
A Sampling-Based Method for Tensor Ring Decomposition

Osman Asif Malik¹ Stephen Becker¹

Abstract

We propose a sampling-based method for computing the tensor ring (TR) decomposition of a data tensor. The method uses leverage score sampled alternating least squares to fit the TR cores in an iterative fashion. By taking advantage of the special structure of TR tensors, we can efficiently estimate the leverage scores and attain a method which has complexity sublinear in the number of input tensor entries. We provide high-probability relative-error guarantees for the sampled least squares problems. We compare our proposal to existing methods in experiments on both synthetic and real data. Our method achieves substantial speedup—sometimes two or three orders of magnitude—over competing methods, while maintaining good accuracy. We also provide an example of how our method can be used for rapid feature extraction.

1. Introduction

Tensor decomposition has recently found many applications in machine learning and related areas. Examples include parameter reduction in neural networks (Novikov et al., 2015; Garipov et al., 2016; Yang et al., 2017; Yu et al., 2017; Ye et al., 2018), understanding the expressivity of deep neural networks (Cohen et al., 2016; Khrulkov et al., 2018), supervised learning (Stoudenmire & Schwab, 2017; Novikov et al., 2017), filter learning (Hazan et al., 2005; Rigamonti et al., 2013), image factor analysis and recognition (Vasilescu & Terzopoulos, 2002; Liu et al., 2019), scaling up Gaussian processes (Izmailov et al., 2018), multimodal feature fusion (Hou et al., 2019), natural language processing (Lei et al., 2014), data mining (Papalexakis et al., 2016), feature extraction (Bengua et al., 2015), and tensor completion (Wang et al., 2017). The CP and Tucker decompositions are the two most popular tensor decompositions

that each have strengths and weaknesses (Kolda & Bader, 2009). The number of parameters in the CP decomposition scales linearly with the tensor dimension, but finding a good decomposition numerically can be challenging. The Tucker decomposition is usually easier to work with numerically, but the number of parameters grows exponentially with tensor dimension. To address these challenges, decompositions based on tensor networks have recently gained in popularity. The tensor train (TT) is one such decomposition (Oseledets, 2011). The number of parameters in a TT scales linearly with the tensor dimension. The tensor ring (TR) decomposition is a generalization of the TT decomposition. It is more expressive than the TT decomposition (see Remark 1), but maintains the linear scaling of the number of parameters with increased tensor dimension. Although the TR decomposition is known to have certain numerical stability issues it usually works well in practice, achieving a better compression ratio and lower storage cost than the TT decomposition (Mickelin & Karaman, 2020).

It is crucial to have efficient algorithms for computing the decomposition of a tensor. This task is especially challenging for large and high-dimensional tensors. In this paper, we present a sampling-based method for TR decomposition which has complexity sublinear in the number of input tensor entries. The method fits the TR core tensors by iteratively updating them in an alternating fashion. Each update requires solving a least squares problem. We use leverage score sampling to speed up the solution of these problems and avoid having to consider all elements in the input tensor for each update. The system matrices in the least squares problems have a particular structure that we take advantage of to efficiently estimate the leverage scores. To summarize, we make the following contributions:

- Propose a sampling-based method for TR decomposition which has complexity sublinear in the number of input tensor entries.
- Provide relative-error guarantees for the sampled least squares problems we use in our algorithm.
- Compare our proposed algorithm to four other methods in experiments on both synthetic and real data.
- Provide an example of how our method can be used for rapid feature extraction.

¹Department of Applied Mathematics, University of Colorado Boulder, Boulder, CO, USA. Correspondence to: Osman Asif Malik <osman.malik@colorado.edu>.

2. Related Work

Leverage score sampling has been used previously for tensor decomposition. Cheng et al. (2016) use it to reduce the size of the least squares problems that come up in the alternating least squares (ALS) algorithm for CP decomposition. They efficiently estimate the leverage scores by exploiting the Khatri–Rao structure of the design matrices in these least squares problems and prove additive-error guarantees. Larsen & Kolda (2020) further improve this method by combining repeatedly sampled rows and by using a hybrid deterministic–random sampling scheme. These techniques make the method faster without sacrificing accuracy. Our paper differs from these previous works since we consider a different kind of tensor decomposition. Since the TR decomposition has core tensors instead of factor matrices, the leverage score estimation formulas used in the previous works cannot be used for TR decomposition.

There are previous works that develop randomized algorithms for TR decomposition. Yuan et al. (2019) develop a method which first applies a randomized variant of the HOSVD Tucker decomposition, followed by a TR decomposition of the core tensor using either a standard SVD- or ALS-based algorithm. Finally, the TR cores are contracted appropriately with the Tucker factor matrices to get a TR decomposition of the original tensor. Ahmadi-Asl et al. (2020) present several TR decomposition algorithms that incorporate randomized SVD. In Section 5, we compare our proposed algorithm to methods by Yuan et al. (2019) and Ahmadi-Asl et al. (2020) in experiments.

Papers that develop randomized algorithms for other tensor decompositions include the works by Wang et al. (2015), Battaglino et al. (2018), Yang et al. (2018) and Aggour et al. (2020) for the CP decomposition; Biagioni et al. (2015) and Malik & Becker (2020) for tensor interpolative decomposition; Drineas & Mahoney (2007), Tsourakakis (2010), da Costa et al. (2016), Malik & Becker (2018), Sun et al. (2020) and Minster et al. (2020) for the Tucker decomposition; Zhang et al. (2018) and Tarzanagh & Michailidis (2018) for t-product-based decompositions; and Huber et al. (2017) and Che & Wei (2019) for the TT decomposition.

Skeleton approximation and other sampling-based methods for TR decomposition include the works by Espig et al. (2012) and Khoo et al. (2019). Espig et al. (2012) propose a skeleton/cross approximation type method for the TR format with complexity linear in the dimensionality of the tensor. Khoo et al. (2019) propose an ALS-based scheme for constructing a TR decomposition based on only a few samples of the input tensor. Papers that use skeleton approximation and sampling techniques for other types of tensor decomposition include those by Mahoney et al. (2008), Oseledets et al. (2008), Oseledets & Tyrtshnikov (2010), Caiafa & Cichocki (2010), and Friedland et al. (2011).

3. Preliminaries

By tensor, we mean a multidimensional array of real numbers. Boldface Euler script letters, e.g. \mathcal{X} , denote tensors of dimension 3 or greater; bold uppercase letters, e.g. \mathbf{X} , denote matrices; bold lowercase letters, e.g. \mathbf{x} , denote vectors; and regular lowercase letters, e.g. x , denote scalars. Elements of tensors, matrices and vectors are denoted in parentheses. For example, $\mathcal{X}(i_1, i_2, i_3)$ is the element on position (i_1, i_2, i_3) in the 3-way tensor \mathcal{X} , and $\mathbf{y}(i_1)$ is the element on position i_1 in the vector \mathbf{y} . A colon is used to denote all elements along a dimension. For example, $\mathbf{X}(i, :)$ denotes the i th row of the matrix \mathbf{X} . For a positive integer I , define $[I] \stackrel{\text{def}}{=} \{1, \dots, I\}$. For indices $i_1 \in [I_1], \dots, i_N \in [I_N]$, the notation $\overline{i_1 i_2 \dots i_N} \stackrel{\text{def}}{=} \sum_{n=1}^N i_n \prod_{j=1}^{n-1} I_j$ will be helpful when working with reshaped tensors. $\|\cdot\|_F$ denotes the Frobenius norm of matrices and tensors, and $\|\cdot\|_2$ denotes the Euclidean norm of vectors.

