed via the matrix multiplication  $B_{(n)} = U A_{(n)}$  followed by a “fold” operation along the  $n$ th mode. The “fold” operation is defined as  $\text{fold}_n(\mathcal{B}_{(n)}) = \mathcal{B}$ . This paper uses

$$\|\mathcal{A}\| = \left( \sum_{i_1, i_2, \dots, i_N} |a_{i_1, i_2, \dots, i_N}|^2 \right)^{1/2}$$

as the Frobenious norm of a tensor and  $|\mathcal{A}| = \sum_{i_1, i_2, \dots, i_N} (|a_{i_1, i_2, \dots, i_N}|)$  as the  $L_1$  norm.

Next, we formally formulate the SNTF problem. Like SNMF, the goal of SNTF is to search for the factorization of a tensor  $\hat{\mathcal{A}}$  to approximate the target tensor  $\mathcal{A}$ . For convenience, let us first consider a simple example where the mode of  $\hat{\mathcal{A}}$  is 2, i.e.  $\hat{\mathcal{A}}$  is a matrix. According to the definition of SNMF,  $\hat{\mathcal{A}}$  can be factored as  $\hat{\mathcal{A}} = MN^T$  or  $\hat{\mathcal{A}} = MIN^T$ . Using the tensor notation,  $\hat{\mathcal{A}}$  is denoted as  $\hat{\mathcal{A}} = I \times_1 M \times_2 N$ . Here, the identity  $I$  plays the role of the core tensor. Through a simple generalization, when the mode of the target tensor is greater than 2, the approximate tensor  $\hat{\mathcal{A}}$  can be factored as

$$\hat{\mathcal{A}} = C \times_1 U_1 \times_2 \dots \times_N U_N,$$

where  $C$  is an identity tensor whose elements are 1 in the diagonal and 0 otherwise. Similar to Eq. (1), the SNTF problem can be described as the following optimization problem:

$$\begin{aligned} \min_{U_1, \dots, U_N} \quad & \frac{1}{2} \|A - C \times_1 U_1 \times_2 \dots \times_N U_N\|^2 + \sum_{1 \leq n \leq N} \lambda_n |U_n| \\ \text{s.t.} \quad & U_n \geq 0, \quad 1 \leq n \leq N, \end{aligned} \quad (3)$$

where  $\hat{A} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  and  $U_n \in \mathbb{R}^{I_n \times J_n}$  for all  $1 \leq n \leq N$ . Since the core tensor  $C$  is assumed to be an identity tensor,  $J_1 = J_2 = \dots = J_N$ . The  $L_1$  norm penalty in the objective function forces  $U_n$  to be sparse.

We propose to solve the optimization problem in Eq. (3) iteratively by updating one part at a time with all other parts fixed. For example, if we fix  $U_1, \dots, U_{n-1}, U_{n+1}, \dots, U_N$  and search for the optimal  $U_n^*$ , we obtain the following optimization sub-problem:

$$\begin{aligned} \min_{U_n} \quad & \frac{1}{2} \|A - C \times_1 U_1 \dots \times_N U_N\|^2 + \lambda_n |U_n| \\ \text{s.t.} \quad & U_n \geq 0. \end{aligned} \quad (4)$$

Since  $\|A - C \times_1 U_1 \dots \times_N U_N\| = \|A_{(n)} - U_n(C \times_1 U_1 \dots \times_{n-1} U_{n-1} \times_{n+1} \dots \times_N U_N)_{(n)}\|$ , the problem is equal to the following one:

$$\begin{aligned} \min_{U_n} \quad & \frac{1}{2} \|A_{(n)} - U_n B_{(n)}\|^2 + \lambda_n |U_n| \\ \text{s.t.} \quad & U_n \geq 0, \end{aligned} \quad (5)$$

where

$$B_{(n)} = (C \times_1 U_1 \dots \times_{n-1} U_{n-1} \times_{n+1} \dots \times_N U_N)_{(n)}.$$

We can further simplify the optimization problem in Eq. (5), by taking the transpose and separating the equations into the  $I_n$  columns of the matrix  $U_n^T$ , resulting in  $I_n$  independent optimization problems:

$$\begin{aligned} \min_{u_i} \quad & \frac{1}{2} \|B_{(n)}^T u_i - a_i\|^2 + \lambda_n |u_i| \\ \text{s.t.} \quad & u_i \geq 0, \end{aligned} \quad (6)$$

where

$$A_{(n)} = [a_1, a_2, \dots, a_{I_n}]^T,$$

$$U_n = [u_1, u_2, \dots, u_{I_n}]^T.$$

For this convex but not differentiable optimization problem, a coordinate descent (CD) method can be applied to find the global minimum [22]. The basic idea of CD is to optimize one element at a time while fixing other elements by decomposing the problem in Eq. (6) into several sub-problems as

$$\begin{aligned} \min_{u_{ij}} \quad & \frac{1}{2} \|B_{(n)}^T u_i - a_i\|^2 + \lambda_n |u_{ij}| \\ \text{s.t.} \quad & u_{ij} \geq 0, \end{aligned} \quad (7)$$

where  $u_i = [u_{i1}, \dots, u_{ij}, \dots, u_{iJ_n}]^T$ . Since the objective function is not differentiable, we need to compute its subdifferential to search for its optimum. The subdifferential of a function  $f(x)$  at a point  $x$  is defined as

$$\partial f(x) = \{d | f(y) \geq f(x) + \langle d, y - x \rangle, \forall y\}.$$

