

A CONSTRUCTIVE ALGORITHM FOR DECOMPOSING A TENSOR INTO A FINITE SUM OF ORTHONORMAL RANK-1 TERMS*

KIM BATSELIER[†], HAOTIAN LIU[†], AND NGAI WONG[†]

Abstract. We propose a constructive algorithm that decomposes an arbitrary real tensor into a finite sum of orthonormal rank-1 outer products. The algorithm, called TTr1SVD, works by converting the tensor into a tensor-train rank-1 (TTr1) series via the singular value decomposition (SVD). TTr1SVD naturally generalizes the SVD to the tensor regime with properties such as uniqueness for a fixed order of indices, orthogonal rank-1 outer product terms, and easy truncation error quantification. Using an outer product column table it also allows, for the first time, a complete characterization of all tensors orthogonal with the original tensor. Incidentally, this leads to a strikingly simple constructive proof showing that the maximum rank of a real $2 \times 2 \times 2$ tensor over the real field is 3. We also derive a conversion of the TTr1 decomposition into a Tucker decomposition with a sparse core tensor. Numerical examples illustrate each of the favorable properties of the TTr1 decomposition.

Key words. tensor decompositions, multiway arrays, singular values, orthogonal rank-1 terms, CANDECOMP/PARAFAC decomposition

AMS subject classifications. 15A69, 15A18, 15A23

DOI. 10.1137/141000658

1. Introduction. There has been a recent surge in the research and utilization of tensors, which are high-order generalization of matrices, and their low-rank approximations [1, 4, 9, 15, 10]. This is due to their natural form to capture high dimensional problems and their efficient compact representation of large-scale data sets.

Among various tensor decompositions, the CANDECOMP/PARAFAC (CP) decomposition¹ [4, 6, 9] has found widespread use. CP expresses a tensor as the sum of a finite number of rank-1 tensors, called outer products, so that the tensor (CP-)rank can be defined as the minimum number of terms in the decomposition. Although CP is regarded as the generalization of the matrix singular value decomposition (SVD) to tensors, unlike matrices, there are no feasible algorithms to determine the rank of a specific tensor. Furthermore, most existing CP algorithms are optimization based, such as the “workhorse” algorithm for CP: the alternating least squares- (ALS-) CP method [4]. ALS-CP minimizes the error between the original tensor and its rank- R approximation (viz., sum of R outer products) in an iterative procedure. The main problem of ALS-CP is that it only works by prescribing the rank R ; therefore the procedure itself does not directly identify the tensor rank. Moreover, the outer products generated by ALS-CP are not orthogonal with each other, unlike the case for matrix singular vectors.

*Received by the editors December 18, 2014; accepted for publication (in revised form) by P. Comon June 17, 2015; published electronically September 10, 2015. This work was supported in part by the Hong Kong Research Grants Council under General Research Fund (GRF) Projects 718213E and 17208514, and the University Research Committee of The University of Hong Kong.

<http://www.siam.org/journals/simax/36-3/100065.html>

[†]Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong, Hong Kong (kim.batselier@gmail.com, htliu@eee.hku.hk, nwong@eee.hku.hk).

¹Originally introduced by Hitchcock [7], the decomposition was rediscovered independently as CANDECOMP (CANonical DECOMPosition) by Carroll and Chang [4] and PARAFAC (PARAllel FACtors) by Harshman [6]. The underlying algorithms are, however, the same.

Other tensor decompositions, for example the Tucker decomposition [4, 20], compress a tensor into a core tensor and several factor matrices. The Tucker decomposition of a tensor is not unique. One of its realizations can be efficiently computed by the higher-order SVD [5]. Each element in its core tensor can be deemed as the weight of a rank-1 factor. In this interpretation, all rank-1 factors of the Tucker decomposition are orthonormal. Nonetheless, the Tucker decomposition is not necessarily canonical and therefore cannot be used to estimate tensor ranks.

To this end, a constructive orthogonal tensor decomposition algorithm, called the tensor-train rank-1 (TTr1) SVD (TTr1SVD), is proposed in this paper. The recent introduction of the tensor-train (TT) decomposition [15] provides a constructive approach to represent and possibly compress tensors. Similar to the TT decomposition, the TTr1 decomposition reshapes and factorizes the tensor in a recursive way. However, unlike the TT decomposition, one needs to progressively reshape and compute the SVD of each singular vector to produce the TTr1 decomposition. The resulting singular values are constructed into a tree structure whereby the product of each branch is the weight of one orthonormal (rank-1) outer product. Most of the main properties and contributions of the TTr1 decomposition are highly reminiscent of the matrix SVD:

1. An arbitrary tensor is for a fixed order of the indices *uniquely* decomposed into a linear combination of orthonormal outer products, each associated with a nonnegative TTr1 singular value.
2. The approximation error of an R -term approximation is easily quantified in terms of the singular values.
3. Numerical stability of the algorithm is due to the use of consecutive SVDs.
4. It characterizes the orthogonal complement tensor space that contains all tensors whose inner product is 0 with the original tensor \mathcal{A} . This orthogonal complement tensor space is, to our knowledge, new in the literature.
5. It is straightforward to convert the TTr1 decomposition into the Tucker format with a sparse core tensor and orthogonal matrix factors.

Having developed TTr1SVD, we found that its core routine turns out to be an independent re-derivation of the PARATREE algorithm [16]. However, TTr1SVD bears the physical insight of enforcing a rank-1 constraint onto the TT decomposition [15]. Such a TT rank-1 perspective provides a much more straightforward appreciation of the favorable properties of this orthogonal SVD-like tensor decomposition. In particular, we provide a significantly more in-depth treatment of TTr1 decomposition than in [16], leading to important new results, such as a perturbation analysis of the singular values, a direct conversion of the TTr1 to the Tucker format featuring a sparse core tensor, and a full characterization of orthogonal complement tensors. Specifically, we introduce a TTr1-based tabulation of all orthogonal outer products that span a tensor \mathcal{A} , as well its orthogonal complement space $\text{opspan}(\mathcal{A})^\perp$ that is proposed for the first time in the literature. This permits, as an immediate application, an elegant and constructive proof that the rank of a real $2 \times 2 \times 2$ tensor over the real field is maximally 3. A MATLAB/Octave implementation of our TTr1SVD algorithm can be freely downloaded and modified from <https://github.com/kbatseli/TTr1SVD>.

The outline of this paper is as follows. First, we introduce some notation and definitions in section 1.1. Section 2 presents a brief overview of the TT decomposition together with a detailed explanation of our TTr1 decomposition. Properties of the TTr1 decomposition such as uniqueness, orthogonality, approximation errors, orthogonal complement tensor space, perturbation of singular values, and Tucker

conversion are discussed in section 3. These properties are illustrated in section 4 by means of several numerical examples. Section 5 concludes and summarizes the contributions.

1.1. Notation and definitions. We will adopt the following notational conventions. A d th-order tensor, assumed real throughout this paper, is a multiway array $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_d}$ with elements $\mathcal{A}_{i_1 i_2 \cdots i_d}$ that can be perceived as an extension of a matrix to its general d th-order, also called d -way, counterpart. We consider only real tensors because we adopt an application point of view. This is, however, without loss of generality; one could easily consider tensors over \mathbb{C} , which would require the replacement of the transpose by the conjugate transpose. Although the words “order” and “dimension” seem to be interchangeable in the tensor community, we prefer to call the number of indices i_k ($k = 1, \dots, d$) the order of the tensor, while the maximal value n_k ($k = 1, \dots, d$) associated with each index is called the dimension. A cubical tensor is a tensor for which $n_1 = n_2 = \cdots = n_d = n$. The k -mode product of a tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_d}$ with a matrix $U \in \mathbb{R}^{p_k \times n_k}$ is defined by

$$(\mathcal{A}_{\times_k} U)_{i_1 \cdots i_{k-1} j_k i_{k+1} \cdots i_d} = \sum_{i_k=1}^{n_k} U_{j_k i_k} \mathcal{A}_{i_1 \cdots i_k \cdots i_d},$$

so that $\mathcal{A}_{\times_k} U \in \mathbb{R}^{n_1 \times \cdots \times n_{k-1} \times p_k \times n_{k+1} \times \cdots \times n_d}$. The inner product between two tensors $\mathcal{A}, \mathcal{B} \in \mathbb{R}^{n_1 \times \cdots \times n_d}$ is defined as

$$\langle \mathcal{A}, \mathcal{B} \rangle = \sum_{i_1, i_2, \dots, i_d} \mathcal{A}_{i_1 i_2 \cdots i_d} \mathcal{B}_{i_1 i_2 \cdots i_d}.$$

The norm of a tensor is taken to be the Frobenius norm $\|\mathcal{A}\|_F = \langle \mathcal{A}, \mathcal{A} \rangle^{1/2}$. The vectorization of a tensor \mathcal{A} , denoted $\text{vec}(\mathcal{A}) \in \mathbb{R}^{n_1 \cdots n_d}$, is the vector obtained from taking all indices together in one mode. A third-order rank-1 tensor can always be written as the outer product [9],

$$\sigma(a \circ b \circ c) \quad \text{with components} \quad \mathcal{A}_{i_1 i_2 i_3} = \sigma a_{i_1} b_{i_2} c_{i_3}$$

with $\sigma \in \mathbb{R}$, whereas a , b , and c are vectors of arbitrary lengths as depicted in Figure 1. Similarly, any d -way rank-1 tensor can be written as an outer product of d vectors. Using the k -mode multiplication, this outer product can also be written as $\sigma_{\times_1} a_{\times_2} b_{\times_3} c$, where σ is now regarded as a $1 \times 1 \times 1$ tensor. In order to facilitate the discussion of the TTr1 decomposition we will make use of a running example tensor $\mathcal{A} \in \mathbb{R}^{3 \times 4 \times 2}$ shown in Figure 2.