3.1. Tensor Ring Decomposition

Let $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ be an N -way tensor. For $n \in [N]$, let $\mathcal{G}^{(n)} \in \mathbb{R}^{R_{n-1} \times I_n \times R_n}$ be 3-way tensors with $R_0 = R_N$. A rank- (R_1, \dots, R_N) TR decomposition of \mathcal{X} takes the form

$$\mathcal{X}(i_1, \dots, i_N) = \sum_{r_1, \dots, r_N} \prod_{n=1}^N \mathcal{G}^{(n)}(r_{n-1}, i_n, r_n), \quad (1)$$

where each r_n in the sum goes from 1 to R_n and $r_0 = r_N$. $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$ are called *core tensors*. We will use $\text{TR}((\mathcal{G}^{(n)})_{n=1}^N)$ to denote a tensor ring with cores $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$. The name of the decomposition comes from the fact that it looks like a ring in tensor network¹ notation; see Figure 1. The TT decomposition is a special case of the TR decomposition with $R_0 = R_N = 1$, which corresponds to severing the connection between $\mathcal{G}^{(1)}$ and $\mathcal{G}^{(5)}$ in Figure 1. The TR decomposition can therefore also be interpreted as a sum of tensor train decompositions with the cores $\mathcal{G}^{(2)}, \dots, \mathcal{G}^{(N-1)}$ in common.

Remark 1 (Expressiveness of TR vs TT). Since the TT decomposition is a special case of the TR decomposition, any TT is also a TR with the same number of parameters. So the TR decomposition is *at least as expressive* as the TT decomposition. To see that the TR decomposition is *strictly more expressive* than the TT decomposition, consider the following example: Let $\mathcal{X} = \text{TR}(\mathcal{G}, \mathcal{G}, \mathcal{G}, \mathcal{G})$ where $\mathcal{G} \in \mathbb{R}^{2 \times 2 \times 2}$ has entries $1, 2, \dots, 8$ ($\mathcal{G} = \text{reshape}(1:8, 2, 2, 2)$ in Matlab). This TR of \mathcal{X} requires $4 \cdot (2 \cdot 2 \cdot 2) = 32$ parameters. The TT ranks of \mathcal{X} are $(2, 4, 2)$ (computed via Theorem 8.8 by Ye & Lim (2018)). An exact TT decomposition of \mathcal{X} therefore requires $2 \cdot (2 \cdot 2) + 2 \cdot (2 \cdot 2 \cdot 4) = 40$ parameters.

¹See Cichocki et al. (2016; 2017) for an introduction to tensor networks.

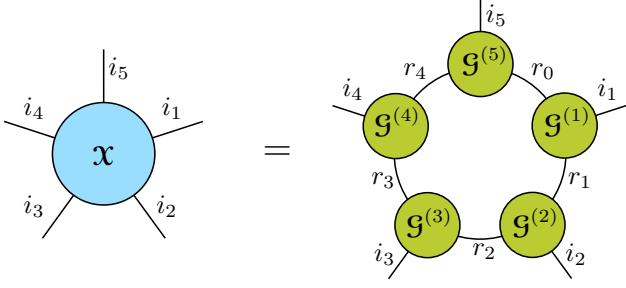


Figure 1. Tensor ring decomposition of 5-way tensor.

The problem of fitting $\text{TR}((\mathcal{G}^{(n)})_{n=1}^N)$ to a data tensor \mathcal{X} can be written as the minimization problem

$$\arg \min_{\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}} \|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N) - \mathcal{X}\|_F, \quad (2)$$

where the size of each core tensor is fixed. There are two common approaches to this fitting problem: one is SVD based and the other uses ALS. We describe the latter below since it is what we use in our work, and refer to Zhao et al. (2016) and Mickelin & Karaman (2020) for a description of the SVD-based approach.

Definition 2. The *classical mode- n unfolding* of \mathcal{X} is the matrix $\mathbf{X}_{(n)} \in \mathbb{R}^{I_n \times \prod_{j \neq n} I_j}$ defined elementwise via

$$\mathbf{X}_{(n)}(i_n, \overline{i_1 \dots i_{n-1} i_{n+1} \dots i_N}) \stackrel{\text{def}}{=} \mathcal{X}(i_1, \dots, i_N).$$

The *mode- n unfolding* of \mathcal{X} is the matrix $\mathbf{X}_{[n]} \in \mathbb{R}^{I_n \times \prod_{j \neq n} I_j}$ defined elementwise via

$$\mathbf{X}_{[n]}(i_n, \overline{i_{n+1} \dots i_N i_1 \dots i_{n-1}}) \stackrel{\text{def}}{=} \mathcal{X}(i_1, \dots, i_N).$$

Definition 3. By merging all cores except the n th, we get a *subchain tensor* $\mathcal{G}^{\neq n} \in \mathbb{R}^{R_n \times \prod_{j \neq n} I_j \times R_{n-1}}$ defined elementwise via

$$\begin{aligned} \mathcal{G}^{\neq n}(r_n, \overline{i_{n+1} \dots i_N i_1 \dots i_{n-1}}, r_{n-1}) \\ \stackrel{\text{def}}{=} \sum_{\substack{r_1, \dots, r_{n-2} \\ r_{n+1}, \dots, r_N}} \prod_{\substack{j=1 \\ j \neq n}}^N \mathcal{G}^{(j)}(r_{j-1}, i_j, r_j). \end{aligned}$$

It follows directly from Theorem 3.5 in Zhao et al. (2016) that the objective in (2) can be rewritten as

$$\|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N) - \mathcal{X}\|_F = \|\mathbf{G}_{[2]}^{\neq n} \mathbf{G}_{(2)}^{(n)\top} - \mathbf{X}_{[n]}^\top\|_F.$$

The ALS approach to finding an approximate solution to (2) is to keep all cores except the n th fixed and solve the least squares problem above with respect to that core, and then repeat this multiple times for each $n \in [N]$ until some termination criterion is met. We summarize the approach

in Algorithm 1. The initialization on line 1 can be done by, e.g., drawing each core entry independently from a standard normal distribution. Notice that the first core does not need to be initialized since it will be immediately updated. Normalization steps can also be added to the algorithm, but such details are omitted here. Throughout this paper we make the reasonable assumption that $\mathcal{G}^{\neq n} \neq \mathbf{0}$ for all $n \in [N]$ during the execution of the ALS-based methods. This is discussed further in Remark S6 in the supplement.

Algorithm 1: TR-ALS (Zhao et al., 2016)

Input: $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$, target ranks (R_1, \dots, R_N)

Output: TR cores $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$

- 1 Initialize cores $\mathcal{G}^{(2)}, \dots, \mathcal{G}^{(N)}$
 - 2 **repeat**
 - 3 **for** $n = 1, \dots, N$ **do**
 - 4 Compute $\mathbf{G}_{[2]}^{\neq n}$ from cores
 - 5 Update $\mathcal{G}^{(n)} = \arg \min_{\mathcal{Z}} \|\mathbf{G}_{[2]}^{\neq n} \mathbf{Z}_{(2)}^\top - \mathbf{X}_{[n]}^\top\|_F$
 - 6 **end**
 - 7 **until** *termination criteria met*
 - 8 **return** $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$
-

3.2. Leverage Score Sampling

Let $\mathbf{A} \in \mathbb{R}^{K \times L}$ and $\mathbf{y} \in \mathbb{R}^K$ where $K \gg L$. Solving the overdetermined least squares problem $\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2$ using standard methods costs $O(KL^2)$. One approach to reducing this cost is to randomly sample and rescale $J \ll K$ of the rows of \mathbf{A} and \mathbf{y} according to a carefully chosen probability distribution and then solve this smaller system instead. This will reduce the solution cost to $O(JL^2)$.