In general, the subdifferential at a particular point is a set rather than a single value. If the function  $f(x)$  is differentiable at the point  $\hat{x}$ , then the differential value  $f'(\hat{x})$  is the unique element in its subdifferential. If  $x^*$  is a global optimum, then  $0 \in \partial f(x^*)$  must be satisfied. For the problem in Eq. (7), we compute the

subdifferential of the objective function as follows:

$$\partial f(u_{ik}) = \begin{cases} \{b_k(B_{(n)}^T u_i - b_k^T a_i) + \lambda_n\}, & u_{ik} > 0, \\ \{b_k(B_{(n)}^T u_i - b_k^T a_i) - \lambda_n\}, & u_{ik} < 0, \\ [b_k(B_{(n)}^T u_i - b_k^T a_i) - \lambda_n, \\ b_k(B_{(n)}^T u_i - b_k^T a_i) + \lambda_n], & u_{ik} = 0, \end{cases} \quad (8)$$

where

$$B_{(n)} = [b_1^T, \dots, b_j^T, \dots, b_{J_n}^T]^T.$$

Eq. (8) distinguishes between three cases. In the first two cases, the function is differentiable at point  $u_{ik}$  and the subdifferential corresponds to the gradient. In the third case the subdifferential is a closed interval. We then look for the optimal  $u_{ij}^*$  which satisfies  $0 \in \partial f(u_{ij}^*)$ . We can obtain the solution in a closed-form as follows:

$$u_{ij}^* = \begin{cases} \frac{t - \lambda_n}{b_k b_j^T}, & t > \lambda_n, \\ \frac{t + \lambda_n}{b_k b_j^T}, & t < \lambda_n, \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

where

$$t = b_j(B_{(n)}^T u_i - b_j^T a_i) - b_j u_{ij}. \quad (10)$$

If we further force the non-negative constraint  $u_{ij} \geq 0$ , the optimum solution can be computed as follows:

$$u_{ij}^* = \begin{cases} \frac{t - \lambda_n}{b_j b_j^T}, & t > \lambda_n, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

The coordinate descent method discussed above optimizes the objective function in Eq. (5) by one element at a time where the formulation for each individual element is given by Eq. (7). This is illustrated in the first row of Fig. 1. A big advantage of our derivation is that several variables are independent and can be updated simultaneously. The optimization formulation in Eq. (6) considers a row of  $U_n$  at the time, but all elements in a row are dependent. Therefore, it seems that no more simplification is possible. However, the trick is to realize that all of the  $n^{\text{th}}$  coordinates in each of the  $I_n$  rows of  $U_n$  are independent. Therefore, we can update one column vector at the time. That means we simultaneously work on  $I_n$  individual elements, one from each of the row equations shown in Eq. (6). We call the algorithm




Fig. 1. Comparison of CD updating (top row) and CCD updating (bottom row).

**Table 1**  
Complexity analysis.

| Method      | Complexity  |
|-------------|---|
| Step 1      | $O(J_n^2 \prod_{i \neq n} I_i)$   |
| Step 2      | $O(J_n \prod_{i=1}^N I_i)$  |
| Step 3      | 0   |
| Step 4      | 0   |
| Step 5–9    | $O(K_I I_n J_n^2)$  |
| Algorithm 1 | $O(K_I I_n J_n^2) + O(J_n \prod_{i=1}^N I_i)$                                   |
| Total       | $O(K_O K_I \sum_{i=1}^N I_i J_i^2) + O(K_O \sum_{i=1}^N I_i \prod_{i=1}^N I_i)$ |

“columnwise coordinate descent” (CCD), which updates all elements in one column together as shown in the second row of Fig. 1. Our empirical results show that CCD is 1–2 orders of magnitude faster than CD. The detailed CCD updating procedure is shown in Algorithm 1.

**Algorithm 1.** Columnwise Coordinate Descent Updating.

**Input:**  $A_{(n)}, B_{(n)}, \lambda_n$

**Output:**  $U_n$

- 1: Set  $M = B_{(n)} B_{(n)}^T$ ;
- 2: Set  $N = A_{(n)} B_{(n)}^T - \lambda_n$ ;
- 3: Set  $D = [M(1,1)M(2,2) \dots M(J_n J_n)]$ ;
- 4: Set  $M(j,j) = 0$  for all  $1 \leq j \leq J_n$ ;
- 5: **while** not convergent **do**
- 6:   **for**  $j = 1$  to  $J_n$  **do**
- 7:     Updating  $U_n(:, j) = \max\left(0, \frac{N(:, j) - U_n M}{D(j)}\right)$
- 8:   **end for**
- 9: **end while**

The computational complexity is presented in Table 1 (considering only the “multiplication” and “addition” operations). We defined  $K_I$  as the maximal iteration number for the iteration in Steps 5–9 in Algorithm 1 and  $K_O$  as the running number of Algorithm 1.