FIG. 1. The outer product of three vectors a, b, c of arbitrary lengths forming a rank-1 outer product.

FIG. 2. A running $3 \times 4 \times 2$ tensor example.

2. TTr1 decomposition.

2.1. TT decomposition. Our decomposition is directly inspired by the TT decomposition [15], which we will succinctly review here. The main idea of the TT decomposition is to re-express a tensor \mathcal{A} as

$$(2.1) \quad \mathcal{A}_{i_1 i_2 \dots i_d} = \mathcal{G}_1(i_1) \mathcal{G}_2(i_2) \cdots \mathcal{G}_d(i_d),$$

where for a fixed i_k each $\mathcal{G}_k(i_k)$ is an $r_{k-1} \times r_k$ matrix, also called the TT core. Note that the subscript k of a core \mathcal{G}_k indicates the k th core of the TT decomposition. The ranks r_k are called the TT ranks. Each core \mathcal{G}_k is in fact a third-order tensor with indices $\alpha_{k-1}, i_k, \alpha_k$ and dimensions r_{k-1}, n_k, r_k , respectively. Since $\mathcal{A}_{i_1 i_2 \dots i_d}$ is a scalar we obviously have that $r_0 = r_d = 1$ and for this reason α_0 and α_d are omitted. Consequently, we can write the elements of \mathcal{A} as

$$(2.2) \quad \mathcal{A}_{i_1 i_2 \dots i_d} = \sum_{\alpha_1, \dots, \alpha_{d-1}} \mathcal{G}_1(i_1, \alpha_1) \mathcal{G}_2(\alpha_1, i_2, \alpha_2) \cdots \mathcal{G}_d(\alpha_{d-1}, i_d),$$

where we always need to sum over the auxiliary indices $\alpha_1, \dots, \alpha_{d-1}$, and therefore (2.2) is equivalent to the matrix product form in (2.1). An approximation of \mathcal{A} is achieved by truncating the α_k indices in (2.2) at values smaller than the TT ranks r_k .

Computing the TT decomposition consists of doing $d - 1$ consecutive reshapings and SVD computations. For our running example $\mathcal{A} \in \mathbb{R}^{3 \times 4 \times 2}$ in Figure 2, this means that the decomposition is computed in two steps. This process is visualized in Figure 3, whereby \mathcal{A} will eventually be converted into its TT format in Figure 4. Referring to Figure 3, the first reshaping of \mathcal{A} into a matrix $\bar{\mathcal{A}}$ that needs to be considered is by grouping the indices i_2, i_3 together. This results in the 3×8 matrix

$$\bar{\mathcal{A}} = \begin{pmatrix} 1 & 4 & 7 & 10 & 13 & 16 & 19 & 22 \\ 2 & 5 & 8 & 11 & 14 & 17 & 20 & 23 \\ 3 & 6 & 9 & 12 & 15 & 18 & 21 & 24 \end{pmatrix}.$$

The “economical” SVD of the 3×8 matrix $\bar{\mathcal{A}}$ is then

$$(2.3) \quad \bar{\mathcal{A}} = U_1 S_1 V_1^T$$

with U_1 a 3×3 matrix and V_1 an 8×3 matrix. In fact, any dyadic decomposition can be used for this step in the TT algorithm, but the SVD is often chosen for its numerical stability. The first TT core \mathcal{G}_1 is given by the 3×3 matrix U_1 indexed by i_1, α_1 . We now form the matrix $Y_1 = S_1 V_1^T$ and reshape it such that its rows are

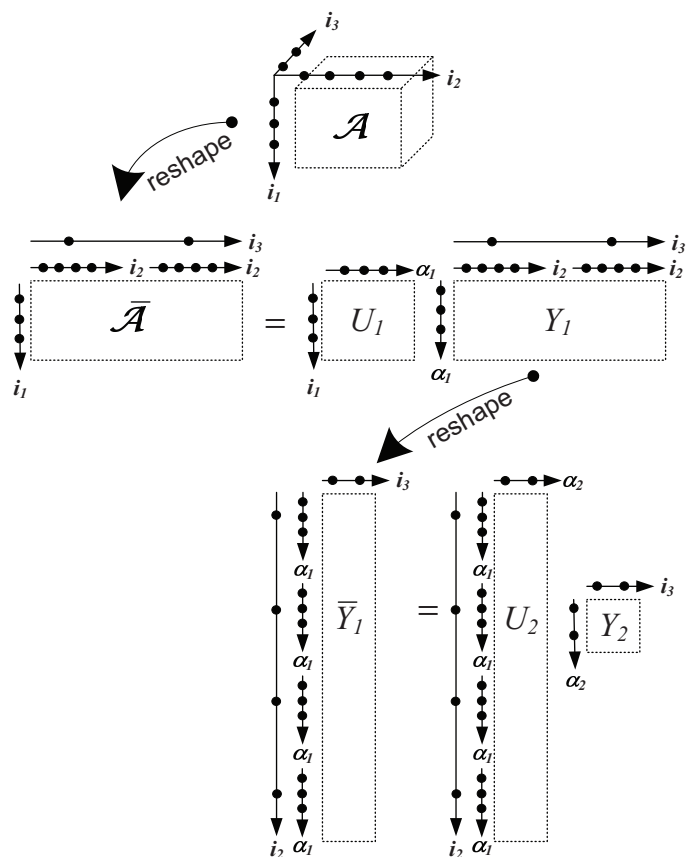


FIG. 3. Computation of the TT decomposition of \mathcal{A} .




FIG. 4. TT decomposition of \mathcal{A} . Each set of orthogonal axes represents a core \mathcal{G}_k of the TT.

indexed by α_1, i_2 and its columns by i_3 . This results in a 12×2 matrix \bar{Y}_1 and its SVD

$$\bar{Y}_1 = U_2 Y_2$$

with $Y_2 = S_2 V_2^T$. The second TT core \mathcal{G}_2 is then given by U_2 , reshaped into a $3 \times 4 \times 2$ tensor. The last TT core \mathcal{G}_3 is then Y_2 , which is a 2×2 matrix indexed by α_2 and i_3 .

We therefore have that

$$\mathcal{A}_{i_1 i_2 i_3} = \mathcal{G}_1(i_1) \mathcal{G}_2(i_2) \mathcal{G}_3(i_3)$$

with $\mathcal{G}_1(i_1)$ a 1×3 row vector, $\mathcal{G}_2(i_2)$ a 3×2 matrix, and $\mathcal{G}_3(i_3)$ a 2×1 column vector for fixed i_1, i_2 , and i_3 , respectively (cf. Figure 4). Observe how the auxiliary indices α_1, α_2 serve as “links” connecting adjacent TT cores. For tensor orders $d \geq 3$, besides the head and tail tensors $\mathcal{G}_1(i_1), \mathcal{G}_d(i_d)$ which are in fact matrices, there will be $(d-2)$ third-order TT cores in between.

2.2. TTr1 decomposition. With the TT decomposition in place, we are now ready to introduce our TTr1 decomposition, which is easily understood from Figure 5. The main idea of the TTr1 decomposition is to force the rank for each auxiliary index α_k link to unity, which gives rise to a linear combination of rank-1 outer products. We go back to the first SVD (2.3) of the TT decomposition algorithm and realize that we can rewrite it as a sum of rank-1 terms

$$(2.4) \quad \bar{\mathcal{A}} = \sum_{i=1}^3 \sigma_{i \times 1} u_{i \times 2} v_i,$$

where each vector v_i is indexed by i_2, i_3 . The next step in the TT decomposition would be to reshape Y_1 and compute its SVD. For the TTr1 decomposition we reshape each v_i into an $i_2 \times i_3$ matrix \bar{v}_i and compute its SVD. This allows us to write \bar{v}_1 also as a sum of rank-1 terms

$$\bar{v}_1 = \sigma_{11 \times 1} u_{11 \times 2} v_{11} + \sigma_{12 \times 1} u_{12 \times 2} v_{12}.$$

The same procedure can be done for v_2 and v_3 : they can also be written as a sum of two rank-1 terms. Combining these six rank-1 terms we can finally write \mathcal{A} as

$$(2.5) \quad \begin{aligned} \mathcal{A} = & \tilde{\sigma}_{1 \times 1} u_{1 \times 2} u_{11 \times 3} v_{11} + \tilde{\sigma}_{2 \times 1} u_{1 \times 2} u_{12 \times 3} v_{12} \\ & + \tilde{\sigma}_{3 \times 1} u_{2 \times 2} u_{21 \times 3} v_{21} + \tilde{\sigma}_{4 \times 1} u_{2 \times 2} u_{22 \times 3} v_{22} \\ & + \tilde{\sigma}_{5 \times 1} u_{3 \times 2} u_{31 \times 3} v_{31} + \tilde{\sigma}_{6 \times 1} u_{3 \times 2} u_{32 \times 3} v_{32} \end{aligned}$$

with $\tilde{\sigma}_1 = \sigma_1 \sigma_{11}, \dots, \tilde{\sigma}_6 = \sigma_3 \sigma_{32}$. Note the similarity of (2.5) with (2.4). The TTr1 decomposition has three main features that render it similar to the matrix SVD:

1. the scalars $\tilde{\sigma}_1, \dots, \tilde{\sigma}_6$ are the weights of the outer products in the decomposition and can therefore be thought of as the singular values of \mathcal{A} ,
2. the outer products affiliated with each singular value are tensors of unit Frobenius norm, since each product vector (or mode vector) is a unit vector, and
3. each outer product in the decomposition is orthogonal to all the others, which we will prove in section 3.

