

## EFFICIENT MATLAB COMPUTATIONS WITH SPARSE AND FACTORED TENSORS\*

BRETT W. BADER<sup>†</sup> AND TAMARA G. KOLDA<sup>‡</sup>

**Abstract.** In this paper, the term tensor refers simply to a multidimensional or  $N$ -way array, and we consider how specially structured tensors allow for efficient storage and computation. First, we study sparse tensors, which have the property that the vast majority of the elements are zero. We propose storing sparse tensors using coordinate format and describe the computational efficiency of this scheme for various mathematical operations, including those typical to tensor decomposition algorithms. Second, we study factored tensors, which have the property that they can be assembled from more basic components. We consider two specific types: A Tucker tensor can be expressed as the product of a core tensor (which itself may be dense, sparse, or factored) and a matrix along each mode, and a Kruskal tensor can be expressed as the sum of rank-1 tensors. We are interested in the case where the storage of the components is less than the storage of the full tensor, and we demonstrate that many elementary operations can be computed using only the components. All of the efficiencies described in this paper are implemented in the Tensor Toolbox for MATLAB.

**Key words.** sparse multidimensional arrays, multilinear algebraic computations, tensor decompositions, Tucker model, parallel factors (PARAFAC) model, MATLAB classes, canonical decomposition (CANDECOMP)

**AMS subject classifications.** 15A69, 68P05, 65F50

**DOI.** 10.1137/060676489

**1. Introduction.** Tensors, by which we mean multidimensional or  $N$ -way arrays, are used today in a wide variety of applications, but many issues of computational efficiency have not yet been addressed. In this article, we consider the problem of efficient computations with sparse and factored tensors, whose dense/unfactored equivalents would require too much memory.

Our particular focus is on the computational efficiency of tensor decompositions, which are being used in an increasing variety of fields in science, engineering, and mathematics. Tensor decompositions date back to the late 1960s with work by Tucker [51], Harshman [19], and Carroll and Chang [9]. Recent decades have seen tremendous growth in this area with a focus towards improved algorithms for computing the decompositions [13, 12, 58, 50]. Many innovations in tensor decompositions have been motivated by applications in chemometrics [3, 31, 8, 44]. More recently, these methods have been applied to signal processing [10, 11], image processing [52, 54, 57, 53], data mining [43, 46, 1, 45], scientific computing [6], and elsewhere [26, 36]. Though this work can be applied in a variety of contexts, we concentrate on operations that are common to tensor decompositions, such as the Tucker [51] and canonical decomposition (CANDECOMP), and parallel factors (PARAFAC) models [9, 19].

---

\*Received by the editors December 1, 2006; accepted for publication (in revised form) July 11, 2007; published electronically December 7, 2007. This work was funded by 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.

<http://www.siam.org/journals/sisc/30-1/67648.html>

<sup>†</sup>Applied Computational Methods Department, Sandia National Laboratories, Albuquerque, NM 87185-1318 (bwbader@sandia.gov).

<sup>‡</sup>Mathematics, Informatics, and Decision Sciences Department, Sandia National Laboratories, Livermore, CA 94551-9159 (tgkolda@sandia.gov).

For the purposes of our introductory discussion, we consider a third-order tensor

$$\mathcal{X} \in \mathbb{R}^{I \times J \times K}.$$

Storing every entry of  $\mathcal{X}$  requires  $IJK$  storage. A sparse tensor is one where the overwhelming majority of the entries are zero. Let  $P$  denote the number of nonzeros in  $\mathcal{X}$ . Then we say  $\mathcal{X}$  is sparse if  $P \ll IJK$ . Typically, only the nonzeros and their indices are stored for a sparse tensor. We discuss several possible storage schemes and select coordinate format as the most suitable for the types of operations required in tensor decompositions. Storing a tensor in coordinate format requires storing  $P$  nonzero values and  $NP$  corresponding integer indices, for a total of  $(N+1)P$  storage.

In addition to sparse tensors, we study two special types of factored tensors that correspond to the Tucker [51] and CANDECOMP/PARAFAC [9, 19] models. The Tucker format stores a tensor as the product of a core tensor and a factor matrix along each mode [25]. For example, if  $\mathcal{X}$  is a third-order tensor that is stored as the product of a core tensor  $\mathcal{G}$  of size  $R \times S \times T$  with corresponding factor matrices, then we express it as

$$\mathcal{X} = \llbracket \mathcal{G}; \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket, \quad \text{which means} \quad x_{ijk} = \sum_{r=1}^R \sum_{s=1}^S \sum_{t=1}^T g_{rst} a_{ir} b_{js} c_{kt} \text{ for all } i, j, k.$$

If  $I, J, K \gg R, S, T$ , then forming  $\mathcal{X}$  explicitly requires more memory than is needed to store only its components. The storage for the factored form with a dense core tensor is  $RST + IR + JS + KT$ . However, the Tucker format is not limited to the case where  $\mathcal{G}$  is dense and smaller than  $\mathcal{X}$ . It could be the case that  $\mathcal{G}$  is a large, sparse tensor so that  $R, S, T \gg I, J, K$  but the total storage is still less than  $IJK$ . Thus, more generally, the storage for a Tucker tensor is  $\text{STORAGE}(\mathcal{G}) + IR + JS + KT$ . The Kruskal format stores a tensor as the sum of rank-1 tensors [25], which corresponds to the CANDECOMP/PARAFAC models. For example, if  $\mathcal{X}$  is a third-order tensor that is stored as the sum of  $R$  rank-1 tensors, then we express it as

$$\mathcal{X} = \llbracket \boldsymbol{\lambda}; \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket, \quad \text{which means} \quad x_{ijk} = \sum_{r=1}^R \lambda_r a_{ir} b_{jr} c_{kr} \text{ for all } i, j, k.$$

As with the Tucker format, when  $I, J, K \gg R$ , forming  $\mathcal{X}$  explicitly requires more memory than storing just its factors, which require only  $(I + J + K + 1)R$  storage.

These storage formats and the techniques in this article are implemented in the MATLAB Tensor Toolbox, version 2.1 [5].

**1.1. Related work and software.** MATLAB (version 2006a) provides dense multidimensional arrays and operations for elementwise and binary operations. Version 1.0 of our MATLAB Tensor Toolbox [4] extends MATLAB's core capabilities to support operations such as tensor multiplication and matricization. The previous version of the toolbox also included objects for storing Tucker and Kruskal factored tensors but did not support mathematical operations on them beyond conversion to unfactored format. MATLAB cannot store sparse tensors except for sparse matrices which are stored in compressed sparse column (CSC) format [16]. Mathematica, an alternative to MATLAB, also supports multidimensional arrays, and there is a Mathematica package for working with tensors that accompanies the book [41]. In terms of sparse arrays, Mathematica stores its `SparseArray`'s in compressed sparse row (CSR) format and claims that its format is general enough to describe arbitrary

order tensors.<sup>1</sup> Maple has the capacity to work with sparse tensors using the `array` command and supports mathematical operations for manipulating tensors that arise in the context of physics and general relativity. Both Mathematica and Maple are geared to symbolic computation, and so their multidimensional array support is also in that vein. MATLAB, on the other hand, is generally considered more suited to numerical computations.

There are three well known packages for (dense) tensor decompositions. The  $N$ -way toolbox for MATLAB by Andersson and Bro [2] provides a suite of efficient functions and alternating least squares algorithms for decomposing dense tensors into a variety of models including Tucker and CANDECOMP/PARAFAC. The Multilinear Engine by Paatero [38] is a FORTRAN code based on the conjugate gradient algorithm that also computes a variety of multilinear models. The commercial PLS\_Toolbox [55] provides a number of multidimensional models with an emphasis towards their application in chemometrics. All three packages can handle missing data and constraints (e.g., nonnegativity) on the models.

A few other software packages for tensors are available that do not explicitly target tensor decompositions. A collection of highly optimized, template-based tensor classes in C++ for general relativity applications has been written by Landry [30] and supports functions such as binary operations and internal and external contractions. The tensors are assumed to be dense, though symmetries are exploited to optimize storage. The most closely related work to this article is the HUJI Tensor Library (HTL) by Zass [56], a C++ library for dealing with tensors using templates. HTL includes a `SparseTensor` class that stores index/value pairs using a standard template library (STL) map. HTL addresses the problem of how to optimally sort the elements of the sparse tensor (discussed in more detail in section 3.1) by letting the user specify how the subscripts should be sorted. It does not appear that HTL supports general tensor multiplication, but it does support inner product, addition, elementwise multiplication, and more. We also briefly mention `MultiArray` [15], which provides a general array class template that supports multiarray abstractions and can be used to store dense tensors.

Because it directly informs our proposed data structure, related work on storage formats for sparse matrices and tensors is deferred to section 3.1.

**1.2. Outline of article.** In section 2, we review notation and matrix and tensor operations that are needed in the paper. In section 3, we consider sparse tensors, motivate our choice of coordinate format, and describe how to make operations with sparse tensors efficient. In section 4, we describe the properties of the Tucker tensor and demonstrate how they can be used for efficient computations. In section 5, we do the same for the Kruskal tensor. In section 6, we discuss inner products and elementwise multiplication between the different types of tensors. Finally, in section 7, we conclude with a discussion on the Tensor Toolbox, our implementation of these concepts in MATLAB.

**2. Notation and background.** We follow the notation of Kiers [23], except that tensors are denoted by boldface Euler script letters, e.g.,  $\mathcal{X}$ , rather than using underlined boldface  $\underline{\mathbf{X}}$ . Matrices are denoted by boldface capital letters, e.g.,  $\mathbf{A}$ ; vectors are denoted by boldface lowercase letters, e.g.,  $\mathbf{a}$ ; and scalars are denoted by lowercase letters, e.g.,  $a$ . MATLAB-like notation specifies subarrays. For example, let  $\mathcal{X}$  be a third-order tensor. Then  $\mathbf{X}_{i::}$ ,  $\mathbf{X}_{:j:}$ , and  $\mathbf{X}_{::k}$  denote the horizontal, lateral,

<sup>1</sup>Visit the Mathematica website ([www.wolfram.com](http://www.wolfram.com)) and search on “SparseArray Data Format.”

and frontal slices, respectively. Likewise,  $\mathbf{x}_{:jk}$ ,  $\mathbf{x}_{i:k}$ , and  $\mathbf{x}_{ij}$  denote the column, row, and tube fibers, respectively. A single element is denoted by  $x_{ijk}$ . As an exception, provided that there is no possibility for confusion, the  $r$ th column of a matrix  $\mathbf{A}$  is denoted as  $\mathbf{a}_r$ . Generally, indices are taken to run from 1 to their capital version, i.e.,  $i = 1, \dots, I$ . All of the concepts in this section are discussed at greater length by Kolda [25]. For sets we use calligraphic font, e.g.,  $\mathcal{R} = \{r_1, r_2, \dots, r_P\}$ . We denote a set of indices by  $I_{\mathcal{R}} = \{I_{r_1}, I_{r_2}, \dots, I_{r_P}\}$ .

**2.1. Standard matrix operations.** The Kronecker product of matrices  $\mathbf{A} \in \mathbb{R}^{I \times J}$  and  $\mathbf{B} \in \mathbb{R}^{K \times L}$  is

$$\mathbf{A} \otimes \mathbf{B} \equiv \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \cdots & a_{1J}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \cdots & a_{2J}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{I1}\mathbf{B} & a_{I2}\mathbf{B} & \cdots & a_{IJ}\mathbf{B} \end{bmatrix} \in \mathbb{R}^{IK \times JL}.$$

The Khatri–Rao product [35, 40, 8, 44] of matrices  $\mathbf{A} \in \mathbb{R}^{I \times K}$  and  $\mathbf{B} \in \mathbb{R}^{J \times K}$  is

$$\mathbf{A} \odot \mathbf{B} \equiv [\mathbf{a}_1 \otimes \mathbf{b}_1 \quad \mathbf{a}_2 \otimes \mathbf{b}_2 \quad \cdots \quad \mathbf{a}_K \otimes \mathbf{b}_K] \in \mathbb{R}^{IJ \times K}.$$

The Hadamard (elementwise) product of matrices  $\mathbf{A}$  and  $\mathbf{B}$  is denoted by  $\mathbf{A} * \mathbf{B}$ . See, e.g., [44] for properties of these operators.

**2.2. Vector outer product.** The symbol  $\circ$  denotes the vector outer product. Let  $\mathbf{a}^{(n)} \in \mathbb{R}^{I_n}$  for all  $n = 1, \dots, N$ . Then the outer product of these  $N$  vectors is an  $N$ -way tensor, defined elementwise as

$$\left( \mathbf{a}^{(1)} \circ \mathbf{a}^{(2)} \circ \cdots \circ \mathbf{a}^{(N)} \right)_{i_1 i_2 \dots i_N} = a_{i_1}^{(1)} a_{i_2}^{(2)} \cdots a_{i_N}^{(N)} \text{ for } 1 \leq i_n \leq I_n, n \in \mathcal{N}.$$

Sometimes the symbol  $\otimes$  is used rather than the symbol  $\circ$  (see, e.g., [24]).