In our experience,  $K_I \leq 100$  and  $K_O \leq 300$  in general. One can see that if the mode number  $N$  is large,  $\prod_{i \neq n} I_i \gg K_I I_n$  such that the main computational load lies in Steps 1 and 2. The computational complexity in Steps 1 and 2 has the same asymptotic complexity as an iteration in the state-of-the-art tensor/matrix factorization algorithms [18,14]. However, our algorithm converges much faster than the state-of-the-art tensor/matrix factorization algorithms (see the empirical comparison in Section 4). That is why our algorithm is more efficient especially in the large scale data.

### 3.3. Non-identity core tensor

In the exposition above we assume that the core tensor is an identity following the SNMF structure. However, forcing the core tensor to be an identity may not work in some situations. For instance, when the target tensor is very unbalanced, say, its mode is  $1000 \times 1000 \times 15 \times 3$  in size, then choosing an identity core tensor will create a dilemma. A high-mode core tensor (say,  $20 \times 20 \times 20 \times 20$ ) obviously leads to redundant computing; a low-mode core tensor (say,  $3 \times 3 \times 3 \times 3$ ) may result in a high error. Thus, it is essential to construct a non-identity core tensor to deal with unbalanced tensors. In the following, we describe how to establish a non-identity core tensor.

We set out to create a core tensor using the requirements described below. This will lead to a core tensor that is similar to an identity matrix.

1. It consists of “0” and “1” elements.
2. Any two slices of the core tensor along any mode must be orthogonal. In other word, all rows of  $C_{(n)}$  must be orthogonal.



**Fig. 2.** Illustration of C-filling and P-filling.

3. It does not decrease the rank of  $\hat{A}(n)$ . In other word, for any  $n$ ,  $C_{(n)}$  is of full row rank:  $\text{rank}(C_{(n)}) = J_n$ . Here, we assume that  $J_n \leq \prod_{k \neq n} J_k$  for any  $n$ .

It is clear that a tensor designed using the guidelines above is not unique. We propose to automatically generate a core tensor fulfilling all three conditions above.

Since the core tensor contains two types of elements: “1” and “0” only, we use a matrix to store all locations of 1’s. The matrix can be expressed as  $S = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N]$  or  $[\mathbf{r}_1^T, \mathbf{r}_2^T, \dots, \mathbf{r}_M^T]^T$ .

Each row of the matrix  $S$  corresponds to an element in the core tensor which has a value of “1”. For example, if there is a row of the matrix given by  $\mathbf{r}_m = [j_1 j_2 \dots j_N]$ , then the  $(j_1 j_2 \dots j_N)$ -th element of the core tensor is 1. The condition 2 requires that any two rows of the matrix  $S$  have at least two different elements. To fulfill the condition 3, the  $n$ th column  $\mathbf{c}_n$  of the matrix  $S$  contains all integers from 1 to  $J_n$ . The pseudo-code for the core tensor estimation is given in Algorithm 2.

**Algorithm 2.** Core Tensor Estimation.

**Input:**  $J_1 J_2 \dots J_N$

**Output:**  $S$

- 1: Sort  $J_1 J_2 \dots J_N$  in increasing order. Without loss of generality, we assume  $J_1 \leq J_2 \leq \dots \leq J_N$ ;
- 2: Set  $S = [0]_{J_N \times N}$ ;
- 3: Set  $S(:, j) = 1 : J_n$ ;
- 4: Find  $n$  satisfying  $M = \prod_{i=1}^{n-1} J_i \leq J_N \leq \prod_{i=1}^n J_i$ ;
- 5: **if**  $J_n \geq M$  **then**
- 6:   Fill  $S(:, n)$  using the C-filling rule
- 7:   Fill  $S(:, 1 : n-1)$  using the P-filling rule
- 8: **else**
- 9:   Fill  $S(:, n)$  using the P-filling rule
- 10:   Fill  $S(:, 1 : n-1)$  using the C-filling rule
- 11: **end if**
- 12: Fill  $S(:, n+1 : N)$  using the C-filling rule

Under the C-filling rule, all combinations fill the matrix in a circular fashion, while under the P-filling rule, all combinations fill the matrix in a piecewise fashion. To illustrate the algorithm, We show an example in Fig. 2. The mode size is given as  $2 \times 3 \times 4 \times 8$ . Since  $2 \times 3 \leq 8 \leq 2 \times 3 \times 4$ , the first two columns should be packed together. Since  $2 \times 3 \geq 4$ , the first two columns are filled using the C-filling rule and the third column follows the P-filling rule.

## 4. Empirical evaluation

In this section, we evaluate the proposed algorithms using three different types of higher-mode images including brain MRI images, gene expression images, and hyperspectral images. All



**Table 2**

Comparison of Parafac, Tucker, and T-CCD on the brain MRI data: Iter. = iteration number; Comp. = compression ratio; Spar. = sparseness ratio, indicating the percentage of zero values in the factors; time = running time; error = error ratio defined as  $\|A - \hat{A}\|/\|A\|$ .

