

# Numerical optimization for symmetric tensor decomposition

Tamara G. Kolda<sup>1</sup>

Received: 1 October 2014 / Accepted: 17 February 2015 / Published online: 11 April 2015  
© Springer-Verlag Berlin Heidelberg and Mathematical Optimization Society 2015

**Abstract** We consider the problem of decomposing a real-valued symmetric tensor as the sum of outer products of real-valued vectors. Algebraic methods exist for computing complex-valued decompositions of symmetric tensors, but here we focus on real-valued decompositions, both unconstrained and nonnegative, for problems with low-rank structure. We discuss when solutions exist and how to formulate the mathematical program. Numerical results show the properties of the proposed formulations (including one that ignores symmetry) on a set of test problems and illustrate that these straightforward formulations can be effective even though the problem is nonconvex.

**Keywords** Symmetric · Outer product · Canonical polyadic · Tensor decomposition · Completely positive · Nonnegative

**Mathematics Subject Classification** 90C30 · 15A69

## 1 Introduction

We consider the problem of decomposing a real-valued symmetric tensor as the sum of outer products of real-valued vectors. Let  $\mathcal{A}$  represent an  $m$ -way,  $n$ -dimension

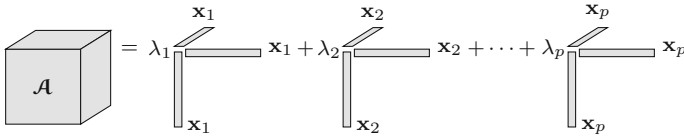
---

This material is based upon work supported by the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

---

✉ Tamara G. Kolda  
tgkolda@sandia.gov

<sup>1</sup> Sandia National Laboratories, Livermore, CA, USA



**Fig. 1** Symmetric tensor factorization for  $m = 3$

symmetric tensor. Given a real-valued vector  $\mathbf{x}$  of length  $n$ , we let  $\mathbf{x}^m$  denote the  $m$ -way,  $n$ -dimensional symmetric outer product tensor such that  $(\mathbf{x}^m)_{i_1 i_2 \dots i_m} = x_{i_1} x_{i_2} \dots x_{i_m}$ . Comon et al. [15] showed that any real-valued symmetric tensor  $\mathcal{A}$  can be decomposed as

$$\mathcal{A} = \sum_{k=1}^p \lambda_k \mathbf{x}_k^m, \tag{1}$$

with  $\lambda_k \in \mathbb{R}$  and  $\mathbf{x}_k \in \mathbb{R}^n$  for  $k = 1, \dots, p$ ; see the illustration in Fig. 1. We assume that the tensor is low-rank, i.e.,  $p$  is small relative to the typical rank of a random tensor. We survey the methods that have been proposed for related problems and discuss several optimization formulations, including a surprisingly effective method that *ignores* the symmetry.

We also consider the related problem of decomposing a real-valued nonnegative symmetric tensor as the sum of outer products of real-valued nonnegative vectors. Let  $\mathcal{A} \geq 0$  represent an  $m$ -way,  $n$ -dimension nonnegative symmetric tensor. In this case, the goal is a factorization of the form

$$\mathcal{A} = \sum_{k=1}^p \mathbf{x}_k^m \quad \text{with} \quad \mathbf{x}_k \geq 0. \tag{2}$$

If such a factorization exists, we say that  $\mathcal{A}$  is completely positive [39]. If such a factorization does not exist, then we propose to solve a “best fit” problem instead.

The paper is structured as follows. Section 2 provides notation and background material. Related decompositions, including the best symmetric rank-1 approximation, the symmetric Tucker decomposition, partially symmetric decompositions, and the complex-valued canonical decompositions are discussed in Sect. 3. We describe two optimization formulations for symmetric decomposition in Sect. 4, and a mathematical program for the nonnegative problem in Sect. 5. Numerical results, including the methodology for generating challenging problems, is presented in Sect. 6. Finally, Sect. 7 discusses our findings and future challenges.

## 2 Background

### 2.1 Notation and preliminaries

A tensor is a multidimensional array. The number of ways or modes is called the *order* of a tensor. For example, a matrix is a tensor of order two. Tensors of order three or greater are called *higher-order* tensors.

Let  $n_1 \times n_2 \times \dots \times n_m$  denote the size of an  $m$ -way tensor. We say that the tensor is *cubic* if all the modes have the same size, i.e.,  $n = n_1 = n_2 \dots = n_m$ . In other words, “cubic” is the tensor generalization of “square.” In this case, we refer to  $n$  as the *dimension* of the tensor. We let  $\mathbb{R}^{[m,n]}$  denote the space of all cubic real-valued tensors of order  $m$  and dimension  $n$ . As appropriate, we use *multiindex* notation to compactly index tensors so that  $\mathbf{i} = (i_1, i_2, \dots, i_m)$ . Thus,  $a_{\mathbf{i}}$  denotes  $a_{i_1 i_2 \dots i_m}$ .

The norm of a tensor  $\mathcal{A} \in \mathbb{R}^{[m,n]}$  is the square root of the sum of squares of its elements, i.e.,

$$\|\mathcal{A}\| = \sqrt{\sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_m=1}^n a_{\mathbf{i}}^2}.$$

Unless otherwise noted, all norms are the (elementwise)  $\ell_2$ -norm.

### 2.2 Symmetric tensors

A tensor is symmetric if its entries do not change under permutation of the indices. Formally, we let  $\pi(m)$  denote the set of permutations of length  $m$ . For instance,

$$\pi(3) = \{(1, 2, 3), (1, 3, 2), (2, 1, 3), (2, 3, 1), (3, 1, 2), (3, 2, 1)\}.$$

It is well known that  $|\pi(m)| = m!$ . We say a real-valued  $m$ -way  $n$ -dimensional tensor  $\mathcal{A}$  is *symmetric* [15] if

$$a_{i_{p(1)} \dots i_{p(m)}} = a_{i_1 \dots i_m} \quad \text{for all } i_1, \dots, i_m \in \{1, \dots, n\} \text{ and } p \in \pi(m).$$

Such tensors are also sometimes referred to as *supersymmetric*. For a 3-way tensor  $\mathcal{A}$  of dimension  $n$ , symmetry means

$$a_{ijk} = a_{ikj} = a_{jik} = a_{kij} = a_{kji} \quad \text{for all } i, j, k \in \{1, \dots, n\}.$$

We let  $\mathbb{S}^{[m,n]} \subset \mathbb{R}^{[m,n]}$  denote the subspace of all symmetric tensors.

### 2.3 Symmetric outer product tensors

A tensor in  $\mathbb{S}^{[m,n]}$  is called *rank one* if it has the form  $\lambda \mathbf{x}^m$  where  $\lambda \in \mathbb{R}$  and  $\mathbf{x} \in \mathbb{R}^n$ . If  $m$  is odd or  $\lambda > 0$ , then the  $m$ th real root of  $\lambda$  always exists, so we can rewrite the tensor as

$$\lambda \mathbf{x}^m = \mathbf{y}^m \quad \text{where } \mathbf{y} = \left(\sqrt[m]{\lambda}\right) \mathbf{x}. \tag{3}$$

If  $m$  is even, however, the  $m$ th real root does not exist if  $\lambda < 0$ , so the scalar cannot be absorbed as in (3).

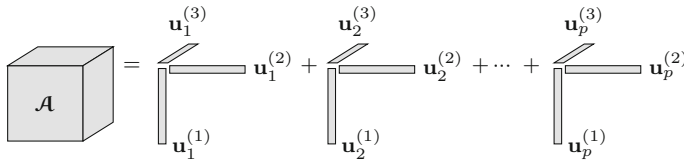


Fig. 2 CP tensor factorization for  $m = 3$

### 2.4 Model parameters

For the symmetric decomposition, we let  $\lambda$  denote the vector of weights and  $\mathbf{X}$  denote the matrix of component vectors, i.e.,

$$\lambda = [\lambda_1 \ \lambda_2 \ \dots \ \lambda_p]^T \quad \text{and} \quad \mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_p].$$

The notation  $x_{ik}$  refers to the  $i$ th entry in the  $k$ th column, so recalling the multiindex notation  $\mathbf{i} = (i_1, \dots, i_m)$ , we have

$$(\mathbf{x}_k^m)_{\mathbf{i}} = x_{i_1 k} x_{i_2 k} \dots x_{i_m k}.$$

## 3 Related problems

### 3.1 Canonical polyadic tensor decomposition

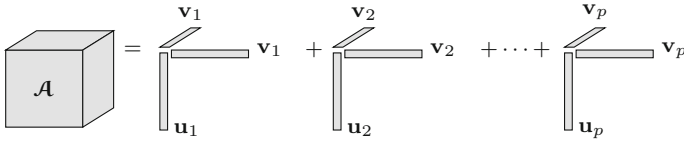
Canonical polyadic (CP) tensor decomposition has been known since 1927 [25, 26]. It is known under several names, two of the most prominent being CANDECOMP as proposed by Carroll and Chang [14] and PARAFAC by Harshman [24]. Originally, the term CP was proposed as a combination of these two names [29], but more recently has been re-purposed to mean “canonical polyadic.” For details on CP, we refer the reader to the survey [31]. Here, we describe the problem in the case of a cubic tensor  $\mathcal{A} \in \mathbb{R}^{[m, n]}$ . Our goal is to discover a decomposition of the form

$$\mathcal{A} = \sum_{k=1}^p \mathbf{u}_k^{(1)} \circ \mathbf{u}_k^{(2)} \circ \dots \circ \mathbf{u}_k^{(m)}. \tag{4}$$

The circle denotes the vector outer product so the  $\mathbf{i} = (i_1, i_2, \dots, i_m)$  entry is

$$(\mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \dots \circ \mathbf{u}^{(m)})_{\mathbf{i}} = u_{i_1}^{(1)} u_{i_2}^{(2)} \dots u_{i_m}^{(m)}.$$

Each summand is called a *component*. An illustration is shown in Fig. 2. One of the most effective methods for this problem is alternating least squares (ALS). We solve for each *factor* matrix



**Fig. 3** INDSCAL tensor factorization for  $m = 3$

$$\mathbf{U}^{(j)} = \left[ \mathbf{u}_1^{(j)} \quad \mathbf{u}_2^{(j)} \quad \dots \quad \mathbf{u}_p^{(j)} \right],$$

in turn by solving a linear least squares problem, cycling through all modes (i.e.,  $j = 1, \dots, m$ ) repeatedly until convergence. See, e.g., [31, Figure 3.3] for details.