2.3. TTr1SVD algorithm. As was shown in the previous subsection, computing the TTr1 decomposition requires recursively reshaping the obtained v vectors and computing their SVDs. This recursive procedure gives rise to the formation of a tree, where each SVD generates additional branches of the tree. The tree for the TTr1 decomposition of \mathcal{A} in (2.5) is shown in Figure 6. As denoted in the figure, we will call a row in the tree a level. Level 0 corresponds to the SVD of $\bar{\mathcal{A}}$ and generates the first level of singular values. This is graphically represented by the node at level 0 branching off into three additional nodes at level 1. The reshaping and SVD of the different v vectors at level 1 then generates level 2 and so forth. Observe how the




FIG. 5. Computation of the TTr1 decomposition of \mathcal{A} .




FIG. 6. Tree representation of TTr1 decomposition, where $\tilde{\sigma}_i$ is the product of all nodes down a branch.

total number of subscript indices of the singular values are equal to the level at which these singular values occur. For example, σ_2 occurs at level 1 and σ_{21} occurs at level 2. The number of levels for the TTr1 decomposition of an arbitrary d -way tensor is $d - 1$. The final singular values $\tilde{\sigma}_i$'s are the product of all σ 's along a branch.

The total number of terms in the decomposition is the total number of leaves. This number is easily determined. Indeed, each node at level k of the tree branches off into

$$r_k \triangleq \min \left(n_{k+1}, \prod_{i=k+2}^d n_i \right) \quad (k = 0, \dots, d-2)$$

nodes. Hence, the total number of leaves or terms N in the TTr1 decomposition is given by

$$N = \prod_{k=0}^{d-2} r_k.$$

The algorithm to compute the TTr1 decomposition is presented in pseudocode in Algorithm 1. First the tensor \mathcal{A} is reshaped into an $n_1 \times \prod_{i=2}^d n_i$ matrix and its SVD is computed. The computational complexity for this first step is approximately $14n_1^2 \prod_{i=2}^d n_i + 8n_1^3$ flops. Observe that the computation of the TT or Tucker decomposition has a computational complexity of the same order of magnitude. Then for all remaining nodes in the tree, except for the leaves, the resulting v_i vectors are reshaped into a matrix and their SVDs are also computed. The U, S, V matrices for each of these SVDs are stored. Note that for levels 0 up to $d-2$ the v vectors do not need to be stored. From the tree it is also easy to determine the total number of SVDs required to do the full TTr1 decomposition. This is simply the total number of nodes in the tree from level 0 up to $d-2$ and equals

$$1 + \sum_{i=0}^{d-3} \prod_{k=0}^i r_k.$$

Assuming that $r_k = n_k$ for all k , then the total number of SVDs required for computing the TTr1 decomposition of a cubical tensor is

$$1 + n + n^2 + \cdots + n^{d-2} = \frac{1 - n^{d-1}}{1 - n}.$$

This exponential dependence on the order of the tensor and the computational complexity of $O(n_1^2 \prod_{i=2}^d n_i)$ for the first SVD are the two major limiting factors to compute the TTr1 decomposition. Note, however, that the tree structure is perfectly suited to do all SVD computations that generate the next level in parallel, and in that case the runtime is linearly proportional to the number of levels. However, such an implementation requires an exponential growing number of computational units.

Input: arbitrary tensor \mathcal{A}
Output: U, S, V matrices of each SVD
 $\bar{\mathcal{A}} \leftarrow$ reshape \mathcal{A} into an $n_1 \times \prod_{i=2}^d n_i$ matrix
 $U_1, S_1, V_1 \leftarrow \text{SVD}(\bar{\mathcal{A}})$
for all remaining nodes in the tree except the leaves **do**
 $\bar{v}_i \leftarrow$ reshape v_i
 $U_k, S_k, V_k \leftarrow \text{SVD}(\bar{v}_i)$
 add U_k, S_k, V_k to U, S, V
end for

ALGORITHM 1. TTr1SVD.

3. Properties. We now discuss many attractive properties of the TTr1 decomposition. Most of these properties are also shared with the matrix SVD and it is in this sense that the TTr1SVD is a natural generalization of the SVD for tensors.

3.1. Uniqueness. A first attractive feature of the TTr1 decomposition is that it is uniquely determined for a fixed order of indices. This means that for any given arbitrary tensor \mathcal{A} its TTr1 decomposition will always be the same. Indeed, Algorithm 1 consists of a sequence of SVD computations so the uniqueness of the TTr1 decomposition follows trivially from the fact that each of the SVDs in Algorithm 1 is unique up to sign. Although the singular values and vectors of a matrix A and its transpose A^T are the same, this is not the case for the TTr1SVD. Indeed, applying a permutation of the indices $\pi(i_1, \dots, i_n)$ will generally result in a different TTr1SVD, which we illustrate in Example 1 in section 4.1. Once the indices are fixed, however, the TTr1SVD algorithm will always return the same decomposition, which is not the case for conventional iterative optimization-based methods.

3.2. Orthogonality of outer products. Any two rank-1 terms $\tilde{\sigma}_i \mathcal{T}_i$ and $\tilde{\sigma}_j \mathcal{T}_j$ of the TTr1 decomposition are orthogonal with respect to one another, which means that $\langle \mathcal{T}_i, \mathcal{T}_j \rangle = 0$. We will use our running example to show why this is so. Let us take two terms of (2.5), for example, $\mathcal{T}_1 = 1_{\times_1} u_{1 \times_2} u_{11 \times_3} v_{11}$ and $\mathcal{T}_2 = 1_{\times_1} u_{1 \times_2} u_{12 \times_3} v_{12}$. Another way of writing $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ is

$$\langle \mathcal{T}_1, \mathcal{T}_2 \rangle = (v_{11} \otimes u_{11} \otimes u_1)^T (v_{12} \otimes u_{12} \otimes u_1),$$

where \otimes denotes the Kronecker product. These Kronecker products generate the vectorization of each of the rank-1 tensors $\mathcal{T}_1, \mathcal{T}_2$, which allows us to easily write their inner product as an inner product between two mode vectors. Applying properties of the Kronecker product we can now write

$$\begin{aligned} (v_{11} \otimes u_{11} \otimes u_1)^T (v_{12} \otimes u_{12} \otimes u_1) &= (v_{11}^T \otimes u_{11}^T \otimes u_1^T) (v_{12} \otimes u_{12} \otimes u_1), \\ &= (v_{11}^T v_{12} \otimes u_{11}^T u_{12} \otimes u_1^T u_1), \end{aligned}$$

where it is clear that the right-hand side vanishes due to the orthogonality $v_{11}^T v_{12} = u_{11}^T u_{12} = 0$. This property generalizes to any tensor \mathcal{A} . Indeed, if any two rank-1 terms do not originate from the same node at level 1, then their respective u_i, u_j vectors are orthogonal and ensure that their inner product vanishes. If the two rank-1 terms do originate from the same node at level 1 but from different nodes at level 2, then their u_{ij}, u_{ik} vectors are orthogonal and again the inner product will vanish. This reasoning extends up to level $d-1$. If any two terms have their first separate nodes at level $k \in [1, d-1]$, then their corresponding u vectors at level k will also be orthogonal. The tree structure, together with the orthogonality of all u vectors that share a same parent node, hence guarantees that any two rank-1 outer factors in the TTr1 decomposition are orthogonal. The TTr1 decomposition is hence an orthogonal decomposition as defined in [8].

3.3. Upper bound on the orthogonal tensor rank. The (CP-)rank of an arbitrary d -way tensor \mathcal{A} is usually defined similarly to the matrix case as the minimum number of rank-1 terms that \mathcal{A} decomposes into.