| Method   | Iter.     | Comp.         | Spar.         | Time           | Error         |
|--|-----------|---------------|---------------|----------------|---------------|
| Mode size = $20 \times 20 \times 20$ $\lambda_1 = \lambda_2 = \lambda_3 = 0.5$ |           |               |               |                |               |
| Parafac  | 211       | <b>0.0016</b> | 0.18532       | 192.875        | 0.1451        |
| Tucker   | 721       | 0.0027        | 0.4632        | 4327.75        | <b>0.1034</b> |
| T-CCD  | <b>49</b> | <b>0.0016</b> | <b>0.5080</b> | <b>94.0803</b> | 0.1402        |
| Mode size = $40 \times 40 \times 40$ $\lambda_1 = \lambda_2 = \lambda_3 = 0.5$ |           |               |               |                |               |
| Parafac  | 227       | <b>0.0033</b> | 0.1856        | 334.145        | 0.0877        |
| Tucker   | –         | 0.0123        | –             | > 6000         | –             |
| T-CCD  | <b>54</b> | <b>0.0033</b> | <b>0.5365</b> | <b>136.876</b> | <b>0.0809</b> |
| Mode size = $60 \times 60 \times 60$ $\lambda_1 = \lambda_2 = \lambda_3 = 0.5$ |           |               |               |                |               |
| Parafac  | 233       | <b>0.0049</b> | 0.2539        | 440.370        | 0.0748        |
| Tucker   | –         | 0.0353        | –             | > 6000         | –             |
| T-CCD  | <b>61</b> | <b>0.0049</b> | <b>0.5844</b> | <b>198.394</b> | <b>0.0651</b> |

algorithms were implemented in Matlab version 7.6.0 and all tests were performed on an Intel Core 2 2.0 Hz and 3 GB RAM computer.

#### 4.1. Comparison with several recent algorithms

We compare our proposed sparse non-negative tensor/matrix factorization using columnwise coordinate descent (called “T-CCD”/“M-CCD”) against three existing methods, including Parafac [18], Tucker [17], and ALS (alternating least squares) [10]. Note that we name the first two algorithms after the well established tensor structure they use, but all three optimization algorithms stem from recent papers.

Our first experiment compares the proposed tensor factorization algorithm T-CCD against the Parafac and Tucker algorithms. We use a Brain MRI image of size  $181 \times 217 \times 181$ . We use a core tensor of size  $20 \times 20 \times 20$ ,  $40 \times 40 \times 40$ , and  $60 \times 60 \times 60$ , respectively. We apply the same initialization and the same sparseness coefficient (“ $\lambda$ ” value in Table 2) for all three methods. The tolerance value is fixed at  $10^{-5} \times \|A\|^2$ . To make a fair comparison, we report the average results over five runs with different initializations. We compare the algorithms in the following aspects:

- running time,
- compression ratio: the memory size of storing the factorization results over the original size of this data,
- sparseness ratio: the percentage of non-zero entries in the factorization result, and
- error ratio: the reconstruction error defined by  $\|A - \hat{A}\|/\|A\|$ .

We also compare three algorithms on an open EEG (electroencephalography) data set<sup>1</sup> used in the Parafac and Tucker algorithms [17,18]. This data of size  $64 \times 512 \times 72$  is linearly normalized into the range [0,1]. We report the performance comparison in Table 3. Note that different from Table 2, we fix the computation time to about 30 s for three algorithms in terms of “Error ratio” and “Sparseness ratio”.

We can observe from Tables 2 and 3 that the proposed method outperforms the other two competing algorithms in terms of the convergence speed and the sparseness ratio. A visual comparison of the convergence speed is shown in Fig. 3. The proposed

**Table 3**

Comparison of Parafac, Tucker, and T-CCD on the EEG data: Iter. = iteration number; Comp. = compression ratio; Spar. = sparseness ratio, indicating the percentage of zero values in the factors; time = running time; error = error ratio defined as  $\|A - \hat{A}\|/\|A\|$ .

| Method  | Iter. | Comp.          | Spar.         | Time    | Error         |
|---|-------|----------------|---------------|---------|---------------|
| Core tensor size = $5 \times 5 \times 5$ $\lambda_1 = \lambda_2 = \lambda_3 = 0.1$    |       |                |               |         |               |
| Parafac   | 330   | <b>0.00141</b> | <b>0.0802</b> | 29.3842 | 0.3955        |
| Tucker  | 46    | 0.00145        | 0.0081        | 29.7751 | 0.4152        |
| T-CCD   | 143   | <b>0.00141</b> | 0.0791        | 29.3618 | <b>0.3935</b> |
| Core tensor size = $10 \times 10 \times 10$ $\lambda_1 = \lambda_2 = \lambda_3 = 0.1$ |       |                |               |         |               |
| Parafac   | 330   | <b>0.0027</b>  | 0.2392        | 30.8170 | 0.3471        |
| Tucker  | 45    | 0.0032         | 0.0567        | 29.4274 | 0.3868        |
| T-CCD   | 130   | <b>0.0027</b>  | <b>0.2423</b> | 29.6574 | <b>0.3425</b> |
| Core tensor size = $15 \times 15 \times 15$ $\lambda_1 = \lambda_2 = \lambda_3 = 0.1$ |       |                |               |         |               |
| Parafac   | 235   | <b>0.0041</b>  | 0.1702        | 29.8741 | 0.2997        |
| Tucker  | 42    | 0.0056         | 0.0097        | 30.6174 | 0.3672        |
| T-CCD   | 113   | <b>0.0041</b>  | <b>0.1822</b> | 29.4654 | <b>0.2943</b> |

algorithm is 1–2 orders of magnitude faster than other methods to reach a certain objective value. Although the Tucker algorithm can achieve a lower error ratio, its computing time and additional storage requirement for the core tensor makes it not competitive for larger data sets. Note that our algorithm can outperform the Tucker algorithm in terms of the error ratio if both algorithms employ the same compression ratio, especially when the the core tensor is unbalanced as verified in the experiment shown in next subsection.