Definition 4. For a matrix $\mathbf{A} \in \mathbb{R}^{K \times L}$, let $\mathbf{U} \in \mathbb{R}^{K \times \text{rank}(\mathbf{A})}$ contain the left singular vectors of \mathbf{A} . The *i th leverage score* of \mathbf{A} is defined as $\ell_i(\mathbf{A}) \stackrel{\text{def}}{=} \|\mathbf{U}(i, :)\|_2^2$ for $i \in [K]$.

Definition 5. Let $\mathbf{q} \in \mathbb{R}^K$ be a probability distribution on $[K]$, and let $\mathbf{v} \in [K]^J$ be a random vector with independent elements satisfying $\mathbb{P}(\mathbf{v}(j) = i) = \mathbf{q}(i)$ for all $(i, j) \in [K] \times [J]$. Let $\mathbf{\Omega} \in \mathbb{R}^{J \times K}$ and $\mathbf{R} \in \mathbb{R}^{J \times J}$ be a random sampling matrix and a diagonal rescaling matrix, respectively, defined as $\mathbf{\Omega}(j, :) \stackrel{\text{def}}{=} \mathbf{e}_{\mathbf{v}(j)}^\top$ and $\mathbf{R}(j, j) \stackrel{\text{def}}{=} 1/\sqrt{J\mathbf{q}(\mathbf{v}(j))}$ for each $j \in [J]$, where \mathbf{e}_i is the i th column of the $K \times K$ identity matrix. We say that $\mathbf{S} \stackrel{\text{def}}{=} \mathbf{R}\mathbf{\Omega} \in \mathbb{R}^{J \times K}$ is a *sampling matrix with parameters* (J, \mathbf{q}) , or $\mathbf{S} \sim \mathcal{D}(J, \mathbf{q})$ for short. Let $\mathbf{A} \in \mathbb{R}^{K \times L}$ be a nonzero matrix, let $\mathbf{p} \in \mathbb{R}^K$ be a probability distribution on $[K]$ with entries $\mathbf{p}(i) \stackrel{\text{def}}{=} \ell_i(\mathbf{A})/\text{rank}(\mathbf{A})$, and suppose $\beta \in (0, 1]$. We say that $\mathbf{S} \sim \mathcal{D}(J, \mathbf{q})$ is a *leverage score sampling matrix for* (\mathbf{A}, β) if $\mathbf{q}(i) \geq \beta\mathbf{p}(i)$ for all $i \in [K]$.

Notice that for any \mathbf{A} we have $\sum_i \ell_i(\mathbf{A}) = \text{rank}(\mathbf{A})$ and

therefore $\sum_i p(i) = 1$ as desired for a probability distribution. One can show that if $\mathbf{S} \sim \mathcal{D}(J, \mathbf{q})$ is a leverage score sampling matrix for (\mathbf{A}, β) with J sufficiently large, then $\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{y}\|_2 \leq (1 + \varepsilon) \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2$ with high probability, where $\tilde{\mathbf{x}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{x}} \|\mathbf{S}\mathbf{A}\mathbf{x} - \mathbf{S}\mathbf{y}\|_2$. How large to make J depends on the probability of success and the parameters ε and β . For further details on leverage score sampling, see Mahoney (2011) and Woodruff (2014).

4. Tensor Ring Decomposition via Sampling

We propose using leverage score sampling to reduce the size of the least squares problem on line 5 in Algorithm 1. A challenge with using leverage score sampling is computing a sampling distribution \mathbf{q} which satisfies the conditions in Definition 5. One option is to use $\mathbf{q} = \mathbf{p}$ and $\beta = 1$, with \mathbf{p} defined as in Definition 5. However, computing \mathbf{p} requires finding the left singular vectors of the design matrix which costs the same as solving the original least squares problem.

When the design matrix is $\mathbf{G}_{[2]}^{\neq n}$, a sampling distribution \mathbf{q} can be computed much more efficiently directly from the cores $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(n-1)}, \mathcal{G}^{(n+1)}, \dots, \mathcal{G}^{(N)}$ without explicitly forming $\mathbf{G}_{[2]}^{\neq n}$. For each $n \in [N]$, let $\mathbf{p}^{(n)} \in \mathbb{R}^{I_n}$ be a probability distribution on $[I_n]$ defined elementwise via²

$$\mathbf{p}^{(n)}(i_n) \stackrel{\text{def}}{=} \frac{\ell_{i_n}(\mathbf{G}_{(2)}^{(n)})}{\text{rank}(\mathbf{G}_{(2)}^{(n)})}. \quad (3)$$

Furthermore, let $\mathbf{q}^{\neq n} \in \mathbb{R}^{\prod_{j \neq n} I_j}$ be a vector defined elementwise via

$$\mathbf{q}^{\neq n}(i_{n+1} \cdots i_N i_1 \cdots i_{n-1}) \stackrel{\text{def}}{=} \prod_{\substack{j=1 \\ j \neq n}}^N \mathbf{p}^{(j)}(i_j). \quad (4)$$

Lemma 6. Let β_n be defined as

$$\beta_n \stackrel{\text{def}}{=} \left(R_{n-1} R_n \prod_{\substack{j=1 \\ j \notin \{n-1, n\}}}^N R_j^2 \right)^{-1}.$$

For each $n \in [N]$, the vector $\mathbf{q}^{\neq n}$ is a probability distribution on $[\prod_{j \neq n} I_j]$, and $\mathbf{S} \sim \mathcal{D}(J, \mathbf{q}^{\neq n})$ is a leverage score sampling matrix for $(\mathbf{G}_{[2]}^{\neq n}, \beta_n)$.

This lemma can now be used to prove the following guarantees for a sampled variant of the least squares problem on line 5 of Algorithm 1.

Theorem 7. Let $\mathbf{S} \sim \mathcal{D}(J, \mathbf{q}^{\neq n})$, $\varepsilon \in (0, 1)$, $\delta \in (0, 1)$

² $\text{rank}(\mathbf{G}_{(2)}^{(n)}) \geq 1$ due to our assumption that $\mathcal{G}^{\neq n} \neq \mathbf{0}$ for all $n \in [N]$, so (3) is well-defined; see Remark S6 in the supplement.

and $\tilde{\mathbf{z}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z}} \|\mathbf{S}\mathbf{G}_{[2]}^{\neq n} \mathbf{z}_{(2)}^\top - \mathbf{S}\mathbf{X}_{[n]}^\top\|_F$. If

$$J > \left(\prod_{j=1}^N R_j^2 \right) \max \left(\frac{16}{3(\sqrt{2}-1)^2} \ln \left(\frac{4R_{n-1}R_n}{\delta} \right), \frac{4}{\varepsilon\delta} \right), \quad (5)$$

then the following holds with probability at least $1 - \delta$:

$$\|\mathbf{G}_{[2]}^{\neq n} \tilde{\mathbf{z}}_{(2)}^\top - \mathbf{X}_{[n]}^\top\|_F \leq (1 + \varepsilon) \min_{\mathbf{z}} \|\mathbf{G}_{[2]}^{\neq n} \mathbf{z}_{(2)}^\top - \mathbf{X}_{[n]}^\top\|_F.$$

Proofs of Lemma 6 and Theorem 7 are given in Section S1 of the supplement. In Algorithm 2 we present our proposed TR decomposition algorithm.