**2.3. Matricization of a tensor.** Matricization, sometimes called unfolding or flattening, is the rearrangement of the elements of a tensor into a matrix. Let  $\mathbf{X} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$  be an order- $N$  tensor. The modes  $\mathcal{N} = \{1, \dots, N\}$  are partitioned into  $\mathcal{R} = \{r_1, \dots, r_L\}$ , the modes that are mapped to the rows, and  $\mathcal{C} = \{c_1, \dots, c_M\}$ , the remaining modes that are mapped to the columns. Recall that  $I_{\mathcal{N}}$  denotes the set  $\{I_1, \dots, I_N\}$ . Then the matricized tensor is specified by

$$\mathbf{X}_{(\mathcal{R} \times \mathcal{C} : I_{\mathcal{N}})} \in \mathbb{R}^{J \times K}, \quad \text{with } J = \prod_{n \in \mathcal{R}} I_n \quad \text{and} \quad K = \prod_{n \in \mathcal{C}} I_n.$$

Specifically,  $(\mathbf{X}_{(\mathcal{R} \times \mathcal{C} : I_{\mathcal{N}})})_{jk} = x_{i_1 i_2 \dots i_N}$ , with

$$j = 1 + \sum_{\ell=1}^L \left[ (i_{r_\ell} - 1) \prod_{\ell'=1}^{\ell-1} I_{r_{\ell'}} \right] \quad \text{and} \quad k = 1 + \sum_{m=1}^M \left[ (i_{c_m} - 1) \prod_{m'=1}^{m-1} I_{c_{m'}} \right].$$

Other notation is used in the literature. For example,  $\mathbf{X}_{(\{1,2\} \times \{3, \dots, N\} : I_{\mathcal{N}})}$  is more typically written as

$$\mathbf{X}^{I_1 I_2 \times I_3 I_4 \dots I_N} \quad \text{or} \quad \mathbf{X}_{(I_1 I_2 \times I_3 I_4 \dots I_N)}.$$

The main nuance in our notation is that we explicitly indicate the tensor dimensions  $I_{\mathcal{N}}$ . This matters in some situations; see, e.g., (4.3).

Two special cases have their own notation. If  $\mathcal{R}$  is a singleton, then the fibers of mode  $n$  are aligned as the columns of the resulting matrix; this is called the mode- $n$  matricization. The result is denoted by

$$(2.1) \quad \mathbf{X}_{(n)} \equiv \mathbf{X}_{(\mathcal{R} \times \mathcal{C}: I_N)}, \text{ with } \mathcal{R} = \{n\} \text{ and } \mathcal{C} = \{1, \dots, n-1, n+1, \dots, N\}.$$

Different authors use different orderings for  $\mathcal{C}$ ; see, e.g., [12] versus [23]. If  $\mathcal{R} = \mathcal{N}$ , the result is a vector and is denoted by

$$(2.2) \quad \text{vec}(\mathbf{X}) \equiv \mathbf{X}_{(\mathcal{N} \times \emptyset: I_N)}.$$

Just as there is row and column rank for matrices, it is possible to define the mode- $n$  rank for a tensor, which is called the  $n$ -rank [12]. The  $n$ -rank of a tensor  $\mathbf{X}$  is defined as

$$\text{rank}_n(\mathbf{X}) = \text{rank}(\mathbf{X}_{(n)}).$$

This is not to be confused with the notion of tensor rank, which is defined in section 2.6.

**2.4. Norm and inner product of a tensor.** The inner (or scalar) product of two tensors  $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  is defined as

$$\langle \mathbf{X}, \mathbf{Y} \rangle \equiv \text{vec}(\mathbf{X})^\top \text{vec}(\mathbf{Y}) = \sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \dots \sum_{i_N=1}^{I_N} x_{i_1 i_2 \dots i_N} y_{i_1 i_2 \dots i_N},$$

and the Frobenius norm is defined as usual:  $\|\mathbf{X}\|^2 = \langle \mathbf{X}, \mathbf{X} \rangle$ .

**2.5. Tensor multiplication.** The  $n$ -mode matrix product [12] defines multiplication of a tensor with a matrix in mode  $n$  of the tensor. Let  $\mathbf{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  and  $\mathbf{A} \in \mathbb{R}^{J \times I_n}$ . Then

$$\mathbf{Y} = \mathbf{X} \times_n \mathbf{A} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N}$$

is defined most easily in terms of the mode- $n$  unfolding:

$$(2.3) \quad \mathbf{Y}_{(n)} = \mathbf{A} \mathbf{X}_{(n)}.$$

The  $n$ -mode vector product defines multiplication of a tensor with a vector in mode  $n$ . Let  $\mathbf{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  and  $\mathbf{a} \in \mathbb{R}^{I_n}$ . Then

$$\mathbf{Y} = \mathbf{X} \bar{\times}_n \mathbf{a} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times I_{n+1} \times \dots \times I_N}$$

is a tensor of order  $(N - 1)$ , defined elementwise as

$$(\mathbf{Y})_{i_1 \dots i_{n-1} i_{n+1} \dots i_N} = \sum_{i_n=1}^{I_n} x_{i_1 i_2 \dots i_N} a_{i_n}.$$

More general concepts of tensor multiplication can be defined; see [4].

**2.6. Tensor decompositions.** As mentioned in the introduction, there are two standard tensor decompositions that are considered in this paper. Let  $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ . The Tucker decomposition [51] approximates  $\mathcal{X}$  as

$$(2.4) \quad \mathcal{X} \approx \mathcal{G} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_N \mathbf{U}^{(N)},$$

where  $\mathcal{G} \in \mathbb{R}^{J_1 \times J_2 \times \dots \times J_N}$  and  $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times J_n}$  for all  $n = 1, \dots, N$ . If  $J_n = \text{rank}_n(\mathcal{X})$  for all  $n$ , then the approximation is exact and the computation is trivial. More typically, an alternating least squares (ALS) approach is used for the computation; see [27, 47, 13]. The Tucker decomposition is not unique, but measures can be taken to correct this [20, 21, 22, 48]. Observe that the right-hand side of (2.4) is a Tucker tensor, to be discussed in more detail in section 4.

The CANDECOMP/PARAFAC decomposition was simultaneously developed as the canonical decomposition of Carroll and Chang [9] and the parallel factors model of Harshman [19]; it is henceforth referred to as CP per Kiers [23]. It approximates the tensor  $\mathcal{X}$  as

$$(2.5) \quad \mathcal{X} \approx \sum_{r=1}^R \lambda_r \mathbf{v}_r^{(1)} \circ \mathbf{v}_r^{(2)} \circ \dots \circ \mathbf{v}_r^{(N)}$$

for some integer  $R > 0$ , with, for  $r = 1, \dots, R$ ,  $\lambda_r \in \mathbb{R}$  and  $\mathbf{v}_r^{(n)} \in \mathbb{R}^{I_n}$  for  $n = 1, \dots, N$ . The scalar multiplier  $\lambda_r$  is optional and can be absorbed into one of the factors, e.g.,  $\mathbf{v}_1^{(r)}$ . The rank of  $\mathcal{X}$  is defined as the minimal  $R$  such that  $\mathcal{X}$  can be exactly reproduced [28]. The right-hand side of (2.5) is a Kruskal tensor, which is discussed in more detail in section 5.

The CP decomposition can also be computed via an ALS algorithm; see, e.g., [44, 50]. Here we briefly discuss a critical part of the CP-ALS computation that can and should be specialized to sparse and factored tensors. Without loss of generality, we assume  $\lambda_r = 1$  for all  $r = 1, \dots, R$ . The CP model can be expressed in matrix form as

$$\mathbf{X}_{(n)} = \mathbf{V}^{(n)} \underbrace{\left( \mathbf{V}^{(N)} \circ \dots \circ \mathbf{V}^{(n+1)} \circ \mathbf{V}^{(n-1)} \circ \dots \circ \mathbf{V}^{(1)} \right)^\top}_{\mathbf{W}},$$

where  $\mathbf{V}^{(n)} = [\mathbf{v}_1^{(n)} \ \dots \ \mathbf{v}_R^{(n)}]$  for  $n = 1, \dots, N$ . If we fix everything but  $\mathbf{V}^{(n)}$ , then solving for it is a linear least squares problem. The pseudoinverse of the Khatri–Rao product  $\mathbf{W}$  has special structure [7, 49]:

$$\mathbf{W}^\dagger = \left( \mathbf{V}^{(N)} \circ \dots \circ \mathbf{V}^{(n+1)} \circ \mathbf{V}^{(n-1)} \circ \dots \circ \mathbf{V}^{(1)} \right) \mathbf{Z}^\dagger, \quad \text{where}$$

$$\mathbf{Z} = \left( \mathbf{V}^{(1)\top} \mathbf{V}^{(1)} \right) * \dots * \left( \mathbf{V}^{(n-1)\top} \mathbf{V}^{(n-1)} \right) * \left( \mathbf{V}^{(n+1)\top} \mathbf{V}^{(n+1)} \right) * \dots * \left( \mathbf{V}^{(N)\top} \mathbf{V}^{(N)} \right).$$

The least squares solution is given by  $\mathbf{V}^{(n)} = \mathbf{Y} \mathbf{Z}^\dagger$ , where  $\mathbf{Y} \in \mathbb{R}^{I_n \times R}$  is defined as

$$(2.6) \quad \mathbf{Y} = \mathbf{X}_{(n)} \left( \mathbf{V}^{(N)} \circ \dots \circ \mathbf{V}^{(n+1)} \circ \mathbf{V}^{(n-1)} \circ \dots \circ \mathbf{V}^{(1)} \right).$$

For CP-ALS on large-scale tensors, the calculation of  $\mathbf{Y}$  is an expensive operation and needs to be specialized. We refer to (2.6) as the “matricized-tensor-times-Khatri–Rao product,” or MTTKRP for short.

**2.7. MATLAB details.** Here we briefly describe the MATLAB code for the functions discussed in this section. The Kronecker and Hadamard matrix products are called by `kron(A,B)` and `A.*B`, respectively. The Khatri–Rao product is provided by the Tensor Toolbox and called by `khatri rao(A,B)`.

Higher-order outer products are not directly supported in MATLAB but can be implemented. For instance,  $\mathbf{X} = \mathbf{a} \circ \mathbf{b} \circ \mathbf{c}$  can be computed with standard functions via

$$\mathbf{X} = \text{reshape}(\text{kron}(\text{kron}(\mathbf{c}, \mathbf{b}), \mathbf{a}), \mathbf{I}, \mathbf{J}, \mathbf{K}),$$

where  $I$ ,  $J$ , and  $K$  are the lengths of the vectors  $\mathbf{a}$ ,  $\mathbf{b}$ , and  $\mathbf{c}$ , respectively. Using the Tensor Toolbox and the properties of the Kruskal tensor, this can be done via

$$\mathbf{X} = \text{full}(\text{ktensor}(\mathbf{a}, \mathbf{b}, \mathbf{c})).$$

Tensor  $n$ -mode multiplication is implemented in the Tensor Toolbox via the `ttm` and `ttv` commands for matrices and vectors, respectively. Implementations for dense tensors were already available in the previous version of the toolbox as discussed in [4]. We describe implementations for sparse and factored forms in this paper.

Matricization of a tensor is accomplished by permuting and reshaping the elements of the tensor. Consider the example below.

```

X = rand(5,6,4,2); R = [2 3]; C = [4 1];
I = size(X); J = prod(I(R)); K = prod(I(C));
Y = reshape(permute(X, [R C]), J, K); % convert X to matrix Y
Z = ipermute(reshape(Y, [I(R) I(C)]), [R C]); % convert back to tensor
```

In the Tensor Toolbox, this functionality is supported transparently via the `tenmat` class, which is a generalization of a MATLAB matrix. The class stores additional information to support conversion back to a `tensor` object as well as to support multiplication with another `tenmat` object for subsequent conversion back into a `tensor` object. These features are fundamental to supporting tensor multiplication. Suppose that a tensor  $\mathbf{X}$  is stored as a `tensor` object. To compute  $\mathbf{A} = \mathbf{X}_{(\mathcal{R} \times \mathcal{C} : I_N)}$ , use `A = tenmat(X,R,C)`; to compute  $\mathbf{A} = \mathbf{X}_{(n)}$ , use `A = tenmat(X,n)`; and to compute  $\mathbf{A} = \text{vec}(\mathbf{X})$ , use `A = tenmat(X, [1:N])`, where  $N$  is the number of dimensions of the tensor  $\mathbf{X}$ . The `tenmat` class was already implemented in the previous version of the toolbox under the name `tensor_as_matrix` and is described in detail in [4]. Support for sparse matricization is handled with `sptenmat`, which is described in section 3.3.

In the Tensor Toolbox, the inner product and norm functions are called via `innerprod(X,Y)` and `norm(X)`, respectively. Efficient implementations for the sparse and factored versions are discussed in the sections that follow.

The MTTKRP in (2.6) is computed via `mttkrp(X, {V1,...,VN}, n)`, where  $n$  is a scalar that indicates in which mode to matricize  $\mathbf{X}$  and which matrix to skip, i.e.,  $\mathbf{V}^{(n)}$ . If  $\mathbf{X}$  is dense, the tensor is matricized, the Khatri–Rao product is formed explicitly, and the two are multiplied together. Efficient implementations for the sparse and factored versions are discussed in the sections that follow.

**3. Sparse tensors.** A sparse tensor is a tensor where most of the elements are zero; in other words, it is a tensor where efficiency in storage and computation can be realized by storing and working with only the nonzeros. We consider storage in section 3.1, operations in section 3.2, and MATLAB details in section 3.3.

**3.1. Sparse tensor storage.** We consider the question of how to efficiently store sparse tensors. As background, we review the closely related topic of sparse matrix storage in section 3.1.1. We then consider two paradigms for storing a tensor: compressed storage in section 3.1.2 and coordinate storage in section 3.1.3.

**3.1.1. Review of sparse matrix storage.** Sparse matrices frequently arise in scientific computing, and numerous data structures have been studied for memory and computational efficiency, in serial and parallel. See [39] for an early survey of sparse matrix indexing schemes; a contemporary reference is [42, section 3.4]. Here we focus on two storage formats that can extend to higher dimensions.