Next, we evaluate the performance of the proposed algorithm for matrix factorization. We compare the algorithms on a *Drosophila* gene expression pattern image data set from the BDGP database.<sup>2</sup> The data matrix is of size  $10,240 \times 1000$ . The M-CCD algorithm is compared with the ALS algorithm [10]. Since ALS employs both  $L_1$  norm and  $L_2$  norm in their objective function, it is difficult to compare the objective function directly. Our comparison will focus on the main aspects of the SNMF algorithm including the sparseness ratio, the running time, and the error ratio. The results are summarized in Table 4. We can observe from Table 4 that (1) the performance of both methods are similar, if the mode size is low; (2) when the mode size become larger, our algorithm is significantly faster than ALS. We present in Fig. 4 a visual comparison between these two algorithms.

#### 4.2. Unbalanced core tensor

When the modes of the target tensor are very different, e.g.,  $\max(I_1, I_2, \dots, I_N)/\min(I_1, I_2, \dots, I_N) \gg 1$ ,

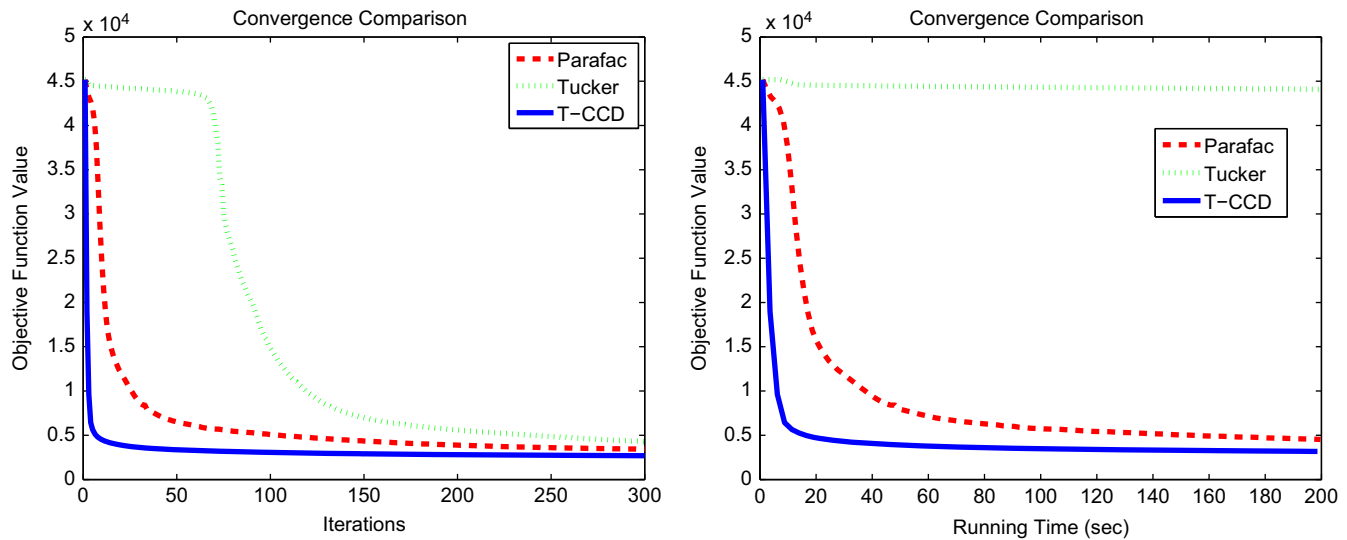
the proposed algorithm may not perform well if we force all modes of the tensor to a common size. We evaluate the effectiveness of the proposed unbalanced core tensor using the MRI brain image and the “lunar lake” hyperspectral image data.<sup>3</sup> We sub-sample them into images of size  $256 \times 207 \times 11$  and  $181 \times 217 \times 21$ , respectively, as the target tensor.

For the first data set, we test two different mode sizes: a balanced size  $29 \times 29 \times 29$  and an unbalanced size  $30 \times 30 \times 10$ . For the second data set, two different mode sizes:  $48 \times 48 \times 48$  and  $50 \times 50 \times 20$  are used. The results in Table 5 show that using an unbalanced mode size requires less running time and iteration number than using a balanced one when achieving a similar compression ratio, sparseness ratio, and error ratio.

<sup>1</sup> <http://www.erpwavelab.org/>

<sup>2</sup> <http://www.fruitfly.org/cgi-bin/ex/insitu.pl>

<sup>3</sup> <http://aviris.jpl.nasa.gov/html/aviris.freedata.html>



**Fig. 3.** Comparison of all three methods in terms of the decrease of the objective function over iterations (left figure) and time (right figure). The size of the core tensor is set to be  $40 \times 40 \times 40$ , and  $\lambda_1 = \lambda_2 = \lambda_3 = 0.5$ . We can observe from the figures that the Tucker algorithm will take orders of magnitude longer to achieve the same value of the objective function.

