

Scalable tensor factorizations for incomplete data[☆]

Evrin Acar^a, Daniel M. Dunlavy^c, Tamara G. Kolda^{b,*}, Morten Mørup^d

^a TUBITAK-UEKAE, Gebze, Turkey

^b Sandia National Laboratories, Livermore, CA 94551-9159, United States

^c Sandia National Laboratories, Albuquerque, NM 87123-1318, United States

^d Technical University of Denmark, 2800 Kgs. Lyngby, Denmark

ARTICLE INFO

Article history:

Received 1 December 2009

Received in revised form 28 July 2010

Accepted 4 August 2010

Available online 12 August 2010

Keywords:

Missing data

Incomplete data

Tensor factorization

CANDECOMP

PARAFAC

Optimization

ABSTRACT

The problem of incomplete data – i.e., data with missing or unknown values – in multi-way arrays is ubiquitous in biomedical signal processing, network traffic analysis, bibliometrics, social network analysis, chemometrics, computer vision, communication networks, etc. We consider the problem of how to factorize data sets with missing values with the goal of capturing the underlying latent structure of the data and possibly reconstructing missing values (i.e., tensor completion). We focus on one of the most well-known tensor factorizations that captures multi-linear structure, CANDECOMP/PARAFAC (CP). In the presence of missing data, CP can be formulated as a weighted least squares problem that models *only* the known entries. We develop an algorithm called CP-WOPT (CP Weighted OPTimization) that uses a first-order optimization approach to solve the weighted least squares problem. Based on extensive numerical experiments, our algorithm is shown to successfully factorize tensors with noise and up to 99% missing data. A unique aspect of our approach is that it scales to sparse large-scale data, e.g., $1000 \times 1000 \times 1000$ with five million known entries (0.5% dense). We further demonstrate the usefulness of CP-WOPT on two real-world applications: a novel EEG (electroencephalogram) application where missing data is frequently encountered due to disconnections of electrodes and the problem of modeling computer network traffic where data may be absent due to the expense of the data collection process.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Missing data can arise in a variety of settings due to loss of information, errors in the data collection process, or costly experiments. For instance, in biomedical signal processing, missing data can be encountered during EEG analysis, where multiple electrodes are used to collect the electrical activity along the scalp. If one of the electrodes becomes loose or disconnected, the signal is either lost or discarded due to contamination with high amounts of mechanical noise. We also encounter the missing data problem in other areas of data mining, such as packet losses in network traffic analysis [2] and occlusions in images in computer vision [3]. Many real-world data with missing entries are ignored because they are deemed unsuitable for analysis, but this work contributes to the growing evidence that such data can be analyzed.

Unlike most previous studies on missing data which have only considered matrices, we focus here on the problem of missing data in *tensors* because it has been shown increasingly that data often have

more than two modes of variation and are therefore best represented as multi-way arrays (i.e., tensors) [4,5]. For instance, in EEG data each signal from an electrode can be represented as a time-frequency matrix; thus, data from multiple channels is three-dimensional (temporal, spectral, and spatial) and forms a three-way array [6]. Social network data, network traffic data, and bibliometric data are of interest to many applications such as community detection, link mining, and more; these data can have multiple dimensions/modalities, are often extremely large, and generally have at least some missing data. These are just a few of the many data analysis applications where one needs to deal with large multi-way arrays with missing entries. Other examples of multi-way arrays with missing entries from different disciplines have also been studied in the literature [7–9]. For instance, [7] shows that, in spectroscopy, intermittent machine failures or different sampling frequencies may result in tensors with missing fibers (i.e., the higher-order analogues of matrix rows or columns, see Fig. 1). Similarly, missing fibers are encountered in multidimensional NMR (Nuclear Magnetic Resonance) analysis, where sparse sampling is used in order to reduce the experimental time [8].

Our goal is to capture the latent structure of the data via a higher-order factorization, even in the presence of missing data. Handling missing data in the context of matrix factorizations, e.g., the widely-used principal component analysis, has long been studied [10,11] (see

[☆] A preliminary conference version of this paper has appeared as [1].

* Corresponding author.

E-mail addresses: evrim.acar@bte.tubitak.gov.tr (E. Acar), dmdunla@sandia.gov (D.M. Dunlavy), tgkolda@sandia.gov (T.G. Kolda), mm@imm.dtu.dk (M. Mørup).

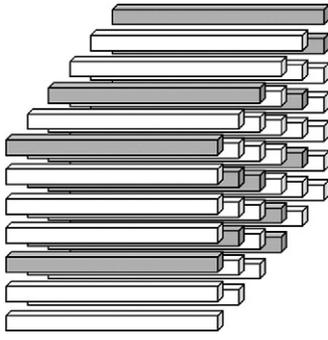


Fig. 1. A 3-way tensor with missing row fibers (in gray).

[3] for a review). It is also closely related to the matrix completion problem, where the goal is to recover the missing entries [12,13] (see Section 3 for more discussion). Higher-order factorizations, i.e., tensor factorizations, have emerged as an important method for information analysis [4,5]. Instead of flattening (unfolding) multi-way arrays into matrices and using matrix factorization techniques, tensor models preserve the multi-way nature of the data and extract the underlying factors in each mode (dimension) of a higher-order array.

We focus here on the CANDECOMP/PARAFAC (CP) tensor decomposition [14,15], which is a tensor model commonly used in various applications [6,16–19]. To illustrate differences between matrix and tensor factorizations, we introduce the CP decomposition for three-way tensors; discussion of the CP decomposition for general N -way tensors can be found in Section 4. Let \mathcal{X} be a three-way tensor of size $I \times J \times K$, and assume its rank is R (see [5] for a detailed discussion on tensor rank). With perfect data, the CP decomposition is defined by factor matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} of sizes $I \times R$, $J \times R$, and $K \times R$, respectively, such that

$$x_{ijk} = \sum_{r=1}^R a_{ir} b_{jr} c_{kr}, \quad \text{for all } i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K.$$

In the presence of noise, the true \mathcal{X} is not observable and we cannot expect equality. Instead, the CP decomposition should minimize the error function

$$f(\mathbf{A}, \mathbf{B}, \mathbf{C}) = \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^K \left(x_{ijk} - \sum_{r=1}^R a_{ir} b_{jr} c_{kr} \right)^2. \quad (1)$$

An illustration of CP for third-order tensors is given in Fig. 2. The CP decomposition is extensible to N -way tensors for $N \geq 3$, and there are numerous methods for computing it [20].

In the case of incomplete data, a standard practice is to impute the missing values in some fashion (e.g., replacing the missing entries using average values along a particular mode). Imputation can be useful as long as the amount of missing data is small; however, performance degrades for large amounts of missing data [1,10]. As a better alternative, factorizations of the data with imputed values for missing entries can be used to re-impute the missing values and the procedure can be repeated to iteratively determine suitable values for the missing entries. Such a procedure is an example of the expectation maximization (EM) algorithm [21]. Computing CP decompositions by combining the alternating least squares method, which computes the factor matrices one at a time, and iterative imputation (denoted EM-ALS in this paper) has been shown to be quite effective and has the advantage of often being simple and fast. Nevertheless, as the amount of missing data increases, the performance of the algorithm may suffer since the initialization and the intermediate models used to impute the missing values will increase the risk of converging to a less than optimal factorization [7]. Also, the poor convergence of alternating

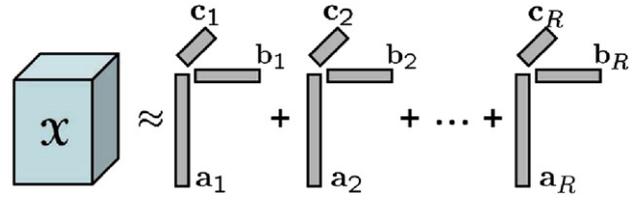


Fig. 2. Illustration of an R -component CP model for a third-order tensor \mathcal{X} .

methods due to their vulnerability to flatlining, i.e., stagnation, is noted in [3].

In this paper, though, we focus on using a weighted version of the error function to ignore missing data and model only the known entries. In that case, nonlinear optimization can be used to directly solve the weighted least squares problem for the CP model. The weighted version of Eq. (1) is

$$f_w(\mathbf{A}, \mathbf{B}, \mathbf{C}) = \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^K \left\{ w_{ijk} \left(x_{ijk} - \sum_{r=1}^R a_{ir} b_{jr} c_{kr} \right) \right\}^2, \quad (2)$$

where w , which is the same size as \mathcal{X} , is a nonnegative weight tensor defined as

$$w_{ijk} = \begin{cases} 1 & \text{if } x_{ijk} \text{ is known,} \\ 0 & \text{if } x_{ijk} \text{ is missing,} \end{cases} \quad \text{for all } i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K.$$

Our contributions in this paper are summarized as follows. (a) We develop a scalable algorithm called CP-WOPT (CP Weighted OPTimization) for tensor factorizations in the presence of missing data. CP-WOPT uses first-order optimization to solve the weighted least squares objective function over all the factor matrices simultaneously. (b) We show that CP-WOPT can scale to sparse, large-scale data using specialized sparse data structures, significantly reducing the storage and computation costs. (c) Using extensive numerical experiments on simulated data sets, we show that CP-WOPT can successfully factor tensors with noise and up to 99% missing data. In many cases, CP-WOPT is significantly faster than the best published direct optimization method in the literature [7]. (d) We demonstrate the applicability of the proposed algorithm on a real data set in a novel EEG application where data is incomplete due to failures of particular electrodes. This is a common occurrence in practice, and our experiments show that even if signals from almost half of the channels are missing, underlying brain activities can still be captured using the CP-WOPT algorithm, illustrating the usefulness of our proposed method. (e) In addition to tensor factorizations, we also show that CP-WOPT can be used to address the tensor completion problem in the context of network traffic analysis. We use the factors captured by the CP-WOPT algorithm to reconstruct the tensor and illustrate that even if there is a large amount of missing data, the algorithm is able to keep the relative error in the missing entries close to the modeling error.

The paper is organized as follows. We introduce the notation used throughout the paper in Section 2. In Section 3, we discuss related work in matrix and tensor factorizations. The computation of the function and gradient values for the general N -way weighted version of the error function and the presentation of the CP-WOPT method are given in Section 4. Numerical results on both simulated and real data are given in Section 5. Conclusions and future work are discussed in Section 6.

2. Notation

Tensors of order $N \geq 3$ are denoted by Euler script letters ($\mathcal{X}, \mathcal{Y}, \mathcal{Z}$), matrices are denoted by boldface capital letters ($\mathbf{A}, \mathbf{B}, \mathbf{C}$), vectors are

denoted by boldface lowercase letters ($\mathbf{a}, \mathbf{b}, \mathbf{c}$), and scalars are denoted by lowercase letters (a, b, c). Columns of a matrix are denoted by boldface lower letters with a subscript ($\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are the first three columns of \mathbf{A}). Entries of a matrix or a tensor are denoted by lowercase letters with subscripts, i.e., the (i_1, i_2, \dots, i_N) entry of an N -way tensor \mathcal{X} is denoted by $x_{i_1 i_2 \dots i_N}$.

An N -way tensor can be rearranged as a matrix; this is called *matricization*, also known as *unfolding* or *flattening*. The mode- n matricization of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is denoted by $\mathbf{X}_{(n)}$ and arranges the mode- n one-dimensional “fibers” to be the columns of the resulting matrix; see [1,5] for details.

Given two tensors \mathcal{X} and \mathcal{Y} of equal size $I_1 \times I_2 \times \dots \times I_N$, their Hadamard (elementwise) product is denoted by $\mathcal{X} * \mathcal{Y}$ and defined as

$$(\mathcal{X} * \mathcal{Y})_{i_1 i_2 \dots i_N} = x_{i_1 i_2 \dots i_N} y_{i_1 i_2 \dots i_N} \quad \text{for all } i_n \in \{1, \dots, I_n\} \text{ and } n \in \{1, \dots, N\}$$

The *inner product* of two same-sized tensors $\mathcal{X}, \mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is the sum of the products of their entries, i.e.,

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \dots \sum_{i_N=1}^{I_N} x_{i_1 i_2 \dots i_N} y_{i_1 i_2 \dots i_N}.$$

For a tensor \mathcal{X} of size $I_1 \times I_2 \times \dots \times I_N$, its *norm* is $\|\mathcal{X}\| = \sqrt{\langle \mathcal{X}, \mathcal{X} \rangle}$. For matrices and vectors, $\|\cdot\|$ refers to the analogous Frobenius and two-norm, respectively. We can also define a weighted norm as follows. Let \mathcal{X} and \mathcal{W} be two tensors of size $I_1 \times I_2 \times \dots \times I_N$. Then the \mathcal{W} -weighted norm of \mathcal{X} is

$$\|\mathcal{X}\|_{\mathcal{W}} = \|\mathcal{W} * \mathcal{X}\|.$$

Given a sequence of matrices $\mathbf{A}^{(n)}$ of size $I_n \times R$ for $n = 1, \dots, N$, the notation $\llbracket \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(N)} \rrbracket$ defines an $I_1 \times I_2 \times \dots \times I_N$ tensor whose elements are given by

$$\left(\llbracket \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(N)} \rrbracket \right)_{i_1 i_2 \dots i_N} = \sum_{r=1}^R \prod_{n=1}^N a_{i_n r}^{(n)},$$

for $i_n \in \{1, \dots, I_n\}, n \in \{1, \dots, N\}$.