The simplest storage format is coordinate format, which stores each nonzero along with its row and column index in three separate one-dimensional arrays, which Duff and Reid [14] called “parallel arrays.” For a matrix  $\mathbf{A}$  of size  $I \times J$  with  $\text{nnz}(\mathbf{A})$  nonzeros, the total storage is  $3 \cdot \text{nnz}(\mathbf{A})$ , and the indices are not necessarily presorted.

More common are CSR and CSC formats, which appear to have originated in [18]. The CSR format stores three one-dimensional arrays: an array of length  $\text{nnz}(\mathbf{A})$  with the nonzero values (sorted by row), an array of length  $\text{nnz}(\mathbf{A})$  with corresponding column indices, and an array of length  $I + 1$  that stores the beginning (and end) of each row in the other two arrays. The total storage for CSR is  $2 \cdot \text{nnz}(\mathbf{A}) + I + 1$ . The CSC format, also known as the Harwell–Boeing format, is analogous except that rows and columns are swapped; this is the format used by MATLAB [16].<sup>2</sup> The CSR/CSC formats are often cited for their storage efficiency, but our opinion is that the minor reduction of storage is of secondary importance. The main advantage of CSR/CSC formats is that the nonzeros are necessarily grouped by row/column, which means that operations that focus on rows/columns are more efficient while other operations become more expensive, such as element insertion and matrix transpose.

**3.1.2. Compressed sparse tensor storage.** Numerous higher-order analogues of CSR and CSC exist for tensors. Just as in the matrix case, the idea is that the indices are somehow sorted by a particular mode (or modes).

For a third-order tensor  $\mathbf{X}$  of size  $I \times J \times K$ , one straightforward idea is to store each frontal slice  $\mathbf{X}_{::k}$  as a sparse matrix in, say, CSC format. The entries are consequently sorted first by the third index and then by the second index.

Another idea, proposed by Lin, Liu, and Chung [34, 33], is to use extended Karnaugh map representation (EKMR). In this case, a three- or four-dimensional tensor is converted to a matrix (see section 2.3) and then stored using a standard sparse matrix scheme, such as CSR or CSC. For example, if  $\mathbf{X}$  is a three-way tensor of size  $I \times J \times K$ , then the EKMR scheme stores  $\mathbf{X}_{\{1\} \times \{2,3\}}$ , which is a sparse matrix of size  $I \times JK$ . EKMR stores a fourth-order tensor as  $\mathbf{X}_{\{1,4\} \times \{2,3\}}$ . Higher-order tensors are stored as a one-dimensional array (which encodes indices from the leading  $n - 4$  dimensions using a Karnaugh map) pointing to  $n - 4$  sparse four-dimensional tensors.

Lin, Chung, and Liu [33] compare the EKMR scheme to the method described above, i.e., storing two-dimensional slices of the tensor in CSR or CSC format. They consider two operations for the comparison: tensor addition and slice multiplication. The latter operation is multiplying subtensors (matrices) of two tensors  $\mathcal{A}$  and  $\mathcal{B}$ , such that  $\mathbf{C}_{::k} = \mathbf{A}_{::k} \mathbf{B}_{::k}$ , which is matrix-matrix multiplication on the horizontal slices. In this comparison, the EKMR scheme is more efficient.

<sup>2</sup>Search on “sparse matrix storage” in MATLAB Help or at the website [www.mathworks.com](http://www.mathworks.com).

Despite these promising results, our opinion is that compressed storage is, in general, not the best option for storing sparse tensors. First, consider the problem of choosing the sort order for the indices, which is really what a compressed format boils down to. For matrices, there are only two cases: rowwise or columnwise. For an  $N$ -way tensor, however, there are  $N!$  possible orderings on the modes. Second, the code complexity grows with the number of dimensions. It is well known that CSC/CSR formats require special code to handle rowwise and columnwise operations; for example, two distinct codes are needed to calculate  $\mathbf{A}\mathbf{x}$  and  $\mathbf{A}^\top\mathbf{x}$ . The analogue for an  $N$ th-order tensor would be a different code for  $\mathbf{A} \bar{\times}_n \mathbf{n}$  for  $n = 1, \dots, N$ . General tensor-tensor multiplication (see [4] for details) would be hard to handle. Third, we face the potential of integer overflow if we compress a tensor in a way that leads to one dimension being too big. For example, in MATLAB, indices are signed 32-bit integers, and so the largest such number is  $2^{31} - 1$ . Storing a tensor  $\mathcal{X}$  of size  $2048 \times 2048 \times 2048 \times 2048$  as the (unfolded) sparse matrix  $\mathbf{X}_{(1)}$  means that the number of columns is  $2^{33}$  and consequently too large to be indexed within MATLAB. Finally, as a general rule, the idea that the data are sorted by a particular mode becomes less and less useful as the number of modes increases. Consequently, we opt for coordinate storage format, discussed in more detail below.

Before moving on, we note that there are many cases where specialized storage formats such as EKMR can be quite useful. In particular, if the number of tensor modes is relatively small (third- or fourth-order) and the operations are specific, e.g., only operations on frontal slices, then formats such as EKMR are likely a good choice.

**3.1.3. Coordinate sparse tensor storage.** As mentioned previously, we focus on coordinate storage in this paper. For a sparse tensor  $\mathcal{X}$  of size  $I_1 \times I_2 \times \dots \times I_N$  with  $\text{nnz}(\mathcal{X})$  nonzeros, this means storing each nonzero along with its corresponding index. The nonzeros are stored in a real array of length  $\text{nnz}(\mathcal{X})$ , and the indices are stored in an integer matrix with  $\text{nnz}(\mathcal{X})$  rows and  $N$  columns (one per mode). The total storage is  $(N + 1) \cdot \text{nnz}(\mathcal{X})$ . We make no assumption on how the nonzeros are sorted. To the contrary, in section 3.2, we show that for certain operations we can entirely avoid sorting the nonzeros.

The advantage of coordinate format is its simplicity and flexibility. For instance, the operation of insertion costs  $O(\text{nnz}(\mathcal{X}))$ , which is due to the need to check whether the element already exists or not. Moreover, the operations are independent of how the nonzeros are sorted, meaning that the functions need not be specialized for different mode orderings.

**3.2. Operations on sparse tensors.** As motivated in the previous section, we consider only the case of a sparse tensor stored in coordinate format. We consider a sparse tensor

$$(3.1) \quad \mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N} \quad \text{stored as} \quad \mathbf{v} \in \mathbb{R}^P \quad \text{and} \quad \mathbf{S} \in \mathbb{R}^{P \times N},$$

where  $P = \text{nnz}(\mathcal{X})$ ,  $\mathbf{v}$  is a vector storing the nonzero values of  $\mathcal{X}$ , and  $\mathbf{S}$  stores the subscripts corresponding to the  $p$ th nonzero as its  $p$ th row. For convenience, the subscript of the  $p$ th nonzero in dimension  $n$  is denoted by  $s_{p_n}$ . In other words, the  $p$ th nonzero is

$$x_{s_{p_1}, s_{p_2}, \dots, s_{p_N}} = v_p.$$

Duplicate subscripts are not allowed.

**3.2.1. Assembling a sparse tensor.** To assemble a sparse tensor, we require a list of nonzero values and the corresponding subscripts as input. Here we consider the issue of resolving duplicate subscripts in that list. Typically, we simply sum the values at duplicate subscripts; for example,

$$\begin{array}{ll} (2, 3, 4, 5) & 3.4 \\ (2, 3, 5, 5) & 4.7 \\ (2, 3, 4, 5) & 1.1 \end{array} \quad \rightarrow \quad \begin{array}{ll} (2, 3, 4, 5) & 4.5 \\ (2, 3, 5, 5) & 4.7. \end{array}$$

If any subscript resolves to a value of zero, then that value and its corresponding subscript are removed.

Summation is not the only option for handling duplicate subscripts on input. We can use any rule to combine a list of values associated with a single subscript, such as max, mean, standard deviation, or even the ordinal count, as shown here:

$$\begin{array}{ll} (2, 3, 4, 5) & 3.4 \\ (2, 3, 5, 5) & 4.7 \\ (2, 3, 4, 5) & 1.1 \end{array} \quad \rightarrow \quad \begin{array}{ll} (2, 3, 4, 5) & 2 \\ (2, 3, 5, 5) & 1. \end{array}$$

Overall, the work of assembling a tensor reduces to finding all of the unique subscripts and applying a reduction function (to resolve duplicate subscripts). The amount of work for this computation depends on the implementation but is no worse than the cost of sorting all of the subscripts, i.e.,  $O(P \log P)$ , where  $P = \text{nnz}(\mathbf{X})$ .

**3.2.2. Arithmetic on sparse tensors.** Consider two same-sized sparse tensors  $\mathbf{X}$  and  $\mathbf{Y}$ , stored as  $(\mathbf{v}_\mathbf{X}, \mathbf{S}_\mathbf{X})$  and  $(\mathbf{v}_\mathbf{Y}, \mathbf{S}_\mathbf{Y})$  as defined in (3.1). To compute  $\mathbf{Z} = \mathbf{X} + \mathbf{Y}$ , we create

$$\mathbf{v}_\mathbf{Z} = \begin{bmatrix} \mathbf{v}_\mathbf{X} \\ \mathbf{v}_\mathbf{Y} \end{bmatrix} \quad \text{and} \quad \mathbf{S}_\mathbf{Z} = \begin{bmatrix} \mathbf{S}_\mathbf{X} \\ \mathbf{S}_\mathbf{Y} \end{bmatrix}.$$

To produce  $\mathbf{Z}$ , the nonzero values  $\mathbf{v}_\mathbf{Z}$  and corresponding subscripts  $\mathbf{S}_\mathbf{Z}$  are assembled by summing duplicates (see section 3.2.1). Clearly,  $\text{nnz}(\mathbf{Z}) \leq \text{nnz}(\mathbf{X}) + \text{nnz}(\mathbf{Y})$ . In fact,  $\text{nnz}(\mathbf{Z}) = 0$  if  $\mathbf{Y} = -\mathbf{X}$ .

It is possible to perform logical operations on sparse tensors in a similar fashion. For example, computing  $\mathbf{Z} = \mathbf{X} \wedge \mathbf{Y}$  (“logical and”) reduces to finding the intersection of the nonzero indices for  $\mathbf{X}$  and  $\mathbf{Y}$ . In this case, the reduction formula is that the final value is 1 (true) only if the number of elements is at least two; for example, suppose the tensor  $\mathbf{X}$  has the first two elements in the list below and the tensor  $\mathbf{Y}$  has the third, then

$$\begin{array}{ll} (2, 3, 4, 5) & 3.4 \\ (2, 3, 5, 5) & 4.7 \\ (2, 3, 4, 5) & 1.1 \end{array} \quad \rightarrow \quad (2, 3, 4, 5) \quad 1 \text{ (true)}.$$

For “logical and,”  $\text{nnz}(\mathbf{Z}) \leq \text{nnz}(\mathbf{X}) + \text{nnz}(\mathbf{Y})$ . Some logical operations, however, do not produce sparse results. For example,  $\mathbf{Z} = \neg \mathbf{X}$  (“logical not”) has nonzeros everywhere that  $\mathbf{X}$  has a zero.

Comparisons can also produce dense or sparse results. For instance, if  $\mathbf{X}$  and  $\mathbf{Y}$  have the same sparsity pattern, then  $\mathbf{Z} = (\mathbf{X} < \mathbf{Y})$  is such that  $\text{nnz}(\mathbf{Z}) \leq \text{nnz}(\mathbf{X})$ . Comparison against a scalar can produce a dense or sparse result. For example,  $\mathbf{Z} = (\mathbf{X} > 1)$  has no more nonzeros than  $\mathbf{X}$ , whereas  $\mathbf{Z} = (\mathbf{X} > -1)$  has nonzeros everywhere that  $\mathbf{X}$  has a zero.

**3.2.3. Norm and inner product for a sparse tensor.** Consider a sparse tensor  $\mathcal{X}$  as in (3.1) with  $P = \text{nnz}(\mathcal{X})$ . The work to compute the norm is  $O(P)$  and does not involve any data movement:

$$\|\mathcal{X}\| = \sqrt{\sum_{p=1}^P v_p^2} .$$

The inner product of two same-sized sparse tensors  $\mathcal{X}$  and  $\mathcal{Y}$  involves finding duplicates in their subscripts, similar to the problem of assembly (see section 3.2.1). The cost is no worse than the cost of sorting all of the subscripts, i.e.,  $O(P \log P)$ , where  $P = \text{nnz}(\mathcal{X}) + \text{nnz}(\mathcal{Y})$ .

**3.2.4.  $n$ -mode vector multiplication for a sparse tensor.** Coordinate storage format is amenable to the computation of a tensor times a vector in mode  $n$ . We can do this computation in  $O(\text{nnz}(\mathcal{X}))$  time, though this does not account for the cost of data movement, which is generally the most time-consuming part of this operation. (The same is true for sparse matrix-vector multiplication.)

Consider

$$\mathcal{Y} = \mathcal{X} \bar{\times}_n \mathbf{a},$$

where  $\mathcal{X}$  is as defined in (3.1) and the vector  $\mathbf{a}$  is of length  $I_n$ . For each  $p = 1, \dots, P$ , nonzero  $v_p$  is multiplied by  $a_{s_{p_n}}$  and added to the  $(s_{p_1}, \dots, s_{p_{n-1}}, s_{p_{n+1}}, \dots, s_{p_N})$  element of  $\mathcal{Y}$ . Stated another way, we can convert  $\mathbf{a}$  to an “expanded” vector  $\mathbf{b} \in \mathbb{R}^P$  such that

$$b_p = a_{s_{n_p}} \text{ for } p = 1, \dots, P.$$

Next we can calculate a vector of values  $\hat{\mathbf{v}} \in \mathbb{R}^P$  so that

$$\hat{\mathbf{v}} = \mathbf{v} * \mathbf{b}.$$

We create a matrix  $\hat{\mathbf{S}}$  that is equal to  $\mathbf{S}$  with the  $n$ th column removed. Then the nonzeros  $\hat{\mathbf{v}}$  and subscripts  $\hat{\mathbf{S}}$  can be assembled (summing duplicates) to create  $\mathcal{Y}$ . Observe that  $\text{nnz}(\mathcal{Y}) \leq \text{nnz}(\mathcal{X})$ , but the number of dimensions is also reduced by one, meaning that the final result is not necessarily sparse even though the number of nonzeros cannot increase.