**Table 4**

Comparison between ALS and M-CCD using *Drosophila* gene expression pattern images.

| Method                            | Iter.     | Spar.         | Time           | Error         |
|-----------------------------------|-----------|---------------|----------------|---------------|
| Core tensor size = $20 \times 20$ |           |               |                |               |
| ALS                               | <b>42</b> | 0.2345        | <b>76.4423</b> | 0.1641        |
| M-CCD                             | 47        | <b>0.3541</b> | 81.3976        | <b>0.1603</b> |
| Core tensor size = $40 \times 40$ |           |               |                |               |
| ALS                               | <b>53</b> | 0.3281        | 312.395        | 0.1349        |
| M-CCD                             | 59        | <b>0.4219</b> | <b>126.812</b> | <b>0.1323</b> |
| Core tensor size = $80 \times 80$ |           |               |                |               |
| ALS                               | <b>73</b> | 0.4530        | 1475.39        | 0.1074        |
| M-CCD                             | 75        | <b>0.4913</b> | <b>306.123</b> | <b>0.1072</b> |

**Table 5**

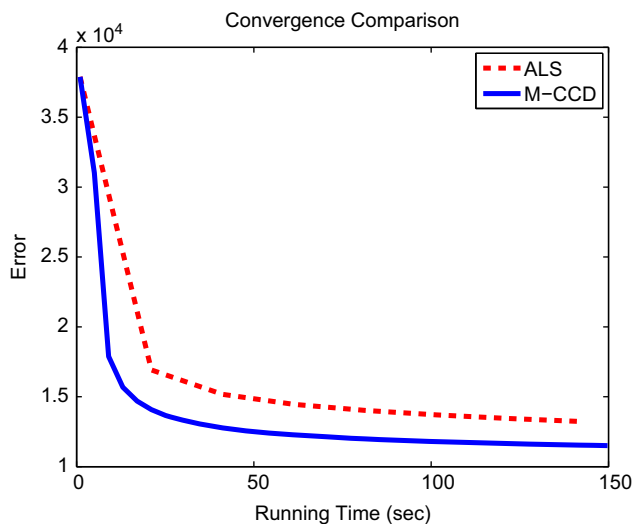
Comparison between Balanced Mode and Unbalanced Mode.

| Mode  | Iter. | Comp.  | Spar.  | Time           | Error  |
|---|-------|--------|--------|----------------|--------|
| Data: hyperspectral image size = $256 \times 307 \times 11$ |       |        |        |                |        |
| $29 \times 29 \times 29$                                    | 64    | 0.0193 | 0.4968 | 87.3137        | 0.0892 |
| $30 \times 30 \times 10$                                    | 53    | 0.0197 | 0.4897 | <b>56.9032</b> | 0.0884 |
| Data: brain image size = $181 \times 217 \times 20$         |       |        |        |                |        |
| $48 \times 48 \times 48$                                    | 77    | 0.0246 | 0.5301 | 95.4573        | 0.0765 |
| $50 \times 50 \times 20$                                    | 65    | 0.0244 | 0.5398 | <b>69.3374</b> | 0.0772 |

**Table 6**

Comparison between Tucker and T-CCD when using the same compression ratio.

| Method  | Mode                     | Comp. | Spar.  | Time         | Error         |
|---|--------------------------|-------|--------|--------------|---------------|
| Data: hyperspectral image size = $256 \times 307 \times 11$ |                          |       |        |              |               |
| Tucker  | $30 \times 30 \times 10$ | 0.030 | 0.5049 | 1250         | 0.0852        |
| T-CCD   | $46 \times 46 \times 10$ | 0.030 | 0.4917 | <b>67.32</b> | <b>0.0813</b> |
| Data: brain MRI image size = $181 \times 217 \times 61$     |                          |       |        |              |               |
| Tucker  | $30 \times 30 \times 10$ | 0.009 | 0.5187 | 1834         | 0.0972        |
| T-CCD   | $50 \times 50 \times 30$ | 0.009 | 0.5297 | <b>105.3</b> | <b>0.0779</b> |



**Fig. 4.** Comparison between M-CCD and ALS algorithm for matrix factorization. The mode size is set to be  $80 \times 80$ .

Next, we compare the proposed unbalanced core tensor with the one automatically learnt from the Tucker model. Note that the learning of the core tensor can further decrease the error ratio, however, additional time and space are required for the core

tensor optimization. In this experiment, we compare T-CCD with the core tensor derived from **Algorithm 2** with the Tucker model in terms of the error ratio when using the same compression ratio. The results are summarized in **Table 6**. We can observe from the table that the proposed algorithm can achieve a lower error ratio than the Tucker model when using the same compression ratio.

#### 4.3. An application example on biological images

In this experiment, we employ the proposed algorithm (M-CCD) on 5 groups of biological images. The target data in each group consists of 1000 *Drosophila* gene expression pattern images from the BDGP database. The *Drosophila* gene expression pattern images [21] document the spatial and temporal dynamics of gene expression and provide valuable resources for explicating the gene functions, interactions, and networks during *Drosophila* embryogenesis. The images of 5 groups are from stage ranges

4–6, 7–8, 9–10, 11–12, 13–16 of embryo development, respectively. Each image is of size  $64 \times 160$  and is unfolded into a column vector. The resulting target tensor (matrix) for each group is of size  $10,240 \times 1000$ . We apply M-CCD to extract a set of basis images by setting the mode size to be  $100 \times 100$ . We show some sample basis images of stage range 11–12 in Fig. 5. The complete 100 basis images are not shown due to the space constraint. In Fig. 6, we show sample images on the decomposition of an image as a linear combination of several basis images. We are currently




Fig. 5. Sample basis images at stage range 11–12 learnt by M-CCD.