Algorithm 2: TR-ALS-Sampled (proposal)

Input: $\mathcal{X} \in \mathbb{R}^{I_1 \times \cdots \times I_N}$, target ranks (R_1, \dots, R_N)

Output: TR cores $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$

```

1 Initialize cores  $\mathcal{G}^{(2)}, \dots, \mathcal{G}^{(N)}$ 
2 Compute distributions  $\mathbf{p}^{(2)}, \dots, \mathbf{p}^{(N)}$  via (3)
3 repeat
4   for  $n = 1, \dots, N$  do
5     Set sample size  $J$ 
6     Draw sampling matrix  $\mathbf{S} \sim \mathcal{D}(J, \mathbf{q}^{\neq n})$ 
7     Compute  $\tilde{\mathbf{G}}_{[2]}^{\neq n} = \mathbf{S}\mathbf{G}_{[2]}^{\neq n}$  from cores
8     Compute  $\tilde{\mathbf{X}}_{[n]}^\top = \mathbf{S}\mathbf{X}_{[n]}^\top$ 
9     Update  $\mathcal{G}^{(n)} = \arg \min_{\mathbf{z}} \|\tilde{\mathbf{G}}_{[2]}^{\neq n} \mathbf{z}_{(2)}^\top - \tilde{\mathbf{X}}_{[n]}^\top\|_F$ 
10    Update  $n$ th distribution  $\mathbf{p}^{(n)}$  via (3)
11  end
12 until termination criteria met
13 return  $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(N)}$ 

```

We draw each core entry independently from a standard normal distribution during the initialization on line 1 in our experiments. The sample size on line 5 can be provided as a user input or be determined adaptively; this is discussed further in Section 4.2. The sampling matrix \mathbf{S} is never explicitly constructed. Instead, on line 6, we draw and store a realization of the vector $\mathbf{v} \in [\prod_{j \neq n} I_j]^J$ in Definition 5. Entries of \mathbf{v} can be drawn efficiently using Algorithm 3.

Algorithm 3: Sampling from $\mathbf{q}^{\neq n}$ (proposal)

Input: Distributions $\{\mathbf{p}^{(k)}\}_{k \neq n}$

Output: Sample $i \sim \mathbf{q}^{\neq n}$

```

1 for  $j \in [N] \setminus n$  do draw realization  $i_j \sim \mathbf{p}^{(j)}$ 
2 return  $i = i_{n+1} \cdots i_N i_1 \cdots i_{n-1}$ 

```

The sampling on line 7 in Algorithm 2 can be done efficiently directly from the cores without forming $\mathbf{G}_{[2]}^{\neq n}$. To see this, note that the matrix row $\mathbf{G}_{[2]}^{\neq n}(i, :)$ is the vectorization of the tensor slice $\mathcal{G}^{\neq n}(:, i, :)$ due to Definition 2.

From Definition 3, this tensor slice in turn is given by the sequence of matrix multiplications

$$\mathcal{G}^{\neq n}(:, i, :) = \mathcal{G}^{(n+1)}(:, i_{n+1}, :) \cdots \mathcal{G}^{(N)}(:, i_N, :) \\ \cdot \mathcal{G}^{(1)}(:, i_1, :) \cdots \mathcal{G}^{(n-1)}(:, i_{n-1}, :),$$

where $i = \overline{i_{n+1} \cdots i_N i_1 \cdots i_{n-1}}$ and each $\mathcal{G}^{(j)}(:, i_j, :)$ is treated as an $R_{j-1} \times R_j$ matrix. For a realization of $\mathbf{v} \in [\prod_{j \neq n} I_j]^J$, it is therefore sufficient to extract J lateral slices from each of the cores $\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(n-1)}, \mathcal{G}^{(n+1)}, \dots, \mathcal{G}^{(N)}$ and then multiply together the j th slice from each core, for each $j \in [J]$. This is illustrated in Figure 2. An algorithm for computing $\tilde{\mathcal{G}}_{[2]}^{\neq n}$ in this way is given in Algorithm 4.

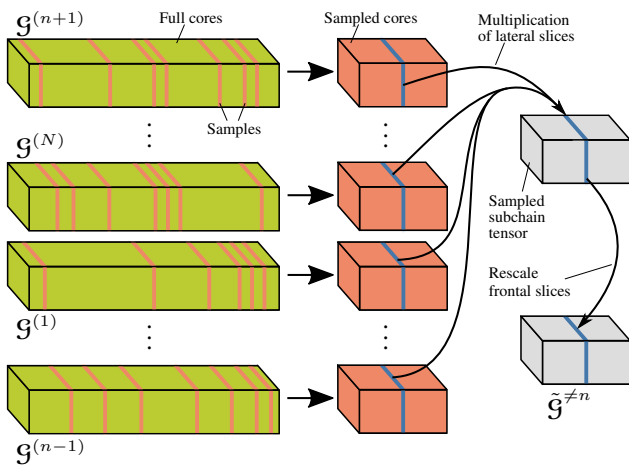


Figure 2. Illustration of how to efficiently construct $\tilde{\mathcal{G}}_{[2]}^{\neq n}$ by sampling the core tensors.

Algorithm 4: Computation of $\tilde{\mathcal{G}}_{[2]}^{\neq n}$ (proposal)

Input: Samples \mathbf{v} , distributions $\{\mathbf{p}^{(k)}\}_{k \neq n}$

Output: Sketch $\tilde{\mathcal{G}}_{[2]}^{\neq n} = \mathbf{S}\mathcal{G}_{[2]}^{\neq n}$

- 1 Initialize $\tilde{\mathcal{G}}_{[2]}^{\neq n} = \mathbf{0} \in \mathbb{R}^{R_n \times J \times R_{n-1}}$
 - 2 **for** $j \in [J]$ **do**
 - 3 $i = \mathbf{v}(j)$ // Where $i = \overline{i_{n+1} \cdots i_N i_1 \cdots i_{n-1}}$
 - 4 $c = \sqrt{J \prod_{k \neq n} \mathbf{p}(i_k)}$
 - 5 $\tilde{\mathcal{G}}_{[2]}^{\neq n}(:, j, :) = \mathcal{G}^{(n+1)}(:, i_{n+1}, :) \cdots \mathcal{G}^{(N)}(:, i_N, :) \\ \cdot \mathcal{G}^{(1)}(:, i_1, :) \cdots \mathcal{G}^{(n-1)}(:, i_{n-1}, :)/c$
 - 6
 - 7 **end**
 - 8 **return** $\tilde{\mathcal{G}}_{[2]}^{\neq n}$
-

The sampling of the data tensor on line 8 in Algorithm 2 only needs to access $J I_n$ entries of \mathcal{X} , or less if \mathcal{X} is sparse. This is why our proposed algorithm has a cost which is sublinear in the number of entries in \mathcal{X} . Due to Theorem 7, each

sampled least squares solve on line 9 has high-probability relative-error guarantees, provided J is large enough.