### 3.2 Canonical decomposition with partial symmetry

Partial symmetry has been considered since the work of Carroll and Chang [14]. At the same time Carroll and Chang [14] introduced CANDECOMP, they also defined INDSCAL which assumes two modes are symmetric. For simplicity of discussion, we assume a cubic tensor  $\mathcal{A} \in \mathbb{R}^{[m,n]}$ . For  $m = 3$  and the last two dimensions being symmetric, this means

$$a_{ijk} = a_{ikj} \quad \text{for all } i, j, k \in \{1, \dots, n\},$$

and the factorization should be of the form

$$\mathcal{A} = \sum_{k=1}^p \mathbf{u}_k \circ \mathbf{v}_k \circ \mathbf{v}_k.$$

In other words, the last two vectors in each component are equal. An illustration is provided in Fig. 3.

Carroll and Chang [14] proposed to use an alternating method that ignores symmetry, with the idea that it will often converge to a symmetric solution (up to diagonal scaling). Later work showed that not all KKT points satisfy this condition [18]. In §4.7, we show how a generalization of this method can be surprisingly effective for symmetric tensor decomposition and provide some motivation for why this might be the case.

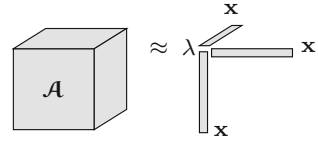
We also note that the methods proposed in this manuscript can be extended to partial symmetries.

### 3.3 Best symmetric rank-1 approximation

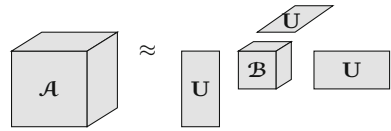
The best symmetric rank-1 approximation problem is

$$\min \|\mathcal{A} - \lambda \mathbf{x}^m\|^2 \quad \text{subject to } \lambda \in \mathbb{R}, \mathbf{x} \in \mathbb{R}^n. \tag{5}$$

**Fig. 4** Best symmetric rank-1 decomposition for  $m = 3$



**Fig. 5** Symmetric Tucker decomposition for  $m = 3$



An illustration is shown in Fig. 4. This problem was first considered in De Lathauwer et al. [17], but their proposed symmetric higher-order power method was not convergent. The power method has been improved so that it is convergent in subsequent work [30,32,33,41].

This problem is directly related to the problem of computing tensor Z-eigenpairs. A pair  $(\lambda, \mathbf{x})$  is a Z-eigenpair [34,38] of a tensor  $\mathcal{A} \in \mathbb{S}^{[m,n]}$  if

$$\mathcal{A}\mathbf{x}^{m-1} = \lambda\mathbf{x} \quad \text{and} \quad \|\mathbf{x}\| = 1,$$

where  $\mathcal{A}\mathbf{x}^{m-1}$  denotes a vector in  $\mathbb{R}^n$  such that

$$\left(\mathcal{A}\mathbf{x}^{m-1}\right)_{i_1} = \sum_{i_2} \cdots \sum_{i_m} a_{i_1 i_2 \cdots i_m} x_{i_2} \cdots x_{i_m} \quad \text{for } i_1 \in \{1, \dots, n\}.$$

The problems are related because any Karush-Kuhn-Tucker (KKT) point of (5) is a Z-eigenpair of  $\mathcal{A}$ ; see, e.g., [32].

Han [23] has considered an unconstrained optimization formulation of the problem (5). Cui et al. [16] use Jacobian SDP relaxations in polynomial optimization to find *all* real eigenvalues sequentially, from the largest to the smallest. Nie and Wang [36] consider semidefinite relaxations.

### 3.4 Symmetric Tucker decomposition

A related problem is symmetric Tucker decomposition. Here the goal is to find an orthogonal matrix  $\mathbf{U} \in \mathbb{R}^{n \times p}$  and a symmetric tensor  $\mathcal{B} \in \mathbb{S}^{[m,p]}$  that solves

$$\min \|\mathcal{A} - \hat{\mathcal{A}}\|^2 \quad \text{subject to } \hat{\mathbf{a}}_{\mathbf{i}} = \sum_{j_1=1}^p \sum_{j_2=1}^p \cdots \sum_{j_m=1}^p b_{j_1 j_2 \cdots j_m} u_{i_1 j_1} u_{i_2 j_2} \cdots u_{i_m j_m}.$$

An illustration is shown in Fig. 5. This topic has been considered in [13,27,40] and is useful for compression and signal processing applications. Alas, the computational techniques are quite different, so we do not consider them further in this work.

### 3.5 Complex-valued symmetric tensor decomposition

An alternative version of the problem allows a complex decomposition, i.e.,

$$\mathcal{A} = \sum_{k=1}^p \mathbf{x}_k^m \quad \text{with } \mathbf{x}_k \in \mathbb{C}^n \text{ for } k = 1, \dots, p. \tag{6}$$

Techniques from algebraic geometry have been proposed to solve (6) in [10–12,37]. More recently, Nie [35] devised has a combination of algebraic and numerical approaches for solving this problem. Generally, these approaches do not scale to large  $n$ , though Nie’s numerical method scales much better than previous approaches.

In the complex case, the typical rank (i.e., with probability one) is given by the theorem below. To the best of our knowledge, for the real case, no analogous results are known [15].

**Theorem 1** (Alexander-Hirschowitz [4,15]) *For  $m > 2$ , the typical symmetric rank (over  $\mathbb{C}$ ) of an order- $m$  symmetric tensor of dimension  $n$  is*

$$\left\lceil \frac{1}{n} \binom{n+k-1}{k} \right\rceil$$

except for  $(m, n) \in \{(3, 5), (4, 3), (4, 4), (4, 5)\}$  where it should be increased by one.

## 4 Optimization formulations for symmetric tensor decomposition

### 4.1 Index multiplicities

A tensor  $\mathcal{A} \in \mathbb{S}^{[m,n]}$  has  $n^m$  entries, but not all are distinct. Let the set of all possible indices be denoted by

$$\mathcal{R} = \{(i_1, \dots, i_m) \mid i_1, \dots, i_m \in \{1, \dots, n\}\}.$$

Clearly,  $|\mathcal{R}| = n^m$ .

Following [9], we define an *index class* as a set of tensor indices such that the corresponding tensor entries all share a value due to symmetry. For example, for  $m = 3$  and  $n = 2$ , the tensor indices  $(1, 1, 2)$  and  $(1, 2, 1)$  are in the same index class since  $a_{112} = a_{121}$ . For each index class, we specify an *index representation* which is an index such that the entries are in nondecreasing order. For instance,  $(1, 1, 2)$  is the index representation for the index class that includes  $a_{121}$ . The set

$$\mathcal{I} = \{(i_1, \dots, i_m) \mid i_1, \dots, i_m \in \{1, \dots, n\} \text{ and } i_1 \leq i_2 \leq \dots \leq i_m\} \subset \mathcal{R}$$

denotes all possible index representations.

Each index class also has a *monomial representation* [9]. For each  $\mathbf{i} \in \mathcal{I}$  there is a corresponding monomial representation  $\mathbf{c}$  such that

$$x_{i_1}x_{i_2}\cdots x_{i_m} = x_1^{c_1}x_2^{c_2}\cdots x_n^{c_n}.$$

Specifically,  $c_j$  represents that number of occurrences of index  $j$  in  $\mathbf{i}$  for  $j = 1, \dots, n$ . Clearly,  $\sum_j c_j = m$ . Conversely, for a given  $\mathbf{c}$ , we build an index  $\mathbf{i}$  with  $c_1$  copies of 1,  $c_2$  copies of 2, etc. This results in an  $m$ -long index representation. The set of monomial representations is denoted by

$$\mathcal{C} = \{(c_1, \dots, c_n) \mid c_1, \dots, c_n \in \{0, \dots, m\} \text{ and } c_1 + \dots + c_n = m\}.$$

From [9], we have that the number of distinct entries of  $\mathcal{A}$  is given by

$$|\mathcal{I}| = |\mathcal{C}| = \binom{m+n-1}{m} = \frac{n^m}{m!} + O(n^{m-1}).$$

It can be shown [9] that the multiplicity of the entry corresponding to a monomial representation  $\mathbf{c}$  is

$$\sigma_{\mathbf{c}} = \binom{m}{c_1, c_2, \dots, c_n} = \frac{m!}{c_1! c_2! \cdots c_n!}. \tag{7}$$

Table 1 shows an example of index and monomial representations for  $\mathbb{S}^{[3,2]}$ , including the multiplicities of each element.