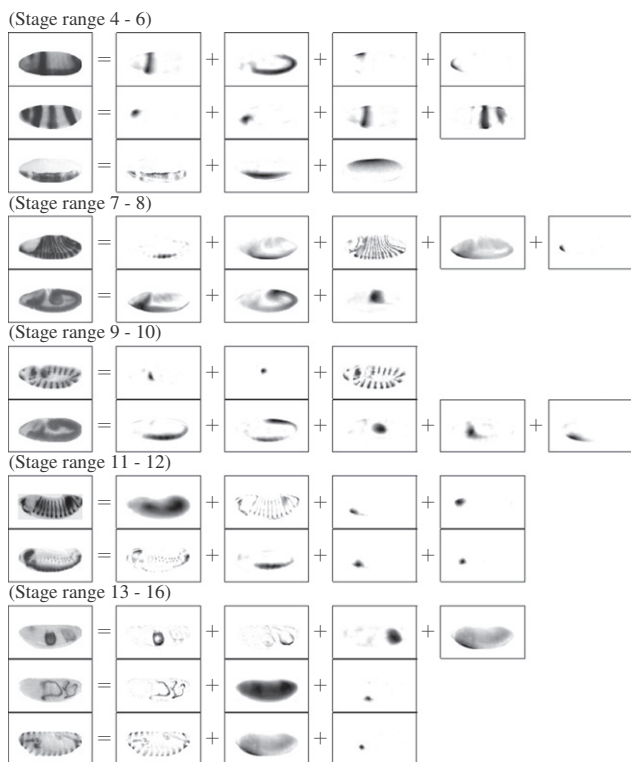


Fig. 6. decomposition of an image as the summation of a set of the learnt basis images.

working with developmental biologists to analyze the biological significance of the learnt basis images.

## 5. Conclusion

In this paper, we propose a fast and flexible algorithm for sparse non-negative tensor factorization (SNTF) based on column-wise coordinate descent (CCD). Different from the traditional coordinate descent, CCD updates one column vector simultaneously, resulting in a significant reduction in the computation time. Our empirical results on brain MRI images, gene expression images, and hyperspectral images show that the proposed algorithm is 1–2 orders of magnitude faster than several state-of-the-art algorithms. In addition, we propose to construct non-identity core tensors. Our experiments show the effectiveness of the proposed unbalanced core tensor especially when the target tensor is very unbalanced.

We have constructed a collection of basis images for *Drosophila* gene expression pattern images from stages 11–12. We plan to analyze the biological significance of the learnt basis images in the future. In addition, we plan to construct and compare the basis images for all stages to study the dynamics of the embryo development. We will explore the automatic estimation of the parameters involved in T-CCD including the size of the core tensor and  $\lambda$  in the future.

## Acknowledgments

This work was supported by NSF IIS-0612069, IIS-0812551, CCF-0811790, NGA HM1582-08-1-0016, NSFC 60905035 and NSFC 61035003.

## References

- [1] M.W. Berry, M. Browne, A.N. Langville, V. Paul Pauca, R.J. Plemmons, Algorithms and applications for approximate nonnegative matrix factorization, *Computational Statistics and Data Analysis* 52 (1) (2007) 155–173.
- [2] R. Bro, C.A. Andersson, The *n*-way toolbox for matlab, *Chemometrics and Intelligent Laboratory Systems* 52 (2000) 1–4.
- [3] J.D. Carroll, J.J. Chang, Analysis of individual differences in multidimensional scaling via an *n*-way generalization of Eckart-Young decomposition, *Psychometrika* 35 (1970) 283–319.
- [4] A. Cichocki, R. Zdunek, S. Amari, in: *Csiszar's divergences for non-negative matrix factorization: family of new algorithms*, 2006, pp. 32–39.
- [5] A. Cichocki, R. Zdunek, S. Choi, R. Plemmons, S. Amari, Nonnegative tensor factorization using alpha and beta divergencies, in: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2007, pp. 1393–1396.
- [6] D. FitzGerald, M. Cranitch, E. Coyle, Non-negative tensor factorisation for sound source separation, *Irish Signals and Systems Conference*, vol. 5, 2006, pp. 8–12.
- [7] E. Gonzalez, Y. Zhang, Accelerating the Lee-Seung algorithm for nonnegative matrix factorization, *Technical Report*, Rice University 5(2), 2005.
- [8] R.A. Harshman, Foundations of the Parafac procedure: models and conditions for an “explanatory” multi-modal factor analysis, *UCLA Working Papers in Phonetics* 16 (1970) 1–84.
- [9] P.O. Hoyer, Nonnegative matrix factorization with sparseness constraints, *Journal of Machine Learning Research* 5 (2004) 1457–1469.
- [10] H. Kim, H. Park, Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis, *Bioinformatics* 23 (12) (2007) 1495–1502.
- [11] H. Kim, H. Park, L. Elden, Non-negative tensor factorization based on alternating large-scale non-negativity-constrained least squares, in: *The Seventh IEEE International Conference on Bioinformatics and Bioengineering*, 2007, pp. 1147–1151.
- [12] L.D. Lathauwer, B.D. Moor, J. Vandewalle, Multilinear singular value decomposition, *SIAM Journal on Matrix Analysis and Applications* C 21 (4) (2000) 1253–1278.
- [13] D.D. Lee, H.S. Seung, Learning the parts of objects by nonnegative matrix factorization, *Nature* 401 (755) (1999) 788–791.
- [14] D.D. Lee, H.S. Seung, Algorithms for non-negative matrix factorization, *Advances in Neural Information Processing Systems* 13 (2001) 556–562.