4.1. Termination Criteria

Both Algorithms 1 and 2 rely on using some termination criterion. One such criterion is to terminate the outer loop when the relative error $\|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N) - \mathcal{X}\|_{\text{F}} / \|\mathcal{X}\|_{\text{F}}$, or its change, is below some threshold. Computing the relative error exactly costs $O(I^N R^2)$ which may be too costly for large tensors. In particular, it would defeat the purpose of the sublinear complexity of our method. An alternative is to terminate when the change in $\|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N)\|_{\text{F}}$ is below some threshold. Such an approach is used for large scale Tucker decomposition by Kolda & Sun (2008) and Malik & Becker (2018). The norm $\|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N)\|_{\text{F}}$ can be computed efficiently using an algorithm by Micklin & Karaman (2020). When $I_n = I$ and $R_n = R$ for all $n \in [N]$ it has complexity $O(NIR^4)$; see Section 3.3 of Micklin & Karaman (2020). Another approach is to estimate the relative error via sampling, which is used by Battagliano et al. (2018) for their randomized CP decomposition method.

4.2. Adaptive Sample Size

Algorithm 2 can produce good results even if the sample size J is smaller than what Theorem 7 suggests. Choosing an appropriate J for a particular problem can therefore be challenging. Aggour et al. (2020) develop an ALS-based algorithm for CP decomposition with an adaptive sketching rate to speed up convergence. Similar ideas can easily be incorporated in our Algorithm 2. Indeed, one of the benefits of our method is that the sample size J can be changed at any point with no overhead cost. By contrast, for rTR-ALS, another randomized method for TR decomposition which we describe in Section 5, the entire compression phase would need to be redone if the sketch rate is updated.

4.3. Complexity Analysis

Table 1. Comparison of leading order computational complexity. “#iter” denotes the number of outer loop iterations in the ALS-based methods.

Method	Complexity
TR-ALS	$NIR^2 + \text{\#iter} \cdot NI^N R^2$
rTR-ALS	$NI^N K + \text{\#iter} \cdot NK^N R^2$
TR-SVD	$I^{N+1} + I^N R^3$
TR-SVD-Rand	$I^N R^2$
TR-ALS-S. (proposal)	$NIR^4 + \text{\#iter} \cdot \frac{NIR^{2(N+1)}}{\epsilon \delta}$

Table 1 compares the leading order complexity of our proposed method to the four other methods we consider in the experiments. The other methods are described in Section 5.

For simplicity, we assume that $I_n = I$ and $R_n = R$ for all $n \in [N]$, and that $N < I$, $R^2 < I$ and $NR < I$. For rTR-ALS, the intermediate N -way Tucker core is assumed to be of size $K \times \dots \times K$; see [Yuan et al. \(2019\)](#) for details. The complexity for our method in the table assumes J satisfies (5), and that ε and δ are small enough so that (5) simplifies to $J > 4R^{2N}/(\varepsilon\delta)$. Any costs for checking convergence criteria are ignored. This is justified since e.g. the termination criterion based on computing $\|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N)\|_{\text{F}}$ described in Section 4.1 would not impact the complexity of TR-ALS or our TR-ALS-Sampled, and would only impact that of rTR-ALS if $K^N < IR^2$. The complexities in Table 1 are derived in Section S2 of the supplement.

The complexity of our method is sublinear in the number of tensor elements, avoiding the dependence on I^N that the other methods have. While there is still an exponential dependence on N , this is still much better than I^N for low-rank decomposition of large tensors in which case $R \ll I$. This is particularly true when the input tensor is too large to store in RAM and accessing its elements dominate the cost of the decomposition. We note that our method typically works well in practice even if (5) is not satisfied.

5. Experiments

5.1. Decomposition of Real and Synthetic Datasets

We compare our proposed TR-ALS-Sampled method (Alg. 2) to four other methods: TR-ALS (Alg. 1), rTR-ALS (Alg. 1 in [Yuan et al. \(2019\)](#)), TR-SVD (Alg. 1 in [Mickelin & Karaman \(2020\)](#)) and a randomized variant of TR-SVD (Alg. 7 in [Ahmadi-Asl et al. \(2020\)](#)) which we refer to as TR-SVD-Rand. TR-SVD takes an error tolerance as input and makes the ranks large enough to achieve this tolerance. To facilitate comparison to the ALS-based methods, we modify TR-SVD so that it takes target ranks as inputs instead. TR-SVD-Rand takes target ranks as inputs by default. For TR-ALS, rTR-ALS and TR-SVD-Rand, we wrote our own implementations since codes were not publicly available. For TR-SVD, we modified the function `TRdecomp.m` provided by [Mickelin & Karaman \(2020\)](#)³. As suggested by [Ahmadi-Asl et al. \(2020\)](#), we use an oversampling parameter of 10 in TR-SVD-Rand. The relative error of a decomposition is computed as $\|\text{TR}((\mathcal{G}^{(n)})_{n=1}^N) - \mathcal{X}\|_{\text{F}}/\|\mathcal{X}\|_{\text{F}}$.

To get a fair comparison between the different methods we first decompose the input tensor with TR-ALS by running it until a convergence criterion has been met or a maximum number of iterations reached. All ALS-based algorithms (including TR-ALS) are then applied to the input tensor, running for twice as many iterations as it took the initial TR-ALS run to terminate, without checking any convergence

criteria. This is done to strip away any run time influence that checking convergence has and only compare those aspects that differ between the methods. TR-ALS-Sampled and rTR-ALS apply sketching in fundamentally different ways. To best understand the trade-off between run time and accuracy, we start out these two algorithm with small sketch rates (J and K , respectively) and increase them until the resulting relative error is less than $(1 + \varepsilon)E$, where E is the relative error achieved by TR-ALS during the second run and $\varepsilon \in (0, 1)$. The relative error and run time for the smallest sketch rate that achieve this relative-error bound are then reported. The exact details on how this is done for each individual experiment is described in Section S4 of the supplement.

All experiments are run in Matlab R2017a with Tensor Toolbox 2.6 ([Bader et al., 2015](#)) on a computer with an Intel Xeon E5-2630 v3 @ 32x 3.2GHz CPU and 32 GB of RAM. All our code used in the experiments is available online⁴.

5.1.1. RANDOMLY GENERATED DATA

Synthetic 3-way tensors are generated by creating 3 cores of size $R' \times I \times R'$ with entries drawn independently from a standard normal distribution. Here, R' is used to denote true rank, as opposed to the rank R used in the decomposition. For each core, one entry is chosen uniformly at random and set to 20. The goal of this is to increase the coherence in the least squares problems, which makes them more challenging for sampling-based methods like ours. The cores are then combined into a full tensor via (1). Finally, Gaussian noise with standard deviation 0.1 is added to all entries in the tensor independently.

In the first synthetic experiment $R' = 10$ and $I \in \{100, 200, \dots, 500\}$. The target rank is also set to $R = 10$. Ideally, decompositions should therefore have low error. Figure 3 shows the relative error and run time for the five different algorithms. All plotted quantities are the median over 10 runs. The ALS-based algorithms reach very low errors. On average, only 21 iterations are needed, and the number of iterations is fairly stable when the tensor size is increased. The SVD-based algorithms, by contrast do very poorly, having relative errors close to 1. Our proposed TR-ALS-Sampled is the fastest method, aside from TR-SVD-Rand which returns very poor results. We achieve a substantial speedup of $668\times$ over TR-ALS and $230\times$ over rTR-ALS for $I = 500$.

In the second synthetic experiment, we set $R' = 20$ but keep the target rank at $R = 10$. We therefore expect the decompositions to have a larger relative error. Figure 4 shows the results for these experiments. All plotted quantities are

³Available at <https://github.com/oscardmickelin/tensor-ring-decomposition>.