For just two matrices, this reduces to familiar expressions: $\llbracket \mathbf{A}, \mathbf{B} \rrbracket = \mathbf{A}\mathbf{B}^T$. Using the notation defined here, Eq. (2) can be rewritten as

$$f_{\mathcal{W}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) = \frac{1}{2} \|\mathcal{X} - \llbracket \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket\|_{\mathcal{W}}^2.$$

3. Related work in factorizations with missing data

In this section, we first review the approaches for handling missing data in matrix factorizations and then discuss how these techniques have been extended to tensor factorizations.

3.1. Matrix factorizations

Matrix factorization in the presence of missing entries is a problem that has been studied for several decades; see, e.g., [10,11]. The problem is typically formulated analogously to Eq. (2) as

$$f_{\mathcal{W}}(\mathbf{A}, \mathbf{B}) = \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{B}^T\|_{\mathcal{W}}^2. \quad (3)$$

A common procedure for solving this problem is EM which combines imputation and alternation [7,22]. In this approach, the

missing values of \mathbf{X} are imputed using the current model, $\hat{\mathbf{X}} = \mathbf{A}\mathbf{B}^T$, as follows:

$$\bar{\mathbf{X}} = \mathbf{W} * \mathbf{X} + (1 - \mathbf{W}) * \hat{\mathbf{X}},$$

where $\mathbf{1}$ is the matrix of all ones. Once $\bar{\mathbf{X}}$ is generated, the matrices \mathbf{A} and \mathbf{B} can then be alternately updated according to the error function $\frac{1}{2} \|\bar{\mathbf{X}} - \mathbf{A}\mathbf{B}^T\|^2$ (e.g., using the linear least squares method). See [22,23] for further discussion of the EM method in the missing data and general weighted case.

Recently, a direct nonlinear optimization approach was proposed for matrix factorization with missing data [3]. In this case, Eq. (3) is solved directly using a 2nd-order damped Newton method. This new method is compared to other standard techniques based on some form of alternation and/or imputation as well as hybrid techniques that combine both approaches. Overall, the conclusion is that nonlinear optimization strategies are key to successful matrix factorization. Moreover, the authors observe that the alternating methods tend to take much longer to converge to the solution even though they make faster progress initially. This work is theoretically the most related to what we propose—the main differences are 1) we focus on tensors rather than matrices, and 2) we use first-order rather than second-order optimization methods (we note that first order methods are mentioned as future work in [3]).

A major difference between matrix and tensor factorizations is worth noting here. In [3,22], the lack of uniqueness in matrix decompositions is discussed. Given any invertible matrix \mathbf{G} , $\llbracket \mathbf{A}, \mathbf{B} \rrbracket = \llbracket \mathbf{A}\mathbf{G}, \mathbf{B}\mathbf{G}^{-T} \rrbracket$. This means that there is an infinite family of equivalent solutions to Eq. (3). In [3], regularization is recommended as a partial solution; however regularization can only control scaling indeterminacies and not rotational freedom. In the case of the CP model, often a unique solution (including trivial indeterminacies of scaling and column permutation) can be recovered exactly; see, e.g., [5] for further discussion on uniqueness of the CP decomposition.

Factorization of matrices with missing entries is also closely related to the matrix completion problem. In matrix completion, one tries to recover the missing matrix entries using the low-rank structure of the matrix. Recent work in this area [12,13] shows that even if a small amount of matrix entries are available and those are corrupted with noise, it is still possible to recover the missing entries up to the level of noise. In [13], it is also discussed how this problem relates to the field of compressive sensing, which exploits structures in data to generate more compact representations of the data. Practically speaking, the difference between completion and factorization is how success is measured. Factorization methods aim to increase accuracy in the factors; in other words, capture the underlying phenomena as well as possible. Completion methods, on the other hand, seek accuracy in filling in the missing data. Obviously, once a factorization has been computed, it can be used to reconstruct the missing entries. In fact, many completion methods use this procedure.

3.2. Tensor factorizations

The EM procedure discussed for matrices has also been widely employed for tensor factorizations with missing data. If the current model is $\llbracket \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket$, then we fill in the missing entries of \mathcal{X} to produce a complete tensor according to

$$\bar{\mathcal{X}} = \mathcal{W} * \mathcal{X} + (1 - \mathcal{W}) * \llbracket \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket,$$

where $\mathbf{1}$ is the tensor of all ones the same size as \mathcal{W} . The factor matrices are then updated using alternating least squares (ALS) as those that best fit $\bar{\mathcal{X}}$. See, e.g., [24,25] for further details. As noted previously, we denote this method as EM-ALS.

Paatero [26] and Tomasi and Bro [7] have investigated direct nonlinear approaches based on Gauss–Newton (GN). The code from

[26] is not widely available; therefore, we focus on [7] and its INDAFAC (INcomplete DATA paraFAC) procedure which specifically uses the Levenberg–Marquardt version of GN for fitting the CP model to data with missing entries. The primary application in [7] is missing data in chemometrics experiments. This approach is compared to EM-ALS with the result being that INDAFAC and EM-ALS perform almost equally well in general with the exception that INDAFAC is more accurate for difficult problems, i.e., higher collinearity and systematically missing patterns of data. In terms of computational efficiency, EM-ALS is usually faster but becomes slower than INDAFAC as the percentage of missing entries increases and also depending on the missing entry patterns.

Both INDAFAC and CP-WOPT address the problem of fitting the CP model to incomplete data sets by solving Eq. (2). The difference is that INDAFAC is based on second-order optimization while CP-WOPT is first-order with a goal of scaling to larger problem sizes.

4. CP-WOPT Algorithm

We consider the general N -way CP factorization problem for tensors with missing entries. Let \mathcal{X} be a real-valued tensor of size $I_1 \times I_2 \times \dots \times I_N$ and assume its rank is known to be R .¹ Define a nonnegative weight tensor \mathcal{W} of the same size as \mathcal{X} such that

$$w_{i_1 i_2 \dots i_N} = \begin{cases} 1 & \text{if } x_{i_1 i_2 \dots i_N} \text{ is known,} \\ 0 & \text{if } x_{i_1 i_2 \dots i_N} \text{ is missing,} \end{cases} \quad (4)$$

for all $i_n \in \{1, \dots, I_n\}$ and $n \in \{1, \dots, N\}$.

The N -way objective function is defined by

$$f_{\mathcal{W}}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(N)}) = \frac{1}{2} \|\mathcal{X} - \llbracket \mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)} \rrbracket\|_{\mathcal{W}}^2. \quad (5)$$

It is a function of $P = R \sum_{n=1}^N I_n$ variables. Due to the well-known indeterminacies of the CP model, it may also be desirable to add regularization to the objective function as in [20], but this has not been necessary thus far in our experiments.

Eq. (5) is equivalent to

$$f_{\mathcal{W}}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(N)}) = \frac{1}{2} \|\boldsymbol{\Psi} - \boldsymbol{\mathcal{Z}}\|^2, \quad (6)$$

where $\boldsymbol{\Psi} = \mathcal{W} * \mathcal{X}$ and $\boldsymbol{\mathcal{Z}} = \mathcal{W} * \llbracket \mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)} \rrbracket$.

The tensor $\boldsymbol{\Psi}$ can be precomputed as neither \mathcal{W} nor \mathcal{X} change during the iterations. We will see later that $\boldsymbol{\mathcal{Z}}$ can be computed efficiently for sparse \mathcal{W} .

Our goal is to find matrices $\mathbf{A}^{(n)} \in \mathbb{R}^{I_n \times R}$ for $n = 1, \dots, N$ that minimize the weighted objective function in Eq. (5). We show how to compute the function and gradient values efficiently when \mathcal{W} is dense in Section 4.1 and when \mathcal{W} is sparse in Section 4.2. Once the function and gradient are known, any gradient-based optimization method [27] can be used to solve the optimization problem.

The derivation of the gradient in the weighted case is given in [1]; here we just report the formula. In matrix notation, using $\boldsymbol{\Psi}$ and $\boldsymbol{\mathcal{Z}}$ from Eq. (6), we can rewrite the gradient equation as

$$\frac{\partial f_{\mathcal{W}}}{\partial \mathbf{A}^{(n)}} = (\mathbf{Z}_{(n)} - \mathbf{Y}_{(n)}) \mathbf{A}^{(-n)}, \quad (7)$$

¹ In practice, the rank is generally not known and is not easily determined. Understanding the performance of the methods under consideration in that scenario is a topic of future work. Results in [20] indicate that direct optimization methods have an advantage over alternating least squares approaches when the rank is over-estimated.

where

$$\mathbf{A}^{(-n)} = \mathbf{A}^{(N)} \odot \dots \odot \mathbf{A}^{(n+1)} \odot \mathbf{A}^{(n-1)} \odot \dots \odot \mathbf{A}^{(1)}$$

for $n = 1, \dots, N$. The symbol \odot denotes the Khatri-Rao product; see, e.g., [1,5] for details.

4.1. Computations with dense \mathcal{W}

Fig. 3a shows the algorithmic steps for computing the function and gradient values. We are optimizing a function of P variables as defined by the factor matrices $\mathbf{A}^{(1)}$ through $\mathbf{A}^{(N)}$. The gradient is computed as a series of matrices $\mathbf{G}^{(n)} \equiv \frac{\partial f}{\partial \mathbf{A}^{(n)}}$ for $n = 1, \dots, N$. Recall that $\mathbf{T}_{(n)}$ in the last step is the unfolding of the tensor \mathbf{T} in mode n .

4.2. Computations with sparse \mathcal{W}

If a large number of entries is missing, then \mathcal{W} is sparse. In this case, there is no need to allocate storage for every entry of the tensor \mathcal{X} . Instead, we can store and work with just the known values, making the method efficient in both storage and time. Fig. 3b shows analogous computations to Fig. 3a in the case that \mathcal{W} is sparse.

Let the ordered set $\mathcal{I} = \{\mathbf{i}^{(1)}, \mathbf{i}^{(2)}, \dots, \mathbf{i}^{(Q)}\}$ be the indices of the known values, i.e., all the locations where $\mathcal{W} = 1$. Each $\mathbf{i}^{(q)}$, $q \in \{1, \dots, Q\}$, is an N -tuple of indices whose n th entry is denoted by $i_n^{(q)}$. The known entries of \mathcal{X} can be stored in an array \mathbf{y} of length Q so that

$$y_q = x_{i_1^{(q)} i_2^{(q)} \dots i_N^{(q)}} \quad \text{for } q = 1, \dots, Q.$$

Observe that $\|\mathcal{W} * \mathcal{X}\|^2 = \|\mathbf{y}\|^2$. Thus, the value of γ is the same in both Fig. 3a and b.

In the dense version (Fig. 3a), we need to compute $\boldsymbol{\mathcal{Z}} = \mathcal{W} * \llbracket \mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)} \rrbracket$. In the sparse version, we observe that $\boldsymbol{\mathcal{Z}}$ is necessarily zero anywhere there is a missing entry in \mathcal{X} . Consequently, it is only necessary to compute the entries of $\boldsymbol{\mathcal{Z}}$ corresponding to known values in \mathcal{I} ; the vector $\boldsymbol{\mathcal{Z}}$ corresponds to these known entries. The computations for Step 2 in Fig. 3b can be done efficiently in MATLAB using the “expanded” vectors technique of ([28], Section 3.2.4). Let the r th summand be denoted by \mathbf{u} , i.e.,

$$u_q = \prod_{n=1}^N a_{i_n^{(q)}}^{(n)} \quad \text{for } q = 1, \dots, Q.$$

Let $\mathbf{v}^{(n)}$ be “expanded” vectors of length Q for $n = 1, \dots, N$ defined by

$$v_q^{(n)} = a_{i_n^{(q)}}^{(n)} \quad \text{for } q = 1, \dots, Q. \quad (8)$$

Then the vector \mathbf{u} can be calculated as

$$\mathbf{u} = \mathbf{v}^{(1)} * \mathbf{v}^{(2)} * \dots * \mathbf{v}^{(N)}.$$

This can naturally be done iteratively (i.e., only one $\mathbf{v}^{(n)}$ is calculated at a time) to reduce storage costs. Likewise, each summand \mathbf{u} can be iteratively added to compute the final answer $\boldsymbol{\mathcal{Z}}$.