Without loss of generality, we exploit the one-to-one correspondence between index and monomial representations to change between them. For example,

$$\|\mathcal{A}\|^2 = \sum_{\mathbf{i} \in \mathcal{R}} a_{\mathbf{i}}^2 = \sum_{\mathbf{i} \in \mathcal{I}} \sigma_{\mathbf{i}} a_{\mathbf{i}}^2 = \sum_{\mathbf{c} \in \mathcal{C}} \sigma_{\mathbf{c}} a_{\mathbf{c}}^2,$$

and

$$(\mathbf{x}_k^m)_{\mathbf{i}} = x_{i_1 k} x_{i_2 k} \cdots x_{i_m k} = (\mathbf{x}_k^m)_{\mathbf{c}} = x_{1k}^{c_1} x_{2k}^{c_2} \cdots x_{nk}^{c_n}.$$

### 4.2 Two formulations

For given  $\mathcal{A} \in \mathbb{S}^{[m,n]}$  and  $p$ , our goal is to find  $\lambda$  and  $\mathbf{X}$  such that (1) is satisfied in a minimization sense. We consider two optimization formulations. The first formulation is the standard least squares formulation, i.e.,

**Table 1** Index and monomial representations for  $\mathbb{S}^{[3,2]}$

Index ( $\mathcal{I}$ )	Monomial ( $\mathcal{C}$ )	Multiplicity ( $\sigma$ )
(1,1,1)	(3,0)	1
(1,1,2)	(2,1)	3
(1,2,2)	(1,2)	3
(2,2,2)	(0,3)	1



$$f_1(\boldsymbol{\lambda}, \mathbf{X}) = \sum_{\mathbf{i} \in \mathcal{R}} \left( a_{\mathbf{i}} - \sum_{k=1}^p \lambda_k (\mathbf{x}_k^m)_{\mathbf{i}} \right)^2 = \sum_{\mathbf{i} \in \mathcal{I}} \sigma_{\mathbf{i}} \left( a_{\mathbf{i}} - \sum_{k=1}^p \lambda_k (\mathbf{x}_k^m)_{\mathbf{i}} \right)^2. \tag{8}$$

Observe that this counts each unique entry multiple times, according to its multiplicity. The second formulation counts each unique entry only once, i.e.,

$$f_2(\boldsymbol{\lambda}, \mathbf{X}) = \sum_{\mathbf{i} \in \mathcal{I}} \left( a_{\mathbf{i}} - \sum_{k=1}^p \lambda_k (\mathbf{x}_k^m)_{\mathbf{i}} \right)^2. \tag{9}$$

Either formulation can be expressed generically as

$$f_{\mathbf{w}}(\boldsymbol{\lambda}, \mathbf{X}) = \sum_{\mathbf{i} \in \mathcal{I}} w_{\mathbf{i}} \left( a_{\mathbf{i}} - \sum_{k=1}^p \lambda_k (\mathbf{x}_k^m)_{\mathbf{i}} \right)^2 = \sum_{\mathbf{i} \in \mathcal{I}} w_{\mathbf{i}} \delta_{\mathbf{i}}^2.$$

Choosing  $w_{\mathbf{i}} = \sigma_{\mathbf{i}}$  yields  $f_1$  whereas  $w_{\mathbf{i}} = 1$  yields  $f_2$ . The value  $\delta_{\mathbf{i}}$  denotes the difference between the model and the tensor at entry  $\mathbf{i}$ . Note that this formulation easily adapts to the case of missing data, i.e., missing data should have weight of zero in the optimization formulation [2,3].

### 4.3 Gradients

Using the generic formulation, the gradients are given by

$$\begin{aligned} \frac{\partial f_{\mathbf{w}}}{\partial \lambda_k} &= -2 \sum_{\mathbf{i} \in \mathcal{I}} w_{\mathbf{i}} \delta_{\mathbf{i}} (\mathbf{x}_k^m)_{\mathbf{i}}, \\ \frac{\partial f_{\mathbf{w}}}{\partial x_{jk}} &= -2\lambda_k \sum_{\mathbf{c} \in \mathcal{C}} c_j w_{\mathbf{c}} \delta_{\mathbf{c}} \left( x_{1k}^{c_1} \cdots x_{jk}^{c_j-1} \cdots x_{nk}^{c_n} \right). \end{aligned} \tag{10}$$

For  $f_1$ , we mention an alternate gradient expression because it is more efficient to compute for larger values  $n$  and  $m$ . The derivation follows [1], and the gradients are given by

$$\begin{aligned} \frac{\partial f_1}{\partial \lambda_k} &= -2\mathcal{A}\mathbf{x}_k^m + 2 \sum_{\ell=1}^p \lambda_{\ell} \left( \mathbf{x}_k^{\top} \mathbf{x}_{\ell} \right)^m, \\ \frac{\partial f_1}{\partial \mathbf{x}_k} &= -2m\lambda_k \mathcal{A}\mathbf{x}_k^{m-1} + 2m\lambda_k \sum_{\ell=1}^p \lambda_{\ell} \left( \mathbf{x}_k^{\top} \mathbf{x}_{\ell} \right)^{m-1} \mathbf{x}_{\ell}. \end{aligned} \tag{11}$$

This formulation does not easily accommodate missing data since  $\mathbf{w}$  is implicit.

### 4.4 Scaling ambiguity

Observe that either objective function suffers from scaling ambiguity. Suppose we have two equivalent models defined by

$$\sum_{k=1}^p \lambda_k \mathbf{x}_k^m = \sum_{k=1}^p \hat{\lambda}_k \hat{\mathbf{x}}_k^m,$$

related by a positive scaling vector  $\boldsymbol{\rho} \in \mathbb{R}_+^p$  such that

$$\hat{\lambda}_k = \rho_k^m \lambda_k \quad \text{and} \quad \hat{\mathbf{x}}_k = \mathbf{x}_k / \rho_k \quad \text{for} \quad k = 1, \dots, p.$$

To avoid this ambiguity, it is convenient to require  $\|\mathbf{x}_k\| = 1$  for all  $k$ . We could enforce this condition as an equality constraint, but instead we treat it as a exact penalty, i.e.,

$$p_\gamma(\mathbf{X}) = \gamma \sum_{k=1}^p \left( \mathbf{x}_k^\top \mathbf{x}_k - 1 \right)^2. \tag{12}$$

It is straightforward to observe that the gradient is given by

$$\frac{\partial p_\gamma}{\partial \mathbf{x}_k} = 4\gamma \left( \mathbf{x}_k^\top \mathbf{x}_k - 1 \right) \mathbf{x}_k.$$

In the experimental results, we see that choosing  $\gamma = 0.1$  appears to be adequate for enforcing the penalty.

### 4.5 Sparse component weights

We assume so far that  $p$  is known, but this is not always the case. One technique to get around this problem is to guess a large value for  $p$  and then add a sparsity penalty on  $\boldsymbol{\lambda}$ , the weight vector. Specifically, we can use an approximate  $\ell_1$  penalty of the form suggested by [42]:

$$p_{\alpha,\beta}(\boldsymbol{\lambda}) = \frac{\beta}{\alpha} \sum_{k=1}^p \log(1 + \exp(-\alpha\lambda_k)) + \log(1 + \exp(\alpha\lambda_k)) \approx \beta \|\boldsymbol{\lambda}\|_1$$

In this case, the gradient is

$$\frac{\partial p_{\alpha,\beta}}{\partial \lambda_k} = \beta \left[ (1 + \exp(-\alpha\lambda_k))^{-1} + (1 + \exp(\alpha\lambda_k))^{-1} \right].$$

Note that the  $\beta$  term is not part of the approximation but rather the weight of the penalization. In our experiments, the results are insensitive to the precise choices of  $\alpha$  and  $\beta$ .

### 4.6 Putting it all together

The final function to be optimized is

$$\hat{f}(\boldsymbol{\lambda}, \mathbf{X}) = f_w(\boldsymbol{\lambda}, \mathbf{X}) + p_\gamma(\mathbf{X}) + p_{\alpha,\beta}(\boldsymbol{\lambda}).$$

The choice of  $\mathbf{w}$  determines the choice of objective function. We can also set  $w_i = 0$  for any missing values. The choice of  $\gamma$  determines the weight of the penalty on the norm of the columns of  $\mathbf{X}$ . Since this constraint is easy to satisfy and mostly convenience, the exact choice of  $\gamma$  is not critical. We later show experiments with  $\gamma = 0$  and  $\gamma = 0.1$ , to contrast the difference between no penalty and a small penalty. (Increasing  $\gamma$  beyond 0.1 did not have any impact on the experiments.) The parameter  $\alpha$  determines the “steepness” of the approximate  $\ell_1$  penalty function, and the choice of  $\beta$  determines the weight of the sparsity-encouraging penalty. In [42], they start with a small value of  $\alpha$  and gradually increase it. In our experiments, we use fixed values  $\alpha = 10$ . The  $\beta$  term is the weight given to the penalty, which is usually determined heuristically; we use  $\beta = 0.1$  in our experiments.

### 4.7 Ignoring symmetry

Another approach to symmetric decomposition is to ignore the symmetry altogether and use a standard CP tensor decomposition method such as ALS [19, 31]; surprisingly, there are situations under which this non-symmetric method yields a symmetric solution.

Under mild conditions, the CP decomposition (4) is unique up to permutation and scaling of the components, i.e., *essentially unique*. Sidiropoulos and Bro [43, Theorem 3] have a general *a posteriori* result on the essential uniqueness of the CP decompositions for tensors. If we specialize this result to the symmetric case by assuming  $\mathbf{U}^{(j)} = \mathbf{X}$  for  $j = 1, \dots, m$ , the result says that a sufficient condition for the uniqueness of (4) is

$$2p + (m - 1) \leq m \text{ k-rank}(\mathbf{X}). \tag{13}$$

Here, the k-rank of the matrix  $\mathbf{X}$  is the largest number  $k$  such that every subset of  $k$  columns of  $\mathbf{X}$  is linearly independent. Table 2 shows sufficient k-rank’s for various values of  $m$  and  $p$ . For instance, if  $m = 3$  and  $p = 25$ , then  $\text{k-rank}(\mathbf{X}) \geq 18$  is sufficient for uniqueness. The table does not directly depend on  $n$ ; however, recall that  $\mathbf{X}$  is an  $n \times p$  matrix, so  $\text{k-rank}(\mathbf{X}) \leq \min\{n, p\}$ .