We can generalize the previous discussion to multiplication by vectors in multiple modes. For example, consider the case of multiplication in every mode:

$$\alpha = \mathcal{X} \bar{\times}_1 \mathbf{a}^{(1)} \dots \bar{\times}_N \mathbf{a}^{(N)}.$$

Define “expanded” vectors  $\mathbf{b}^{(n)} \in \mathbb{R}^P$  for  $n = 1, \dots, N$  such that

$$b_p^{(n)} = a_{s_{n_p}}^{(n)} \text{ for } p = 1, \dots, P.$$

We then calculate  $\mathbf{w} = \mathbf{v} * \mathbf{b}^{(1)} * \dots * \mathbf{b}^{(N)}$ , and the final scalar result is  $\alpha = \sum_{p=1}^P w_p$ . Observe that we calculate all of the  $n$ -mode products simultaneously rather than in sequence. Hence, only one “assembly” of the final result is needed.

**3.2.5.  $n$ -mode matrix multiplication for a sparse tensor.** The computation of a sparse tensor times a matrix in mode  $n$  is straightforward. To compute

$$\mathbf{Y} = \mathbf{X} \times_n \mathbf{A},$$

we use the matricized version in (2.3), storing  $\mathbf{X}_{(n)}$  as a sparse matrix. As one might imagine, CSR format works well for mode- $n$  unfoldings, but CSC format does not because there are so many columns. For CSC, use the transposed version of the equation, i.e.,

$$\mathbf{Y}_{(n)}^\top = \mathbf{X}_{(n)}^\top \mathbf{A}^\top.$$

Unless  $\mathbf{A}$  has special structure (e.g., diagonal), the result is dense. Consequently, this works only for relatively small tensors (and is why we have glossed over the possibility of integer overflow when we convert  $\mathbf{X}$  to  $\mathbf{X}_{(n)}$ ). The cost boils down to that of converting  $\mathbf{X}$  to a sparse matrix, doing a matrix-by-sparse-matrix multiplication, and converting the result into a (dense) tensor  $\mathbf{Y}$ . Multiple  $n$ -mode matrix multiplications are performed sequentially.

**3.2.6. General tensor multiplication for sparse tensors.** For tensor-tensor multiplication, the modes to be multiplied are specified. For example, if we have two tensors  $\mathbf{X} \in \mathbb{R}^{3 \times 4 \times 5}$  and  $\mathbf{Y} \in \mathbb{R}^{4 \times 3 \times 2 \times 2}$ , we can calculate:

$$\mathbf{Z} = \langle \mathbf{X}, \mathbf{Y} \rangle_{\{1,2;2,1\}} \in \mathbb{R}^{5 \times 2 \times 2},$$

which means that we multiply modes 1 and 2 of  $\mathbf{X}$  with modes 2 and 1 of  $\mathbf{Y}$ . Here we refer to the modes that are being multiplied as the “inner” modes and the other modes as the “outer” modes because, in essence, we are taking inner and outer products along these modes. Because it takes several pages to explain tensor-tensor multiplication, we have omitted it from the background material in section 2 and instead refer the interested reader to [4].

In the sparse case, we have to find all of the matches of the inner modes of  $\mathbf{X}$  and  $\mathbf{Y}$ , compute the Kronecker product of the matches, associate each element of the product with a subscript that comes from the outer modes, and then resolve duplicate subscripts by summing the corresponding nonzeros. Depending on the modes specified, the work can be as high as  $O(PQ)$ , where  $P = \text{nnz}(\mathbf{X})$  and  $Q = \text{nnz}(\mathbf{Y})$ , but can be closer to  $O(P \log P + Q \log Q)$  depending on which modes are multiplied and the structure on the nonzeros.

**3.2.7. Matricized sparse tensor times Khatri–Rao product.** Consider the calculation of the matricized tensor times a Khatri–Rao product in (2.6). We compute this indirectly using the  $n$ -mode vector multiplication, which is efficient for large, sparse tensors (see section 3.2.4), by rewriting (2.6) as

$$\mathbf{y}_r = \mathbf{X} \bar{\mathbf{x}}_1 \mathbf{v}_r^{(1)} \cdots \bar{\mathbf{x}}_{n-1} \mathbf{v}_r^{(n-1)} \bar{\mathbf{x}}_{n+1} \mathbf{v}_r^{(n+1)} \cdots \bar{\mathbf{x}}_N \mathbf{v}_r^{(N)} \quad \text{for } r = 1, 2, \dots, R.$$

In other words, the solution  $\mathbf{W}$  is computed column by column. The cost equates to computing the product of the sparse tensor with  $N - 1$  vectors  $R$  times.

**3.2.8. Computing  $\mathbf{X}_{(n)} \mathbf{X}_{(n)}^\top$  for a sparse tensor.** Generally, the product  $\mathbf{Z} = \mathbf{X}_{(n)} \mathbf{X}_{(n)}^\top \in \mathbb{R}^{I_n \times I_n}$  can be computed directly by storing  $\mathbf{X}_{(n)}$  as a sparse matrix. As in section 3.2.5, we must be wary of CSC format, in which case we should actually store  $\mathbf{A} = \mathbf{X}_{(n)}^\top$  and then calculate  $\mathbf{Z} = \mathbf{A}^\top \mathbf{A}$ . The cost is primarily the cost of

converting to a sparse matrix format (e.g., CSC) plus the matrix-matrix multiplication to form the dense matrix  $\mathbf{Z} \in \mathbb{R}^{I_n \times I_n}$ . However, the matrix  $\mathbf{X}_{(n)}$  is of size

$$I_n \times \prod_{\substack{m=1 \\ m \neq n}}^N I_m,$$

which means that its column indices may overflow the integers if the tensor dimensions are very big.

**3.2.9. Collapsing and scaling on sparse tensors.** We present the concepts of collapsing and scaling on tensors to extend well-known (and mostly unnamed) operations on matrices.

For a matrix, one might want to compute the sum of all elements in each row, or the maximum element in each column, or the average of all elements, and so on. To the best of our knowledge, these sorts of operations do not have a name, so we call them *collapse* operations—we are collapsing the object in one or more dimensions to get some statistical information. Conversely, we often want to use the results of a collapse operation to *scale* the elements of a matrix. For example, to convert a matrix  $\mathbf{A}$  to a row-stochastic matrix, we compute the collapsed sum in mode 1 (rowwise) and call it  $\mathbf{z}$  and then scale  $\mathbf{A}$  by multiplying each row in  $\mathbf{A}$  by the corresponding element in the vector  $(1/\mathbf{z})$ .

We can define similar operations in the  $N$ -way context for tensors. For collapsing, we define the modes to be collapsed and the operation (e.g., sum, max, number of elements, etc.). Likewise, scaling can be accomplished by specifying the modes to scale.

Suppose, for example, that we have an  $I \times J \times K$  tensor  $\mathcal{X}$  and want to scale each frontal slice so that its largest entry is one. First, we collapse the tensor in modes 1 and 2 using the max operation. In other words, we compute the maximum of each frontal slice, i.e.,

$$z_k = \max\{x_{ijk} \mid i = 1, \dots, I \text{ and } j = 1, \dots, J\} \quad \text{for } k = 1, \dots, K.$$

This is accomplished in coordinate format by considering only the third subscript corresponding to each nonzero and doing assembly with duplicate resolution (see section 3.2.1) via the appropriate collapse operation (in this case, max). Then the scaled tensor can be computed elementwise by

$$y_{ijk} = \frac{x_{ijk}}{z_k}.$$

This computation can be completed by “expanding”  $\mathbf{z}$  to a vector of length  $\text{nnz}(X)$  as was done for the sparse-tensor-times-vector operation in section 3.2.4.

**3.3. MATLAB details for sparse tensors.** MATLAB does not natively support sparse tensors. In the Tensor Toolbox, sparse tensors are stored in the `sptensor` class, which stores the size as an integer  $N$ -vector along with the vector of nonzero values  $\mathbf{v}$  and corresponding integer matrix of subscripts  $\mathbf{S}$  from (3.1).

We can assemble a sparse tensor from a list of subscripts and corresponding values, as described in section 3.2.1. By default, we sum repeated entries, though we allow the option of using other functions to resolve duplicates. To this end, we rely on the MATLAB `accumarray` function, which takes a list of subscripts, a corresponding list

of values, and a function to resolve the duplicates (sum, by default). To use this with large-scale sparse data is complex. We first calculate a codebook of the  $Q$  unique subscripts (using the MATLAB `unique` function), use the codebook to convert each  $N$ -way subscript to an integer value between 1 and  $Q$ , call `accumarray` with the integer indices, and then use the codebook to map the final result back to the corresponding  $N$ -way subscripts.

MATLAB relies heavily on *linear indices* for any operation that returns a list of subscripts. For example, the `find` command on a sparse matrix returns linear indices (by default) that can be subsequently converted to row and column indices. For tensors, we are wary of linear indices due to the possibility of integer overflow discussed in section 3.1.2. Specifically, linear indices may produce integer overflow if the product of the dimensions of the tensor is greater than or equal to  $2^{32}$ , e.g., a four-way tensor of size  $2048 \times 2048 \times 2048 \times 2048$ . Thus, our versions of subscripted reference (`subsref`) and assignment (`subsasgn`) as well as our version of `find` explicitly use subscripts and do not support linear indices.

We do, however, support the conversion of a sparse tensor to a matrix stored in coordinate format via the class `sptenmat`. This matrix can then be converted into a MATLAB sparse matrix via the command `double`.

All operations are called in the same way for sparse tensors as they are for dense tensor, e.g.,  $\mathbf{Z} = \mathbf{X} + \mathbf{Y}$ . Logical operations always produce `sptensor` results, even if they would be more efficiently stored as dense tensors. To convert to a dense tensor, call `full(X)`.

The three multiplication operations may produce dense results: tensor-times-tensor (`ttt`), tensor-times-matrix (`ttm`), and tensor-times-vector (`ttv`). In the case of `ttm`, since it is often called repeatedly for multiplication in multiple modes, any intermediate product may be dense, and the remaining calls will be to the dense version of `ttm`. For general tensor multiplication, which reduces to sparse matrix-matrix multiplication, we take measures to avoid integer overflow by instead finding the unique subscripts and using only that many rows/columns in the matrices that are multiplied. This is similar to how we use `accumarray` to assemble a tensor.

Generating a random sparse tensor is complicated because it requires generating the locations of the nonzeros as well as the nonzeros. Thus, the Tensor Toolbox provides the command `sptenrand(sz, nnz)` to produce a sparse tensor. It is analogous to the command `sprand` to produce a random sparse matrix in MATLAB with two exceptions. First, the size is passed in as a single (row vector) input. Second, the last argument can be either a percentage (as in `sprand`) or an explicit number of nonzeros desired. We also provide a function `sptendiag` to create a superdiagonal tensor.

**4. Tucker tensors.** Consider a tensor  $\mathbf{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  such that

$$(4.1) \quad \mathbf{X} = \llbracket \mathcal{G}; \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \rrbracket \equiv \mathcal{G} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_N \mathbf{U}^{(N)},$$

where  $\mathcal{G} \in \mathbb{R}^{J_1 \times J_2 \times \dots \times J_N}$  is the core tensor and  $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times J_n}$  for  $n = 1, \dots, N$ . This is the format that results from a Tucker decomposition [51] and is therefore termed a *Tucker tensor*. We use the shorthand notation  $\llbracket \mathcal{G}; \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \rrbracket$  from [25], but other notation can be used. For example, Lim [32] proposes that the covariant aspect of the multiplication be made explicit by expressing (4.1) as

$$\left( \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \right) \cdot \mathcal{G}.$$

As another example, Grigorascu and Regalia [17] emphasize the role of the core tensor in the multiplication by expressing (4.1) as

$$\mathbf{X} = \mathbf{U}^{(1)} \underset{\star}{\mathfrak{G}} \mathbf{U}^{(2)} \underset{\star}{\mathfrak{G}} \dots \underset{\star}{\mathfrak{G}} \mathbf{U}^{(N)},$$

which is called the *weighted Tucker product*; the unweighted version has  $\mathfrak{G} = \mathfrak{J}$ , the identity tensor. Regardless of the notation, the properties of a Tucker tensor are the same.

**4.1. Tucker tensor storage.** Storing  $\mathbf{X}$  as a Tucker tensor can have major advantages in terms of memory requirements. In its explicit form,  $\mathbf{X}$  requires storage of

$$\prod_{n=1}^N I_n \quad \text{versus} \quad \text{STORAGE}(\mathfrak{G}) + \sum_{n=1}^N I_n J_n$$

elements for the factored form. Thus, the Tucker tensor factored format is clearly advantageous if  $\text{STORAGE}(\mathfrak{G})$  is sufficiently small. This certainly is the case if

$$(4.2) \quad \prod_{n=1}^N J_n \ll \prod_{n=1}^N I_n.$$

However, there is no reason to assume that the core tensor  $\mathfrak{G}$  is dense; on the contrary,  $\mathfrak{G}$  might itself be sparse or factored. For instance, the authors of [37, 48] have studied methods to impose structure on  $\mathfrak{G}$ . The next section discusses computations on  $\mathbf{X}$  in its factored form, making minimal assumptions about the format of  $\mathfrak{G}$ .