In Step 3, the function values in Fig. 3a and b are clearly equivalent since \mathbf{y} and $\boldsymbol{\mathcal{Z}}$ contain the nonzero values of $\boldsymbol{\Psi}$ and $\boldsymbol{\mathcal{Z}}$, respectively. Similarly, the vector \mathbf{t} in Step 4 of Fig. 3b represents just the nonzero values of $\boldsymbol{\mathcal{Z}}$ in Fig. 3a.

The computation of the gradients in Step 6 of Fig. 3b performs a matricized-tensor-times-Khatri-Rao-product (mttkrp) calculation, which has been described for the sparse data case in ([28], Section 5.2.6). Here we briefly summarize the methodology. The r th column

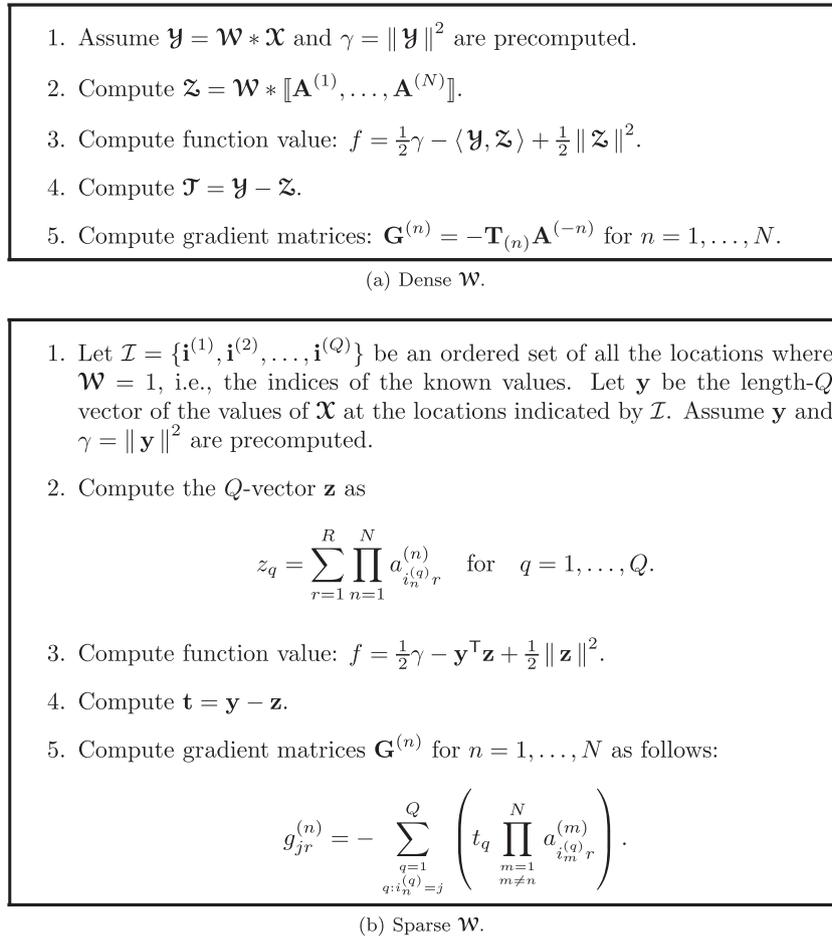


Fig. 3. CP-WOPT computation of function value and gradient.

of $\mathbf{G}^{(n)}$, $\mathbf{g}_r^{(n)}$ is calculated as follows. Let the vectors $\mathbf{v}^{(n)}$ be defined as above in Eq. (8), but define \mathbf{u} instead as

$$\mathbf{u} = \mathbf{t} * \mathbf{v}^{(1)} * \mathbf{v}^{(2)} * \dots * \mathbf{v}^{(n-1)} * \mathbf{v}^{(n+1)} * \dots * \mathbf{v}^{(N)}.$$

Then

$$\left(\mathbf{g}_r^{(n)} \right)_j = \sum_{i_n^{(q)}=j} u_q \quad \text{for } j = 1, \dots, I_n.$$

This can be computed efficiently using the accumarray function in MATLAB, and the code is available in version 2.4 of the Tensor Toolbox for MATLAB [29].

5. Experiments

On both real and simulated three-way data, we assess the performance of the CP-WOPT method. We compare CP-WOPT with other methods and also demonstrate its performance on two applications.

5.1. Computational environment

All experiments were performed using MATLAB 2009b on a Linux Workstation (RedHat 5.2) with 2 Quad-Core Intel Xeon 3.0 GHz processors and 32 GB RAM. Timings were performed using MATLAB's tic and toc functions since cpu time is known to produce inaccurate results for multi-CPU and/or multi-core systems.

CP-WOPT is implemented in the Tensor Toolbox [29]. We consider dense and sparse versions, based on the gradient and function computations shown in Fig. 3a and b, respectively. We use the nonlinear conjugate gradient (NCG) method with Hestenes–Stiefel updates [27] and the Moré–Thuente line search [30] provided in the Poblano Toolbox [31] as the optimization method.

We compare CP-WOPT to two other methods: EM-ALS (implemented in the N-way Toolbox for MATLAB, version 3.10 [32]) and INDAFAC [33], which is a damped Gauss–Newton method proposed by Tomasi and Bro [7]. Previously, Tomasi and Bro showed that INDAFAC converged to solutions in many fewer iterations than EM-ALS.

The stopping conditions are set as follows. All algorithms use the relative change in the function value $f_{\mathbf{W}}$ in Eq. (2) as a stopping condition (set to 10^{-8}). In INDAFAC, the tolerance on the infinity norm of the gradient is set to 10^{-8} and the maximum number of iterations is set to 500. In CP-WOPT, the tolerance on the two-norm of the gradient divided by the number of entries in the gradient is set to 10^{-8} , the maximum number of iterations is set to 500, and the maximum number of function evaluations is set to 10000. In EM-ALS, the maximum number of iterations (equivalent to one function evaluation) is set to 10000.

All the methods under consideration are iterative methods. We used multiple starting points for each randomly generated problem. The first starting point is generated using the n -mode singular vectors² of \mathcal{X} with missing entries replaced by zero. In our preliminary experiments, this

² The n -mode singular vectors are the left singular vectors of $\mathbf{X}_{(n)}$ (\mathcal{X} unfolded in mode n).

starting procedure produced significantly better results than random initializations, even with large amounts of missing data. In order to improve the chances of reaching the global minimum, additional starting points were generated randomly. The same set of starting points was used by all the methods in the same order.

5.2. Validation metrics

If the true factors are known, then we can assess the recovery of the factors via the *factor match score* (FMS) defined as follows. Let the correct and computed factorizations be given by

$$\sum_{r=1}^R \lambda_r \mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)} \circ \dots \circ \mathbf{a}_r^{(N)} \quad \text{and} \quad \sum_{r=1}^{\bar{R}} \bar{\lambda}_r \bar{\mathbf{a}}_r^{(1)} \circ \bar{\mathbf{a}}_r^{(2)} \circ \dots \circ \bar{\mathbf{a}}_r^{(N)},$$

respectively. Without loss of generality, we assume that all the vectors have been scaled to unit length and that the scalars are positive. Further, we assume $\bar{R} \geq R$ so that the computed solution has at least as many components as the true solution. (One could add all-zero components to the computed solution if $\bar{R} < R$.) Recall that there is a permutation ambiguity, so all possible matchings of components between the two solutions must be considered. Under these conditions, the FMS is defined as

$$\text{FMS} = \max_{\sigma \in \Pi(R, \bar{R})} \frac{1}{\bar{R}} \sum_{r=1}^{\bar{R}} \left(1 - \frac{|\lambda_r - \bar{\lambda}_{\sigma(r)}|}{\max\{\lambda_r, \bar{\lambda}_{\sigma(r)}\}} \right) \prod_{n=1}^N |a_r^{(n)T} \bar{a}_{\sigma(r)}^{(n)}|. \quad (9)$$

If $\bar{R} = R$, then the set $\Pi(R, \bar{R})$ is all permutations of 1 to R ; otherwise, it is all possible permutations of all $\binom{\bar{R}}{R}$ mappings of $\{1, \dots, \bar{R}\}$ to $\{1, \dots, R\}$. The FMS can be between 0 and 1, and the best possible FMS is 1. If $\bar{R} \geq R$, some components in the computed solution are completely ignored in the calculation of FMS, but this is not an issue for us because we use $\bar{R} = R$ throughout.

We also consider the problem of recovering missing data. Let \mathcal{X} be the original data and let $\hat{\mathcal{X}}$ be the tensor that is produced by the computed model. Then the *tensor completion score* (TCS) is

$$\text{TCS} = \frac{\|(1-\mathcal{W}) * (\mathcal{X} - \hat{\mathcal{X}})\|}{\|(1-\mathcal{W}) * \mathcal{X}\|}. \quad (10)$$

In other words, the TCS is the relative error in the missing entries. TCS is always nonnegative, and the best possible score is 0.

5.3. Simulated data

We consider the performance of the methods on moderately-sized problems of sizes $50 \times 40 \times 30$, $100 \times 80 \times 60$, and $150 \times 120 \times 90$. For all sizes, we set the number of components in the CP model to be $R = 5$. We test 60%, 70%, 80%, 90%, and 95% missing data. The experiments show that the underlying factors can be captured even if the CP model is fit to a tensor with a significant amount of missing data; this is because the low-rank structure of the tensor is being exploited. A rank- R CP model for a tensor of size $I \times J \times K$ has $R(I+J+K)$ variables. The reason that the factors can be recovered accurately, even with 95% missing data, is that there is still a lot more data than variables, i.e., the size of the data is equal to $0.05IJK$ which is much greater than the $R(I+J+K)$ variables for large values of I, J , and K . Because it is a nonlinear problem, we do not know exactly how many data entries are needed in order to recover a CP model of a low-rank tensor; however, a lower bound for the number of entries needed in the matrix case has been derived in [12].

5.3.1. Generating simulated data

We create the test problems as follows. Assume that the tensor size is $I \times J \times K$ and that the number of factors is R . We generate factor

matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} of sizes $I \times R$, $J \times R$, and $K \times R$, respectively, by randomly choosing each entry from $\mathcal{N}(0, 1)$ and then normalizing every column to unit length. Note that this method of generating the factor matrices ensures that the solution is unique with probability one for the sizes and number of components that we are considering because $R \ll \min\{I, J, K\}$.³

We then create the data tensor as

$$\mathcal{X} = [\mathbf{A}, \mathbf{B}, \mathbf{C}] + \eta \frac{[\mathcal{X}]}{[\mathcal{M}]}.$$

Here \mathcal{M} is a noise tensor (of the same size as \mathcal{X}) with entries randomly selected from $\mathcal{N}(0, 1)$, and η is the noise parameter. We use $\eta = 10\%$ in our experiments.

Finally, we set some entries of each generated tensor to be missing according to an indicator tensor \mathcal{W} . We consider two situations for determining the missing data: randomly missing entries and, in Appendix B, structured missing data in the form of randomly missing fibers. In the case of randomly missing entries, \mathcal{W} is a binary tensor such that exactly $\lfloor MIJK \rfloor$ randomly selected entries are set to zero, where $M \in (0, 1)$ defines the percentage of missing data. We require that every slice of \mathcal{W} (in every direction) have at least one nonzero because otherwise we have no chance of recovering the factor matrices; this is related to the problem of *coherence* in the matrix completion problem where it is well-known that missing an entire row (or a column) of a matrix means that the true factors can never be recovered. For each problem size and missing data percentage, thirty independent test problems are generated.