The importance of essential uniqueness is that the global solution of the unconstrained problem (4) is the same as for the symmetric problem (1) so long as  $\mathbf{X}$  satisfies (13). If we normalize the factors in (4) and, without loss of generality, ignore the permutation ambiguity, then uniqueness implies, for  $k = 1, \dots, p$ ,

**Table 2** Minimal k-rank( $\mathbf{X}$ ) sufficient for uniqueness of symmetric outer product factorization

Order ( $m$ )	Components ( $p$ )							
	2	3	4	5	10	25	50	100
3	2	3	4	4	8	18	34	68
4	2	3	3	4	6	14	26	51
5	2	2	3	3	5	11	21	41
6	2	2	3	3	5	10	18	35

$$\lambda_k = \pm \|\mathbf{u}_1^{(k)}\| \cdots \|\mathbf{u}_m^{(k)}\| \quad \text{and} \quad \mathbf{x}_k = \pm \mathbf{u}_k^{(1)} / \|\mathbf{u}_1^{(k)}\| = \cdots = \pm \mathbf{u}_k^{(m)} / \|\mathbf{u}_m^{(k)}\|$$

A bit of care must be taken to convert from a solution that ignores symmetry since it could be the case, e.g., that  $\mathbf{u}_k^{(1)} = -\mathbf{u}_k^{(2)}$ . Algorithm 1 gives a simple procedure to “symmetrize” a tensor so that the signs align. It also averages the final sign-aligned factor matrices in case they are not exactly equal.

The benefit of ignoring symmetry is that we can use existing software for the CP decomposition. The disadvantage is that it requires  $m$  times as much storage, i.e., it must store the matrices  $\mathbf{U}^{(1)}$  thru  $\mathbf{U}^{(m)}$  rather than just  $\mathbf{X}$ . Moreover, there is no guarantee that the optimization algorithm will find the global minimum.

---

**Algorithm 1** Symmetrize Kruskal tensor

---

```

Input: CP decomposition defined by  $\mathbf{U}^{(1)}, \dots, \mathbf{U}^{(m)}$ 
Output: Symmetric CP decomposition defined by  $\lambda$  and  $\mathbf{X}$ 
1: for  $k = 1, \dots, p$  do
2:    $\lambda_k \leftarrow 1$ 
3:   for  $j = 1, \dots, m$  do
4:      $\eta \leftarrow \|\mathbf{u}_k^{(j)}\|_2$ 
5:      $\lambda_k \leftarrow \eta \lambda_k$  and  $\mathbf{u}_k^{(j)} \leftarrow \mathbf{u}_k^{(j)} / \eta$  ▷ Normalize
6:     if  $j > 1$  and  $\langle \mathbf{u}_k^{(1)}, \mathbf{u}_k^{(j)} \rangle < 0$  then
7:        $\lambda_k \leftarrow -\lambda_k$  and  $\mathbf{u}_k^{(j)} \leftarrow -\mathbf{u}_k^{(j)}$  ▷ Flip  $\mathbf{u}_k^{(j)}$  to align with  $\mathbf{u}_k^{(1)}$ 
8:     end if
9:   end for
10: end for
11:  $\mathbf{X} \leftarrow \sum_j \mathbf{U}^{(j)} / m.$ 

```

---

### 5 Optimization formulation for nonnegative symmetric factorization

The notion of completely positive tensors has been introduced by Qi et al. [39]. It is a natural extension of completely positive matrices. A nonnegative tensor  $\mathcal{A} \in \mathbb{S}^{[m,n]}$  is called completely positive if it has a decomposition of the form in (2).

The formulation is analogous to the unconstrained case, except that there is no  $\lambda$  (or equivalently, we constrain  $\lambda = 1$ ) and we add nonnegativity constraints. For given  $\mathcal{A} \in \mathbb{S}^{[m,n]}$ , our goal is to find  $\mathbf{X}$  such that (2) is satisfied. We again assume  $p$  is known. The mathematical program is given by

$$\min f_+(\mathbf{X}) = \sum_{\mathbf{i} \in \mathcal{I}} w_{\mathbf{i}} \left( a_{\mathbf{i}} - \sum_{k=1}^p (\mathbf{x}_k^m)_{\mathbf{i}} \right)^2 = \sum_{\mathbf{i} \in \mathcal{I}} w_{\mathbf{i}} \delta_{\mathbf{i}}^2 \quad \text{s.t.} \quad \mathbf{X} \geq 0.$$

Choosing  $w_{\mathbf{i}} = \sigma_{\mathbf{i}}$  yields the analogue of  $f_1$  whereas  $w_{\mathbf{i}} = 1$  yields the analogue  $f_2$ . The value  $\delta_{\mathbf{i}}$  is the difference between the model and the tensor at entry  $\mathbf{i}$ .

Using the generic formulation and following (10) without  $\lambda_k$ , the gradients are given by

$$\frac{\partial f_+}{\partial x_{jk}} = -2 \sum_{\mathbf{c} \in \mathcal{C}} c_j w_{\mathbf{c}} \delta_{\mathbf{c}} \left( x_{1k}^{c_1} \cdots x_{jk}^{c_j-1} \cdots x_{nk}^{c_n} \right).$$

Our formulation finds the best nonnegative factorization. Fan and Zhou [20] consider the problem of verifying that a tensor is completely positive.

## 6 Numerical results

For our numerical results, we assume the tensor has underlying low-rank structure, as is typical in comparisons of numerical methods for tensor factorization (see, e.g., [44]). Hence, we assume there is some underlying  $\lambda^* \in \mathbb{R}^p$  and  $\mathbf{X}^* \in \mathbb{R}^{n \times p}$  to be recovered, where  $p$  is lower than the typical rank. The noise-free data tensor is given by

$$\mathcal{A}^* = \sum_{k=1}^p \lambda_k^* (\mathbf{x}_k^*)^m. \tag{14}$$

The data tensor  $\mathcal{A}$  may also be contaminated by noise as controlled by the parameter  $\eta \geq 0$ , i.e.,

$$\mathcal{A} = \mathcal{A}^* + \eta \frac{\|\mathcal{A}^*\|}{\|\mathcal{N}\|} \mathcal{N} \quad \text{where } n_i \sim \mathcal{N}(0, 1). \tag{15}$$

Here  $\mathcal{N}$  is a noise tensor such that each element is drawn from a normal distribution, i.e.,  $n_i \sim \mathcal{N}(0, 1)$ . The parameters  $m, n, p$  control the size of the problem. If the vectors in  $\mathbf{X}^*$  are collinear, then the problem is generally more difficult [28, 44].

For the  $f_1$  objective function in (8), we calculate the gradients as specified in (11). For small problems this may not be as fast as (10), but for larger problems it makes a significant difference in speed, as shown in the results. For  $f_2$ , we precompute the index set  $\mathcal{I}$  as well as the corresponding monomial representations  $\mathcal{C}$  and multiplicities  $\sigma$ . This means that these values need not be computed each time the objective function and gradient are evaluated. The time for this preprocessing is included in the reported runtimes.

All tests were conducted on a laptop with an Intel Dual Core i7-3667U CPU and 8 GB of RAM, using MATLAB R2013a. For the optimization, unless otherwise noted, all tests are based on SNOPT, Version 7.2–9 [21, 22], using the MATLAB MEX interface. SNOPT default parameters were used except for the following: Major iteration limit = 10,000, New superbasics limit/Superbasics limit = 999, Major optimality tolerance = 1e-8. All tensor computations use the Tensor Toolbox for MATLAB, Version 2.5 [6–8] as well as additional codes for symmetric tensors (e.g., to calculate the index sets) that will be included in the next release.

### 6.1 Numerical results on a collection of test problems

We consider the impact of the problem formulation resulting from the choice of objective function and column normalization penalty. The objective function can be *weighted*, based on the standard least squares formulation denoted by  $f_1$  in (8), or *unweighted*, which counts each unique entry only once denoted by  $f_2$  in (9). The column normalization penalty is either  $\gamma = 0$  (no penalty) or  $\gamma = 0.1$ . Higher values of  $\gamma$  did not change the results.

We test the choices for several test problems as follows. We consider four sizes:

- $m = 3, n = 4, p = 2$ ;
- $m = 4, n = 3, p = 5$ ;
- $m = 4, n = 25, p = 3$ ; and
- $m = 6, n = 6, p = 4$ .

In the first case, since  $m$  is odd, we have the option to exclude  $\lambda$  from the optimization, but we include it here for consistency in this set of experiments. For each size, we also consider three noise levels:  $\eta \in \{0, 0.01, 0.1\}$ .

A random instance is created as follows. We generate a *true solution* defined by  $\lambda^* \in \mathbb{R}^p$  and  $\mathbf{X}^* \in \mathbb{R}^{n \times p}$ . The weight vector has entries selected uniformly from  $\{-1, 1\}$ , i.e.,

$$\lambda^* \in \mathbb{R}^p \quad \text{such that} \quad \lambda_k^* \in \mathcal{U}\{-1, 1\}.$$

The factor matrix is computed by first generating a matrix with random values from the normal distribution, i.e.,

$$\hat{\mathbf{X}}^* \in \mathbb{R}^{n \times p} \quad \text{such that} \quad \hat{x}_{ik}^* \in \mathcal{N}(0, 1),$$

and then normalizing each column to length one, i.e.,  $\mathbf{x}_k^* = \hat{\mathbf{x}}_k^* / \|\hat{\mathbf{x}}_k^*\| = 1$ .

Finally, given  $\lambda^*$  and  $\mathbf{X}^*$ , we can compute the tensor  $\mathcal{A}^*$  from (14) and add noise at the level specified by  $\eta$  per (15). For each problem size and noise level, we generate ten instances.

For each problem size, we generate five random starting points by choosing entries of  $\mathbf{X}$  from a Gaussian distribution (no column normalization) and entries of  $\lambda$  uniformly at random from  $\{-1, 1\}$ . The same five starting points are used for all problems of that size.