DEFINITION 3.1. *The rank of an arbitrary d -way tensor \mathcal{A} , denoted $\text{rank}(\mathcal{A})$, is the minimum number of rank-1 tensors that yields \mathcal{A} in a linear combination.*

In [8, 12] the orthogonal rank, $\text{rank}_\perp(\mathcal{A})$, is defined as the minimal number of terms in an orthogonal rank-1 decomposition. Apparently,

$$\text{rank}(\mathcal{A}) \leq \text{rank}_\perp(\mathcal{A}),$$

where strict inequality is possible for tensors of orders $d > 2$. The TTr1 decomposition allows a straightforward determination of an upper bound on $\text{rank}_\perp(\mathcal{A})$. Indeed, this is simply the total number of leaves in the tree and is therefore

$$(3.1) \quad \text{rank}_\perp(\mathcal{A}) \leq N = \prod_{k=0}^{d-2} r_k.$$

Applying (3.1) to our running example $\mathcal{A} \in \mathbb{R}^{3 \times 4 \times 2}$ we obtain

$$\text{rank}_\perp(\mathcal{A}) \leq \min(3, 8) \cdot \min(4, 2) = 3 \cdot 2 = 6.$$

For a cubical tensor with $n_1 = \cdots = n_d = n$, (3.1) then tells us that

$$\text{rank}_\perp(\mathcal{A}) \leq \prod_{k=0}^{d-2} \min(n, n^{d-1}) = \prod_{k=0}^{d-2} n = n^{d-1}.$$

The dependency of the TTr1SVD on the ordering of the indices implies that a permutation of the indices can lead to different upper bounds on the orthogonal rank. Indeed, if we permute the indices of \mathcal{A} to $\{i_2, i_3, i_1\}$ we get

$$\text{rank}_\perp(\mathcal{A}) \leq \min(4, 6) \cdot \min(2, 3) = 4 \cdot 2 = 8.$$

Consequently, there exists the notion of a minimum upper bound on the orthogonal rank of a tensor, obtained from computing the rank upper bounds through all permutations of indices. Whether the TTr1SVD algorithm is able to derive a minimal orthogonal decomposition needs further investigation. Furthermore, we will demonstrate by an example in section 4 that the orthogonality as it occurs in the TTr1 decomposition is not enough to make the problem of computing a low-rank approximation of an arbitrary tensor well-posed. This agrees with [21], in which a necessary condition of pairwise orthogonality of all rank-1 terms in at least two modes is proved.

3.4. Quantifying the approximation error. As soon as the number of levels is large it becomes very cumbersome to write all the different subscript indices of the u and v vectors in the TTr1 decomposition. We therefore introduce a shorter and more convenient notation. Herein, u_{ki} denotes the u vector at level k that contributes to the i th rank-1 term. Similarly, v_i denotes the v vector that contributes to the i th rank-1 term. The TTr1SVD algorithm decomposes an arbitrary tensor \mathcal{A} into a linear combination of N orthogonal rank-1 terms

$$(3.2) \quad \mathcal{A} = \sum_{i=1}^N \tilde{\sigma}_{i \times 1} u_{1i \times 2} u_{2i \times 3} \cdots u_{d-1i \times d} v_i$$

with

$$\|1_{\times 1} u_{1i \times 2} u_{2i \times 3} \cdots u_{d-1i \times d} v_i\|_F = 1 \quad \text{and} \quad N = \prod_{k=0}^{d-2} r_k.$$

Suppose that we have ordered and relabeled the terms such that $\tilde{\sigma}_1 \geq \tilde{\sigma}_2 \geq \cdots \geq \tilde{\sigma}_N$. An R -term approximation is then computed by truncating (3.2) to the first R terms

$$\tilde{\mathcal{A}} = \sum_{i=1}^R \tilde{\sigma}_{i \times 1} u_{1i \times 2} u_{2i \times 3} \cdots u_{d-1i \times d} v_i.$$

The following lemma tells us exactly what the error is when breaking off the summation at R terms.

LEMMA 3.2. *Let $\tilde{\mathcal{A}}$ be the summation of the first R terms in (3.2). Then*

$$\|\mathcal{A} - \tilde{\mathcal{A}}\|_F = \sqrt{\tilde{\sigma}_{R+1}^2 + \cdots + \tilde{\sigma}_N^2}.$$

Proof. Using the fact that $\|1_{\times_1} u_{1i \times_2} u_{2i \times_3} \cdots \times_{d-1} u_{d-1i \times_d} v_i\|_F = 1$ we can write

$$\|\mathcal{A} - \tilde{\mathcal{A}}\|_F = \left\| \sum_{i=R+1}^n \tilde{\sigma}_{i \times_1} u_{1i \times_2} u_{2i \times_3} \cdots \times_{d-1} u_{d-1i \times_d} v_i \right\|_F = \sqrt{\tilde{\sigma}_{R+1}^2 + \cdots + \tilde{\sigma}_N^2}. \quad \square$$

Lemma 3.2 can also be used to determine the lowest number of terms R with a guaranteed accuracy. Indeed, once a tolerance ϵ is chosen such that it is required that

$$\|\mathcal{A} - \tilde{\mathcal{A}}\|_F \leq \epsilon,$$

the minimal number of terms R in the TTr1 decomposition of $\tilde{\mathcal{A}}$ is easily determined by the requirement that

$$\sqrt{\tilde{\sigma}_{R+1}^2 + \cdots + \tilde{\sigma}_N^2} \leq \epsilon \text{ such that } \sqrt{\tilde{\sigma}_R^2 + \cdots + \tilde{\sigma}_N^2} > \epsilon.$$

It is tempting to choose R such that $\tilde{\sigma}_R > \epsilon > \tilde{\sigma}_{R+1}$. However, when the approxi-rank gap, defined as $\tilde{\sigma}_R / \tilde{\sigma}_{R+1}$ [13, p. 920], is not large enough, then there is a possibility that $\sqrt{\tilde{\sigma}_{R+1}^2 + \cdots + \tilde{\sigma}_N^2} \geq \epsilon$ due to the contributions of the smaller singular values. A large approxi-rank gap implies that the number of terms in the approximation $\tilde{\mathcal{A}}$ is relatively insensitive to the given tolerance ϵ . In Example 6 in section 4.6 a tensor is presented for which this is not the case.

3.5. Reducing the number of SVDs. Suppose that an approximation $\tilde{\mathcal{A}}$ of \mathcal{A} is desired such that $\|\mathcal{A} - \tilde{\mathcal{A}}\|_F \leq \epsilon$. Computing the full TTr1 decomposition and applying Lemma 3.2 solves this problem. It is, however, possible to reduce the total number of required SVDs by taking into account that the final singular values $\tilde{\sigma}_i$'s are the product of the singular values along each branch of the TTr1-tree. An important observation is that all singular values $\sigma_{ij \dots m}$ at levels 2 up to $d-1$ satisfy $\sigma_{ij \dots m} \leq 1$. This is easily seen from the fact that they are computed from a reshaped unit vector $v_{ij \dots n}$ at their parent node. Indeed, since $\|v_{ij \dots n}\|_2 = 1$ it follows that $\|\bar{v}_{ij \dots n}\|_F = 1$. This allows us to make an educated guess about the impact of the singular values at level l on the final rank-1 terms. Suppose we have a singular value σ_k at level l , preceded by a product σ_{k0} of parent singular values. An upper bound on the size of the final $\tilde{\sigma}$'s that are descendants from σ_k can be derived by assuming that σ_k is unchanged throughout each branch. Since one node at level l results in $\prod_{i=l}^{d-2} r_i$ rank-1 terms, this then implies that there are $\prod_{i=l}^{d-2} r_i$ rank-1 terms with $\tilde{\sigma} = \sigma_{k0} \sigma_k^{d-l-1}$, so

$$\tilde{\sigma}_1^2 + \cdots + \tilde{\sigma}_N^2 \leq \tilde{\sigma}_1^2 + \cdots + \tilde{\sigma}_r^2 + \prod_{i=l}^{d-2} r_i (\sigma_{k0} \sigma_k^{d-l-1})^2.$$

If now

$$(3.3) \quad e_k^2 \triangleq \prod_{i=l}^{d-2} r_i (\sigma_{k0} \sigma_k^{d-l-1})^2 \leq \epsilon^2$$

is satisfied, then removing σ_k at level l produces an approximation $\tilde{\mathcal{A}}$ that is guaranteed to satisfy the approximation error bound. Removing σ_k at level l implies that not a full but a reduced TTr1 decomposition is computed. Indeed, the total number of computed rank-1 terms is effectively lowered by $\prod_{i=l}^{d-2} r_i$ terms, decreasing the total number of required SVDs in the TTr1SVD algorithm. This condition on σ_k is easily extended to m singular values at level l as

$$(3.4) \quad \sum_{j=1}^m e_j^2 \leq \epsilon^2,$$

where we compute an e_j term for each of the m singular values at level l . Checking whether (3.4) holds for m σ 's at level l can be easily implemented in Algorithm 1. As shown in section 4, a rather gradual decrease of σ_k is seen in practice as the level increases. This implies that it might still be possible to find an $\tilde{\mathcal{A}}$ of lower rank that satisfies the approximation error bound from the rank-1 terms of a reduced TTr1 decomposition. Lemma 3.2 can also be used to find the desired $\tilde{\mathcal{A}}$ in this case.