5.3.2. Results

In our experiments, CP-WOPT, INDAFAC, and EM-ALS were indistinguishable in terms of accuracy. Fig. 4 show box plots⁴ of the FMS scores for problems of size $50 \times 40 \times 30$. Results for problems of size $100 \times 80 \times 60$ and $150 \times 120 \times 90$ are provided in Appendix A. The plots summarize the FMS scores of the computed factors for 30 independent problems for each problem size and percentage of missing data: the median values are shown as black dots inside a colored circle, the 25th and 75th percentile values are indicated by the top and bottom edges of the solid bars extending from the median values, and the outliers (corresponding to values more than 2.7 standard deviations from the means) are shown as smaller colored circles. A lack of solid bars extending from the median indicates that the 25th and 75th percentile values are equal to the median. The results shown are cumulative (i.e., the best so far) across the multiple starts, with the first start using the n -mode singular vectors and the remaining starts using random vectors.

For all methods, additional starting points improve performance; however, no one method clearly requires fewer starting points than the others. In general, using two or three starting points suffices to solve all problems to high accuracy in terms of FMS. For problems with up to 90% missing data, this is illustrated by median values which are all very close to one; further, the number of outliers decreases with multiple starts.

Perhaps contrary to intuition, we note that smaller problems are generally more difficult than larger problems for the same percentage of missing data. Thus, we can see in the bottom plot of Fig. 4 that the FMS scores remain low even with multiple starts for problems of size

³ Since each matrix has R linearly independent columns with probability one and the k -rank (i.e., the maximum value k such that any k columns are linearly independent) of each factor matrix is R , we satisfy the necessary conditions defined by Kruskal [34] for uniqueness.

⁴ The box plots were generated using the `boxplot` command in the Statistics Toolbox of MATLAB. The plots shown here use the “compact” plot style for that command. For more details on box plots and the use of this MATLAB command, see <http://www.mathworks.com/access/helpdesk/help/toolbox/stats/boxplot.html>.

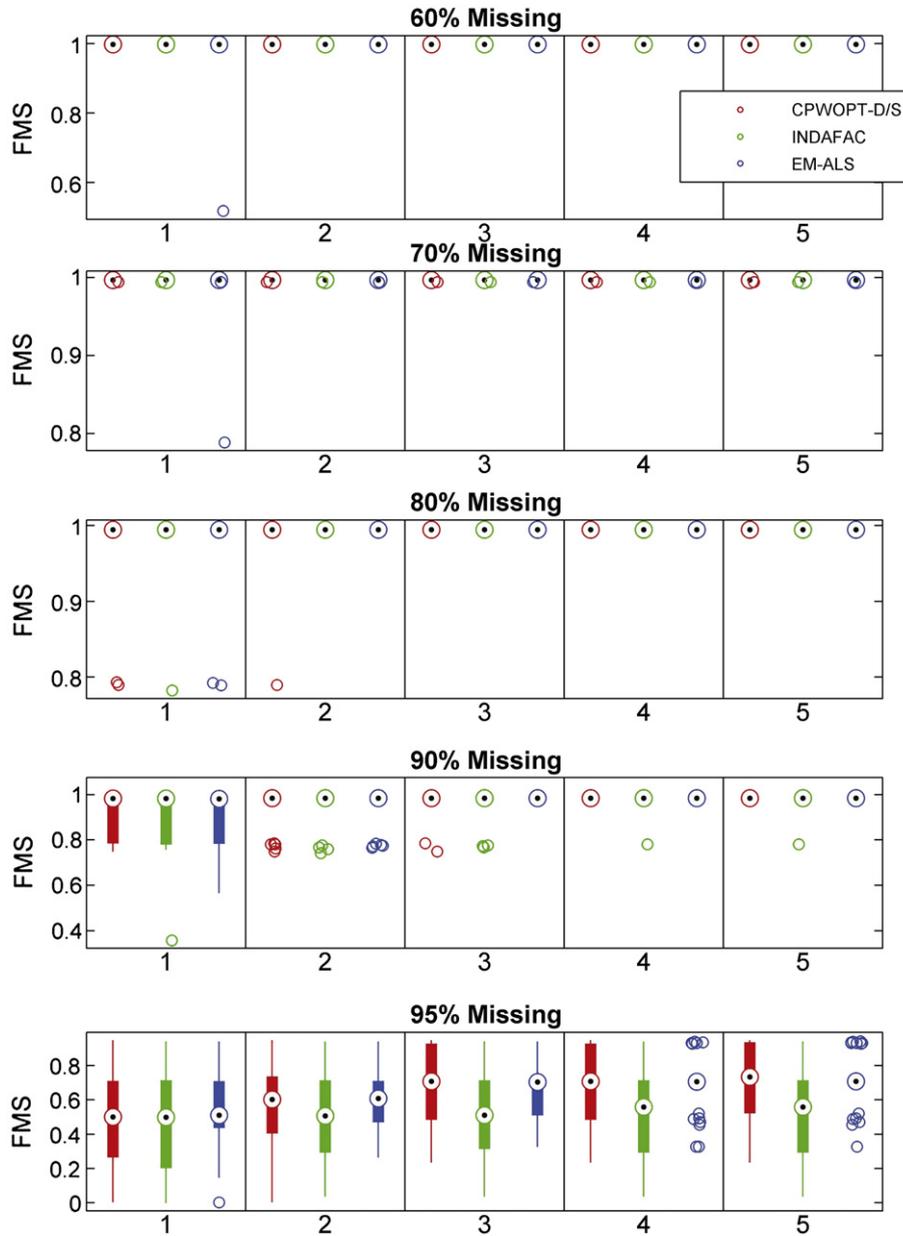


Fig. 4. Comparison of accuracy for problems of size $50 \times 40 \times 30$ for randomly missing entries with different amounts of missing data. The methods are CP-WOPT (red), INDAFAC (green), and EM-ALS (blue). We plot cumulative results for multiple starts. The first start is based on the n -mode singular vectors, and the remaining starts are random.

$50 \times 40 \times 30$ with 95% missing data. It may be that some of these problems have gone below the lower bound on the amount of data needed to find a solution; as mentioned previously, a lower bound on how much data is needed is known for the matrix case [12] but finding such a bound is still an open problem for higher-order tensors. We can define the ratio ρ as

$$\rho = \frac{\text{Number of known tensor entries}}{\text{Number of variables}} = \frac{(1-M)JK}{R(I+J+K)}, \quad (11)$$

where smaller values of ρ indicate more difficult problems. For example, problems with 95% missing data of size $50 \times 40 \times 30$ ($\rho = 5$) are more difficult than those of size $100 \times 80 \times 60$ ($\rho = 20$), as illustrated in Fig. 4 and Fig. A.10, respectively. In the latter figure, corresponding to the larger size (and thus larger value of ρ), nearly all problems are solved to high accuracy by all methods. Since this is a

nonlinear problem, ρ does not tell the entire story, but it is at least a partial indicator of problem difficulty.

The differentiator between the methods is computational time required. Fig. 5 shows a comparison of the sparse and dense versions of CP-WOPT, along with INDAFAC and EM-ALS. The timings reported in the figure are the sum of the times for all starting points. The y-axis is time in seconds on a log scale. We present results for time rather than iterations because the iterations for INDAFAC are much more expensive than those for CP-WOPT or EM-ALS. We make several observations. For missing data levels up to 80%, the dense version of CP-WOPT is faster than INDAFAC, but EM-ALS is the fastest overall. For 90% and 95% missing data, the sparse version of CP-WOPT is fastest (by more than a factor of 10 in some cases) with the exception of the case of 90% missing data for problems of size $150 \times 120 \times 90$. However, for this latter case, the difference is not significant; furthermore, there are two problems where EM-ALS required more than 10 times the median to find a correct factorization.

As demonstrated in earlier studies [1,7], problems with structured missing data are generally more difficult than those with randomly missing values. We present results for problems of structured missing data in Appendix B. The set-up is the same as we have used here, except that we only go up to 90% missing data because the accuracy of all methods degrades significantly after that point. The results indicate that accuracies of all methods start to diminish at lower percentages of structured missing data than for those problems with randomly missing data. Timings of the methods also indicate similar behavior in the case of structure missing data: with 80% or less missing data EM-ALS is fastest, whereas the sparse version of CP-WOPT is in general faster when there is 90% missing data.

Because it ignores missing values (random or structured), a major advantage of CP-WOPT is its ability to perform sparse computations, enabling it to scale to very large problems. Neither INDAFAC nor EM-ALS can do this. The difficulty with EM-ALS is that it must impute all missing values; therefore, it loses any speed or scalability advantage of

sparsity since it ultimately operates on dense data. Although INDAFAC also ignores missing values, its scalability is limited due to the expense of solving the Newton equation. Even though the Newton system is extremely sparse, we can see that the method is expensive even for moderate-sized problems such as those discussed in this section. In the next section, we consider very large problems, demonstrating this strength of CP-WOPT.

5.4. Large-scale simulated data

A unique feature of CP-WOPT is that it can be applied to problems that are too large to fit in memory if dense storage were used. We consider two situations:

- (a) $500 \times 500 \times 500$ with 99% missing data (1.25 million known values), and

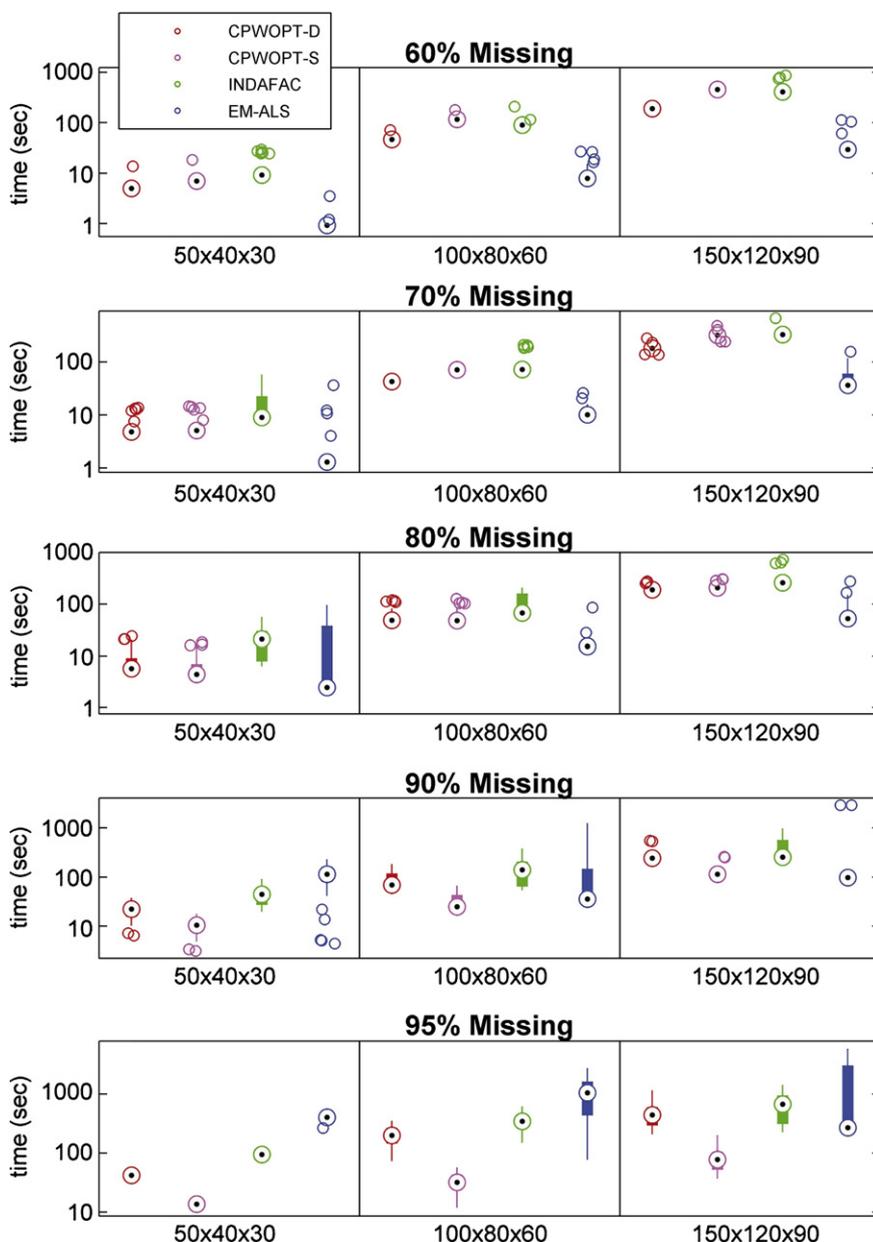


Fig. 5. Comparison of timings for randomly missing entries for the total time for all starting points for different problem sizes and different amounts of missing data. The methods are CP-WOPT-Dense (red), CP-WOPT-Sparse (magenta), INDAFAC (green), and EM-ALS (blue).