For each problem formulation corresponding to a choice for objective function and for normalization penalty, we do fifty runs, i.e., ten instances with five random starts each. The same instances and starting points are used across all formulations. The output of each run is a weight vector  $\lambda$  and a matrix  $\mathbf{X}$ . Table 3a compares the relative error which measures the proportion of the observed data that is explained by the model, i.e.,

$$\text{relative error} = \left\| \mathcal{A} - \sum_{k=1}^p \lambda_k \mathbf{x}_k^m \right\| / \|\mathcal{A}\|.$$

In the case of no noise, the ideal relative error is zero; otherwise, we hope for something near the noise level, i.e.,  $\eta$ . In our experiments, we say a run or instance is *successful* if the relative error is  $\leq 0.1$ . For each formulation, three values are reported. The first value is the number of successful runs. Since we are using five starting points per instance, the second value is the number of instances such that at least one starting point is successful. Finally, the last value is the median relative error across all fifty runs. Summary totals are provide in the last line for the 600 runs and 150 instances. Clearly,  $\gamma = 0.1$  is superior to  $\gamma = 0$  in terms of number of successful runs and instances. The comparison of unweighted ( $f_2$ ) and

**Table 3** Results of different formulations for a set of test problems

Size			Noise	Unweighted $f_2$		Weighted $f_1$	
$m$	$n$	$p$	$\eta$	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0$	$\gamma = 0.1$
3	4	2	0.00	48/10/9e-07	50/ 10/1e-06	35/10/7e-07	42/ 10/4e-07
			0.01	43/10/8e-03	46/ 10/8e-03	32/ 9/8e-03	39/ 10/8e-03
			0.10	48/10/8e-02	48/ 10/8e-02	39/10/8e-02	41/ 10/8e-02
4	3	5	0.00	34/ 9/5e-02	38/ 10/6e-03	27/10/9e-02	37/ 10/4e-02
			0.01	31/ 9/5e-02	39/ 9/6e-03	29/10/7e-02	39/ 10/9e-03
			0.10	36/10/6e-02	39/ 9/4e-02	38/10/5e-02	40/ 10/4e-02
4	25	3	0.00	6/ 5/7e-01	40/ 10/2e-05	4/ 4/6e-01	16/ 9/6e-01
			0.01	10/ 7/7e-01	44/ 10/1e-02	10/ 8/6e-01	19/ 10/6e-01
			0.10	11/ 7/7e-01	44/ 10/1e-01	11/ 6/6e-01	26/ 10/1e-01
6	6	4	0.00	23/10/4e-01	39/ 10/2e-05	7/ 5/5e-01	18/ 9/4e-01
			0.01	15/ 9/5e-01	40/ 10/1e-02	9/ 8/5e-01	25/ 10/1e-01
			0.10	1/ 1/5e-01	5/ 1/1e-01	7/ 7/5e-01	18/ 10/3e-01
Total				306/97/1e-01	472/109/1e-02	248/97/3e-01	360/118/9e-02

**(a) Relative error:** runs  $\leq 0.1$ , instances  $\leq 0.1$ , and median.

Size			Noise	Unweighted $f_2$		Weighted $f_1$	
$m$	$n$	$p$	$\eta$	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0$	$\gamma = 0.1$
3	4	2	0.00	48/10/ 1.00	50/10/ 1.00	35/10/ 1.00	42/10/ 1.00
			0.01	43/10/ 0.99	46/10/ 1.00	32/ 9/ 0.99	39/10/ 1.00
			0.10	34/ 7/ 0.97	34/ 7/ 0.97	30/ 8/ 0.93	32/ 8/ 0.93
4	3	5	0.00	3/ 2/ 0.43	5/ 3/ 0.64	0/ 0/ 0.42	6/ 4/ 0.60
			0.01	0/ 0/ 0.43	1/ 1/ 0.54	0/ 0/ 0.37	2/ 2/ 0.52
			0.10	0/ 0/ 0.25	0/ 0/ 0.45	0/ 0/ 0.32	1/ 1/ 0.51
4	25	3	0.00	7/ 6/ 0.55	40/10/ 1.00	9/ 7/ 0.66	16/ 9/ 0.67
			0.01	10/ 7/ 0.53	44/10/ 1.00	11/ 8/ 0.67	19/10/ 0.67
			0.10	13/ 7/ 0.51	44/10/ 1.00	16/ 9/ 0.67	26/10/ 1.00
6	6	4	0.00	21/10/ 0.72	38/10/ 1.00	6/ 4/ 0.72	15/ 8/ 0.75
			0.01	15/ 9/ 0.73	40/10/ 1.00	9/ 8/ 0.67	25/10/ 0.87
			0.10	18/ 8/ 0.72	32/10/ 0.98	7/ 7/ 0.73	18/10/ 0.74
Total				212/76/7e-01	374/91/1e+00	155/70/6e-01	241/92/7e-01

**(b) Solution score:** runs  $\geq 0.9$ , instances  $\geq 0.9$ , and median.

Size			Noise	Unweighted $f_2$		Weighted $f_1$	
$m$	$n$	$p$	$\eta$	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0$	$\gamma = 0.1$
3	4	2	0.00	0.10 $\pm$ 0.02	0.13 $\pm$ 0.03	0.71 $\pm$ 0.99	0.69 $\pm$ 0.63
			0.01	0.31 $\pm$ 0.74	0.21 $\pm$ 0.27	0.92 $\pm$ 1.09	1.03 $\pm$ 1.19
			0.10	0.12 $\pm$ 0.08	0.17 $\pm$ 0.28	0.92 $\pm$ 1.47	1.02 $\pm$ 1.48
4	3	5	0.00	1.19 $\pm$ 1.81	0.75 $\pm$ 0.87	4.67 $\pm$ 3.29	6.28 $\pm$ 5.44
			0.01	0.96 $\pm$ 1.25	0.77 $\pm$ 0.72	5.60 $\pm$ 3.88	7.07 $\pm$ 5.16
			0.10	1.43 $\pm$ 1.48	0.98 $\pm$ 0.78	6.35 $\pm$ 4.60	6.05 $\pm$ 4.85
4	25	3	0.00	41.38 $\pm$ 20.20	54.90 $\pm$ 13.09	3.80 $\pm$ 2.03	5.02 $\pm$ 1.45
			0.01	44.96 $\pm$ 25.17	55.69 $\pm$ 21.31	4.42 $\pm$ 2.34	5.47 $\pm$ 1.99
			0.10	44.32 $\pm$ 24.24	56.11 $\pm$ 13.29	5.01 $\pm$ 2.65	8.95 $\pm$ 3.42
6	6	4	0.00	1.79 $\pm$ 1.31	1.64 $\pm$ 0.48	4.91 $\pm$ 2.76	6.55 $\pm$ 2.29
			0.01	1.52 $\pm$ 0.89	1.57 $\pm$ 0.41	7.20 $\pm$ 5.17	8.09 $\pm$ 2.88
			0.10	1.57 $\pm$ 0.74	1.76 $\pm$ 0.73	6.16 $\pm$ 2.88	8.56 $\pm$ 4.08

**(c) Run time:** mean and standard deviation.

For each size and noise combination, the number of runs is fifty and the number of instances is ten (five random starts per instance)

weighted ( $f_1$ ) is less clear cut—the unweighted formulation is successful for many more runs overall, but the weighted formulation is successful for more instances overall.

Table 3b compares the solution scores which is a measure of how accurately  $\lambda$  and  $\mathbf{X}$  are as compared to  $\lambda^*$  and  $\mathbf{X}^*$ . Without loss of generality, we assume both  $\mathbf{X}$  and  $\mathbf{X}^*$  have normalized columns. (If  $\|\mathbf{x}_k\|_2 \neq 1$ , then we rescale  $\lambda_k = \lambda_k \sqrt{\|\mathbf{x}_k\|}$  and  $\mathbf{x}_k = \mathbf{x}_k / \|\mathbf{x}_k\|$ .) There is a permutation ambiguity, but we permute the computed solution so as to maximize the following score:

$$\text{solution score} = \frac{1}{p} \sum_{k=1}^p \left( 1 - \frac{|\lambda_k - \lambda_k^*|}{\max\{|\lambda_k|, |\lambda_k^*|\}} \right) \left| \mathbf{x}_k^T \mathbf{x}_k^* \right|.$$

A solution score of 1 indicates a perfect match, and we say a run or instance is *successful* if its solution score is  $\geq 0.9$ . As with the relative error, we report three values. The first value is the number of runs out of fifty that are successful, the second value is the number of instances out of ten that are successful (i.e., at least one starting point was successful), and the third value is the median solution score. We also report totals for each formulation across the 600 runs and 150 instances. Consistent with Table 3a, using  $\gamma = 0.1$  is more successful than  $\gamma = 0$ . The unweighted is once again successful for more runs, but the two methods are nearly tied in terms of number of instances.

Observe in Table 3b that the second size ( $m = 4, n = 3, p = 5$ ) has very low solution scores despite having good performance in terms of relative error. This is because the solution may not be unique, i.e., the k-rank of  $\mathbf{X}^*$  is no more than three, but the minimum k-rank that is sufficient for uniqueness is four per Table 2. If the solution is not unique, then multiple solutions exist and there is no reason to expect that the particular solution we find will be that one. For example, a particular instance for  $m = 4, n = 3, p = 5$  with  $\eta = 0$  is defined by

$$\lambda^* = \begin{bmatrix} 1 \\ 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} \quad \text{and} \quad \mathbf{X}^* = \begin{bmatrix} -0.3859 & -0.9285 & 0.4922 & -0.1094 & 0.4107 \\ 0.8403 & -0.1678 & -0.6975 & 0.8395 & 0.0308 \\ 0.3807 & 0.3313 & -0.5208 & -0.5322 & 0.9112 \end{bmatrix}.$$

The alternate model given by

$$\lambda = \begin{bmatrix} 1 \\ 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} \quad \text{and} \quad \mathbf{X} = \begin{bmatrix} -0.7872 & 0.5136 & -0.7809 & -0.1081 & 0.4157 \\ -0.1928 & -0.9150 & -0.0704 & 0.8249 & 0.0387 \\ 0.2039 & -0.5355 & 0.3678 & -0.5477 & 0.9065 \end{bmatrix}$$

has a relative error  $< 10^{-6}$ . The last two columns generally agree, but the first three do not and the solution score is only 0.65. It may be interesting to know that in the matrix



**Table 4** Constraint violation runs  $\leq 0.01$  and mean

Size			Noise	Unweighted $f_2$		Weighted $f_1$	
$m$	$n$	$p$	$\eta$	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0$	$\gamma = 0.1$
3	4	2	0.00	0/6.73e+00	50/1.17e-06	1/3.54e+02	50/1.87e-06
			0.01	1/3.76e+01	50/9.11e-05	1/2.04e+02	49/9.47e-04
			0.10	0/4.32e+01	50/9.32e-07	0/3.16e+02	48/6.36e-04
4	3	5	0.00	0/7.99e+01	46/7.34e-03	0/9.24e+01	40/9.76e-03
			0.01	0/6.56e+01	46/2.40e-03	0/ 6.47e+01	39/1.09e-02
			0.10	0/2.76e+02	44/3.20e-03	0/1.05e+02	39/1.71e-02
4	25	3	0.00	0/1.70e+03	50/2.63e-06	0/8.36e+02	49/3.65e-02
			0.01	0/1.76e+03	50/4.35e-06	0/6.52e+02	50/1.17e-06
			0.10	0/1.53e+03	50/2.99e-06	0/1.13e+03	49/6.29e-02
6	6	4	0.00	1/2.44e+01	50/6.04e-05	0/6.13e+00	50/5.46e-05
			0.01	1/2.45e+01	50/1.74e-05	0/2.11e+01	50/3.29e-04
			0.10	0/3.12e+01	50/3.34e-05	0/4.42e+01	49/4.75e-04

case ( $m = 2$ ), we would never compare the computed solution without imposing additional constraints such as orthogonality.