3.6. Orthogonal complement tensors. We can consider the vectorization of \mathcal{A} as a vector living in an $(n_1 \cdots n_d)$ -dimensional vector space. Naturally, there must be an $(n_1 \cdots n_d - 1)$ -dimensional vector space $\text{opspan}(\mathcal{A})^\perp$ of tensors that are orthogonal to \mathcal{A} . Note that each basis vector of $\text{opspan}(\mathcal{A})^\perp$ is required to be the vectorization of an outer product of vectors. The TTr1 decomposition allows us to easily find an orthogonal basis for $\text{opspan}(\mathcal{A})^\perp$. We will illustrate how this comes about using tensor \mathcal{A} from Figure 2 and notions in (2.5). Recall from section 2.2 that the first step in the TTr1SVD algorithm was the economical SVD of the 3×8 matrix $\bar{\mathcal{A}} = USV^T$. Each of the v_i vectors was then reshaped into a 4×2 matrix \bar{v}_i . Now consider a full SVD of each of these \bar{v}_i matrices,

$$(3.5) \quad \bar{v}_i = \begin{pmatrix} u_{i1} & u_{i2} & u_{i3} & u_{i4} \end{pmatrix} \begin{pmatrix} \sigma_{i1} & 0 \\ 0 & \sigma_{i2} \\ 0 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} v_{i1}^T \\ v_{i2}^T \end{pmatrix},$$

which is a sum of eight orthogonal rank-1 outer products with only two nonzero σ 's. There are hence six additional outer product terms $u_i \circ u_{ij} \circ v_{ij}$ with a zero singular value, orthogonal to the outer product terms of the economical TTr1 decomposition (2.5). It is easily seen that the rank-1 terms obtained from the zero entries of the S matrix in (3.5) are orthogonal to \mathcal{A} and are therefore basis vectors of $\text{opspan}(\mathcal{A})^\perp$. Table 1 lists all eight orthogonal rank-1 outer product terms that are obtained for the u_1 branch in the TTr1-tree. Each rank-1 term can be read off from Table 1 by starting from the top row and going down along a particular branch of the TTr1-tree. For example, the fourth rank-1 term is given by $\tilde{\sigma}_2 u_1 \circ u_{12} \circ v_{12}$ and the seventh rank-1 term by $0 u_1 \circ u_{14} \circ v_{11}$. We call such a table that exhibits the TTr1-tree structure and allows us to reconstruct all rank-1 terms an outer product column table. The extra orthogonal terms for the u_2, u_3 branches are completely analogous to the u_1 branch. Note that the economical TTr1 decomposition described in section 2.2 computes only the $\tilde{\sigma}_1 u_1 \circ u_{11} \circ v_{11}$ and $\tilde{\sigma}_2 u_1 \circ u_{12} \circ v_{12}$ terms.

The full TTr1 decomposition therefore consists of $8 \times 3 = 24$ orthogonal terms and can be written in vectorized form as

TABLE 1

All eight orthogonal rank-1 outer products in the u_1 branch of the TTr1-tree.

u_1							
u_{11}		u_{12}		u_{13}		u_{14}	
v_{11}	v_{12}	v_{11}	v_{12}	v_{11}	v_{12}	v_{11}	v_{12}
$\tilde{\sigma}_1$	0	0	$\tilde{\sigma}_2$	0	0	0	0

$$\text{vec}(\mathcal{A}) = (\text{vec}(\mathcal{T}_1) \quad \cdots \quad \text{vec}(\mathcal{T}_6) \quad \text{vec}(\mathcal{T}_7) \quad \cdots \text{vec}(\mathcal{T}_{24})) \begin{pmatrix} \tilde{\sigma}_1 \\ \vdots \\ \tilde{\sigma}_6 \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

where $\text{vec}(\mathcal{T}_1), \dots, \text{vec}(\mathcal{T}_6)$ are the orthogonal terms computed in section 2.2 and $\text{vec}(\mathcal{T}_7), \dots, \text{vec}(\mathcal{T}_{24})$ are the orthogonal terms that partly span $\text{opspan}(\mathcal{A})^\perp$. Note that we have found only 18 basis vectors for $\text{opspan}(\mathcal{A})^\perp$. The remaining 5 basis vectors are to be found as the following linear combinations of $\text{vec}(\mathcal{T}_1), \dots, \text{vec}(\mathcal{T}_6)$,

$$(\text{vec}(\mathcal{T}_1) \quad \cdots \quad \text{vec}(\mathcal{T}_6)) S^\perp,$$

where S^\perp is the 6×5 matrix orthogonal to $(\tilde{\sigma}_1 \quad \cdots \quad \tilde{\sigma}_6)$. The property that for every tensor $\mathcal{B} \in \text{opspan}(\mathcal{A})^\perp$ we have that $\langle \mathcal{A}, \mathcal{B} \rangle = 0$ allows us to interpret $\text{opspan}(\mathcal{A})^\perp$ as the orthogonal complement of $\text{vec}(\mathcal{A})$.

3.7. Constructive proof of the maximal CP-rank of a $2 \times 2 \times 2$ tensor.

As an application of the outer product column table, we show how it leads to an elegant proof of the maximal CP-rank of a real $2 \times 2 \times 2$ tensor over \mathbb{R} . It is known that the maximum rank of a real $2 \times 2 \times 2$ tensor over \mathbb{R} is 3 (i.e., any such tensor can be expressed as the sum of at most three real outer products [9]), for which rather complicated proofs were given in [11, 19]. Incidentally, we show that the TTr1 decomposition allows us to formulate a remarkably simpler proof. As in section 3.6, we first consider all orthogonal outer products that span $\mathbb{R}^{2 \times 2 \times 2}$ in the outer product table, Table 2.

TABLE 2

Outer product table for a general $2 \times 2 \times 2$ tensor.

u_1				u_2			
u_{11}		u_{12}		u_{21}		u_{22}	
v_{11}	v_{12}	v_{11}	v_{12}	v_{21}	v_{22}	v_{21}	v_{22}
$\tilde{\sigma}_1$	0	0	$\tilde{\sigma}_2$	$\tilde{\sigma}_3$	0	0	$\tilde{\sigma}_4$

The columns in Table 2 with nonzero singular values are the “active” columns in the TTr1 decomposition of a random real tensor $\mathcal{A} \in \mathbb{R}^{2 \times 2 \times 2}$. The “inactive” (orthogonal) columns carry zero weights but are crucial for proving the maximum rank-3 property of \mathcal{A} .

We first enumerate two important yet straightforward properties for the columns in Table 2, ignoring the bottom row for the time being. First, scaling a column can be regarded as multiplying a scalar onto the whole outer product or absorbing it into

any one of the mode vectors. Taking the first column and a scalar $\alpha \in \mathbb{R}$ for instance, this means that

$$\alpha(u_1 \circ u_{11} \circ v_{11}) = (\alpha u_1) \circ u_{11} \circ v_{11} = u_1 \circ (\alpha u_{11}) \circ v_{11} = u_1 \circ u_{11} \circ (\alpha v_{11}).$$

In other words, the scalar is “mobile” across the various modes. The second property is that any two columns differing in only one mode can be added to form a new rank-1 outer product [8]. We list two examples showing the rank-1 outer products resulting from the linear combinations of columns 1 and 3, and columns 3 and 4, respectively,

$$\begin{aligned}\alpha(u_1 \circ u_{11} \circ v_{11}) + \beta(u_1 \circ u_{12} \circ v_{11}) &= u_1 \circ (\alpha u_{11} + \beta u_{12}) \circ v_{11}, \\ \alpha(u_1 \circ u_{12} \circ v_{11}) + \beta(u_1 \circ u_{12} \circ v_{12}) &= u_1 \circ u_{12} \circ (\alpha v_{11} + \beta v_{12}).\end{aligned}$$

Now to prove the maximum rank-3 property of \mathcal{A} then is to show that the four active columns of the outer product column in Table 2 can always be “merged” into three. To begin, it is readily seen that if we add any nonzero multiple of column 3 to column 1, and then subtract the same multiple of column 3 from column 4, the overall tensor by summing all columns in Table 2 remains unchanged. Our final goal is to merge columns 1 and 5 into one outer product by making two of their modes the same (up to a scalar factor). This is done by appropriately adding column 3 to column 1 such that the second mode vectors of columns 1 and 5 align, while adding column 6 to column 5 such that the third mode vectors of columns 1 and 5 align. Of course, subtractions of column 3 from column 4, and column 6 from column 8, respectively, are necessary to offset the addition. This intermediate step is summarized in Table 3, where the four intermediate columns are now shown individually with the $\tilde{\sigma}_i$ ’s absorbed into the mode vectors.

TABLE 3
The four intermediate outer products that can be merged into three outer products.

u_1	u_1	u_2	u_2
$\tilde{\sigma}_1 u_{11}$			$\tilde{\sigma}_4 u_{22}$
$+\alpha u_{12}$	u_{11}	u_{21}	$-\gamma u_{21}$
$(= \beta u_{21})$			
	$-\alpha v_{11}$	$\tilde{\sigma}_3 v_{21}$	
v_{11}	$+\tilde{\sigma}_2 v_{12}$	$+\gamma v_{22}$	v_{22}
		$(= \delta v_{11})$	

The two linear equations that need to be solved in this process are

$$\begin{bmatrix} -u_{12} & u_{21} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \tilde{\sigma}_1 u_{11} \quad \text{and} \quad \begin{bmatrix} -v_{22} & v_{11} \end{bmatrix} \begin{bmatrix} \gamma \\ \delta \end{bmatrix} = \tilde{\sigma}_3 v_{21}.$$

It is not hard to see that columns 1 and 3 of Table 3 can now be merged into one outer product,

$$(\beta u_1 + \delta u_2) \circ u_{21} \circ v_{11},$$

due to two of their mode vectors now being parallel. Hence an overall rank-3 representation for the original tensor \mathcal{A} is obtained from its TTr1 decomposition.