(b) $1000 \times 1000 \times 1000$ with 99.5% missing data (5 million known values).

5.4.1. Generating large-scale simulated data

The method for generating test problems of this size is necessarily different than the one for smaller problems presented in the previous section because we cannot, for example, generate a full noise tensor. In fact, the tensors $\mathcal{Y} = [\mathbf{A}, \mathbf{B}, \mathbf{C}]$ (the noise-free data) and \mathcal{N} (the noise) are never explicitly fully formed for the large problems studied here.

Let the tensor size be $I \times J \times K$ and the missing value rate be M . We generate the factor matrices as described in Section 5.3, using $R=5$ components. Next, we create the set \mathcal{I} with $(1-M)IJK$ randomly generated indices; this set represents the indices of the known (non-missing) values in our tests. The binary indicator tensor \mathcal{W} is stored as a sparse tensor [28] and defined by

$$w_{ijk} = \begin{cases} 1 & \text{if } (i,j,k) \in \mathcal{I}, \\ 0 & \text{otherwise.} \end{cases}$$

Rather than explicitly forming all of $\mathcal{Y} = [\mathbf{A}, \mathbf{B}, \mathbf{C}]$, we only calculate its values for those indices in \mathcal{I} . This is analogous to the calculations described for Step 2 of Fig. 3b. All missing entries of \mathcal{Y} are set to zero, and \mathcal{Y} is stored as a sparse tensor. Finally, we set

$$\mathcal{X} = \mathcal{Y} + \eta \frac{[\mathcal{X}]}{[\mathcal{N}]} \mathcal{N},$$

where \mathcal{N} is a sparse noise tensor such that

$$w_{ijk} = \begin{cases} \mathcal{N}(0, 1) & \text{if } (i,j,k) \in \mathcal{I}, \\ 0 & \text{otherwise.} \end{cases}$$

For each problem size, ten independent test problems are generated. The initial guess is generated by computing the n -mode singular vectors of the tensor with missing values filled in by zeros.

5.4.2. Results for CP-WOPT on large-scale data

The computational set-up was the same as that described in Section 5.1 except that the tolerance for the two-norm of the gradient divided by the number of entries in the gradient was set to 10^{-10} (previously 10^{-8}). The results across all ten runs for each problem size are shown in Fig. 6.

In the $500 \times 500 \times 500$ case, storing a dense version of the tensor (assuming 8 bytes per entry) would require exactly 1 GB. Storing just 1% of the data (1.25 M entries) in sparse format (assuming 32 bytes per entry, 8 for the value and 8 for each of the indices) requires only 40 MB. In our randomly generated experiments, all ten problems were solved with an FMS score greater than 0.99, with solve times ranging between 200 and 1000 s.

In the $1000 \times 1000 \times 1000$ case, storing a dense version of the tensor would require 8 GB. Storing just 0.5% of the data as a sparse tensor (5 M entries) requires 160 MB. In our randomly generated experiments, only nine of the ten problems were solved with an FMS score greater than 0.99. The solve times ranged from 2000 to 10,000 s, approximately 10 times slower than the $500 \times 500 \times 500$ case, which had half as many variables and 1/4 the nonzero entries. In the failed case, the optimization method exited because the relative change in the function tolerance was smaller than the tolerance. We restarted the optimization method with no changes except that we use the solution that had been computed previously as the initial guess. After 985 s of additional work, an answer was found with an FMS score of 0.9999.

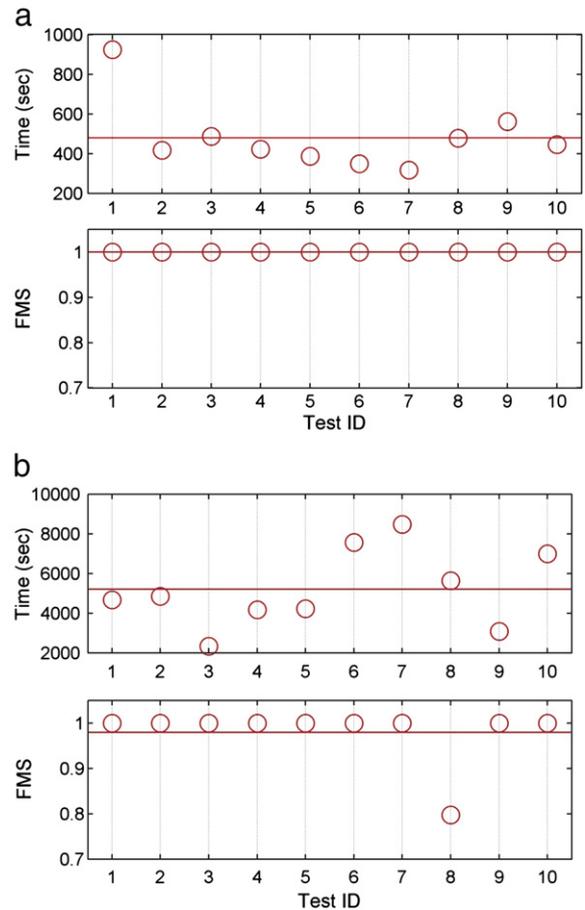


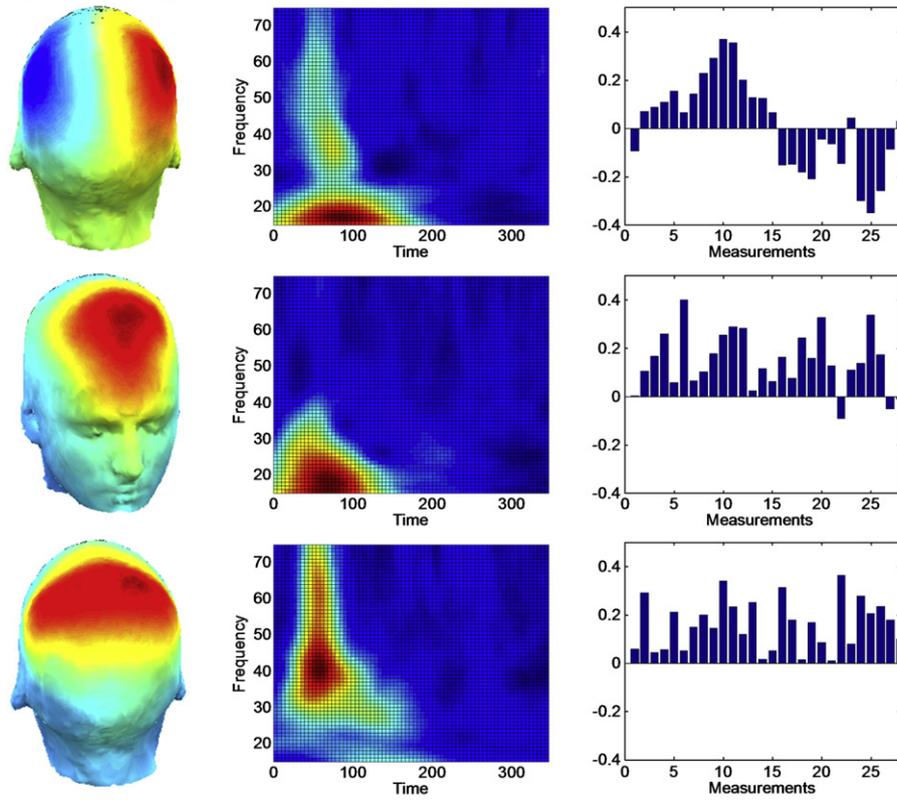
Fig. 6. Results of CP-WOPT-Sparse for large-scale problems with at least 99% missing data. The means are shown as solid lines.

5.5. EEG data

In this section, we demonstrate the use of CP-WOPT in multi-channel EEG (electroencephalogram) analysis by capturing the underlying brain dynamics even in the presence of missing signals. We use an EEG data set collected to observe the gamma activation during proprioceptive stimuli of left and right hand [19]. The data set contains multi-channel signals (64 channels) recorded from 14 subjects during stimulation of left and right hand (i.e., 28 measurements in total). For each measurement, the signal from each channel is represented in both time and frequency domains using a continuous wavelet transform and then vectorized (forming a vector of length 4392); in other words, each measurement is represented by a channels by time-frequency matrix. The data for all measurements are then arranged as a channels by time-frequency by measurements tensor of size $64 \times 4392 \times 28$. For details about the data, see [19].

We model the data using a CP model with $R=3$, denoting \mathbf{A} , \mathbf{B} , and \mathbf{C} as the extracted factor matrices corresponding to the channels, time-frequency, and measurements modes, respectively. We demonstrate the columns of the factor matrices in each mode in Fig. 7a. The 3-D head plots correspond to the columns of \mathbf{A} , i.e., coefficients corresponding to the channels ranging from low values in blue to high values in red. The time-frequency domain representations correspond to the columns of \mathbf{B} rearranged as a matrix and again ranging from low values in blue to high values in red. The bar plots represent the columns of \mathbf{C} . Note that the three rows of images in Fig. 7a (3-D head plot, matrix plot, bar plot) correspond to columns $r=1, 2, 3$ of the factor matrices (\mathbf{A} , \mathbf{B} , \mathbf{C}), respectively. Observe that the first row of images highlights the differences between left and right hand

a) No missing entries



b) 30 channels missing per measurement

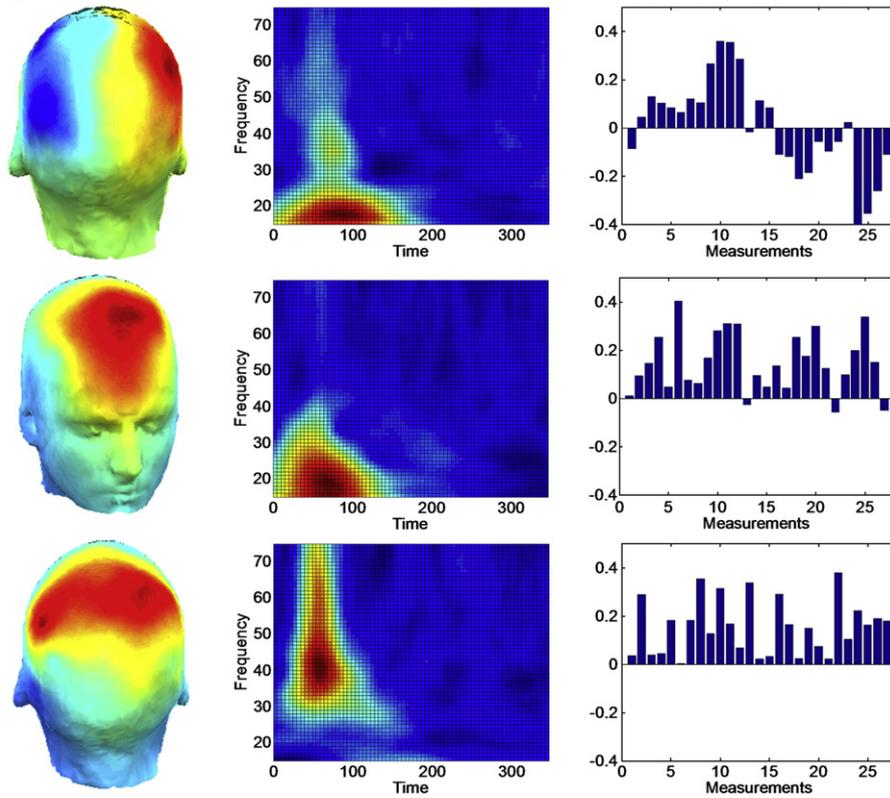


Fig. 7. Columns of the CP factor matrices (A, B and C with $R=3$) extracted from the EEG data arranged as a channels by time-frequency by measurements tensor. The 3-D head images were drawn using EEGLab [35].

Table 1

EEG analysis with incomplete data: the similarity between the columns of **A**, **B**, **C** and the columns of factor matrices extracted by CP-WOPT. The similarity is measured in terms of FMS defined in (9).

Number of missing channels	CP-WOPT
1	0.9959
10	0.9780
20	0.9478
30	0.8949
40	0.6459

stimulation while the second and third rows of images pertain to frontal and parietal activations that are shared by the stimuli. Unlike [19], we do not use nonnegativity constraints; we convert tensor entries from complex to real values by using the absolute values of the entries and center the data across the channels mode before the analysis.