Table 3c compares the total runtimes for each method. As with any nonconvex optimization problem, there is significant variation from run to run, but we can gain a sense of the general expense for each method. As a reminder, we computed the gradient in the weighted case as shown in (11). If we compute it instead using (10), the runtimes for the weighted and unweighted methods are roughly the same. For size  $m = 4, n = 25, p = 3$ , the computation in (11) yields a 5-15X speed improvement because  $n$  is large; otherwise for smaller  $n$ , the computation in (10) will generally be faster.

Finally, we briefly consider the impact on  $\gamma$  with respect to the *constraint violation* from (12), i.e.,

$$\text{constraint violation} = \sum_{k=1}^p \left( \mathbf{x}_k^T \mathbf{x}_k - 1 \right)^2 .$$

In Table 4, we report the number of runs where the constraint violation is  $\leq 0.01$  and the mean value. Recall that the addition of the constraint violation is mainly a convenience, but it does improve the formulation by eliminating a manifold of equivalent solutions.

Table 5 shows results for more difficult test problems where  $\mathbf{X}^*$  has collinear normalized columns, i.e.,  $(\mathbf{x}_k^*)^T \mathbf{x}_\ell^* = 0.9$  for all  $k \neq \ell$  with  $k, \ell \in \{1, \dots, p\}$ . The procedure for generating the collinear columns is described by Tomasi and Bro [44]. The setup is the same as in the previous subsection except for the change in how we generate  $\mathbf{X}^*$  and the omission of size  $m = 4, n = 3, p = 5$  (since the procedure we are using does not allow  $p > n$ ). The results in Table 5 are analogous to those in Table 3. We omit the runtimes since they are similar. Although fewer runs are successful, the number of instances solved is similar.

**Table 5** Results of different formulations for “collinear test problems

Size			Noise	Unweighted $f_2$		Weighted $f_1$	
$m$	$n$	$p$	$\eta$	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0$	$\gamma = 0.1$
3	4	2	0.00	48/10/9e-07	50/10/1e-06	35/10/7e-07	42/10/4e-07
			0.01	43/10/8e-03	46/10/8e-03	32/ 9/8e-03	39/10/8e-03
			0.10	48/10/8e-02	48/10/8e-02	39/10/8e-02	41/10/8e-02
4	25	3	0.00	8/ 7/7e-01	40/10/2e-05	4/ 4/6e-01	16/ 9/6e-01
			0.01	9/ 6/7e-01	44/10/1e-02	10/ 8/6e-01	19/10/6e-01
			0.10	11/ 6/7e-01	44/10/1e-01	11/ 6/6e-01	26/10/1e-01
6	6	4	0.00	11/ 6/3e-01	26/ 9/2e-02	2/ 2/3e-01	6/ 5/2e-01
			0.01	17/10/2e-01	30/ 9/1e-02	4/ 3/3e-01	16/10/2e-01
			0.10	3/ 2/2e-01	8/ 2/1e-01	11/ 8/2e-01	19/10/1e-01

**(a) Relative error:** runs  $\leq 0.1$ , instances  $\leq 0.1$ , and median.

Size			Noise	Unweighted $f_2$		Weighted $f_1$	
$m$	$n$	$p$	$\eta$	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0$	$\gamma = 0.1$
3	4	2	0.00	48/10/9e-07	50/10/1e-06	35/10/7e-07	42/10/4e-07
			0.01	43/10/8e-03	46/10/8e-03	32/ 9/8e-03	39/10/8e-03
			0.10	48/10/8e-02	48/10/8e-02	39/10/8e-02	41/10/8e-02
4	25	3	0.00	8/ 7/7e-01	40/10/2e-05	4/ 4/6e-01	16/ 9/6e-01
			0.01	9/ 6/7e-01	44/10/1e-02	10/ 8/6e-01	19/10/6e-01
			0.10	11/ 6/7e-01	44/10/1e-01	11/ 6/6e-01	26/10/1e-01
6	6	4	0.00	11/ 6/3e-01	26/ 9/2e-02	2/ 2/3e-01	6/ 5/2e-01
			0.01	17/10/2e-01	30/ 9/1e-02	4/ 3/3e-01	16/10/2e-01
			0.10	3/ 2/2e-01	8/ 2/1e-01	11/ 8/2e-01	19/10/1e-01

**(b) Solution score:** runs  $\geq 0.9$ , instances  $\geq 0.9$ , and median.

For each size and noise combination, the number of runs is fifty and the number of instances is ten (five random starts per instance)

From these results, we have a sense that the symmetric factorization problem can be solved using standard optimization techniques. Because the problems are nonconvex, multiple starting points are needed to improve the odds of finding a global minimizer. Our results also indicate that it is helpful to add a penalty to remove the scaling ambiguity; otherwise, with no penalty, the Jacobian at the solution is singular which seems to have a negative impact on the solution quality.

### 6.2 Ignoring symmetry

As noted previously, Carroll and Chang [14] ignored symmetry with the idea that it may not be required. Ideally, the solution that is computed by a standard method, like CP-ALS [19,31] or CP-OPT [1], will be symmetric up to scaling.

Using the same problems from Table 3, we apply CP-ALS (as implemented in the Tensor Toolbox), followed by Algorithm 1 to symmetrize the solution. Three of the four sizes generically satisfy the sufficient uniqueness condition in (13).

- For  $m = 3$  and  $p = 2$ , we require  $k\text{-rank}(\mathbf{X}^*) \geq 2$ . Since  $\mathbf{X}^*$  is an  $n \times p$  matrix with  $n = 4$  whose columns are randomly generated,  $k\text{-rank}(\mathbf{X}^*) = 2$  with probability 1.

**Table 6** Results of CP-ALS plus symmetrization on test problems from Table 3

Size			Noise	CP-ALS + Symmetrization				
<i>m</i>	<i>n</i>	<i>p</i>	$\eta$	Relative Error		Soln. Score		Runtime
3	4	2	0.00	44/	10/2e-04/2e-04	44/10/	1.00	0.07 ± 0.05
			0.01	42/	10/8e-03/8e-03	40/10/	0.99	0.06 ± 0.04
			0.10	47/	10/8e-02/8e-02	40/ 9/	0.97	0.04 ± 0.04
4	3	5	0.00	39/	10/3e-02/3e-02	3/ 3/	0.63	0.23 ± 0.08
			0.01	36/	10/4e-02/2e-02	1/ 1/	0.57	0.22 ± 0.10
			0.10	37/	10/4e-02/4e-02	0/ 0/	0.59	0.21 ± 0.09
4	25	3	0.00	37/	9/9e-07/1e-06	37/ 9/	1.00	0.07 ± 0.04
			0.01	44/	10/1e-02/1e-02	44/10/	1.00	0.07 ± 0.03
			0.10	46/	10/1e-01/1e-01	46/10/	1.00	0.07 ± 0.03
6	6	4	0.00	29/	9/3e-04/3e-04	26/ 8/	1.00	0.13 ± 0.11
			0.01	18/	8/5e-01/5e-01	18/ 8/	0.73	0.08 ± 0.06
			0.10	23/	10/4e-01/4e-01	23/10/	0.74	0.09 ± 0.07
Total				442/116/5e-02/4e-02		322/88/1e+00		

Relative error runs  $\leq 0.1$ , instances  $\leq 0.1$ , median symmetrized, and median unsymmetrized  
 Solution score runs  $\geq 0.9$ , instances  $\geq 0.9$ , and median Runtime mean and standard deviation

- For  $m = 4$  and  $p = 5$ , we require  $\text{k-rank}(\mathbf{X}^*) \geq 4$ . Since  $\mathbf{X}^*$  is an  $n \times p$  matrix with  $n = 3$ , it *cannot* satisfy the condition because  $\text{k-rank}(\mathbf{X}^*) \leq \text{rank}(\mathbf{X}^*) \leq \min\{n, p\} = 3$ . Hence, the solutions may not be unique, and an example of a non-unique solution is provided in the previous subsection.
- For  $m = 4$  and  $p = 3$ , we require  $\text{k-rank}(\mathbf{X}^*) \geq 3$ . Since  $\mathbf{X}^*$  is an  $n \times p$  matrix with  $n = 25$  whose columns are randomly generated,  $\text{k-rank}(\mathbf{X}^*) = 3$  with probability 1.
- For  $m = 6$  and  $p = 6$ , we require  $\text{k-rank}(\mathbf{X}^*) \geq 3$ . Since  $\mathbf{X}^*$  is an  $n \times p$  matrix with  $n = 4$  whose columns are randomly generated,  $\text{k-rank}(\mathbf{X}^*) = 4$  with probability 1.