Obviously, this rank-3 representation is not unique since alternatively we can first align the third mode of columns 1 and 5, followed by their second mode. Furthermore, instead of columns 1 and 3, we can also merge columns 2 and 4, etc.

Details are omitted as they are all based on the same idea of merging columns. Another big advantage of our rank-3 construction is that the relative numerical error $\|\mathcal{A} - \hat{\mathcal{A}}\|_F / \|\mathcal{A}\|_F < 10^{-15}$, whereas the CP rank-3 decomposition has a median error of $\approx 10^{-6}$ over a 100 trials of arbitrary $2 \times 2 \times 2$ tensors.

3.8. Perturbations of singular values. When an $m \times n$ matrix A is additively perturbed by a matrix E to form $\hat{A} = A + E$, then Weyl's theorem [18] bounds the absolute perturbations of the corresponding singular values by

$$|\sigma_i - \hat{\sigma}_i| \leq \|E\|_2,$$

where the $\hat{\sigma}_i$'s are the singular values of \hat{A} . It is possible to extend Weyl's theorem to the TTr1 decomposition of the perturbed tensor $\hat{\mathcal{A}} = \mathcal{A} + \mathcal{E}$. Suppose we want to determine an upper bound for the perturbation of one of the singular values $\tilde{\sigma}$. We first introduce the simpler notation

$$\tilde{\sigma} = \sigma_1 \sigma_2 \cdots \sigma_{d-2},$$

where σ_k ($k = 1, \dots, d-2$) denotes the singular value at level k in the branch of the TTr1-tree corresponding with $\tilde{\sigma}$. Applying Weyl's theorem to the first factor gives

$$|\sigma_1 - \hat{\sigma}_1| \leq \|\bar{\mathcal{E}}\|_2,$$

which we can rewrite into

$$(3.6) \quad |\hat{\sigma}_1| \leq |\sigma_1| + \|\bar{\mathcal{E}}\|_2.$$

Each of the remaining factors $\sigma_2, \dots, \sigma_{d-2}$ are the singular values of a reshaped right singular vector v_1, \dots, v_{d-3} . Again, application of Weyl's theorem allows us to write

$$(3.7) \quad \begin{aligned} |\hat{\sigma}_k| &\leq |\sigma_k| + \|\Delta \bar{v}_{k-1}\|_2 \quad (k = 2, \dots, d-2), \\ &\leq |\sigma_k| + \|\Delta \bar{v}_{k-1}\|_F = |\sigma_k| + \|\Delta v_{k-1}\|_2 \quad (k = 2, \dots, d-2). \end{aligned}$$

An upper bound for the $\|\Delta \bar{v}_{k-1}\|_2$ term is difficult to derive. Fortunately, it is possible to replace the $\|\Delta \bar{v}_{k-1}\|_2$ term by $\|\Delta \bar{v}_{k-1}\|_F = \|\Delta v_{k-1}\|_2$, for which first-order approximations exist [14]. Multiplying (3.6) with (3.7) over all k we obtain

$$|\hat{\sigma}_1| \cdots |\hat{\sigma}_{d-2}| \leq (|\sigma_1| + \|\bar{\mathcal{E}}\|_2) \cdots (|\sigma_k| + \|\Delta v_{k-1}\|_2),$$

which can be simplified by ignoring higher-order terms to

$$(3.8) \quad |\hat{\sigma}_1| \cdots |\hat{\sigma}_{d-2}| \leq |\sigma_1| \cdots |\sigma_{d-2}| + |\sigma_2| \cdots |\sigma_{d-2}| \|\bar{\mathcal{E}}\|_2 + \sum_{k=2}^{d-2} \left(\prod_{i \neq k} |\sigma_i| \right) \|\Delta v_{k-1}\|_2.$$

The maximal value for $\sigma_2, \dots, \sigma_{d-2}$ is 1 and hence (3.8) can be written as

$$|\hat{\sigma}_1| \cdots |\hat{\sigma}_{d-2}| \leq |\sigma_1| \cdots |\sigma_{d-2}| + \|\bar{\mathcal{E}}\|_2 + \sigma_1 \left(\sum_{k=1}^{d-3} \|\Delta v_k\|_2 \right).$$

Hence we arrive at the expression

$$(3.9) \quad |\tilde{\sigma} - \hat{\sigma}| \leq \|\bar{\mathcal{E}}\|_2 + \sigma_1 \left(\sum_{k=1}^{d-3} \|\Delta v_k\|_2 \right),$$

which generalizes Weyl's theorem to the TTr1 decomposition by the addition of a correction term $\sigma_1 (\sum_{k=1}^{d-3} \|\Delta v_k\|_2)$. This correction term depends on the largest singular value of the first level and the perturbations on the right singular vectors for levels 1 up to $d-3$.

3.9. Conversion to the Tucker decomposition. It is possible to convert the sum of orthogonal rank-1 terms obtained from a (truncated) TTr1 decomposition into the Tucker decomposition,

$$\mathcal{A} = \mathcal{S}_{\times_1} U_1 \times_2 U_2 \cdots \times_d U_d,$$

where $\mathcal{S} \in \mathbb{R}^{r_1 \times r_2 \times \cdots \times r_d}$ is called a core tensor and $U_k \in \mathbb{R}^{n_k \times r_k}$ are orthogonal factor matrices. In this way it becomes relatively easy to compute an approximation of a tensor with a known approximation error in the Tucker format with orthogonal factor matrices. This conversion is easily achieved using simple matrix operations. To avoid notational clumsiness, we illustrate the conversion from the TTr1 representation into the Tucker form through the specific TTr1 decomposition in (2.5). Suppose only alternate terms in the equation are significant and we therefore keep only the three terms associated with $\sigma_1 \sigma_{11}$, $\sigma_2 \sigma_{21}$, and $\sigma_3 \sigma_{31}$. Then, the mode vectors of these outer products are collected and subjected to economic QR factorization. We note that $U_1 = [u_1 \ u_2 \ u_3] \in \mathbb{R}^{3 \times 3}$ is already orthogonal and does not need to go through a QR factorization, while

$$\begin{aligned} \begin{bmatrix} u_{11} & u_{21} & u_{31} \end{bmatrix} &= \underbrace{\begin{bmatrix} q_{21} & q_{22} & q_{23} \end{bmatrix}}_{U_2} \begin{bmatrix} 1 & \alpha_{12} & \alpha_{13} \\ 0 & \alpha_{22} & \alpha_{23} \\ 0 & 0 & \alpha_{33} \end{bmatrix}, \\ \begin{bmatrix} v_{11} & v_{21} & v_{31} \end{bmatrix} &= \underbrace{\begin{bmatrix} q_{31} & q_{32} \end{bmatrix}}_{U_3} \begin{bmatrix} 1 & \beta_{12} & \beta_{13} \\ 0 & \beta_{22} & \beta_{23} \end{bmatrix}. \end{aligned}$$

Consequently, the truncated TTr1SVD of (2.5) reads

$$\begin{aligned} \mathcal{A} &\approx \sigma_1 \sigma_{11} (u_1 \circ u_{11} \circ v_{11}) + \sigma_2 \sigma_{21} (u_2 \circ u_{21} \circ v_{21}) + \sigma_3 \sigma_{31} (u_3 \circ u_{31} \circ v_{31}) \\ &= \sigma_1 \sigma_{11} (u_1 \circ q_{21} \circ q_{31}) + \sigma_2 \sigma_{21} (u_2 \circ (\alpha_{12} q_{21} + \alpha_{22} q_{22}) \circ (\beta_{12} q_{31} + \beta_{22} q_{32})) \\ &\quad + \sigma_3 \sigma_{31} (u_3 \circ (\alpha_{13} q_{21} + \alpha_{23} q_{22} + \alpha_{33} q_{23}) \circ (\beta_{13} q_{31} + \beta_{23} q_{32})) \\ (3.10) \quad &= \mathcal{S}_{\times_1} U_1 \times_2 U_2 \times_3 U_3, \end{aligned}$$

where the core tensor $\mathcal{S} \in \mathbb{R}^{3 \times 3 \times 2}$ is filled with coefficients found through expanding the outer products and collecting terms in (3.10). Observe that the dimensions of the core tensor \mathcal{S} are completely determined by the ranks of the orthogonal factor matrices. From practical examples we observe that the Tucker core obtained in this way is more sparse compared to the Tucker core computed from the ALS algorithm [2, 3]. If we take, for example, a random tensor in $\mathbb{R}^{4 \times 3 \times 15}$, then its TTr1 decomposition consists of 12 terms. The ranks of the orthogonal factor matrices U_1, U_2, U_3 are then 4, 3, 12, respectively. Consequently, we have that $\mathcal{S} \in \mathbb{R}^{4 \times 3 \times 12}$ with $3 + 3 \cdot 9 = 30$ nonzero entries. In contrast, computing the Tucker decomposition using the ALS method results in a maximally dense core tensor of 144 nonzero entries.