It is not uncommon in EEG analysis that the signals from some channels are ignored due to malfunctioning of the electrodes. This corresponds to missing fibers in the tensor (as in Fig. 1) when we arrange the data as described above. To reflect such cases of missing data, we randomly set data for one or more of the 64 channels for each measurement to be missing, center the tensor across the channels mode ignoring the missing entries, and then fit a CP model with $R=3$ to the resulting data using CP-WOPT. Let $\bar{\mathbf{A}}$, $\bar{\mathbf{B}}$, $\bar{\mathbf{C}}$ be the factor matrices extracted from a tensor with missing entries using the CP-WOPT algorithm.

Table 1 presents the average FMS scores (over 50 instances of the data with randomly missing channels) for different numbers of missing channels per measurement. As the number of missing channels increases, the average FMS scores decrease as expected. However, even with up to 30 missing channels per measurement, or about 47% of the data, the extracted factor matrices match with the original factor matrices well, with average FMS scores around 0.90. Fig. 7b presents the images for $\bar{\mathbf{A}}$, $\bar{\mathbf{B}}$ and $\bar{\mathbf{C}}$ corresponding to those for **A**, **B**, **C** in Fig. 7a. We see that the underlying brain dynamics are still captured even when 30 channels per measurement are missing; only slight local differences can be observed between the original data and the model generated using CP-WOPT.

It can be argued that the activations of the electrodes are highly correlated and even if some of the electrodes are removed, the underlying brain dynamics can still be captured. However, in these experiments we do not set the same channels to missing for each measurement; the channels are randomly missing from each measurement. There are cases when a CP model may not be able to

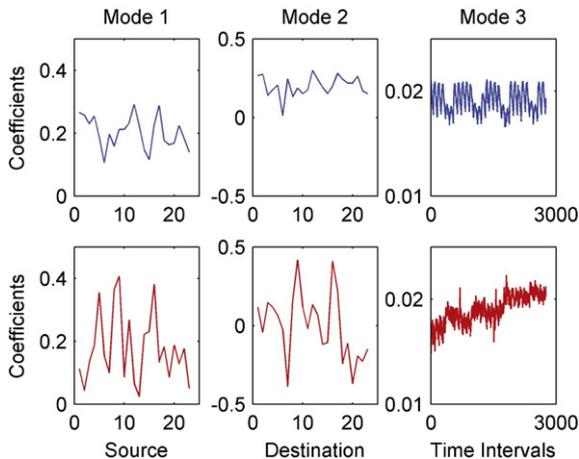


Fig. 8. Factor matrices extracted from Géant data using a 2-component CP model. The first row illustrates the first column of the factor matrices in each mode and the second row shows the second column of the factor matrices.

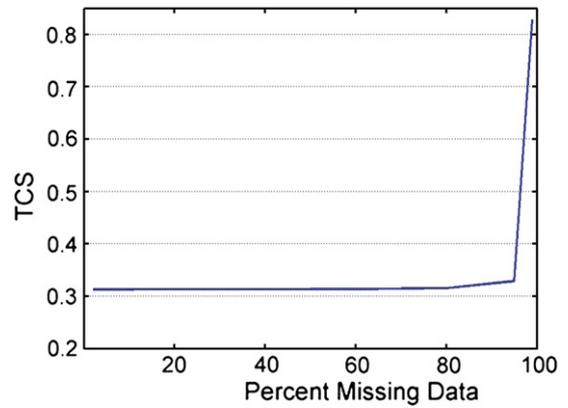


Fig. 9. Tensor completion score for different amounts of missing data for Géant data when a 2-component CP model fitted using CP-WOPT is used to recover the missing entries.

recover factors when data has certain patterns of missing values, e.g., missing the signals from the same side of the brain for all measurements. However, such data is not typical in practice.

Table 1 shows the FMS scores for varying amounts of missing data. The solution does not seriously degrade until 40 of 64 channels are removed.

5.6. Network traffic data

In the previous section, we were interested in tensor factorizations in the presence of missing data. In this section, we address the problem of recovering missing entries of a tensor; in other words, the *tensor completion problem*. One application domain where this problem is frequently encountered is computer network traffic analysis. Network traffic data consists of traffic matrices (TM), which record the amount of network data exchanged between source and destination pairs (computers, routers, etc.). Since TMs evolve over time, network data can be represented as a tensor. For instance, Géant data [36] records the traffic exchanged between 23 routers over a period of 4 months collected using 15-minute intervals. This data set forms a third-order tensor with *source routers*, *destination routers* and *time* modes, and each entry indicates the amount of traffic sent from a source to a destination during a particular time interval. Missing data arises due to the expense of the data collection process.

In our study, we used the Géant data collected in April⁵ stored as a tensor of size $23 \times 23 \times 2756$. Let \mathcal{X} represent this raw data tensor. In order to adjust for scaling bias in the amount of traffic, we preprocess \mathcal{X} as $\mathcal{X} = \log(\mathcal{X} + 1)$. Fig. 8 presents the results of 2-component CP model (i.e., $R=2$); the first and second row correspond the first and second column of the factor matrices of the model, respectively. This CP model fits the data well but not perfectly and there is some unexplained variation left in the residuals. Let $\hat{\mathcal{X}}$ be the tensor constructed using the computed factor matrices. The modeling error, defined as $\frac{\|\mathcal{X} - \hat{\mathcal{X}}\|}{\|\mathcal{X}\|}$, is approximately 0.31 for the 2-component CP model computed. Even though extracting more components slightly lowers the modeling error, models with more components do not look appropriate for the data.

In order to assess the performance of CP-WOPT in terms of recovering missing data, randomly chosen entries of \mathcal{X} are set to missing and a 2-component CP model is fit to the altered data. The extracted factor matrices are then used to reconstruct the data and fill in the missing entries. These recovered values are compared with the

⁵ The data was incomplete for the full four-month period represented in the data, e.g., 6 days of data was missing at the end of February. We used a subset of the data corresponding to a period with no missing time slices.

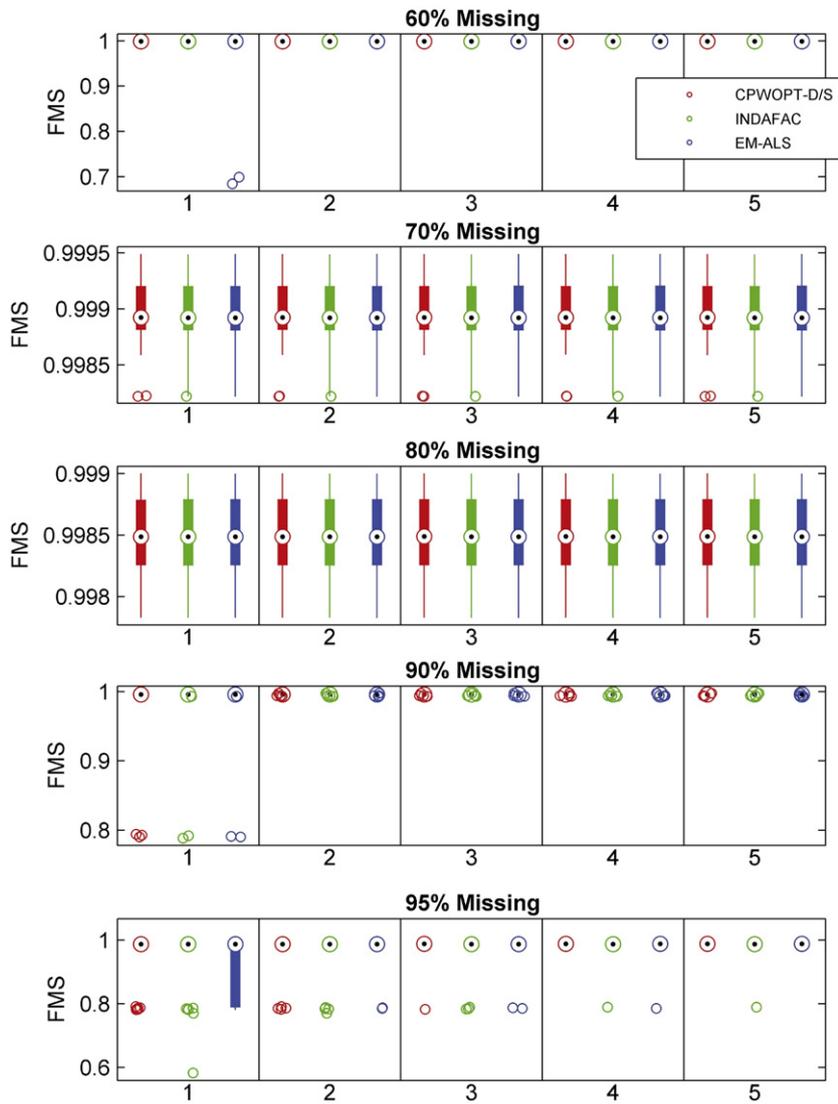


Fig. A.10. Comparison of accuracy for problems of size $100 \times 80 \times 60$ for randomly missing entries with different amounts of missing data. The methods are CP-WOPT (red), INDAFAC (green), and EM-ALS (blue). We plot cumulative results for multiple starts. The first start is based on the n -mode singular vectors, and the remaining starts are random.

actual values based on the tensor completion score (TCS) defined in Eq. (10). Fig. 9 presents the average TCS (across 30 instances) for different amounts of missing data. We observe that the average TCS is around 0.31 when there is little missing data and increases very slowly as we increase the amount of missing data. The average TCS is only slightly higher, approximately 0.33, even when 95% of the entries are missing. However, the average TCS increases sharply when the amount of missing data is 99%. Note that even if there is no missing data, there will be completion error due to the modeling error, i.e., 0.31. Fig. 9 demonstrates the robustness of CP-WOPT, illustrating that the TCS can be kept close to the level of the modeling error using this method, even when the amount of missing data is high.

We have addressed the tensor completion problem using a CP model, which gives easily-interpretable factors in each mode. However, for the tensor completion problem, the recovery of the missing entries is more important than the underlying factors and their interpretability. Therefore, modeling the data using a restricted model like CP may not be the best approach in this case. It may be possible to achieve lower modeling error using a more flexible model such as a Tucker model [37]. If this model is fit to the data using algorithms akin to CP-WOPT, missing entries may be recovered more accurately. This is a topic of future research.

6. Conclusions

The closely related problems of matrix factorizations with missing data and matrix completion have recently been receiving a lot of attention. In this paper, we consider the more general problem of tensor factorization in the presence of missing data, formulating the canonical tensor decomposition for incomplete tensors as a weighted least squares problem. Unlike imputation-based techniques, this formulation ignores the missing entries and models only the known data entries. We develop a scalable algorithm called CP-WOPT using gradient-based optimization to solve the weighted least squares formulation of the CP problem. The code has been released in version 2.4 of the Tensor Toolbox for MATLAB [29].