Table 6 shows the results, which are analogous to those in Table 3. CP-ALS with symmetrization is highly competitive. In terms of the relative error, its total number of 442 successful runs is near the high of 472 for the symmetric optimization methods; likewise, it has 116 successful instances versus 118 for symmetric optimization. Its scores are not as impressive in terms of the solution score, though this is mainly a problem for the size  $m = 4, n = 3, p = 5$ , as expected due to lack of symmetry. The major advantage of CP-ALS is runtime, where it is typically ten times faster or more. Despite the fact that CP-ALS may not find a symmetric solution, using a standard CP solution procedure followed by symmetrization is indeed an effective approach in many situations.

### 6.3 Sparsity penalty for rank determination

In Example 5.5(i) of, Nie [35] considers an method for determining the rank of a tensor. The example tensor is of order  $m = 4$  and defined by

$$\lambda^* = \begin{bmatrix} 676 \\ 196 \end{bmatrix} \quad \text{and} \quad \mathbf{X}^* = \begin{bmatrix} 0 & 3/\sqrt{14} \\ 1/\sqrt{26} & 2/\sqrt{14} \\ -5/\sqrt{26} & -1/\sqrt{14} \end{bmatrix} = \begin{bmatrix} 0.00 & 0.80 \\ 0.20 & 0.53 \\ -0.98 & -0.27 \end{bmatrix}.$$

**Table 7** Impact of sparsity penalty for problems of size  $m = 4, n = 3,$  and  $p = 2$  with a solution that has  $p = 3$

Noise	Solution score $\geq 0.9$		Median	Runtime
	Instances	Runs	Relative error	(Mean $\pm$ SD)
$\eta = 0.00$	10	50	3.4e-04	0.60 $\pm$ 0.20
$\eta = 0.01$	9	45	6.9e-03	0.54 $\pm$ 0.19
$\eta = 0.10$	4	19	6.4e-02	0.47 $\pm$ 0.13

Using our optimization approach with  $f_1$  and  $\gamma = 0.1,$  we impose the approximate  $\ell_1$  penalty of the form suggested by [42], using  $\alpha = 10$  and  $\beta = 0.1$  to arrive at the following result:

$$\lambda = \begin{bmatrix} 675.998 \\ 195.965 \\ 0.001 \\ 0.001 \\ 0.001 \\ 0.001 \end{bmatrix} \quad \text{and} \quad \mathbf{X} = \begin{bmatrix} -0.00 & 0.80 & -0.80 & 0.80 & -0.79 & -0.02 \\ -0.20 & 0.53 & -0.53 & 0.54 & -0.55 & -0.26 \\ 0.98 & -0.27 & 0.27 & -0.25 & 0.27 & 0.97 \end{bmatrix}.$$

We calculate the similarity score as described previously, selecting the two components that yield the best match for a score of 0.999865. The calculation takes approximately 2 sec. Using  $\alpha = 1000$  causes numerical blow-up, but  $\alpha = 100$  or  $\alpha = 1$  work nearly as well as  $\alpha = 10,$  i.e., the solution score is 0.9998 (with  $\beta = 0.1$ ). Likewise, varying  $\beta$  has little impact on the solution quality (with  $\alpha = 10$ ).

Using the same penalty parameters ( $\alpha = 10$  and  $\beta = 0.1$ ), we construct ten instances of problems of size  $m = 4, n = 3,$  and  $p = 2$  for each noise level  $\eta \in \{0, 0.01, 0.1\}.$  We use a solution with three components but once again apply the sparsity penalty, using the same parameters as above. We use five random starts per instance. The results as shown in Table 7. The second column shows the number of instances (out of 10) where the solution score was  $\geq 0.9,$  and the third column is the total number of runs that are successful (out of 50) for which this condition was satisfied. The fourth column shows that median relative error, and the last column shows the mean and standard deviation of the runtime. In the noise-free case, the correct solution is found in every run. For  $\eta = 0.01,$  the correct solution is obtained for 9 out of 10 instances. For  $\eta = 0.1,$  the problem is only solved to the desired accuracy in 4 out of 10 instances.

### 7 Conclusions and future challenges

Alas, the penalty approach is a heuristic; forthcoming work [5] will use statistical validation to select the rank.

**Table 8** Results of nonnegative optimization on test problems

Size	$m$	$n$	$p$	Noise		Relative error		Solution score		Runtime	
				$\eta$		Unweighted	Weighted	Unweighted	Weighted	Unweighted	Weighted
3	4	4	2	0.00		50/10/4e-07	49/10/1e-07	48/10/1.00	49/10/1.00	0.08	0.48
				0.01		50/10/8e-03	50/10/8e-03	30/6/0.95	29/6/0.93	0.11	0.53
				0.10		50/10/8e-02	50/10/7e-02	10/2/0.71	10/2/0.77	0.07	0.31
4	3	5	5	0.00		50/10/1e-04	50/10/3e-05	2/1/0.72	8/6/0.76	0.30	4.64
				0.01		50/10/3e-03	50/10/3e-03	0/0/0.56	0/0/0.60	0.40	4.72
				0.10		50/10/5e-02	50/10/4e-02	0/0/0.58	0/0/0.56	0.21	2.53
4	25	3	3	0.00		50/10/1e-08	50/10/1e-08	50/10/1.00	50/10/1.00	20.28	2.30
				0.01		50/10/1e-02	50/10/1e-02	50/10/1.00	50/10/1.00	21.59	2.38
				0.10		50/10/1e-01	50/10/1e-01	50/10/1.00	50/10/1.00	21.58	2.27
6	6	4	4	0.00		50/10/2e-07	50/10/1e-08	50/10/1.00	50/10/1.00	0.91	3.17
				0.01		50/10/1e-02	50/10/1e-02	50/10/0.99	50/10/0.99	0.88	2.70
				0.10		8/2/1e-01	47/10/1e-01	18/4/0.77	15/3/0.84	0.86	2.64

Relative error runs  $\leq 0.1$ , instances  $\leq 0.1$ , median

Solution score runs  $\geq 0.9$ , instances  $\geq 0.9$ , and median

Runtime mean

## 7.1 Nonnegative factorization

Lastly, we consider the problem of nonnegative factorization. We use the same problem setup as in §6.1 with the exception that we set all entries  $\lambda^*$  equal to one and choose entries of  $\mathbf{X}^*$  to be uniform on  $[0, 1]$ , i.e.,  $x_{ij}^* \in \mathcal{U}[0, 1]$ . The optimization formulation excludes  $\lambda$ , so there is no penalty on the columns norms of  $\mathbf{X}$  ( $\gamma = 0$ ). We add bound constraints that all entries of  $\mathbf{X}$  are nonnegative. We compare only the weighted and unweighted formulations. Table 8 shows the results, which are analogous to Table 3. There is little difference between the two formulations, except runtimes as discussed previously.

We consider straightforward optimization formulations for real-valued symmetric and nonnegative symmetric tensor decompositions. These methods can be used as baselines for comparison as new methods are developed. In particular, these methods should be useful for larger problems with inherent low-rank structure. For instance, the size  $m = 4$  and  $n = 25$  is larger in terms of dimension than most other symmetric tensor decomposition problems in the literature, though other works consider larger values of  $p$  [35]. Furthermore, we consider noise-contaminated problems, which may be problematic for algebraic methods.

Although the symmetric and nonnegative symmetric tensor decomposition problems are nonconvex, these numerical optimization approaches are effective at recovering the known solution in our experiments, especially when we use multiple random starting points. These optimization formulations can be adapted to the case of partial symmetries. Moreover, we show that if the solution is essentially unique (and the optimization method finds a global minima), then symmetry need not be directly enforced by the optimization method. In this case, efficient tools for the nonsymmetric CP problem may be employed directly.

We expect many further improvements, including different optimization formulations that exploit structure and consideration of other optimization methods.

**Acknowledgments** The anonymous referees provided extremely useful feedback that has greatly improved the manuscript. This material is based upon work supported by the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

## References

1. Acar, E., Dunlavy, D.M., Kolda, T.G.: A scalable optimization approach for fitting canonical tensor decompositions. *J. Chemom.* **25**(2), 67–86 (2011). doi:[10.1002/cem.1335](https://doi.org/10.1002/cem.1335)
2. Acar, E., Dunlavy, D.M., Kolda, T.G., Mørup, M.: Scalable tensor factorizations with missing data. In: *SDM10 Proceedings of the 2010 SIAM International Conference on Data Mining*, pp. 701–712 (2010). doi:[10.1137/1.9781611972801.61](https://doi.org/10.1137/1.9781611972801.61)
3. Acar, E., Dunlavy, D.M., Kolda, T.G., Mørup, M.: Scalable tensor factorizations for incomplete data. *Chemom. Intell. Lab. Syst.* **106**(1), 41–56 (2011). doi:[10.1016/j.chemolab.2010.08.004](https://doi.org/10.1016/j.chemolab.2010.08.004)
4. Alexander, J., Hirschowitz, A.: Polynomial interpolation in several variables. *J. Algebra. Geom.* **4**, 201–222 (1995)
5. Austin, W., Kolda, T.G., Plantenga, T.: Tensor rank prediction via cross validation. In preparation (2015)