**4.2. Tucker tensor properties.** It is common knowledge (dating back to [51]) that matricized versions of the Tucker tensor (4.1) have a special form; specifically,

$$(4.3) \quad \mathbf{X}_{(\mathcal{R} \times \mathcal{C} : J_N)} = \left( \mathbf{U}^{(r_L)} \otimes \dots \otimes \mathbf{U}^{(r_1)} \right) \mathbf{G}_{(\mathcal{R} \times \mathcal{C} : I_N)} \left( \mathbf{U}^{(c_M)} \otimes \dots \otimes \mathbf{U}^{(c_1)} \right)^\top,$$

where  $\mathcal{R} = \{r_1, \dots, r_L\}$  and  $\mathcal{C} = \{c_1, \dots, c_M\}$ . Note that the order of the indices in  $\mathcal{R}$  and  $\mathcal{C}$  does matter, and reversing the order of the indices is a frequent source of coding errors. For the special case of mode- $n$  matricization (2.1), we have

$$(4.4) \quad \mathbf{X}_{(n)} = \mathbf{U}^{(n)} \mathbf{G}_{(n)} \left( \mathbf{U}^{(N)} \otimes \dots \otimes \mathbf{U}^{(n+1)} \otimes \mathbf{U}^{(n-1)} \otimes \dots \otimes \mathbf{U}^{(1)} \right)^\top.$$

Likewise, for the vectorized version (2.2), we have

$$(4.5) \quad \text{vec}(\mathbf{X}) = \left( \mathbf{U}^{(N)} \otimes \dots \otimes \mathbf{U}^{(1)} \right) \text{vec}(\mathfrak{G}).$$

**4.2.1.  $n$ -mode matrix multiplication for a Tucker tensor.** Multiplying a Tucker tensor times a matrix in mode  $n$  reduces to multiplying its  $n$ th factor matrix; in other words, the result retains the factored Tucker tensor structure. Let  $\mathbf{X}$  be as in (4.1) and  $\mathbf{V}$  be a matrix of size  $K \times I_n$ . Then from (2.3) and (4.4) we have

$$\mathbf{X} \times_n \mathbf{V} = \llbracket \mathfrak{G} ; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(n-1)}, \mathbf{V} \mathbf{U}^{(n)}, \mathbf{U}^{(n+1)}, \dots, \mathbf{U}^{(N)} \rrbracket.$$

The cost is that of the matrix-matrix multiplication, that is,  $O(I_n J_n K)$ . More generally, let  $\mathbf{V}^{(n)}$  be of size  $K_n \times I_n$  for  $n = 1, \dots, N$ . Then

$$\llbracket \mathbf{X} ; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket = \llbracket \mathfrak{G} ; \mathbf{V}^{(1)} \mathbf{U}^{(1)}, \dots, \mathbf{V}^{(N)} \mathbf{U}^{(N)} \rrbracket.$$

The cost here is the cost of  $N$  matrix-matrix multiplications, for a total of  $O(\sum_n I_n J_n K_n)$ , and the Tucker tensor structure is retained. As an aside, if  $\mathbf{U}^{(n)}$  has full column rank and  $\mathbf{V}^{(n)} = \mathbf{U}^{(n)\dagger}$  for  $n = 1, \dots, N$ , then  $\mathfrak{G} = \llbracket \mathbf{X} ; \mathbf{U}^{(1)\dagger}, \dots, \mathbf{U}^{(N)\dagger} \rrbracket$ .

**4.2.2.  $n$ -mode vector multiplication for a Tucker tensor.** Multiplication of a Tucker tensor by a vector follows logic similar to the matrix case except that the  $n$ th factor matrix necessarily disappears and the problem reduces to  $n$ -mode vector multiplication with the core. Let  $\mathcal{X}$  be a Tucker tensor as in (4.1) and  $\mathbf{v}$  be a vector of size  $I_n$ ; then

$$\mathcal{X} \bar{\times}_n \mathbf{v} = \llbracket \mathcal{G} \bar{\times}_n \mathbf{w} ; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(n-1)}, \mathbf{U}^{(n+1)}, \dots, \mathbf{U}^{(N)} \rrbracket, \quad \text{where } \mathbf{w} = \mathbf{U}^{(n)\top} \mathbf{v}.$$

The cost here is that of multiplying a matrix times a vector,  $O(I_n J_n)$ , plus the cost of multiplying the core (which could be dense, sparse, or factored) times a vector. The Tucker tensor structure is retained but with one less factor matrix. More generally, multiplying a Tucker tensor by a vector in every mode converts to the problem of multiplying its core by a vector in every mode. Let  $\mathbf{V}^{(n)}$  be of size  $I_n$  for  $n = 1, \dots, N$ ; then

$$\mathcal{X} \bar{\times}_1 \mathbf{v}^{(1)} \dots \bar{\times}_N \mathbf{v}^{(N)} = \mathcal{G} \bar{\times}_1 \mathbf{w}^{(1)} \dots \bar{\times}_N \mathbf{w}^{(N)},$$

where  $\mathbf{w}^{(n)} = \mathbf{U}^{(n)\top} \mathbf{v}^{(n)}$  for all  $n = 1, \dots, N$ .

In this case, the work is the cost of  $N$  matrix-vector multiplications,  $O(\sum_n I_n J_n)$ , plus the cost of multiplying the core by a vector in each mode. If  $\mathcal{G}$  is dense, the total cost is

$$O\left(\sum_{n=1}^N \left(I_n J_n + \prod_{m=n}^N J_m\right)\right).$$

Further gains in efficiency are possible by doing the multiplications in order of largest to smallest  $J_n$ . The Tucker tensor structure is clearly not retained for all-mode vector multiplication.

**4.2.3. Inner product.** Let  $\mathcal{X}$  be a Tucker tensor as in (4.1), and let  $\mathcal{Y}$  be a Tucker tensor of the same size, with

$$\mathcal{Y} = \llbracket \mathcal{H} ; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket,$$

with  $\mathcal{H} \in \mathbb{R}^{K_1 \times K_2 \times \dots \times K_N}$  and  $\mathbf{V}^{(n)} \in \mathbb{R}^{I_n \times K_n}$  for  $n = 1, \dots, N$ . If the cores are small in relation to the overall tensor size, we can realize computational savings as follows. Without loss of generality, assume  $\mathcal{G}$  is smaller than (or at least no larger than)  $\mathcal{H}$ , e.g.,  $J_n \leq K_n$  for all  $n$ . Then

$$\begin{aligned} \langle \mathcal{X}, \mathcal{Y} \rangle &= \langle \llbracket \mathcal{G} ; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket, \llbracket \mathcal{H} ; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket \rangle \\ &= \langle \mathcal{G}, \llbracket \mathcal{H} ; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket \times_1 \mathbf{U}^{(1)\top} \dots \times_N \mathbf{U}^{(N)\top} \rangle \\ &= \langle \mathcal{G}, \llbracket \mathcal{H} ; \mathbf{U}^{(1)\top} \mathbf{V}^{(1)}, \dots, \mathbf{U}^{(N)\top} \mathbf{V}^{(N)} \rrbracket \rangle \\ &= \langle \mathcal{G}, \mathcal{F} \rangle, \quad \text{with } \mathcal{F} = \llbracket \mathcal{H} ; \mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)} \rrbracket, \quad \text{and } \mathbf{W}^{(n)} = \mathbf{U}^{(n)\top} \mathbf{V}^{(n)} \text{ for all } n. \end{aligned}$$

Each  $\mathbf{W}^{(n)}$  is of size  $J_n \times K_n$  and costs  $O(I_n J_n K_n)$  to compute. Then, to compute  $\mathcal{F}$ , we do a tensor times matrix in all modes with the tensor  $\mathcal{H}$  (the cost varies depending on the tensor type), followed by an inner product between two tensors of size  $J_1 \times J_2 \times \dots \times J_N$ . If  $\mathcal{G}$  and  $\mathcal{H}$  are dense, then the total cost is

$$O\left(\sum_{n=1}^N I_n J_n K_n + \sum_{n=1}^N \left(\prod_{p=n}^N K_p \prod_{q=1}^n J_q\right) + \prod_{n=1}^N J_n\right).$$

**4.2.4. Norm of a Tucker tensor.** From the previous discussion, it is clear that the norm can also be calculated efficiently if the core tensor is small in relation to the overall tensor, e.g.,  $J_n < I_n$  for all  $n$ . Let  $\mathbf{X}$  be a Tucker tensor as in (4.1). From section 4.2.3, we have

$$(4.6) \quad \|\mathbf{X}\|^2 = \langle \mathbf{X}, \mathbf{X} \rangle = \langle \mathbf{G}, \mathcal{F} \rangle,$$

with  $\mathcal{F} = \llbracket \mathbf{G}; \mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)} \rrbracket$  and  $\mathbf{W}^{(n)} = \mathbf{U}^{(n)\top} \mathbf{U}^{(n)}$  for all  $n$ .

Forming all of the  $\mathbf{W}^{(n)}$  matrices costs  $O(\sum_n I_n J_n^2)$ . To compute  $\mathcal{F}$ , we have to do a tensor times matrix in all  $N$  modes, and if  $\mathbf{G}$  is dense, for example, the cost is  $O(\prod_n J_n \cdot \sum_n J_n)$ . Finally, we compute an inner product of two tensors of size  $J_1 \times J_2 \times \dots \times J_n$ , which costs  $O(\prod_n J_n)$  if both tensors are dense. Moreover, if the matrices  $\mathbf{U}^{(n)}$  are columnwise orthonormal, as is often the case, then they entirely disappear, leading to even greater computational savings.

**4.2.5. Matricized Tucker tensor times Khatri–Rao product.** As noted in section 2.6, a common operation is to calculate a particular matricized tensor times a special Khatri–Rao product (2.6). In the case of a Tucker tensor, we can reduce this to an equivalent operation on the core tensor. Let  $\mathbf{X}$  be a Tucker tensor as in (4.1), and let  $\mathbf{V}^{(m)}$  be a matrix of size  $I_m \times R$  for all  $m \neq n$ . The goal is to calculate

$$\begin{aligned} \mathbf{W} &= \mathbf{X}_{(n)} \left( \mathbf{V}^{(N)} \odot \dots \odot \mathbf{V}^{(n+1)} \odot \mathbf{V}^{(n-1)} \odot \dots \odot \mathbf{V}^{(1)} \right) \\ &= \mathbf{U}^{(n)} \mathbf{G}_{(n)} \left( \mathbf{U}^{(N)} \otimes \dots \otimes \mathbf{U}^{(n+1)} \otimes \mathbf{U}^{(n-1)} \otimes \dots \otimes \mathbf{U}^{(1)} \right)^\top \\ &\quad \left( \mathbf{V}^{(N)} \odot \dots \odot \mathbf{V}^{(n+1)} \odot \mathbf{V}^{(n-1)} \odot \dots \odot \mathbf{V}^{(1)} \right). \end{aligned}$$

Using the properties of the Khatri–Rao product [44] and setting  $\mathbf{W}^{(m)} = \mathbf{U}^{(m)\top} \mathbf{V}^{(m)}$  for  $m \neq n$ , we have

$$\mathbf{W} = \mathbf{U}^{(n)} \underbrace{\mathbf{G}_{(n)} \left( \mathbf{W}^{(N)} \odot \dots \odot \mathbf{W}^{(n+1)} \odot \mathbf{W}^{(n-1)} \odot \dots \odot \mathbf{W}^{(1)} \right)}_{\text{matricized core tensor } \mathcal{G} \text{ times Khatri–Rao product}}.$$

Thus, this requires  $(N - 1)$  matrix-matrix products to form the matrices  $\mathbf{W}^{(m)}$  of size  $J_m \times R$ , each of which costs  $O(I_m J_m R)$ . Then we calculate the MTTKRP with  $\mathcal{G}$ , and the cost is  $O(R \prod_n J_n)$  if  $\mathcal{G}$  is dense. The final matrix-matrix multiplication costs  $O(I_n J_n R)$ . If  $\mathcal{G}$  is dense, the total cost is

$$O \left( R \left( \sum_{n=1}^N I_n J_n + \prod_{n=1}^N J_n \right) \right).$$

**4.2.6. Computing  $\mathbf{X}_{(n)} \mathbf{X}_{(n)}^\top$  for a Tucker tensor.** To compute  $\text{rank}_n(\mathbf{X})$ , we need  $\mathbf{Z} = \mathbf{X}_{(n)} \mathbf{X}_{(n)}^\top$ . Let  $\mathbf{X}$  be a Tucker tensor as in (4.1); then

$$\begin{aligned} \mathbf{Z} &= \mathbf{U}^{(n)} \mathbf{G}_{(n)} \left( \mathbf{U}^{(N)} \otimes \dots \otimes \mathbf{U}^{(n+1)} \otimes \mathbf{U}^{(n-1)} \otimes \dots \otimes \mathbf{U}^{(1)} \right)^\top \\ &\quad \left( \mathbf{U}^{(N)} \otimes \dots \otimes \mathbf{U}^{(n+1)} \otimes \mathbf{U}^{(n-1)} \otimes \dots \otimes \mathbf{U}^{(1)} \right) \mathbf{G}_{(n)}^\top \mathbf{U}^{(n)\top}. \end{aligned}$$

Using the properties of the Kronecker product, this reduces to

$$\mathbf{Z} = \mathbf{U}^{(n)} \mathbf{G}_{(n)} \left( \mathbf{V}^{(N)} \otimes \dots \otimes \mathbf{V}^{(n+1)} \otimes \mathbf{V}^{(n-1)} \otimes \dots \otimes \mathbf{V}^{(1)} \right) \mathbf{G}_{(n)}^T \mathbf{U}^{(n)T},$$

where  $\mathbf{V}^{(m)} = \mathbf{U}^{(m)T} \mathbf{U}^{(m)} \in \mathbb{R}^{J_m \times J_m}$  for all  $m \neq n$  at a cost of  $O(I_m J_m^2)$ . Finally, this becomes the product of two matricized tensors and a third matrix:

$$\mathbf{Z} = \mathbf{H}_{(n)} \mathbf{G}_{(n)}^T \mathbf{U}^{(n)T}, \quad \text{where} \\ \mathcal{H} = \llbracket \mathcal{G}; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(n-1)}, \mathbf{U}^{(n)}, \mathbf{V}^{(n+1)}, \dots, \mathbf{V}^{(N)} \rrbracket.$$

If  $\mathcal{G}$  is dense, forming  $\mathcal{H}$  costs

$$O \left( \prod_{m=1}^N J_m \cdot \left( I_n + \sum_{\substack{m=1 \\ m \neq n}}^N J_m \right) \right).$$

The final multiplication of the three matrices costs  $O(I_n \prod_{m=1}^N J_m + I_n^2 J_n)$ .