Our numerical studies suggest that the proposed CP-WOPT approach is accurate and scalable. CP-WOPT can recover the underlying factors successfully with large amounts of missing data, e.g., 90% missing entries for tensors of size $50 \times 40 \times 30$. We have also studied how CP-WOPT can scale to problems of larger sizes, e.g., $1000 \times 1000 \times 1000$, and recover CP factors from large, sparse tensors with 99.5% missing data. Moreover, through the use of both dense and sparse implementations of the algorithm, we have shown that CP-WOPT was always faster in our studies when compared to the

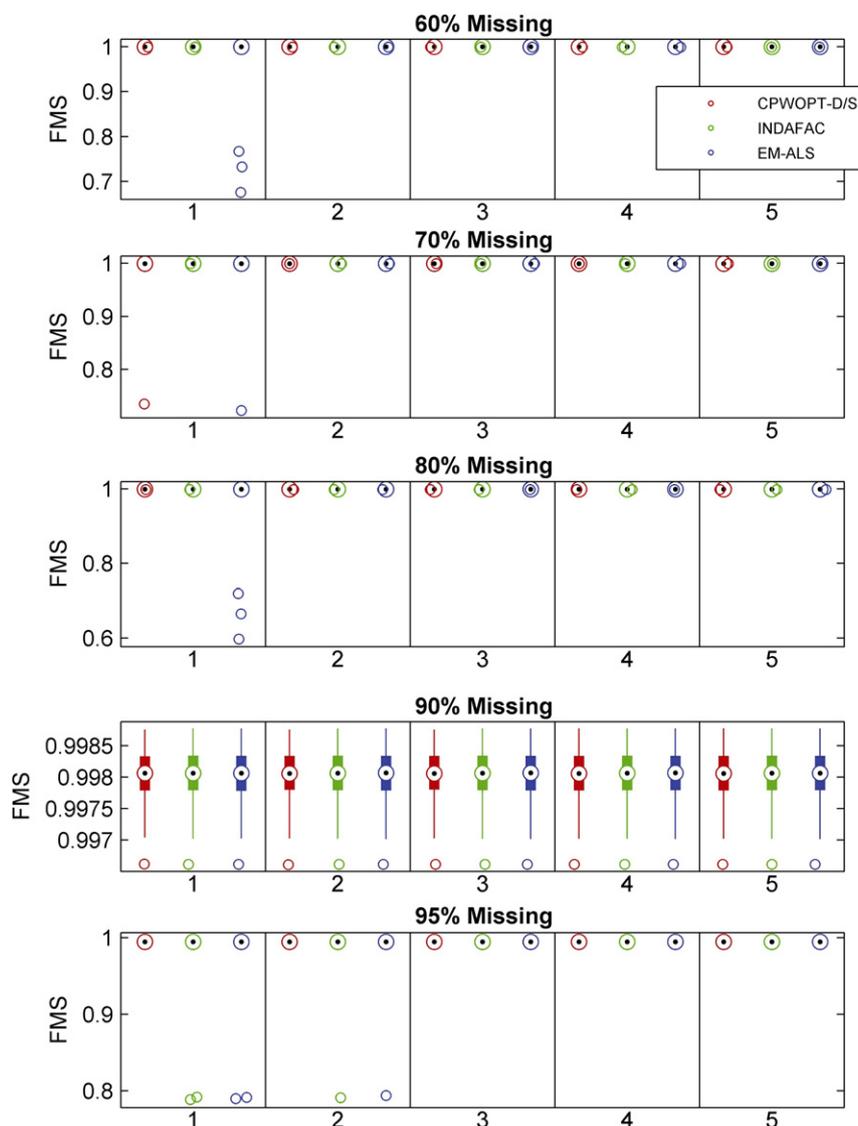


Fig. A.11. Comparison of accuracy for problems of size $150 \times 120 \times 90$ for randomly missing entries with different amounts of missing data. The methods are CP-WOPT (red), INDAFAC (green), and EM-ALS (blue). We plot cumulative results for multiple starts. The first start is based on the n -mode singular vectors, and the remaining starts are random.

best alternative approach based on second-order optimization (INDAFAC) and even faster than EM-ALS for high percentages of missing data.

We consider the practical use of CP-WOPT algorithm in two different applications which demonstrate the effectiveness of CP-WOPT even when there may be low correlation with a multi-linear model. In multi-channel EEG analysis, the factors extracted by the CP-WOPT algorithm can capture brain dynamics even if signals from some channels are missing, suggesting that practitioners can now make better use of incomplete data in their analyses. We note that the EEG data was centered by using the means of only the known entries; however, robust techniques for centering incomplete data are needed and is a topic of future investigation. In network traffic analysis, CP-WOPT algorithm can be used in the context of tensor completion and recover the missing network traffic data.

Although 90% or more missing data may not seem practical, there are situations where it is useful to purposely omit data. For example, if the data is very large, excluding some data will speed up the computation and enable it to fit in memory (for large-scale data like the example of a $1000 \times 1000 \times 1000$ tensor that would normally require 8 GB of storage). It is also common to leave out data for the purpose of rank determination or model assessment via cross-

validation. For these reasons, it is useful to consider high degrees of missing data.

In future studies, we plan to extend our results in several directions. We will include constraints such as non-negativity and penalties to encourage sparsity, which enable us to find more meaningful latent factors from large-scale sparse data. We will also consider the scalability of criss-cross regression [11,23] for fitting the CP model. Finally, we will consider the problem of collective factorizations with missing data, where we are jointly factoring multiple tensors with shared factors.

Acknowledgments

We thank David Gleich and the anonymous referees for helpful comments which greatly improved the presentation of this manuscript.

This work was funded by the Laboratory Directed Research & Development (LDRD) program at Sandia National Laboratories, a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

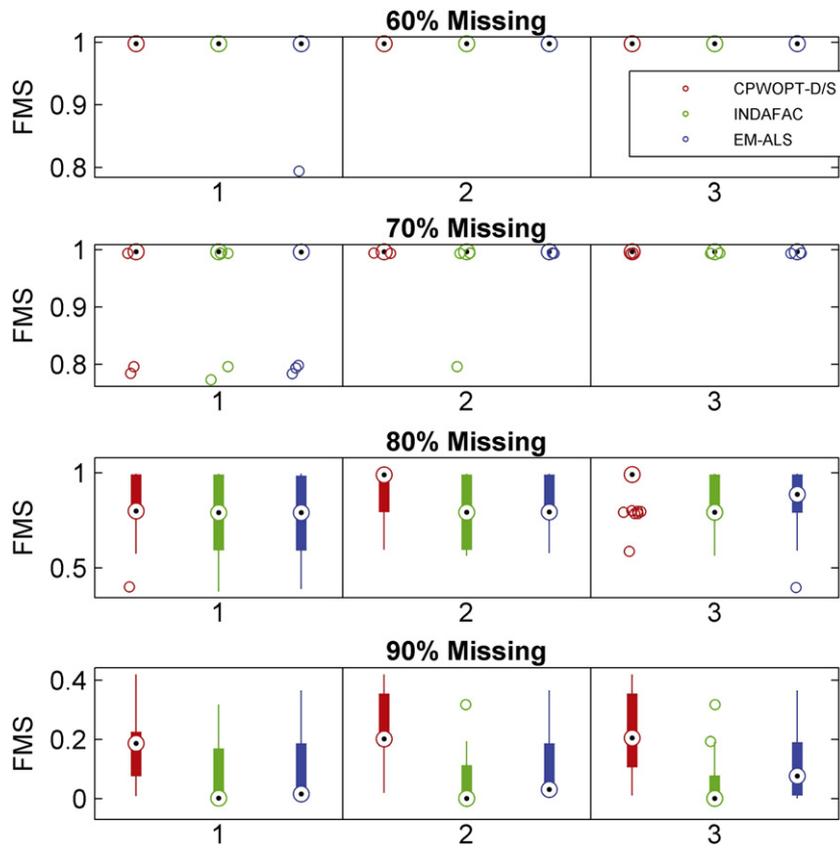


Fig. B.12. Comparison of accuracy for problems of size $50 \times 40 \times 30$ for randomly missing fibers with different amounts of missing data. The methods are CP-WOPT (red), INDAFAC (green), and EM-ALS (blue). We plot cumulative results for multiple starts. The first start is based on the n -mode singular vectors, and the remaining starts are random.

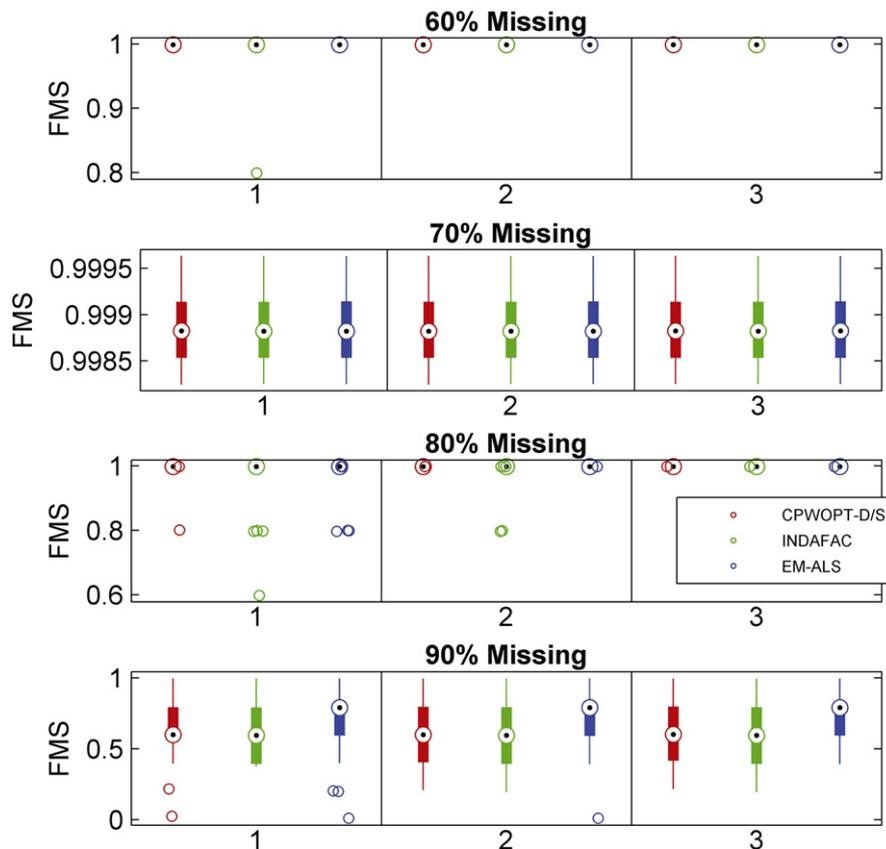


Fig. B.13. Comparison of accuracy for problems of size $100 \times 80 \times 60$ for randomly missing fibers with different amounts of missing data. The methods are CP-WOPT (red), INDAFAC (green), and EM-ALS (blue). We plot cumulative results for multiple starts. The first start is based on the n -mode singular vectors, and the remaining starts are random.

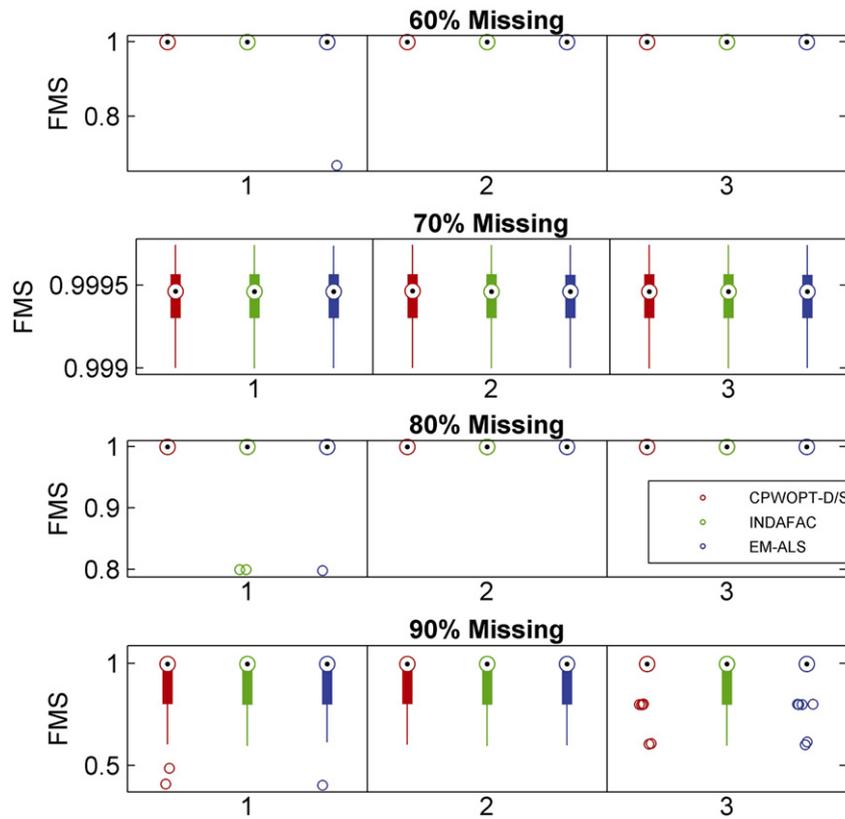


Fig. B.14. Comparison of accuracy for problems of size $150 \times 120 \times 90$ for randomly missing fibers with different amounts of missing data. The methods are CP-WOPT (red), INDAFAC (green), and EM-ALS (blue). We plot cumulative results for multiple starts. The first start is based on the n-mode singular vectors, and the remaining starts are random.