- [15] C.J. Lin, Projected gradient methods for nonnegative matrix factorization, *Neural Computation* 19 (10) (2007) 2756–2779.
- [16] M. Mørup, L.K. Hansen, S.M. Arnfred, Erpwavelab a toolbox for multi-channel analysis of time–frequency transformed event related potentials, *Neuroscience Methods* 2 (2007) 361–368.
- [17] M. Mørup, L.K. Hansen, S.M. Arnfred, Algorithms for sparse nonnegative Tucker decomposition, *Neural Computation* 20 (8) (2008) 2112–2131.
- [18] M. Mørup, L.K. Hansen, J. Parnas, S.M. Arnfred, Decomposing the time–frequency representation of EEG using non-negative matrix and multi-way factorization, Technical University of Denmark Technical Report, 2006.
- [19] T. Murakami, P.M. Kroonenberg, Three-mode models and individual differences in semantic differential data, *Multivariate Behavioral Research* 38 (2) (2003) 247–283.
- [20] M.R. Parry, I. Essa, Estimating the spatial position of spectral components in audio, *Lecture Notes in Computer Science* vol. 3889 (2006) 666–673.
- [21] P. Tomancak, A. Beaton, R. Weiszmman, E. Kwan, S.Q. Shu, S.E. Lewis, S. Richards, M. Ashburner, V. Hartenstein, S.E. Celniker, G. Rubin, Systematic determination of patterns of gene expression during *Drosophila* embryogenesis, *Genome Biology* 3 (12) (2002).
- [22] P. Tseng, Convergence of block coordinate descent method for nondifferentiable minimization, *Journal of Optimization Theory and Applications* 109 (2001) 474–494.
- [23] L.R. Tucker, Some mathematical notes on three-mode factor analysis, *Psychometrika* 31 (1966) 279–311.
- [24] M.A.O. Vasilescu, D. Terzopoulos, Multilinear independent components analysis, The 2005 IEEE Conference on Computer Vision and Pattern Recognition, vol. 1, 2005, pp. 547–553.
- [25] H. Wang, N. Ahuja, Facial expression decomposition, The Ninth IEEE International Conference on Computer Vision, vol. 2, 2003, pp. 958–965.
- [26] M. Welling, M. Weber, Positive tensor factorization, *Pattern Recognition Letters* 22 (12) (2001) 1255–1261.
- [27] S. Wild, J. Curry, A. Dougherty, Improving non-negative matrix factorizations through structured initialization, *Pattern Recognition* 37 (11) (2004) 2217–2232.
- [28] R. Zdunek, A. Cichocki, Non negative matrix factorization with quasi Newton optimization, The Eighth International Conference on Artificial Intelligence and Soft Computing, vol. 4029, 2006, pp. 870–879.

**Ji Liu** received his master degree in Computer Science at the Arizona State University, and bachelor degree in Automation at the University of Science and Technology of China. Now he is a Ph.D. student in the Department of Computer Sciences at the University of Wisconsin-Madison. His research interests include optimization, machine learning and the application in computer vision and graphics.

**Jun Liu** received his B.S. degree from Nantong Institute of Technology (now Nantong University) in 2002, and his Ph.D. degree from Nanjing University of Aeronautics and Astronautics (NUAA) in November, 2007. He joined the Department of Computer Science & Engineering, NUAA, as a Lecturer in 2007. He is currently a Postdoc in the Biodesign Institute of Arizona State University. His research interests include Sparse Learning, Machine Learning and Computer vision, and he has authored or coauthored over 20 scientific papers.

**Peter Wonka** received the M.S. degree in Urban Planning and the Doctorate in Computer Science from the Technical University of Vienna. He is currently with Arizona State University (ASU). Prior to coming to ASU, he was a Postdoctorate Researcher at the Georgia Institute of Technology for two years. His research interests include various topics in computer graphics, visualization, and image processing. He is a member of the IEEE.

**Jieping Ye** received the Ph.D. degree in Computer Science from the University of Minnesota-Twin Cities in 2005. He is currently an Assistant Professor in the Department of Computer Science and Engineering, Arizona State University. He has been a core faculty member of the Center for Evolutionary Functional Genomics, Bio-design Institute, Arizona State University, since August 2005. His research interests include machine learning, data mining, and bioinformatics. He has published extensively in these areas. He received the Guidant Fellowship in 2004 and 2005. In 2004, his paper on generalized low-rank approximations of matrices won the outstanding student paper award at the 21st International Conference on Machine Learning. He is a member of the IEEE and the ACM.