**4.3. MATLAB details for Tucker tensors.** A Tucker tensor  $\mathcal{X}$  is constructed in MATLAB by passing in the core array  $\mathcal{G}$  and factor matrices  $\mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)}$  using  $\mathbf{X} = \mathbf{ttensor}(\mathcal{G}, \{\mathbf{U}1, \dots, \mathbf{U}N\})$ . In version 1.0 of the Tensor Toolbox, the `ttensor` class was instead called `tucker_tensor` [4]. The core tensor can be any of the four classes of tensors supported by the Tensor Toolbox.

A Tucker tensor can be converted to a standard `tensor` by calling `full(X)`. Subscripted reference and assignment can be done only on the factors, not elementwise. For example, it is possible to change the (1, 1) element of  $\mathbf{U}^{(2)}$  but not the (1, 1, 1) element of a three-way Tucker tensor  $\mathcal{X}$ . Scalar multiplication is supported, i.e.,  $\mathbf{X} * 5$ .

The  $n$ -mode product of a Tucker tensor with one or more matrices (section 4.2.1) or vectors (section 4.2.2) is implemented in `ttm` and `ttv`, respectively. The inner product (section 4.2.3 and also section 6) is called `innerprod`, and the norm of a Tucker tensor is called `norm`. The function `mttkrp` computes the matricized-tensor-times-Khatri-Rao-product as described in section 4.2.5. The function `nvecs(X, n)` computes the leading mode- $n$  eigenvectors for  $\mathbf{X}_{(n)} \mathbf{X}_{(n)}^T$  and relies on the efficiencies described in section 4.2.6.

**5. Kruskal tensors.** Consider a tensor  $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  that can be written as a sum of  $R$  rank-1 tensors (with no assumption that  $R$  is minimal), i.e.,

$$\mathcal{X} = \sum_{r=1}^R \lambda_r \mathbf{u}_r^{(1)} \circ \dots \circ \mathbf{u}_r^{(N)},$$

where  $\boldsymbol{\lambda} = [\lambda_1 \ \dots \ \lambda_R]^T \in \mathbb{R}^R$  and  $\mathbf{U}^{(n)} = [\mathbf{u}_1^{(n)} \ \dots \ \mathbf{u}_R^{(n)}] \in \mathbb{R}^{I_n \times R}$ . This is the format that results from a PARAFAC decomposition [19, 9], and we refer to it as a *Kruskal tensor* due to the work of Kruskal on tensors of this format [28, 29]. We use the shorthand notation from [25]:

$$(5.1) \quad \mathcal{X} = \llbracket \boldsymbol{\lambda}; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket.$$

In some cases, the weights  $\lambda_r$  are not explicit, and we write  $\mathcal{X} = \llbracket \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket$ . Other notation can be used. For instance, Kruskal [28] uses

$$\mathcal{X} = \left( \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \right).$$

**5.1. Kruskal tensor storage.** Storing  $\mathcal{X}$  as a Kruskal tensor is efficient in terms of storage. In its explicit form,  $\mathcal{X}$  requires storage of

$$\prod_{n=1}^N I_n \quad \text{versus} \quad R \left( 1 + \sum_{n=1}^N I_n \right)$$

elements for the factored form. We do not assume that  $R$  is minimal.

**5.2. Kruskal tensor properties.** The Kruskal tensor is a special case of the Tucker tensor where the core tensor  $\mathcal{G}$  is an  $R \times R \times \dots \times R$  superdiagonal tensor and all of the factor matrices  $\mathbf{U}^{(n)}$  have  $R$  columns.

It is well known that matricized versions of the Kruskal tensor (5.1) have a special form; namely,

$$\mathbf{X}_{(\mathcal{X} \times e : I_N)} = \left( \mathbf{U}^{(r_L)} \odot \dots \odot \mathbf{U}^{(r_1)} \right) \mathbf{\Lambda} \left( \mathbf{U}^{(c_M)} \odot \dots \odot \mathbf{U}^{(c_1)} \right)^\top,$$

where  $\mathbf{\Lambda} = \text{diag}(\lambda)$ . For the special case of mode- $n$  matricization, this reduces to

$$(5.2) \quad \mathbf{X}_{(n)} = \mathbf{U}^{(n)} \mathbf{\Lambda} \left( \mathbf{U}^{(N)} \odot \dots \odot \mathbf{U}^{(n+1)} \odot \mathbf{U}^{(n-1)} \odot \dots \odot \mathbf{U}^{(1)} \right)^\top.$$

Finally, the vectorized version is

$$(5.3) \quad \text{vec}(\mathcal{X}) = \left( \mathbf{U}^{(N)} \odot \dots \odot \mathbf{U}^{(1)} \right) \boldsymbol{\lambda}.$$

**5.2.1. Adding two Kruskal tensors.** Because the Kruskal tensor is a sum of rank-1 tensors, adding two Kruskal tensors together can be viewed as extending that summation over both sets of terms. For instance, consider Kruskal tensors  $\mathcal{X}$  and  $\mathcal{Y}$  of the same size given by:

$$\mathcal{X} = \llbracket \boldsymbol{\lambda}; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket \quad \text{and} \quad \mathcal{Y} = \llbracket \boldsymbol{\sigma}; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket.$$

Adding  $\mathcal{X}$  and  $\mathcal{Y}$  yields

$$\mathcal{X} + \mathcal{Y} = \sum_{r=1}^R \lambda_r \mathbf{u}_r^{(1)} \circ \dots \circ \mathbf{u}_r^{(N)} + \sum_{p=1}^P \sigma_p \mathbf{v}_p^{(1)} \circ \dots \circ \mathbf{v}_p^{(N)}$$

or, alternatively,

$$\mathcal{X} + \mathcal{Y} = \llbracket \begin{bmatrix} \boldsymbol{\lambda} \\ \boldsymbol{\sigma} \end{bmatrix}; [\mathbf{U}^{(1)} \mathbf{V}^{(1)}], \dots, [\mathbf{U}^{(N)} \mathbf{V}^{(N)}] \rrbracket.$$

The work for this is  $O(1)$ .

**5.2.2. Mode- $n$  matrix multiplication for a Kruskal tensor.** Let  $\mathcal{X}$  be a Kruskal tensor as in (5.1) and  $\mathbf{V}$  be a matrix of size  $J \times I_n$ . From the definition of mode- $n$  matrix multiplication and (5.2), we have

$$\mathcal{X} \times_n \mathbf{V} = \llbracket \boldsymbol{\lambda}; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(n-1)}, \mathbf{V} \mathbf{U}^{(n)}, \mathbf{U}^{(n+1)}, \dots, \mathbf{U}^{(N)} \rrbracket.$$

In other words, mode- $n$  matrix multiplication just modifies the  $n$ th factor matrix in the Kruskal tensor. The work is just a matrix-matrix multiplication,  $O(RI_nJ)$ . More generally, if  $\mathbf{V}^{(n)}$  is of size  $J_n \times I_n$  for  $n = 1, \dots, N$ , then

$$\llbracket \mathcal{X}; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket = \llbracket \boldsymbol{\lambda}; \mathbf{V}^{(1)} \mathbf{U}^{(1)}, \dots, \mathbf{V}^{(N)} \mathbf{U}^{(N)} \rrbracket$$

retains the Kruskal tensor format and the work is  $N$  matrix-matrix multiplications for  $O(R \sum_n I_n J_n)$ .

**5.2.3. Mode- $n$  vector multiplication for a Kruskal tensor.** In multiplication of a Kruskal tensor by a vector, the  $n$ th factor matrix necessarily disappears and is absorbed into the weights. Let  $\mathbf{v} \in \mathbb{R}^{I_n}$ , then

$$\mathcal{X} \bar{\times}_n \mathbf{v} = \llbracket \boldsymbol{\lambda} * \mathbf{w} ; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(n-1)}, \mathbf{U}^{(n+1)}, \dots, \mathbf{U}^{(N)} \rrbracket, \quad \text{where } \mathbf{w} = \mathbf{U}^{(n)\top} \mathbf{v}.$$

This operation retains the Kruskal tensor structure (though its order is reduced), and the work is multiplying a matrix times a vector and then a Hadamard product of two vectors, i.e.,  $O(RI_n)$ . More generally, multiplying a Kruskal tensor by a vector  $\mathbf{v}^{(n)} \in \mathbb{R}^{I_n}$  in every mode yields:

$$\begin{aligned} \mathcal{X} \bar{\times}_1 \mathbf{v}^{(1)} \bar{\times}_2 \mathbf{v}^{(2)} \dots \bar{\times}_N \mathbf{v}^{(N)} &= \boldsymbol{\lambda}^\top \left( \mathbf{w}^{(1)} * \mathbf{w}^{(2)} * \dots * \mathbf{w}^{(N)} \right), \\ &\text{where } \mathbf{w}^{(n)} = \mathbf{U}^{(n)\top} \mathbf{v}^{(n)} \text{ for all } n = 1, \dots, N. \end{aligned}$$

Here the final result is a scalar, which is computed by  $N$  matrix-vector products,  $N$  vector Hadamard products, and one vector dot-product, for total work of  $O(R \sum_n I_n)$ .

**5.2.4. Inner product of two Kruskal tensors.** Consider Kruskal tensors  $\mathcal{X}$  and  $\mathcal{Y}$ , both of size  $I_1 \times I_2 \times \dots \times I_N$ , given by:

$$\mathcal{X} = \llbracket \boldsymbol{\lambda} ; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket \quad \text{and} \quad \mathcal{Y} = \llbracket \boldsymbol{\sigma} ; \mathbf{V}^{(1)}, \dots, \mathbf{V}^{(N)} \rrbracket.$$

Assume that  $\mathcal{X}$  has  $R$  rank-1 factors and  $\mathcal{Y}$  has  $S$ . From (5.3), we have

$$\begin{aligned} \langle \mathcal{X}, \mathcal{Y} \rangle &= \langle \text{vec}(\mathcal{X}), \text{vec}(\mathcal{Y}) \rangle \\ &= \boldsymbol{\lambda}^\top \left( \mathbf{U}^{(N)} \odot \dots \odot \mathbf{U}^{(1)} \right)^\top \left( \mathbf{V}^{(N)} \odot \dots \odot \mathbf{V}^{(1)} \right) \boldsymbol{\sigma} \\ &= \boldsymbol{\lambda}^\top \left( \mathbf{U}^{(N)\top} \mathbf{V}^{(N)} * \dots * \mathbf{U}^{(1)\top} \mathbf{V}^{(1)} \right) \boldsymbol{\sigma}. \end{aligned}$$

Note that this does not require that the number of rank-1 factors in  $\mathcal{X}$  and  $\mathcal{Y}$  to be the same. The work is  $N$  matrix-matrix multiplications, plus  $N$  Hadamard products, and a final vector-matrix-vector product. The total work is  $O(RS \sum_n I_n)$ .

**5.2.5. Norm of a Kruskal tensor.** Let  $\mathcal{X}$  be a Kruskal tensor as defined in (5.1). From section 5.2.4, it follows directly that

$$\|\mathcal{X}\|^2 = \langle \mathcal{X}, \mathcal{X} \rangle = \boldsymbol{\lambda}^\top \left( \mathbf{U}^{(N)\top} \mathbf{U}^{(N)} * \dots * \mathbf{U}^{(1)\top} \mathbf{U}^{(1)} \right) \boldsymbol{\lambda},$$

and the total work is  $O(R^2 \sum_n I_n)$ .

**5.2.6. Matricized Kruskal tensor times Khatri–Rao product.** As noted in section 2.6, a common operation is to calculate (2.6). Let  $\mathcal{X}$  be a Kruskal tensor as in (5.1). Also, let  $\mathbf{V}^{(m)}$  be of size  $I_m \times S$  for  $m \neq n$ . In the case of a Kruskal tensor, the operation simplifies to:

$$\begin{aligned} \mathbf{W} &= \mathcal{X} \left( \mathbf{V}^{(N)} \odot \dots \odot \mathbf{V}^{(n+1)} \odot \mathbf{V}^{(n-1)} \odot \dots \odot \mathbf{V}^{(1)} \right) \\ &= \mathbf{U}^{(n)} \boldsymbol{\Lambda} \left( \mathbf{U}^{(N)} \odot \dots \odot \mathbf{U}^{(n+1)} \odot \mathbf{U}^{(n-1)} \odot \dots \odot \mathbf{U}^{(1)} \right)^\top \\ &\quad \left( \mathbf{V}^{(N)} \odot \dots \odot \mathbf{V}^{(n+1)} \odot \mathbf{V}^{(n-1)} \odot \dots \odot \mathbf{V}^{(1)} \right). \end{aligned}$$

Using the properties of the Khatri–Rao product [44] and setting  $\mathbf{A}^{(m)} = \mathbf{U}^{(m)\top} \mathbf{V}^{(m)} \in \mathbb{R}^{R \times S}$  for all  $m \neq n$ , we have

$$\mathbf{W} = \mathbf{U}^{(n)} \mathbf{\Lambda} \left( \mathbf{A}^{(N)} * \dots * \mathbf{A}^{(n+1)} * \mathbf{A}^{(n-1)} * \dots * \mathbf{A}^{(1)} \right).$$

Computing each  $\mathbf{A}^{(m)}$  requires a matrix-matrix product for a cost of  $O(RSI_m)$  for each  $m = 1, \dots, n-1, n+1, \dots, N$ . There is also a sequence of  $N - 1$  Hadamard products of  $R \times S$  matrices, multiplication with an  $R \times R$  diagonal matrix, and finally matrix-matrix multiplication that costs  $O(RSI_n)$ . Thus, the total cost is  $O(RS \sum_n I_n)$ .