4. Numerical examples. In this section, we demonstrate some of the properties of the TTr1 decomposition and compare it with the ALS-CP and Tucker decompositions by means of numerical examples. All experiments are done in MATLAB on a desktop computer. A MATLAB/Octave implementation of the TTr1SVD algorithm can be freely downloaded and modified from <https://github.com/kbatseli/TTr1SVD>. The ALS-CP and Tucker decompositions are computed by the ALS optimization tool provided in the MATLAB Tensor Toolbox [2, 3]. All ALS procedures are fed by random initial guesses; therefore their errors are defined as the average error over multiple executions.

TABLE 4
 $\tilde{\sigma}$'s of TTr1 decomposition for \mathcal{A} .

Order of indices	$\{i_1, i_2, i_3\}$	$\{i_2, i_3, i_1\}$	$\{i_3, i_1, i_2\}$
$\tilde{\sigma}_1$	69.6306	69.6306	69.6306
$\tilde{\sigma}_2$	6.9190	6.9551	6.9567
$\tilde{\sigma}_3$	1.8036	1.6108	1.6010
$\tilde{\sigma}_4$	0.6729	0.7781	0.7840
$\tilde{\sigma}_5$	6.7e-15	1.9e-15	4.3e-15
$\tilde{\sigma}_6$	1.3e-15	1.9e-15	1.4e-15
$\tilde{\sigma}_7$	NA	7.3e-16	NA
$\tilde{\sigma}_8$	NA	5.3e-16	NA


4.1. Example 1: Singular values and permutation of indices. We start with tensor \mathcal{A} in Figure 2. Since it is discussed in section 3.3 that TTr1 decomposition depends on the ordering of the indices, we demonstrate the TTr1 decomposition with different permutations of the indices. For a three-way tensor, the order of indices can be $\{i_1, i_2, i_3\}$, $\{i_1, i_3, i_2\}$, $\{i_2, i_1, i_3\}$, $\{i_2, i_3, i_1\}$, $\{i_3, i_1, i_2\}$, or $\{i_3, i_2, i_1\}$. Since the order of the last two indices will not affect the σ 's in TTr1 decomposition, we only list the σ 's under the permutations $\{i_1, i_2, i_3\}$, $\{i_2, i_3, i_1\}$, and $\{i_3, i_1, i_2\}$ in Table 4, in descending order. As a result, although permutations of indices may give different upper bounds on the rank, TTr1 decomposition still outputs the same $\text{rank}(\mathcal{A}) = 4$ in all permutations. The largest (dominant) singular value (69.6306) differs only slightly with respect to the permutations, which is also the general observation. Note that the singular values are quite similar over all permutations for this example. It can also be seen that some singular values are numerically zero. The same threshold commonly used to determine the numerical rank of a matrix using the SVD can also be used for the TTr1 decomposition. An interesting consequence of the rank-deficiency of \mathcal{A} is that we can interpret the rank-1 terms corresponding with the very small numerical singular values as being related to $\text{opspan}(\mathcal{A})^\perp$. For the $\{i_1, i_2, i_3\}, \{i_3, i_1, i_2\}$ permutations there are two extra orthogonal complement tensors, while for the $\{i_2, i_3, i_1\}$ there are four extra orthogonal complement tensors. Figure 7 shows the similar singular value curves for all six permutations of a random $3 \times 4 \times 2$ tensor, where it can also be seen that there are basically three distinct permutations and that the largest singular values differ only slightly over all permutations.

4.2. Example 2: Comparison with ALS-CP and Tucker decomposition.

Next, ALS-CP is applied on \mathcal{A} . To begin, we compute the best rank-1 approximation of \mathcal{A} . ALS-CP gives the same weight 69.6306 as the TTr1 decomposition, implying that both decompositions result in the same approximation in terms of the Frobenius norm. The errors between \mathcal{A} and its approximations $\tilde{\mathcal{A}}$, computed using the ALS-CP and TTr1SVD method, are listed in Table 5 for increasing rank.

Table 5 confirms that the $\text{rank}_\perp(\mathcal{A}) = \text{rank}(\mathcal{A}) = 4$. It also indicates that as an optimization approach, ALS-CP itself cannot determine the rank, only the best rank- R approximation for a specific R . Furthermore, it should be noticed that the TTr1 decomposition can always give an R -rank-1-term approximation with orthonormal outer products, while ALS-CP cannot ensure this property. Finally, a Tucker decomposition with a core size $(2, 2, 2)$ is applied on \mathcal{A} . The resulting dense core tensor $\mathcal{S} \in \mathbb{R}^{2 \times 2 \times 2}$ is given by

$$\mathcal{S}_{i_1 i_2 1} = \begin{pmatrix} 69.6306 & -0.0181 \\ -0.0701 & -0.7840 \end{pmatrix}, \quad \mathcal{S}_{i_1 i_2 2} = \begin{pmatrix} -0.0113 & -6.9190 \\ -1.6108 & -0.7010 \end{pmatrix}.$$

FIG. 7. Singular value curves for each of the six permutations of a random $3 \times 4 \times 2$ tensor.

The rank-1 outer factors obtained from the Tucker decomposition are also orthonormal. However, compared to the TTr1 decomposition, the Tucker format needs twice the number of rank-1 terms than TTr1.

TABLE 5
Errors $\|\mathcal{A} - \tilde{\mathcal{A}}\|_F$ of ALS-CP and TTr1SVD for increasing rank R .

Rank	1	2	3	4	5
TTr1SVD	7.2	1.9	0.7	6.8e-15	1.3e-15
ALS-CP	7.2	0.8	3.6e-2	1.4e-10	4.7e-11

4.3. Example 3: Rank behavior under largest rank-1 term subtraction.

In this example we investigate the behavior of the singular value curves and the rank when the largest rank-1 term obtained from the TTr1 decomposition is consecutively subtracted. This means that we start with \mathcal{A} from Figure 2, compute its largest orthogonal rank-1 term \mathcal{T}_1 from the TTr1 decomposition, and subtract it to obtain $\mathcal{A} - \mathcal{T}_1$, after which the procedure is repeated. Figure 8 shows the singular value curves from the TTr1 decompositions obtained for each of the iterations, where it is easily seen that each curve gets shifted to the left with each iteration. In other words, the largest singular value in the next iteration is the second largest singular value of the previous iteration, etc. It is also clear that subtracting the largest orthogonal rank-1 term does not necessarily decrease the rank as also described in [17]. Indeed, using a numerical threshold of $\max(4, 8) \cdot 1.11 \times 10^{-16} \cdot \sigma_1 = 1.13 \times 10^{-13}$ on the singular values obtained in the first iteration will still return a numerical orthogonal rank of 4.

In addition, the rank of the obtained tensors in each iteration was also determined from the CP-decomposition. From Table 5, a numerical threshold of 10^{-10} was set to the absolute error $\|\mathcal{A} - \tilde{\mathcal{A}}\|_F$ to determine the CP-rank. In Table 6, both the CP-rank and the orthogonal rank from the TTr1 decomposition are compared. It can be seen that the rank determined from ALS-CP increases while the orthogonal




FIG. 8. Singular value curves for consecutive subtractions of the largest rank-1 term.

TABLE 6

Ranks determined by ALS-CP and TTr1 decomposition for consecutive subtraction of the largest rank-1 term.

Iteration	0	1	2	3	4	5
CP-rank	4	4	4	5	3	2
TTr1 rank	4	4	4	4	3	2

rank monotonically decreases. In this sense, the orthogonal rank appears to be more robust under the largest rank-1 term subtraction.

4.4. Example 4: Perturbation of the singular values. In this example we illustrate the robustness of the computed singular values of our running example tensor \mathcal{A} when it is subjected to additive perturbations. We construct a perturbation tensor $\mathcal{E} \in \mathbb{R}^{3 \times 4 \times 2}$ where each entry is drawn from a zero mean Gaussian distribution with variance 10^{-6} . We then compute the following two norms of \mathcal{E} and $\bar{\mathcal{E}}$:

$$\|\mathcal{E}\|_F = 5.48 \times 10^{-6} \quad \text{and} \quad \|\bar{\mathcal{E}}\|_2 = 4.13 \times 10^{-6},$$

where $\bar{\mathcal{E}}$ is \mathcal{E} reshaped into a 3×8 matrix. Comparing the perturbed singular values $\bar{\sigma}_1, \dots, \bar{\sigma}_6$ of $\mathcal{A} + \mathcal{E}$ with the singular values $\tilde{\sigma}_1, \dots, \tilde{\sigma}_6$ then shows that

$$(4.1) \quad \sqrt{(\bar{\sigma}_1 - \tilde{\sigma}_1)^2 + \dots + (\bar{\sigma}_6 - \tilde{\sigma}_6)^2} = 3.78 \times 10^{-6} < \|\mathcal{E}\|_F$$

and

$$(4.2) \quad |\bar{\sigma}_i - \tilde{\sigma}_i| \leq \|\bar{\mathcal{E}}\|_2 \quad (i = 1, \dots, 6).$$

These two inequalities (4.1) and (4.2) are very reminiscent of Mirsky's and Weyl's theorems [18], respectively, for the perturbation of singular values for matrices.