⁴Available at <https://github.com/OsmanMalik/tr-als-sampled>.

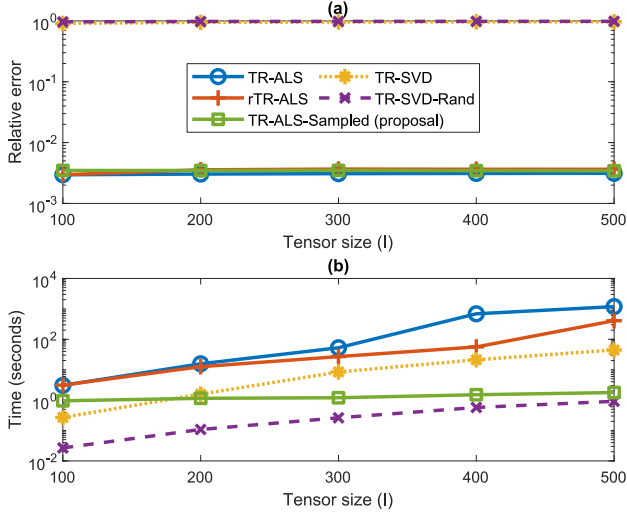


Figure 3. Synthetic experiment with true and target ranks $R' = R = 10$. (a) Relative error and (b) run time for each of the five methods. Average number of ALS iterations used is 21.

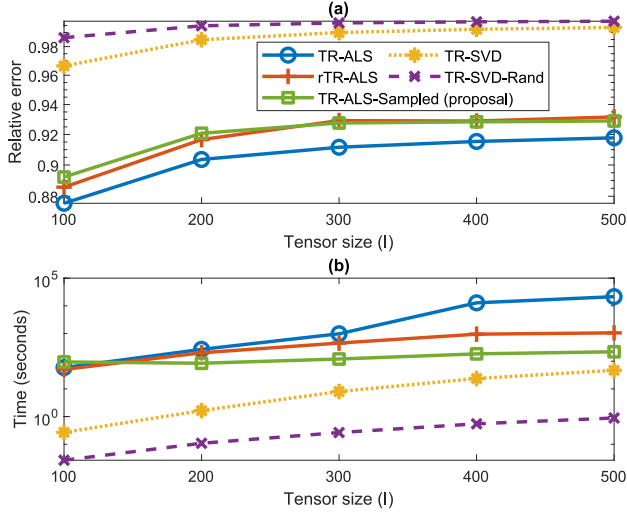


Figure 4. Synthetic experiment with true rank $R' = 20$ and target rank $R = 10$. (a) Relative error and (b) run time for each of the five methods. Average number of ALS iterations used is 417.

the median over 10 runs. The ALS-based methods still outperform the SVD-based ones. The ALS-based methods now require more iterations, on average running for 417 iterations, with greater variability between runs. Our proposed TR-ALS-Sampled remains faster than the other ALS-based methods, although with a smaller difference, achieving a speedup of $98\times$ over TR-ALS and $5\times$ over rTR-ALS for $l = 500$. Although the SVD-based methods are faster, especially the randomized variant, they are not useful since they give such a high error. See Remarks S9 and S10 in the supplement for further discussion of the results in Figures 3 and 4.

The signal-to-noise ratio (SNR) is about 50 dB and 59 dB, respectively, in the two experiments above. To validate the robustness of those results to higher levels of noise we run additional experiments on synthetic $300 \times 300 \times 300$ tensors with added Gaussian noise with standard deviation 1 and 10 which corresponds to an SNR of 30 dB and 10 dB, respectively. All other settings are kept the same as in the first synthetic experiment. The results, shown in Figure 5, indicate that the earlier experiment results are robust to higher noise levels.

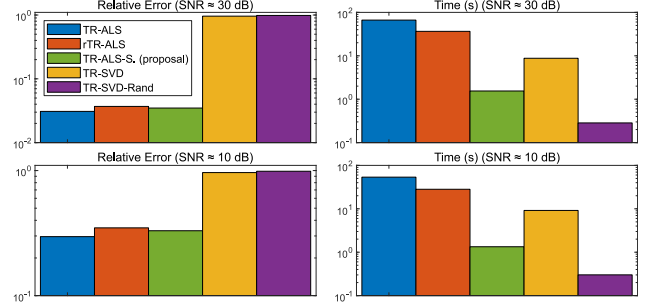


Figure 5. Comparison of methods for higher levels of noise.

5.1.2. HIGHLY OSCILLATORY FUNCTIONS

Next, we decompose the highly oscillatory functions considered by Zhao et al. (2016). The functions are illustrate in Figure 6. Each function is evaluated at 4^{10} points. These values are then reshaped into 10-way tensors of size $4 \times \dots \times 4$ which are then decomposed. Table 2 shows the results for when all target ranks are $R = 10$. The SVD-based methods can not be applied in this experiment since they require $R_0 R_1 \leq I_1$ which is not satisfied.

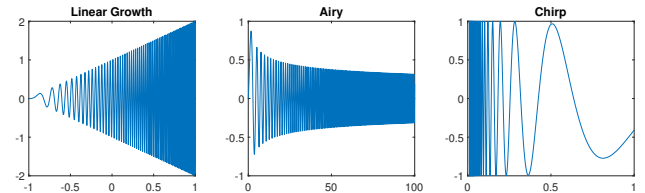


Figure 6. The functions are defined as follows. Linear growth: $(x + 1) \sin(100(x + 1)^2)$. Airy: $x^{-1/4} \sin(2x^{3/2}/3)$. Chirp: $\sin(4/x) \cos(x^2)$.

All algorithms achieve a relative error of only 1%–2%. This is impressive, especially considering that the choice of $R_n = 10$ correspond to a $262\times$ compression rate. With an error very similar to that of TR-ALS, our proposal TR-ALS-Sampled achieves speedups of $108\times$, $193\times$ and $17\times$ on the three tensors over TR-ALS. Our method achieves speedups of $112\times$, $190\times$ and $18\times$ over rTR-ALS. Notice that rTR-ALS yields no speedup over TR-ALS in Table 2. The reason for this is that each $I_n = 4$ is so small. For rTR-

Table 2. Decomposition results for highly oscillatory functions with target rank $R = 10$.

Method	Linear Growth		Airy		Chirp	
	Error	Time (s)	Error	Time (s)	Error	Time (s)
TR-ALS	0.010	370.3	0.020	531.7	0.020	613.1
rTR-ALS	0.011	382.8	0.019	521.6	0.019	645.9
TR-ALS-Sampled (proposal)	0.011	3.4	0.021	2.7	0.021	36.2

ALS to have acceptable accuracy, K must be chosen to be $K = 4$, meaning that rTR-ALS is just as slow as TR-ALS.

5.1.3. IMAGE AND VIDEO DATA

We consider five real image and video datasets⁵ (sizes in parentheses): Pavia University ($610 \times 340 \times 103$) and DC Mall ($1280 \times 307 \times 191$) are 3-way tensors containing hyperspectral images. The first two dimensions are the image height and width, and the third dimension is the number of spectral bands. Park Bench ($1080 \times 1920 \times 364$) and Tabby Cat ($720 \times 1280 \times 286$) are 3-way tensors representing grayscale videos of a man sitting on a park bench and a cat, respectively. The first two dimension are frame height and width, and the third dimension is the number of frames. Red Truck ($128 \times 128 \times 3 \times 72$) is a 4-way tensor consisting of 72 color images of size 128 by 128 pixels depicting a red truck from different angles. It is a subset of the COIL-100 dataset (Nene et al., 1996).