**5.2.7. Computing  $\mathbf{X}_{(n)} \mathbf{X}_{(n)}^\top$ .** Let  $\mathcal{X}$  be a Kruskal tensor as in (5.1). We can use the properties of the Khatri–Rao product to efficiently compute

$$\mathbf{Z} = \mathbf{X}^{(n)} \mathbf{X}^{(n)\top} \in \mathbb{R}^{I_n \times I_n}.$$

From (4.4),

$$\begin{aligned} \mathbf{Z} &= \mathbf{U}^{(n)} \mathbf{\Lambda} \left( \mathbf{U}^{(N)} \odot \dots \odot \mathbf{U}^{(n+1)} \odot \mathbf{U}^{(n-1)} \odot \dots \odot \mathbf{U}^{(1)} \right)^\top \\ &\quad \left( \mathbf{U}^{(N)} \odot \dots \odot \mathbf{U}^{(n+1)} \odot \mathbf{U}^{(n-1)} \odot \dots \odot \mathbf{U}^{(1)} \right) \mathbf{\Lambda} \mathbf{U}^{(n)\top}. \end{aligned}$$

This reduces to

$$\mathbf{Z} = \mathbf{U}^{(n)} \mathbf{\Lambda} \left( \mathbf{V}^{(N)} * \dots * \mathbf{V}^{(n+1)} * \mathbf{V}^{(n-1)} * \dots * \mathbf{V}^{(1)} \right) \mathbf{\Lambda} \mathbf{U}^{(n)\top},$$

where  $\mathbf{V}^{(m)} = \mathbf{U}^{(m)\top} \mathbf{U}^{(m)} \in \mathbb{R}^{R \times R}$  for all  $m \neq n$  and costs  $O(R^2 I_m)$ . This is followed by  $(N - 1)$   $R \times R$  matrix Hadamard products and two matrix multiplications. The total work in  $O(R^2 \sum_n I_n)$ .

**5.3. MATLAB details for Kruskal tensors.** A Kruskal tensor  $\mathcal{X}$  from (5.1) is constructed in MATLAB by passing in the matrices  $\mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)}$  and the weighting vector  $\boldsymbol{\lambda}$  using `X = ktensor(lambda, {U1,U2,U3})`. If all of the  $\boldsymbol{\lambda}$ -values are one, then the shortcut `X = ktensor({U1,U2,U3})` can be used instead. In version 1.0 of the Tensor Toolbox, this object was called the `cp_tensor` [4].

A Kruskal tensor can be converted to a standard `tensor` by calling `full(X)`. Subscripted reference and assignment can be done only on the component matrices not elementwise. For example, it is possible to change the 4th element of  $\boldsymbol{\lambda}$  but not the  $(1, 1, 1)$  element of a three-way Kruskal tensor  $\mathcal{X}$ . Scalar multiplication is supported, i.e., `X*5`. It is also possible to add to Kruskal tensors  $(\mathbf{X}+\mathbf{Y}$  or  $\mathbf{X}-\mathbf{Y})$  as described in section 5.2.1.

The  $n$ -mode product of a Kruskal tensor with one or more matrices (section 5.2.2) or vectors (section 5.2.3) is implemented in `ttm` and `ttv`, respectively. The inner product (section 5.2.4 and also section 6) is called via `innerprod`. The norm of a Kruskal tensor (section 5.2.5) is computed by calling `norm`. The function `mttkrp` computes the matricized-tensor-times-Khatri–Rao-product as described in section 5.2.6. The function `nvecs(X,n)` computes the leading mode- $n$  eigenvectors for  $\mathbf{X}_{(n)} \mathbf{X}_{(n)}^\top$  as described in section 5.2.7.

**6. Operations that combine different types of tensors.** Here we consider two operations that combine different types of tensors. Throughout, we work with the following tensors:

- $\mathcal{D}$  is a dense tensor of size  $I_1 \times I_2 \times \cdots \times I_N$ .
- $\mathcal{S}$  is a sparse tensor of size  $I_1 \times I_2 \times \cdots \times I_N$ , and  $\mathbf{v} \in \mathbb{R}^P$  contains its nonzeros.
- $\mathcal{T} = \llbracket \mathcal{G}; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket$  is a Tucker tensor of size  $I_1 \times I_2 \times \cdots \times I_N$  with a core of size  $\mathcal{G} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$  and factor matrices  $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times J_n}$  for all  $n$ .
- $\mathcal{K} = \llbracket \boldsymbol{\lambda}; \mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)} \rrbracket$  is a Kruskal tensor of size  $I_1 \times I_2 \times \cdots \times I_N$  and  $R$  factor matrices  $\mathbf{W}^{(n)} \in \mathbb{R}^{I_n \times R}$ .

**6.1. Inner product.** Here we discuss how to compute the inner product between any pair of tensors of different types.

For a sparse and dense tensor, we have  $\langle \mathcal{D}, \mathcal{S} \rangle = \mathbf{v}^\top \mathbf{z}$ , where  $\mathbf{z}$  is the vector extracted from  $\mathcal{D}$  using the indices of the nonzeros in the sparse tensor  $\mathcal{S}$ .

For a Tucker and dense tensor, if the core of the Tucker tensor is small, we can compute

$$\langle \mathcal{T}, \mathcal{D} \rangle = \langle \mathcal{G}, \hat{\mathcal{D}} \rangle, \quad \text{where } \hat{\mathcal{D}} = \mathcal{D} \times_1 \mathbf{U}^{(1)\top} \cdots \times_n \mathbf{U}^{(N)\top}.$$

Computing  $\hat{\mathcal{D}}$  and its inner product with a dense  $\mathcal{G}$  costs

$$O\left(\sum_{n=1}^N \left(\prod_{p=n}^N I_p \prod_{q=1}^n J_q\right) + \prod_{n=1}^N J_n\right).$$

The procedure is the same for a Tucker tensor and a sparse tensor, i.e.,  $\langle \mathcal{T}, \mathcal{S} \rangle$ , though the cost is different (see section 3.2.5).

For the inner product of a Kruskal tensor and a dense tensor, we have

$$\langle \mathcal{D}, \mathcal{K} \rangle = \text{vec}(\mathcal{D})^\top \left( \mathbf{U}^{(N)} \odot \cdots \odot \mathbf{U}^{(1)} \right) \boldsymbol{\lambda}.$$

The cost of forming the Khatri–Rao product dominates:  $O(R \prod_n I_n)$ .

The inner product of a Kruskal tensor and a sparse tensor can be written as

$$\langle \mathcal{S}, \mathcal{K} \rangle = \sum_{r=1}^R \lambda_r (\mathcal{S} \bar{\times}_1 \mathbf{w}_r^{(1)} \cdots \bar{\times}_N \mathbf{w}_r^{(N)}).$$

Consequently, the cost is equivalent to doing  $R$  tensor-times-vector products with  $N$  vectors each, i.e.,  $O(RN \cdot \text{nnz}(\mathcal{S}))$ . The same reasoning applies to the inner product of Tucker and Kruskal tensors  $\langle \mathcal{T}, \mathcal{K} \rangle$ .

**6.2. Hadamard product.** We consider the Hadamard product of a sparse tensor with dense and Kruskal tensors.

The product  $\mathcal{Y} = \mathcal{D} * \mathcal{S}$  necessarily has zeros everywhere that  $\mathcal{S}$  is zero, so only the potential nonzeros in the result, corresponding to the nonzeros in  $\mathcal{S}$ , need to be computed. The result is assembled from the nonzero subscripts of  $\mathcal{S}$  and  $\mathbf{v} * \mathbf{z}$ , where  $\mathbf{z}$  is the values of  $\mathcal{D}$  at the nonzero subscripts of  $\mathcal{S}$ . The work is  $O(\text{nnz}(\mathcal{S}))$ .

Once again,  $\mathcal{Y} = \mathcal{S} * \mathcal{K}$  can have nonzeros only where  $\mathcal{S}$  has nonzeros. Let  $\mathbf{z} \in \mathbb{R}^P$  be the vector of possible nonzeros for  $\mathcal{Y}$  corresponding to the locations of the nonzeros in  $\mathcal{S}$ . Observe that

$$z_p = v_p \left( \sum_{r=1}^R \lambda_r w_{r,s_{1p}}^{(1)} w_{r,s_{2p}}^{(2)} \cdots w_{r,s_{Np}}^{(N)} \right).$$

This means that we can compute it vectorwise by a sum of a series of vector Hadamard products with “expanded” vectors as in section 3.2.4, for example. The work is  $O(N \cdot \text{nnz}(\mathbf{S}))$ .

**7. Conclusions.** In this article, we considered the question of how to deal with potentially large-scale tensors stored in sparse or factored (Tucker or Kruskal) form. The Tucker and Kruskal formats can be used, for example, to store the results of a Tucker or CANDECOMP/PARAFAC decomposition of a large, sparse tensor. We demonstrated relevant mathematical properties of structured tensors that simplify common operations appearing in tensor decomposition algorithms, such as mode- $n$  matrix/vector multiplication, inner product, and collapsing/scaling. For many functions, we are able to realize substantial computational efficiencies as compared to working with the tensors in dense/unfactored form.

The Tensor Toolbox provides an extension to MATLAB by adding the ability to work with sparse multidimensional arrays, not to mention the specialized factored tensors. Moreover, relatively few packages in any language have the ability to work with sparse tensors, and our investigations have not revealed any others that have the variety of capabilities available in the Tensor Toolbox. A complete listing of functions for dense (`tensor`), sparse (`sptensor`), Tucker (`ttensor`), and Kruskal (`ktensor`) tensors is provided in Table 7.1. In general, Tensor Toolbox objects work the same as MATLAB arrays. For example, for a 3-way tensor  $\mathcal{A}$  in any format (`tensor`, `sptensor`, `ktensor`, `ttensor`), it is possible to call functions such as `size(A)`, `ndims(A)`, `permute(A, [3 2 1])`, `-A`, `2*A`, and `norm(A)` (always the Frobenius norm for tensors). A major difference between Tensor Toolbox objects and MATLAB arrays is that the tensor classes support subscript indexing (i.e., passing in a matrix of subscripts) and do not support linear indexing. This avoids possible complications with integer overflow for large-scale arrays; see section 3.3.

Due to their structure, factored tensors cannot support every operation that is supported for dense and sparse tensors. For instance, most element-level operations are prohibited, such as subscripted reference/assignment, logical operations/comparisons, etc. In these cases, memory permitting, the factored tensors can be converted to dense tensors by calling `full`. However, there are certain operations that can be adapted to the structure. For example, it is possible to add two Kruskal tensors, as described in section 5.2.1, and it is possible to do tensor multiplication and inner products involving Kruskal tensors; see section 6.

A major feature of the Tensor Toolbox is that it defines multiplication on tensor objects. For example, generalized tensor-tensor multiplication and contraction is supported for dense and sparse tensors. The specialized operations of  $n$ -mode multiplication of a tensor by a matrix or a vector is supported for dense, sparse, and factored tensors. Likewise, inner products, even between tensors of different types, and norms are supported across the board.

The Tensor Toolbox also includes specialized functions, such as `collapse` and `scale` (see section 3.2.9), the matricized-tensor-times-Khatri-Rao-product (see section 2.6), the computation of the leading mode- $n$  singular vectors (equivalent to the leading eigenvectors of  $\mathbf{X}_{(n)}\mathbf{X}_{(n)}^T$ ), and conversion of a tensor to a matrix.