4.5. Example 5: Gradual decrease of intermediate singular value products. In the discussion on reducing the total number of required SVDs it was shown that the product of the singular values along a branch becomes smaller and smaller for every additional level. In this example we demonstrate this gradual decrease for a random $2 \times 2 \times 2 \times 2 \times 2$ tensor where each entry is drawn from a zero mean Gaussian distribution with variance 1. The TTr1 decomposition always has 16 rank-1 terms. Figure 9 shows the intermediate singular value products $\sigma_i \sigma_{ij} \cdots \sigma_{ij \dots m}$ as a function of the level for $\tilde{\sigma}_1, \tilde{\sigma}_8, \tilde{\sigma}_{13}, \tilde{\sigma}_{16}$. The figure shows that the intermediate singular value products indeed decrease as the level increases. The TTr1-tree for this tensor is a binary tree. Each SVD of a v vector therefore produces 2 singular values. It is consistently observed that of the two singular values of \bar{v} , one is very close to unity, with values around 0.8 or 0.9. The other singular value typically has values around 0.5. Branches of the tree that mostly choose the singular value close to unity therefore exhibit a very slow decrease, while branches that predominantly choose the smaller singular value decrease faster. This is seen in Figure 9 as a bigger descent of the intermediate products of $\tilde{\sigma}_8, \tilde{\sigma}_{16}$ compared to $\tilde{\sigma}_1, \tilde{\sigma}_{13}$.




FIG. 9. Gradual decrease of intermediate singular value products as a function of the level.

4.6. Example 6: Exponential decaying singular values. In this example we illustrate the computation of an approximation $\tilde{\mathcal{A}}$ using Lemma 3.2 when the singular values $\tilde{\sigma}$ decay exponentially. Consider the tensor $\mathcal{A} \in \mathbb{R}^{5 \times 5 \times 5}$ with

$$\mathcal{A}_{i_1 i_2 i_3} = \frac{1}{i_1 + i_2 + i_3},$$

which has very smoothly decaying singular values, as shown in Figure 10. There are a total number of 25 rank-1 terms in the TTr1 decomposition. Suppose we are interested in obtaining an approximation $\tilde{\mathcal{A}}$ such that $\|\mathcal{A} - \tilde{\mathcal{A}}\|_F \leq 10^{-6}$. Sorting the

rank-1 terms by descending singular values and using Lemma 3.2, the approximation would then consist of 17 terms since

$$\tilde{\sigma}_{16} = 3.18 \times 10^{-6} \geq \tilde{\sigma}_{17} = 1.18 \times 10^{-6} \geq 1.00 \times 10^{-6} \geq \tilde{\sigma}_{18} = 9.00 \times 10^{-7}.$$

The approxi-rank gap $\tilde{\sigma}_{17}/\tilde{\sigma}_{18} = 1.31$, which indicates that there is no clear “gap” between $\tilde{\sigma}_{17}$ and $\tilde{\sigma}_{18}$. In contrast, the tensor in Example 1 (section 4.1) has an approxi-rank gap of $\tilde{\sigma}_4/\tilde{\sigma}_5 \approx 10^{14}$. Also note that it is not possible to reduce the number of SVDs during execution of the TTr1SVD algorithm since none of the first five computed singular values $\sigma_1, \dots, \sigma_5$ satisfy condition (3.3), with $\sigma_{k0} = 1$. Next, the approximations obtained from the TTr1 decomposition and CP decomposition of this tensor are compared for increasing rank. The CP decomposition was computed over 10 trials with different initial guesses using the CP-ALS method. The absolute errors in terms of the rank are listed in Table 7. For the CP decomposition case the mean absolute error over the 10 trials is reported. From Table 7 it is seen that the errors are almost identical up to the first 24 terms. Since the TTr1 decomposition consists of 25 terms, the error drops at that term to the order of the machine precision, while the ALS-CP method fails to produce any significant improvement in the error. Even when a CP decomposition of 100 rank-1 terms are computed, the average absolute error is around 10^{-12} .




FIG. 10. Singular value decay of the function-generated tensor $\mathcal{A}_{i_1 i_2 i_3} = 1/(i_1 + i_2 + i_3)$.

TABLE 7
Errors $\|\mathcal{A} - \tilde{\mathcal{A}}\|_F$ of ALS-CP and TTr1SVD for increasing rank R .

Rank	1	5	10	15	20	25
TTr1SVD	9.5e-2	2.6e-3	7.6e-5	3.6e-6	2.2e-7	9.9e-16
ALS-CP	9.5e-2	1.6e-3	3.6e-5	4.5e-6	1.2e-7	1.6e-7

5. Conclusion. In this paper, a constructive TTr1 decomposition algorithm, TTr1SVD, was proposed to decompose high-order real tensors into a finite sum of real orthogonal rank-1 outer products. Compared to existing CP approaches, the TTr1 decomposition has many favorable properties, such as uniqueness, easy quantification of the approximation error, and an easy conversion to the Tucker format with a sparse core tensor. A complete characterization of all tensors orthogonal to the original tensor was also provided for the first time, which is readily available via the TTr1SVD and easily visualized by an outer product column table. As an application example, this outer product column table was used to provide an elegant constructive proof of the maximum rank-3 property of $2 \times 2 \times 2$ tensors over the real field. Numerical examples verified and demonstrated the favorable properties of TTr1SVD in decomposing and analyzing real tensors.

Acknowledgments. The authors would like to thank the associate editor, Pierre Comon, and the anonymous referees for the many constructive comments.

REFERENCES

- [1] 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.
- [2] B. W. BADER AND T. G. KOLDA, *Efficient MATLAB computations with sparse and factored tensors*, SIAM J. Sci. Comput., 30 (2007), pp. 205–231.
- [3] B. W. BADER, T. G. KOLDA, ET AL., *MATLAB Tensor Toolbox Version 2.5*, www.sandia.gov/~tgkolda/TensorToolbox/index-2.5.html (2012).
- [4] 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.
- [5] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1253–1278.
- [6] R. A. HARSHMAN, *Foundations of the PARAFAC Procedure: Models and Conditions for an “Explanatory” Multi-Modal Factor Analysis*, UCLA Working Papers in Phonetics 16, University of California, 1970.
- [7] F. L. HITCHCOCK, *The expression of a tensor or a polyadic as a sum of products*, J. Math. Phys., 6 (1927), pp. 164–189.
- [8] T. G. KOLDA, *Orthogonal tensor decompositions*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 243–255.
- [9] T. G. KOLDA AND B.W. BADER, *Tensor decompositions and applications*, SIAM Rev., 51 (2009), pp. 455–500.
- [10] T. G. KOLDA, *Multilinear Operators for Higher-Order Decompositions*, Tech. report SAND2006-2081, Sandia National Laboratories, 2006.
- [11] 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, the Netherlands, 1989, pp. 7–18.
- [12] D. LEIBOVICI AND R. SABATIER, *A singular value decomposition of a k -way array for a principal component analysis of multiway data*, PTA- k , Linear Algebra Appl., 269 (1998), pp. 307–329.
- [13] T. Y. LI AND Z. ZENG, *A rank-revealing method with updating, downdating, and applications*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 918–946.
- [14] J. LIU, X. LIU, AND X. MA, *First-order perturbation analysis of singular vectors in singular value decomposition*, IEEE Trans. Signal Process., 56 (2008), pp. 3044–3049.
- [15] I. OSELEDETS, *Tensor-train decomposition*, SIAM J. Sci. Comput., 33 (2011), pp. 2295–2317.
- [16] J. SALMI, A. RICHTER, AND V. KOIVUNEN, *Sequential unfolding SVD for tensors with applications in array signal processing*, IEEE Trans. Signal Process., 57 (2009), pp. 4719–4733.
- [17] A. STEGEMAN AND P. COMON, *Subtracting a best rank-1 approximation may increase tensor rank*, Linear Algebra Appl., 433 (2010), pp. 1276–1300.
- [18] G. W. STEWART, *Perturbation Theory for the Singular Value Decomposition*, in SVD and Signal Processing, II: Algorithms, Analysis and Applications, Elsevier, New York, 1990, pp. 99–109.

- [19] J. M. F. TEN BERGE, *Kruskal's polynomial for $2 \times 2 \times 2$ arrays and a generalization to $2 \times n \times n$ arrays*, Psychometrika, 56 (1991), pp. 631–636.
- [20] L. TUCKER, *Some mathematical notes on three-mode factor analysis*, Psychometrika, 31 (1966), pp. 279–311.
- [21] N. VANNIEUWENHOVEN, J. NICAISE, R. VANDEBRIL, AND K. MEERBERGEN, *On generic nonexistence of the Schmidt–Eckart–Young decomposition for complex tensors*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 886–903.