Tables 3 and 4 show results for when all target ranks are $R = 10$ and $R = 20$, respectively. The SVD-based algorithms require $R_0 R_1 \leq I_1$, so they cannot handle the Red Truck tensor when $R = 20$ since $I_1 = 128$ for that dataset. All results are based on a single trial. Our proposed TR-ALS-Sampled runs faster than the other ALS-based methods on all image and video datasets, achieving up to $1518\times$ speedup over TR-ALS (for the Park Bench tensor with $R = 10$) and a $21\times$ speedup over rTR-ALS (for the Red Truck tensor with $R = 20$). The SVD-based methods perform better than they did on the randomly generated data, but remain worse than the ALS-based methods. TR-SVD-Rand is always the fastest method, but also has a substantially higher error.

A popular preprocessing step used in tensor decomposition is to first reshape a tensor into a tensor with more modes. Next, we therefore try reshaping the five tensors above so that they each have twice as many modes. Some of the dimensions are truncated slightly to allow for this reshaping. Our method yields speedups in the range $94\times$ – $2880\times$ over TR-ALS and $73\times$ – $1484\times$ over rTR-ALS. The SVD-based methods cannot be applied since they require $R_0 R_1 \leq I_1$

⁵Links to the datasets are provided in Section S3 of the supplement.

which is not satisfied for these reshaped tensors. The speed benefit of our method appears to be *even greater* when increasing the number of modes. Please see Section S4.3 of the supplement for further details on this experiment, including detailed results.

An alternative to sampling according to the distribution defined by the leverage scores in (3) and (4) is to simply sample according to the uniform distribution, i.e., $q^{z^n} = 1/\prod_{j \neq n} I_j$. Although such a sampling approach does not come with any guarantees, it is faster than leverage score sampling since it avoids the cost of computing the sampling distribution. In order to investigate the performance of uniform sampling, we repeat the experiments whose results are show in Tables 3 and 4 for TR-ALS-Sampled, but using uniform sampling instead of leverage score sampling. Everything else is kept the same. Table 5 shows the results.

Using uniform instead of leverage score sampling does not appear to impact the error and overall run time much. However, uniform sampling requires 4.6–8.8 times as many samples to reach the target accuracy. Although more samples increase the cost of the least squares solve, the cost of updating the sampling distribution is avoided. It may be possible to combine uniform and leverage score sampling to speed up our method while still retaining some guarantees. Uniform sampling has also yielded promising empirical results for the CP decomposition; see e.g. Battaglini et al. (2018) and Aggour et al. (2020).

5.2. Rapid Feature Extraction for Classification

Inspired by experiments in Zhao et al. (2016), we now give a practical example of how our method can be used for rapid feature extraction for classification. We consider the full COIL-100 dataset which consists of 7200 color images of size 128×128 depicting 100 different objects from 72 different angles each. The images are downsampled to 32×32 and stacked into a tensor \mathcal{X} of size $32 \times 32 \times 3 \times 7200$. We then use the various TR decomposition algorithms to compute a TR decomposition of \mathcal{X} with each rank $R_n = 5$. Additional details on algorithm settings are given in Section S4.4 of the supplement. The TR core $\mathcal{G}^{(4)}$ is of size $5 \times 7200 \times 5$ and contains latent features for the 4th mode of \mathcal{X} . We reshape it into a 7200×25 feature matrix and apply the k -nearest neighbor algorithm with $k = 1$. Ta-

Table 3. Decomposition results for real datasets with target rank $R = 10$. Time is in seconds.

Method	Pavia Uni.		DC Mall		Park Bench		Tabby Cat		Red Truck	
	Error	Time	Error	Time	Error	Time	Error	Time	Error	Time
TR-ALS	0.16	62.1	0.16	801.4	0.07	5384.9	0.14	1769.4	0.18	181.2
rTR-ALS	0.18	16.3	0.17	25.4	0.08	64.8	0.14	9.0	0.19	22.0
TR-ALS-S. (proposal)	0.18	1.7	0.18	3.1	0.08	3.5	0.15	2.4	0.19	1.9
TR-SVD	0.21	2.1	0.20	12.8	0.08	136.5	0.15	37.7	0.26	1.3
TR-SVD-Rand	0.33	0.2	0.30	0.3	0.11	2.8	0.20	1.6	0.29	0.1

Table 4. Decomposition results for real datasets with target rank $R = 20$. The \times signifies that the SVD-based methods cannot handle this case since they require $R_0R_1 \leq I_1$. Time is in seconds.

Method	Pavia Uni.		DC Mall		Park Bench		Tabby Cat		Red Truck	
	Error	Time	Error	Time	Error	Time	Error	Time	Error	Time
TR-ALS	0.06	249.9	0.06	622.8	0.04	2385.0	0.11	1397.2	0.11	1279.1
rTR-ALS	0.07	194.6	0.06	349.6	0.05	531.5	0.12	92.6	0.12	470.6
TR-ALS-S. (proposal)	0.07	13.7	0.06	26.5	0.05	25.0	0.12	19.0	0.12	22.6
TR-SVD	0.10	4.6	0.10	25.4	0.05	272.0	0.13	110.0	\times	\times
TR-SVD-Rand	0.28	0.5	0.25	1.0	0.10	8.5	0.17	4.1	\times	\times

Table 5. Decomposition results for TR-ALS-Sampled with *uniform sampling*.

Dataset	$R = 10$		$R = 20$	
	Error	Time (s)	Error	Time (s)
Pavia Uni.	0.18	1.3	0.07	9.0
DC Mall	0.18	3.7	0.06	22.2
Park Bench	0.08	5.8	0.05	28.0
Tabby Cat	0.15	2.2	0.12	23.1
Red Truck	0.20	1.5	0.12	18.7

ble 6 reports the decomposition error as well as the accuracy achieved when classifying the 7200 images into the 100 different classes. The classification is done using 10-fold cross validation. The ALS-based methods yield a lower error than the SVD-based methods. All methods result in a similar classification accuracy. The higher decomposition errors for the SVD-based methods do not seem to compromise classification accuracy.

6. Conclusion

We have proposed a method for TR decomposition which uses leverage score sampled ALS and has complexity sub-linear in the number of input tensor entries. We also proved high-probability relative-error guarantees for the sampled least squares problems. Our method achieved substantial speedup over competing methods in experiments on both

Table 6. Decomposition error and classification accuracy for rapid feature extraction experiment.

Method	Time (s)	Error	Acc. (%)
TR-ALS	248.7	0.27	99.07
rTR-ALS	3.1	0.29	98.96
TR-ALS-S. (proposal)	12.7	0.28	99.32
TR-SVD	7.3	0.37	99.68
TR-SVD-Rand	0.7	0.33	99.85

synthetic and real data, in some cases by as much as two or three orders of magnitude, while maintaining good accuracy. We also demonstrated how our method can be used for rapid feature extraction.

The datasets in our experiments are dense. Based on limited testing, rTR-ALS seems to do particularly well on sparse datasets, outperforming our proposed method. Our method still yielded a substantial speedup over TR-ALS. Further investigating this, and optimizing the various methods for sparse tensors, is an interesting direction for future research.