While we believe that the Tensor Toolbox is a useful package, we look forward to greater availability of storage formats and increased functionality in software for tensors, especially sparse tensors. For instance, the benefits of storing matrices in sorted order using CSR or CSC format generally outweigh the negatives, and so it

TABLE 7.1  
*Methods in the Tensor Toolbox.*

<i>Functionality</i>	<b>tensor</b>	<b>sptensor</b>	<b>ttensor</b>	<b>ktensor</b>
Sizes ( <b>size</b> , <b>ndims</b> )	✓	✓	✓	✓
Number of nonzeros ( <b>nnz</b> )	✓ <sup>d</sup>	✓	–	–
Permute	✓	✓	✓	✓
Remove singleton dimensions ( <b>squeeze</b> )	✓	✓	–	–
Subscripted reference and assignment	✓ <sup>a</sup>	✓ <sup>a</sup>	✓ <sup>b</sup>	✓ <sup>b</sup>
Unary plus and minus (e.g., $-\mathbf{X}$ )	✓	✓	✓	✓
Plus and minus	✓	✓	–	✓
Logical ( <b>and/or/xor/not</b> )	✓	✓ <sup>d</sup>	–	–
Comparisons ( <b>eq/ne/gt/ge/lt/le</b> )	✓	✓ <sup>d</sup>	–	–
Scalar multiplication ( <b>A*5</b> )	✓	✓	✓	✓
Scalar elementwise power ( <b>A.<sup>5</sup></b> )	✓	✓ <sup>d</sup>	–	✓
Array (Hadamard) multiplication ( <b>A.*B</b> )	✓ <sup>c</sup>	✓ <sup>c</sup>	–	✓ <sup>c</sup>
Array right division ( <b>A./B</b> )	✓ <sup>c</sup>	✓ <sup>c</sup>	–	✓ <sup>c</sup>
Convert to multidimensional array (MDA) ( <b>double</b> )	✓	✓	✓	✓
Convert to dense ( <b>full</b> )	✓	✓	✓	✓
Find subscripts of nonzero elements ( <b>find</b> )	✓	✓	–	–
Apply a function to every element ( <b>tenfun</b> )	✓	–	–	–
Apply a function to every nonzero element ( <b>elemfun</b> )	–	✓	–	–
Tensor times tensor ( <b>ttt</b> )	✓	✓	–	–
Generalized trace ( <b>contract</b> )	✓	✓	–	–
Tensor time matrix ( <b>ttm</b> )	✓	✓	✓	✓
Tensor times vector ( <b>ttv</b> )	✓	✓	✓	✓
Matricized-tensor-times-Khatri-Rao-product ( <b>mttkrp</b> )	✓	✓	✓	✓
Mode- $n$ singular vectors ( <b>nvecs</b> )	✓	✓	✓	✓
Inner product ( <b>innerprod</b> )	✓ <sup>c</sup>	✓ <sup>c</sup>	✓ <sup>c</sup>	✓ <sup>c</sup>
Norm	✓	✓	✓	✓
Collapse along dimensions	✓	✓	–	–
Scale along dimensions	✓	✓	–	–
Matricize	✓	✓	–	–

<sup>a</sup> Multiple subscripts passed explicitly (no linear indices).

<sup>b</sup> Only the factors may be referenced/modified.

<sup>c</sup> Supports combinations of different types of tensors.

<sup>d</sup> New as of version 2.1.

makes sense to seek multidimensional extensions that are both practical and useful, at least for specialized contexts as with the EKMR [33, 34].

Furthermore, extensions to parallel data structures and architectures require further innovation, especially as we hope to leverage existing codes for parallel linear algebra.

**Acknowledgments.** We gratefully acknowledge all of those who have influenced the development of the Tensor Toolbox through their conversations and email exchanges with us—you have helped us to make this a much better package. In particular, we thank Evrim Acar, Rasmus Bro, Jerry Gregoire, Richard Harshman, Morten Mørup, Teresa Seele, and Giorgio Tomasi. We also thank Jimeng Sun for being a beta tester and using the results in [45]. We thank the referees for their constructive comments, which have greatly improved the manuscript.

## REFERENCES

- [1] E. ACAR, S. A. ÇAMTEPE, AND B. YENER, *Collective sampling and analysis of high order tensors for chatroom communications*, in ISI 2006: IEEE International Conference on Intelligence and Security Informatics, Lecture Notes in Comput. Sci. 3975, Springer, Berlin, 2006, pp. 213–224.

- [2] C. A. ANDERSSON AND R. BRO, *The N-way toolbox for MATLAB*, Chemometr. Intell. Lab., 52 (2000), pp. 1–4. See also <http://www.models.kvl.dk/source/nwaytoolbox/>.
- [3] C. J. APPELLOF AND E. R. DAVIDSON, *Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents*, Anal. Chem., 53 (1981), pp. 2053–2056.
- [4] B. W. BADER AND T. G. KOLDA, *Algorithm 862: MATLAB tensor classes for fast algorithm prototyping*, ACM Trans. Math. Software, 32 (2006), pp. 635–653.
- [5] B. W. BADER AND T. G. KOLDA, *Matlab Tensor Toolbox, Version 2.1*, <http://csmr.ca.sandia.gov/~tgkolda/TensorToolbox/> (December 2006).
- [6] G. BEYLKIN AND M. J. MOHLENKAMP, *Algorithms for numerical analysis in high dimensions*, SIAM J. Sci. Comput., 26 (2005), pp. 2133–2159.
- [7] R. BRO, *PARAFAC. Tutorial and applications*, Chemometr. Intell. Lab., 38 (1997), pp. 149–171.
- [8] R. BRO, *Multi-way Analysis in the Food Industry: Models, Algorithms, and Applications*, Ph.D. thesis, University of Amsterdam, 1998. Available at <http://www.models.kvl.dk/research/theses/>.
- [9] J. D. CARROLL AND J. J. CHANG, *Analysis of individual differences in multidimensional scaling via an N-way generalization of ‘Eckart-Young’ decomposition*, Psychometrika, 35 (1970), pp. 283–319.
- [10] B. CHEN, A. PETROPOLU, AND L. DE LATHAUWER, *Blind identification of convolutive MIM systems with 3 sources and 2 sensors*, Applied Signal Processing, (2002), pp. 487–496 (Special Issue on Space-Time Coding and Its Applications, Part II).
- [11] P. COMON, *Tensor decompositions: State of the art and applications*, in Mathematics in Signal Processing V, J. G. McWhirter and I. K. Proudler, eds., Oxford University Press, Oxford, 2001, pp. 1–24.
- [12] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1253–1278.
- [13] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *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 (2000), pp. 1324–1342.
- [14] I. S. DUFF AND J. K. REID, *Some design features of a sparse matrix code*, ACM Trans. Math. Software, 5 (1979), pp. 18–35.
- [15] R. GARCIA AND A. LUMSDAINE, *MultiArray: A C++ library for generic programming with arrays*, Software: Practice and Experience, 35 (2004), pp. 159–188.
- [16] J. R. GILBERT, C. MOLER, AND R. SCHREIBER, *Sparse matrices in MATLAB: Design and implementation*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 333–356.
- [17] V. S. GRIGORASCU AND P. A. REGALIA, *Tensor displacement structures and polyspectral matching*, in Fast Reliable Algorithms for Matrices with Structure, T. Kaliath and A. H. Sayed, eds., SIAM, Philadelphia, 1999, pp. 245–276.
- [18] F. G. GUSTAVSON, *Some basic techniques for solving sparse systems*, in Sparse Matrices and their Applications, D. J. Rose and R. A. Willoughby, eds., Plenum Press, New York, 1972, pp. 41–52.
- [19] 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://publish.uwo.ca/~harshman/wpppfac0.pdf>.
- [20] R. HENRION, *Body diagonalization of core matrices in three-way principal components analysis: Theoretical bounds and simulation*, J. Chemometr., 7 (1993), pp. 477–494.
- [21] R. HENRION, *N-way principal component analysis theory, algorithms and applications*, Chemometr. Intell. Lab., 25 (1994), pp. 1–23.
- [22] H. A. KIERS, *Joint orthomax rotation of the core and component matrices resulting from three-mode principal components analysis*, J. Classification, 15 (1998), pp. 245–263.
- [23] H. A. L. KIERS, *Towards a standardized notation and terminology in multiway analysis*, J. Chemometr., 14 (2000), pp. 105–122.
- [24] T. G. KOLDA, *Orthogonal tensor decompositions*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 243–255.
- [25] T. G. KOLDA, *Multilinear Operators for Higher-Order Decompositions*, Technical report SAND2006-2081, Sandia National Laboratories, Albuquerque, NM and Livermore, CA, 2006.
- [26] P. KROONENBERG, *Applications of three-mode techniques: Overview, problems, and prospects (slides)*, Presentation at the AIM Tensor Decompositions Workshop, Palo Alto, CA, 2004. Available at <http://csmr.ca.sandia.gov/~tgkolda/tdw2004/Kroonenberg%20-%20Talk.pdf>.

- [27] P. M. KROONENBERG AND J. DE LEEUW, *Principal component analysis of three-mode data by means of alternating least squares algorithms*, Psychometrika, 45 (1980), pp. 69–97.
- [28] J. B. KRUSKAL, *Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics*, Linear Algebra Appl., 18 (1977), pp. 95–138.
- [29] J. B. KRUSKAL, *Rank, decomposition, and uniqueness for 3-way and  $N$ -way arrays*, in Multiway Data Analysis, R. Coppi and S. Bolasco, eds., North-Holland, Amsterdam, 1989.
- [30] W. LANDRY, *Implementing a high performance tensor library*, Scientific Programming, 11 (2003), pp. 273–290.
- [31] S. LEURGANS AND R. T. ROSS, *Multilinear models: Applications in spectroscopy*, Statist. Sci., 7 (1992), pp. 289–310.
- [32] L.-H. LIM, *Singular values and eigenvalues of tensors: A variational approach*, in CAMAP2005: 1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, 2005, pp. 129–132.
- [33] C.-Y. LIN, Y.-C. CHUNG, AND J.-S. LIU, *Efficient data compression methods for multidimensional sparse array operations based on the ekmr scheme*, IEEE Trans. Comput., 52 (2003), pp. 1640–1646.
- [34] C.-Y. LIN, J.-S. LIU, AND Y.-C. CHUNG, *Efficient representation scheme for multidimensional array operations*, IEEE Trans. Comput., 51 (2002), pp. 327–345.
- [35] R. P. McDONALD, *A simple comprehensive model for the analysis of covariance structures*, British J. Math. Statist. Psych., 33 (1980), p. 161. Cited in [8].
- [36] M. MØRUP, L. HANSEN, J. PARNAS, AND S. M. ARNFRED, *Decomposing the Time-Frequency Representation of EEG Using Nonnegative Matrix and Multi-way Factorization*, [http://www2.imm.dtu.dk/pubdb/views/edoc\\_download.php/4144/pdf/imm4144.pdf](http://www2.imm.dtu.dk/pubdb/views/edoc_download.php/4144/pdf/imm4144.pdf) (2006).
- [37] T. MURAKAMI, J. M. F. TEN BERGE, AND H. A. L. KIERS, *A case of extreme simplicity of the core matrix in three-mode principal components analysis*, Psychometrika, 63 (1998), pp. 255–261.
- [38] P. PAATERO, *The multilinear engine - a table-driven, least squares program for solving multilinear problems, including the  $n$ -way parallel factor analysis model*, J. Comput. Graph. Statist., 8 (1999), pp. 854–888.
- [39] U. W. POOCH AND A. NIEDER, *A survey of indexing techniques for sparse matrices*, ACM Computing Surveys, 5 (1973), pp. 109–133.
- [40] C. R. RAO AND S. MITRA, *Generalized Inverse of Matrices and its Applications*, Wiley, New York, 1971. Cited in [8].
- [41] J. R. RUÍZ-TOLOSA AND E. CASTILLO, *From Vectors to Tensors*, Universitext, Springer, Berlin, 2005.
- [42] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003.
- [43] B. SAVAS, *Analyses and Tests of Handwritten Digit Recognition Algorithms*, master’s thesis, Linköping University, Sweden, 2003.
- [44] A. SMILDE, R. BRO, AND P. GELADI, *Multi-way Analysis: Applications in the Chemical Sciences*, Wiley, West Sussex, England, 2004.
- [45] J. SUN, D. TAO, AND C. FALOUTSOS, *Beyond streams and graphs: Dynamic tensor analysis*, in KDD ’06: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2006, pp. 374–383.
- [46] J.-T. SUN, H.-J. ZENG, H. LIU, Y. LU, AND Z. CHEN, *CubeSVD: A novel approach to personalized web search*, in WWW 2005: Proceedings of the 14th International Conference on World Wide Web, ACM Press, New York, 2005, pp. 382–390.
- [47] J. TEN BERGE, J. DE LEEUW, AND P. M. KROONENBERG, *Some additional results on principal components analysis of three-mode data by means of alternating least squares algorithms*, Psychometrika, 52 (1987), pp. 183–191.
- [48] J. M. F. TEN BERGE AND H. A. L. KIERS, *Simplicity of core arrays in three-way principal component analysis and the typical rank of  $p \times q \times 2$  arrays*, Linear Algebra Appl., 294 (1999), pp. 169–179.
- [49] G. TOMASI, *Use of the Properties of the Khatri-Rao Product for the Computation of Jacobian, Hessian, and Gradient of the PARAFAC Model under MATLAB*, manuscript, 2005.
- [50] G. TOMASI AND R. BRO, *A comparison of algorithms for fitting the PARAFAC model*, Comput. Statist. Data Anal., 50 (2006), pp. 1700–1734.
- [51] L. R. TUCKER, *Some mathematical notes on three-mode factor analysis*, Psychometrika, 31 (1966), pp. 279–311.
- [52] M. A. O. VASILESCU AND D. TERZOPOULOS, *Multilinear analysis of image ensembles: TensorFaces*, in ECCV 2002: 7th European Conference on Computer Vision, Lecture Notes in Comput. Sci. 2350, Springer, Berlin, 2002, pp. 447–460.

- [53] D. VLASIC, M. BRAND, H. PFISTER, AND J. POPOVIĆ, *Face transfer with multilinear models*, ACM Trans. Graphics, 24 (2005), pp. 426–433.
- [54] H. WANG AND N. AHUJA, *Facial expression decomposition*, in ICCV 2003: 9th IEEE International Conference on Computer Vision, vol. 2, 2003, pp. 958–965.
- [55] B. M. WISE AND N. B. GALLAGHER, *PLS\_Toolbox 4.0*, <http://www.eigenvector.com> (2007).
- [56] R. ZASS, *HUJI Tensor Library*, <http://www.cs.huji.ac.il/~zass/html/> (May 2006).
- [57] E. ZHANG, J. HAYS, AND G. TURK, *Interactive Tensor Field Design and Visualization on Surfaces*, <http://eecs.oregonstate.edu/library/files/2005-106/tenflddesn.pdf> (2005).
- [58] T. ZHANG AND G. H. GOLUB, *Rank-one approximation to high order tensors*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 534–550.