6. Bader, B.W., Kolda, T.G.: Algorithm 862: MATLAB tensor classes for fast algorithm prototyping. *ACM Trans. Math. Soft.* **32**(4), 635–653 (2006). doi:[10.1145/1186785.1186794](https://doi.org/10.1145/1186785.1186794)
7. Bader, B.W., Kolda, T.G.: Efficient MATLAB computations with sparse and factored tensors. *SIAM J. Sci. Comput.* **30**(1), 205–231 (2007). doi:[10.1137/060676489](https://doi.org/10.1137/060676489)
8. Bader, B.W., Kolda, T.G., et al.: Matlab tensor toolbox version 2.5. <http://www.sandia.gov/~tgkolda/TensorToolbox/> (2012)
9. Ballard, G., Kolda, T.G., Plantenga, T.: Efficiently computing tensor eigenvalues on a GPU. In: *IPDPSW'11: Proceedings of the 2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Ph.D. Forum*, pp. 1340–1348. IEEE Computer Society (2011). doi:[10.1109/IPDPS.2011.287](https://doi.org/10.1109/IPDPS.2011.287)
10. Ballico, E., Bernardi, A.: Decomposition of homogeneous polynomials with low rank. *Math. Z.* **271**(3–4), 1141–1149 (2011). doi:[10.1007/s00209-011-0907-6](https://doi.org/10.1007/s00209-011-0907-6)
11. Bernardi, A., Gimigliano, A., Id, M.: Computing symmetric rank for symmetric tensors. *J. Symb. Comput.* **46**(1), 34–53 (2011). doi:[10.1016/j.jsc.2010.08.001](https://doi.org/10.1016/j.jsc.2010.08.001). <http://www.sciencedirect.com/science/article/pii/S0747717110001240>
12. Brachat, J., Comon, P., Mourrain, B., Tsigaridas, E.: Symmetric tensor decomposition. *Linear Algebra Appl.* **433**(11–12), 1851–1872 (2010). doi:[10.1016/j.laa.2010.06.046](https://doi.org/10.1016/j.laa.2010.06.046). <http://www.sciencedirect.com/science/article/B6V0R-50M0TVJ-2/2/060f24da5301d406f2e2504cce6ff9e>
13. Cambre, J., De Lathauwer, L., De Moor, B.: Best rank (R, R, R) super-symmetric tensor approximation—a continuous-time approach. In: *Proceedings of the IEEE 1999 Signal Processing Workshop on Higher-Order Statistics*, pp. 242–246 (1999). doi:[10.1109/HOST.1999.778734](https://doi.org/10.1109/HOST.1999.778734). <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=778734>
14. Carroll, J.D., Chang, J.J.: Analysis of individual differences in multidimensional scaling via an N-way generalization of “Eckart-Young” decomposition. *Psychometrika* **35**, 283–319 (1970). doi:[10.1007/BF02310791](https://doi.org/10.1007/BF02310791)
15. Comon, P., Golub, G., Lim, L.H., Mourrain, B.: Symmetric tensors and symmetric tensor rank. *SIAM J. Matrix Anal. Appl.* **30**(3), 1254–1279 (2008). doi:[10.1137/060661569](https://doi.org/10.1137/060661569)
16. Cui, C.F., Dai, Y.H., Nie, J.: All real eigenvalues of symmetric tensors, [arXiv:1403.3720](https://arxiv.org/abs/1403.3720) (2014)
17. De Lathauwer, L., De Moor, B., Vandewalle, J.: On the best rank-1 and rank- $(R_1, R_2, \dots, R_N)$  approximation of higher-order tensors. *SIAM J. Matrix Anal. Appl.* **21**(4), 1324–1342 (2000). doi:[10.1137/S0895479898346995](https://doi.org/10.1137/S0895479898346995)
18. Dosse, M.B., Ten Berge, J.M.F.: The assumption of proportional components when CANDECOMP is applied to symmetric matrices in the context of INDSCAL. *Psychometrika* **73**(2), 303–307 (2008). doi:[10.1007/s11336-007-9044-x](https://doi.org/10.1007/s11336-007-9044-x). <http://www.springerlink.com/content/1683x3734320025h/fulltext.pdf>
19. Faber, N.K.M., Bro, R., Hopke, P.K.: Recent developments in CANDECOMP/PARAFAC algorithms: a critical review. *Chemom. Intell. Lab. Syst.* **65**(1), 119–137 (2003). doi:[10.1016/S0169-7439\(02\)00089-8](https://doi.org/10.1016/S0169-7439(02)00089-8)
20. Fan, J., Zhou, A.: Completely positive tensor decomposition, [arXiv:1411.5149](https://arxiv.org/abs/1411.5149) (2014)
21. Gill, P.E., Murray, W., Saunders, M.A.: SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Rev.* **47**(1), 99–131 (2005). doi:[10.1137/S0036144504446096](https://doi.org/10.1137/S0036144504446096)
22. Gill, P.E., Murray, W., Saunders, M.A.: User’s guide for SNOPT version 7: software for large-scale nonlinear programming (2008)
23. Han, L.: An unconstrained optimization approach for finding real eigenvalues of even order symmetric tensors. *Numer. Algebra Control Optim. (NACO)* **3**(3), 583–599 (2012). doi:[10.3934/naco.2013.3.583](https://doi.org/10.3934/naco.2013.3.583)
24. Harshman, R.A.: Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis. *UCLA working papers in phonetics* **16**, 1–84 (1970). <http://www.psychology.uwo.ca/faculty/harshman/wpppfac0.pdf>
25. Hitchcock, F.L.: The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys.* **6**(1), 164–189 (1927)
26. Hitchcock, F.L.: Multiple invariants and generalized rank of a p-way matrix or tensor. *J. Math. Phys.* **7**(1), 39–79 (1927)
27. Ishteva, M., Basil, P.A., Van Dooren, P.: Jacobi algorithm for the best low multilinear rank approximation of symmetric tensors. *SIAM J. Matrix Anal. Appl.* **34**(2), 651–672 (2013). doi:[10.1137/11085743X](https://doi.org/10.1137/11085743X)
28. Kiers, H.A.L.: A three-step algorithm for CANDECOMP/PARAFAC analysis of large data sets with multicollinearity. *J. Chemom.* **12**(3), 255–171 (1998). doi:[10.1002/\(SICI\)1099-128X\(199805/06\)12:3<155::AID-CEM502>3.0.CO;2-5](https://doi.org/10.1002/(SICI)1099-128X(199805/06)12:3<155::AID-CEM502>3.0.CO;2-5)

29. Kiers, H.A.L.: Towards a standardized notation and terminology in multiway analysis. *J. Chemom.* **14**(3), 105–122 (2000). doi:[10.1002/1099-128X\(200005/06\)14:3<105::AID-CEM582>3.0.CO;2-I](https://doi.org/10.1002/1099-128X(200005/06)14:3<105::AID-CEM582>3.0.CO;2-I)
30. Kofidis, E., Regalia, P.A.: On the best rank-1 approximation of higher-order supersymmetric tensors. *SIAM J. Matrix Anal. Appl.* **23**(3), 863–884 (2002). doi:[10.1137/S0895479801387413](https://doi.org/10.1137/S0895479801387413). <http://link.aip.org/link/?SML/23/863/1>
31. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM Rev.* **51**(3), 455–500 (2009). doi:[10.1137/07070111X](https://doi.org/10.1137/07070111X)
32. Kolda, T.G., Mayo, J.R.: Shifted power method for computing tensor eigenpairs. *SIAM J. Matrix Anal. Appl.* **32**(4), 1095–1124 (2011). doi:[10.1137/100801482](https://doi.org/10.1137/100801482)
33. Kolda, T.G., Mayo, J.R.: An adaptive shifted power method for computing generalized tensor eigenpairs, [arXiv:1401.1183](https://arxiv.org/abs/1401.1183) (2014)
34. Lim, L.H.: Singular values and eigenvalues of tensors: A variational approach. In: *CAMSAP'05: Proceedings of the IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 129–132 (2005). doi:[10.1109/CAMAP.2005.1574201](https://doi.org/10.1109/CAMAP.2005.1574201)
35. Nie, J.: Generating polynomials and symmetric tensor decompositions, [arXiv:1408.5664](https://arxiv.org/abs/1408.5664) (2014)
36. Nie, J., Wang, L.: Semidefinite relaxations for best rank-1 tensor approximations. *SIAM J. Matrix Anal. Appl.* **35**(3), 1155–1179 (2014). doi:[10.1137/130935112](https://doi.org/10.1137/130935112)
37. Oeding, L., Ottaviani, G.: Eigenvectors of tensors and algorithms for waring decomposition. *J. Symb. Comput.* **54**, 9–35 (2013). doi:[10.1016/j.jsc.2012.11.005](https://doi.org/10.1016/j.jsc.2012.11.005)
38. Qi, L.: Eigenvalues of a real supersymmetric tensor. *J. Symb. Comput.* **40**, 1302–1324 (2005). doi:[10.1016/j.jsc.2005.05.007](https://doi.org/10.1016/j.jsc.2005.05.007)
39. Qi, L., Xu, C., Xu, Y.: Nonnegative tensor factorization, completely positive tensors and an hierarchical elimination algorithm, [arXiv:1305.5344](https://arxiv.org/abs/1305.5344) (2013)
40. Regalia, P.A.: Monotonically convergent algorithms for symmetric tensor approximation. *Linear Algebra Appl.* **438**(2), 875–890 (2013). doi:[10.1016/j.laa.2011.10.033](https://doi.org/10.1016/j.laa.2011.10.033). <http://www.sciencedirect.com/science/article/pii/S0024379511007300>
41. Regalia, P.A., Kofidis, E.: Monotonic convergence of fixed-point algorithms for ICA. *IEEE Trans. Neural Netw.* **14**(4), 943–949 (2003). doi:[10.1109/TNN.2003.813843](https://doi.org/10.1109/TNN.2003.813843)
42. Schmidt, M., Fung, G., Rosales, R.: Fast optimization methods for L1 regularization: A comparative study and two new approaches. *Lecture notes in computer science* pp. 286–297 (2007). doi:[10.1007/978-3-540-74958-5\\_28](https://doi.org/10.1007/978-3-540-74958-5_28)
43. Sidiropoulos, N.D., Bro, R.: On the uniqueness of multilinear decomposition of N-way arrays. *J. Chemom.* **14**(3), 229–239 (2000). doi:[10.1002/1099-128X\(200005/06\)14:3<229::AID-CEM587>3.0.CO;2-N](https://doi.org/10.1002/1099-128X(200005/06)14:3<229::AID-CEM587>3.0.CO;2-N)
44. Tomasi, G., Bro, R.: A comparison of algorithms for fitting the PARAFAC model. *Comput. Stat. Data Anal.* **50**(7), 1700–1734 (2006). doi:[10.1016/j.csda.2004.11.013](https://doi.org/10.1016/j.csda.2004.11.013)