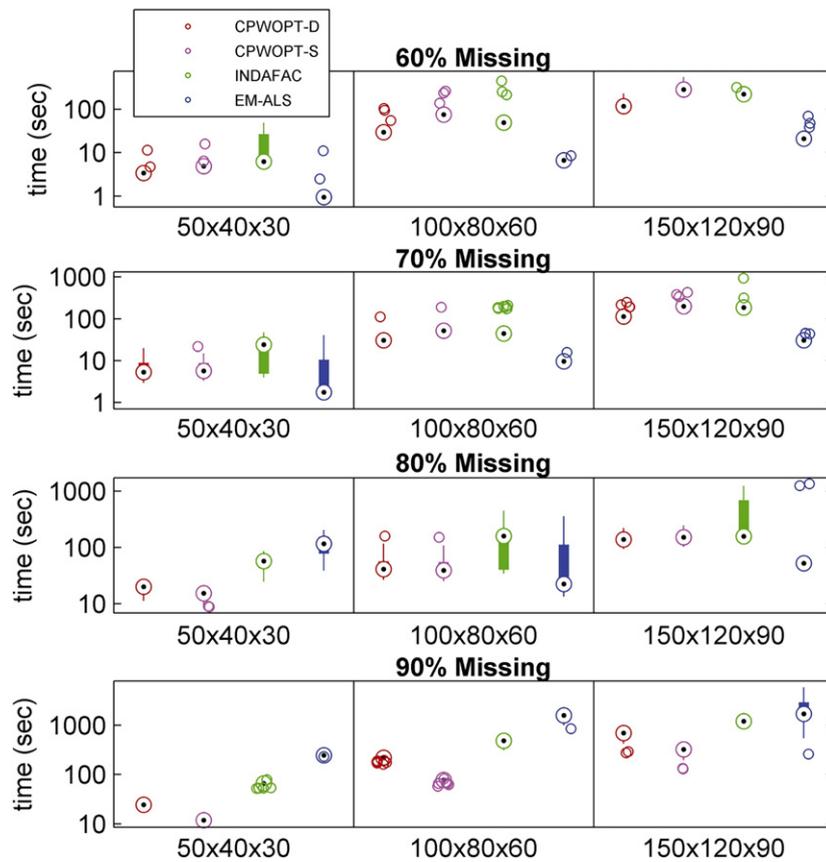


Fig. B.15. Comparison of timings for randomly missing fibers for the total time for all starting points for different problem sizes and different amounts of missing data. The methods are CP-WOPT-Dense (red), CP-WOPT-Sparse (magenta), INDAFAC (green), and EM-ALS (blue).

Appendix A. Additional comparisons for randomly missing data

Figs. A.10 and A.11 show box plots of the FMS scores for problems of size $100 \times 80 \times 60$ and $150 \times 120 \times 90$, respectively. See Section 5.3 and Fig. 4 for a detailed description of the experimental setup and the results for analogous problems of size $50 \times 40 \times 30$. We can see from these figures that the three methods being compared (CP-WOPT, INDAFAC, EM-ALS) are indistinguishable in terms of accuracy, as was the case for the smaller sized problems. Furthermore, comparable improvements using multiple starts can be seen for these larger problems; we are able to solve a majority of the problems with very high accuracy using five or fewer starting points (the n -mode singular vectors used as the basis for the first run and random points used for the others).

Appendix B. Comparisons for structured missing data

We use the same experimental set-up as described in Section 5.3 with the exception that we performed experiments using problems with up to 90% missing data (as opposed to 95%) and we only show results for three starting points (the first is the n -mode singular vectors and the remaining two are random). In the case of randomly missing fibers, without loss of generality, we consider missing fibers in the third mode only. In that case, \mathbf{W} is an $I \times J$ binary matrix with exactly $[MIJ]$ randomly selected entries are set to zero, and the binary tensor \mathcal{W} is created by stacking K copies of \mathbf{W} together. We again require that every slice of \mathcal{W} (in every direction) have at least one nonzero, which is equivalent to requiring that \mathbf{W} has no zero rows or columns. For each problem size and missing data percentage, thirty independent test problems were generated.

The average FMS results are presented in Figs. B.12–B.14 for problems of sizes $50 \times 40 \times 30$, $100 \times 80 \times 60$ and $150 \times 120 \times 90$, respectively. As for the problems with randomly missing data, we see that as the percentage of missing data increases, the average FMS scores tend to decrease. However, one notable difference is that the average FMS scores for structured missing data are in general lower than those for problems with comparable amounts of randomly missing data. This indicates that problems with structured missing data are more difficult than those with randomly missing data.

The computational times required to compute the CP models using the different methods are present in Fig. B.15, where we observe comparable results to those presented in Fig. B.15 for the problems with randomly missing data. With 80% or less structured missing data EM-ALS is fastest, whereas the sparse version of CP-WOPT is in general faster when there is more than 80% missing data.

References

- [1] E. Acar, D.M. Dunlavy, T.G. Kolda, M. Mørup, Scalable Tensor Factorizations with Missing Data, Proceedings of the Tenth SIAM International Conference on Data Mining, SIAM, 2010, pp. 701–712. URL http://www.siam.org/proceedings/data-mining/2010/dm10_061_acare.pdf.
- [2] Y. Zhang, M. Roughan, W. Willinger, L. Qiu, Spatio-temporal compressive sensing and internet traffic matrices, SIGCOMM '09: Proceedings of the ACM SIGCOMM 2009 conference on Data communication, ACM, New York, NY, USA, 2009, pp. 267–278. doi:10.1145/1592568.1592600.
- [3] A.M. Buchanan, A.W. Fitzgibbon, Damped Newton algorithms for matrix factorization with missing data, CVPR'05: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2, IEEE Computer Society, 2005, pp. 316–322. doi:10.1109/CVPR.2005.118.
- [4] E. Acar, B. Yener, Unsupervised multiway data analysis: a literature survey, IEEE Transactions on Knowledge and Data Engineering 21 (1) (2009) 6–20. doi:10.1109/TKDE.2008.112.
- [5] T.G. Kolda, B.W. Bader, Tensor decompositions and applications, SIAM Review 51 (3) (2009) 455–500. doi:10.1137/07070111X.
- [6] F. Miwakeichi, E. Martínez-Montes, P.A. Valdés-Sosa, N. Nishiyama, H. Mizuhara, Y. Yamaguchi, Decomposing EEG data into space-time-frequency components using parallel factor analysis, NeuroImage 22 (3) (2004) 1035–1045. doi:10.1016/j.neuroimage.2004.03.039.
- [7] G. Tomasi, R. Bro, PARAFAC and missing values, Chemometrics and Intelligent Laboratory Systems 75 (2) (2005) 163–180. doi:10.1016/j.chemolab.2004.07.003.
- [8] V.Y. Orekhov, I. Ibragimov, M. Billeter, Optimizing resolution in multidimensional NMR by three-way decomposition, Journal of Biomolecular NMR 27 (2003) 165–173. doi:10.1023/A:1024944720653.
- [9] X. Geng, K. Smith-Miles, Z.-H. Zhou, L. Wang, Face Image Modeling by Multilinear Subspace Analysis with Missing Values, MM '09: Proceedings of the seventeen ACM international conference on Multimedia, ACM, 2009, pp. 629–632. doi:10.1145/1631272.1631373.
- [10] A. Ruhe, Numerical Computation of Principal Components when Several Observations are Missing, Tech. Rep. UMINF-48-74, Department of Information Processing, Institute of Mathematics and Statistics, University of Umea, Umea, Sweden, 1974.
- [11] K.R. Gabriel, S. Zamir, Lower rank approximation of matrices by least squares approximation with any choice of weights, Technometrics 21 (4) (1979) 489–498 URL <http://www.jstor.org/stable/1268288>.
- [12] E. J. Candès, T. Tao, The power of convex relaxation: Near-optimal matrix completion, arXiv:0903.1476v1 (Mar. 2009). URL <http://arxiv.org/abs/0903.1476>.
- [13] E. J. Candès, Y. Plan, Matrix completion with noise, arXiv:0903.3131v1 (Mar. 2009). URL <http://arxiv.org/abs/0903.3131>.
- [14] 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. doi:10.1007/BF02310791.
- [15] 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, pp. 1–84, available at <http://www.psychology.uwo.ca/faculty/harshman/wppfac0.pdf>.
- [16] T.G. Kolda, B.W. Bader, J.P. Kenny, Higher-Order Web Link Analysis Using Multilinear Algebra, ICDM 2005: Proceedings of the 5th IEEE International Conference on Data Mining, IEEE Computer Society, 2005, pp. 242–249. doi:10.1109/ICDM.2005.77.
- [17] R. Bro, Review on multiway analysis in chemistry—2000–2005, Critical Reviews in Analytical Chemistry 36 (3–4) (2006) 279–293. doi:10.1080/10408340600969965.
- [18] E. Acar, C.A. Bingol, H. Bingol, R. Bro, B. Yener, Multiway analysis of epilepsy tensors, Bioinformatics 23 (13) (2007) i10–i18. doi:10.1093/bioinformatics/btm210.
- [19] M. Mørup, L.K. Hansen, S.M. Arnfred, ERPWAVELAB a toolbox for multi-channel analysis of time-frequency transformed event related potentials, Journal of Neuroscience Methods 161 (2) (2007) 361–368. doi:10.1016/j.jneumeth.2006.11.008.
- [20] E. Acar, D.M. Dunlavy, T.G. Kolda, A scalable optimization approach for fitting canonical tensor decompositions, Journal of Chemometrics 25 (2011) 67–86. doi:10.1002/cem.1335.
- [21] A.P. Dempster, N.M. Laird, D.B. Rubin, Maximum likelihood from incomplete data via the em algorithm, Journal of the Royal Statistical Society, Series B 39 (1) (1977) 1–38.
- [22] N. Srebro, T. Jaakkola, Weighted Low-Rank Approximations, IMCL-2003: Proceedings of the Twentieth International Conference on Machine Learning, 2003, pp. 720–727.
- [23] H.A.L. Kiers, Weighted least squares fitting using ordinary least squares algorithms, Psychometrika 62 (2) (1997) 215–266. doi:10.1007/BF02295279.
- [24] R. Bro, Multi-way analysis in the food industry: Models, algorithms, and applications, Ph.D. thesis, University of Amsterdam, available at <http://www.models.kvl.dk/research/theses/> (1998).
- [25] B. Walczak, D.L. Massart, Dealing with missing data: Part I, Chemometrics and Intelligent Laboratory Systems 58 (1) (2001) 15–27. doi:10.1016/S0169-7439(01)00131-9.
- [26] P. Paatero, A weighted non-negative least squares algorithm for three-way “PARAFAC” factor analysis, Chemometrics and Intelligent Laboratory Systems 38 (2) (1997) 223–242. doi:10.1016/S0169-7439(97)00031-2.
- [27] J. Nocedal, S.J. Wright, Numerical Optimization, Springer, 1999.
- [28] B.W. Bader, T.G. Kolda, Efficient MATLAB computations with sparse and factored tensors, SIAM Journal on Scientific Computing 30 (1) (2007) 205–231. doi:10.1137/060676489.
- [29] B. W. Bader, T. G. Kolda, MATLAB tensor toolbox version 2.4, <http://csmr.ca.sandia.gov/tgkolda/TensorToolbox/> (last accessed March, 2010).
- [30] J.J. Moré, D.J. Thuente, Line search algorithms with guaranteed sufficient decrease, ACM Transactions on Mathematical Software 20 (3) (1994) 286–307. doi:10.1145/192115.192132.
- [31] D.M. Dunlavy, T.G. Kolda, E. Acar, Poblano v1.0: A MATLAB Toolbox for Gradient-Based Optimization, Tech. Rep. SAND2010-1422, Sandia National Laboratories, Albuquerque, Mar. 2010, NM and Livermore, CA.
- [32] C.A. Andersson, R. Bro, The N-way toolbox for MATLAB, Chemometrics and Intelligent Laboratory Systems 52 (1) (2000) 1–4 see also <http://www.models.kvl.dk/source/nwaytoolbox/>. doi:10.1016/S0169-7439(00)00071-X.
- [33] G. Tomasi, Incomplete data PARAFAC (INDAFAC), <http://www.models.kvl.dk/source/indafac/index.asp> (last accessed May, 2009).
- [34] J.B. Kruskal, Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics, Linear Algebra and its Applications 18 (2) (1977) 95–138. doi:10.1016/0024-3795(77)90069-6.
- [35] A. Delorme, S. Makeig, EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics, Journal of Neuroscience Methods 134 (2004) 9–21. doi:10.1016/j.jneumeth.2003.10.009.
- [36] S. Uhlig, B. Quoitin, J. Lepropre, S. Balon, Providing public intradomain traffic matrices to the research community, ACM SIGCOMM Computer Communication Review 36 (1) (2006) 83–86. doi:10.1145/1111322.1111341.
- [37] L.R. Tucker, Some mathematical notes on three-mode factor analysis, Psychometrika 31 (1966) 279–311. doi:10.1007/BF02289464.