Acknowledgments

We would like to thank the reviewers for their many helpful comments and suggestions which helped improve this paper. The work of OAM was supported by the AFOSR grant FA9550-20-1-0138.

References

- Aggour, K. S., Gittens, A., and Yener, B. Adaptive sketching for fast and convergent canonical polyadic decomposition. In *International Conference on Machine Learning*. PMLR, 2020.
- Ahmadi-Asl, S., Cichocki, A., Phan, A. H., Asante-Mensah, M. G., Mousavi, F., Oseledets, I., and Tanaka, T. Randomized algorithms for fast computation of low-rank tensor ring model. *Machine Learning: Science and Technology*, 2020.
- Bader, B. W., Kolda, T. G., and others. MATLAB Tensor Toolbox, version 2.6. Available online at <https://www.tensortoolbox.org>, 2015.
- Battaglino, C., Ballard, G., and Kolda, T. G. A practical randomized CP tensor decomposition. *SIAM Journal on Matrix Analysis and Applications*, 39(2):876–901, 2018.
- Bengua, J. A., Phien, H. N., and Tuan, H. D. Optimal feature extraction and classification of tensors via matrix product state decomposition. In *2015 IEEE International Congress on Big Data*, pp. 669–672. IEEE, 2015.
- Biagioni, D. J., Beylkin, D., and Beylkin, G. Randomized interpolative decomposition of separated representations. *Journal of Computational Physics*, 281(C):116–134, January 2015. ISSN 0021-9991. doi: 10.1016/j.jcp.2014.10.009.
- Caiafa, C. F. and Cichocki, A. Generalizing the column–row matrix decomposition to multi-way arrays. *Linear Algebra and its Applications*, 433(3):557–573, 2010.
- Che, M. and Wei, Y. Randomized algorithms for the approximations of Tucker and the tensor train decompositions. *Advances in Computational Mathematics*, 45(1):395–428, 2019.
- Cheng, D., Peng, R., Liu, Y., and Perros, I. SPALS: Fast alternating least squares via implicit leverage scores sampling. In *Advances In Neural Information Processing Systems*, pp. 721–729, 2016.
- Cichocki, A., Lee, N., Oseledets, I., Phan, A.-H., Zhao, Q., and Mandic, D. P. Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions. *Foundations and Trends® in Machine Learning*, 9(4-5):249–429, 2016.
- Cichocki, A., Phan, A.-H., Zhao, Q., Lee, N., Oseledets, I., Sugiyama, M., and Mandic, D. P. Tensor networks for dimensionality reduction and large-scale optimization: Part 2 applications and future perspectives. *Foundations and Trends® in Machine Learning*, 9(6):431–673, 2017.
- Cohen, N., Sharir, O., and Shashua, A. On the expressive power of deep learning: A tensor analysis. In *Conference on Learning Theory*, pp. 698–728, 2016.
- da Costa, M. N., Lopes, R. R., and Romano, J. M. T. Randomized methods for higher-order subspace separation. In *2016 24th European Signal Processing Conference (EUSIPCO)*, pp. 215–219. IEEE, 2016.
- Drineas, P. and Mahoney, M. W. A randomized algorithm for a tensor-based generalization of the singular value decomposition. *Linear algebra and its applications*, 420(2-3):553–571, 2007.
- Espig, M., Naraparaju, K. K., and Schneider, J. A note on tensor chain approximation. *Computing and Visualization in Science*, 15(6):331–344, 2012.
- Friedland, S., Mehrmann, V., Miedlar, A., and Nkengla, M. Fast low rank approximations of matrices and tensors. *Electronic Journal of Linear Algebra*, 22:1031–1048, 2011. ISSN 1081-3810. doi: 10.13001/1081-3810.1489.
- Garipov, T., Podoprikin, D., Novikov, A., and Vetrov, D. Ultimate tensorization: Compressing convolutional and FC layers alike. *arXiv preprint arXiv:1611.03214*, 2016.
- Hazan, T., Polak, S., and Shashua, A. Sparse image coding using a 3D non-negative tensor factorization. In *Tenth IEEE International Conference on Computer Vision (ICCV’05)*, volume 1, pp. 50–57. IEEE, 2005.
- Hou, M., Tang, J., Zhang, J., Kong, W., and Zhao, Q. Deep multimodal multilinear fusion with high-order polynomial pooling. In *Advances in Neural Information Processing Systems*, pp. 12136–12145, 2019.
- Huber, B., Schneider, R., and Wolf, S. A randomized tensor train singular value decomposition. In *Compressed Sensing and Its Applications*, pp. 261–290. Springer, 2017.
- Izmailov, P., Novikov, A., and Kropotov, D. Scalable Gaussian processes with billions of inducing inputs via tensor train decomposition. In *International Conference on Artificial Intelligence and Statistics*, pp. 726–735, 2018.
- Khoo, Y., Lu, J., and Ying, L. Efficient construction of tensor ring representations from sampling. *arXiv preprint arXiv:1711.00954*, 2019.
- Khrulkov, V., Novikov, A., and Oseledets, I. Expressive power of recurrent neural networks. In *International Conference on Learning Representations*, 2018.
- Kolda, T. G. and Bader, B. W. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, August 2009. ISSN 0036-1445. doi: 10.1137/07070111X.

- Kolda, T. G. and Sun, J. Scalable tensor decompositions for multi-aspect data mining. In *2008 Eighth IEEE International Conference on Data Mining*, pp. 363–372. IEEE, 2008.
- Larsen, B. W. and Kolda, T. G. Practical leverage-based sampling for low-rank tensor decomposition. *arXiv preprint arXiv:2006.16438*, 2020.
- Lei, T., Xin, Y., Zhang, Y., Barzilay, R., and Jaakkola, T. Low-rank tensors for scoring dependency structures. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1381–1391, 2014.
- Liu, D., Ran, S.-J., Wittek, P., Peng, C., García, R. B., Su, G., and Lewenstein, M. Machine learning by unitary tensor network of hierarchical tree structure. *New Journal of Physics*, 21(7):073059, 2019.
- Mahoney, M. W. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2): 123–224, 2011.
- Mahoney, M. W., Maggioni, M., and Drineas, P. Tensor-CUR decompositions for tensor-based data. *SIAM Journal on Matrix Analysis and Applications*, 30(3):957–987, 2008.
- Malik, O. A. and Becker, S. Low-rank Tucker decomposition of large tensors using TensorSketch. In *Advances in Neural Information Processing Systems*, pp. 10096–10106, 2018.
- Malik, O. A. and Becker, S. Fast randomized matrix and tensor interpolative decomposition using CountSketch. *Advances in Computational Mathematics*, 46:76, 2020. doi: 10.1007/s10444-020-09816-9.
- Mickelin, O. and Karaman, S. On algorithms for and computing with the tensor ring decomposition. *Numerical Linear Algebra with Applications*, 27(3):e2289, 2020.
- Minster, R., Saibaba, A. K., and Kilmer, M. E. Randomized algorithms for low-rank tensor decompositions in the Tucker format. *SIAM Journal on Mathematics of Data Science*, 2(1):189–215, 2020.
- Nene, S. A., Nayar, S. K., and Murase, H. Columbia Object Image Library (COIL-100). Technical Report CUCS-006-96, Columbia University, 1996.
- Novikov, A., Podoprikhin, D., Osokin, A., and Vetrov, D. P. Tensorizing neural networks. In *Advances in Neural Information Processing Systems*, pp. 442–450, 2015.
- Novikov, A., Trofimov, M., and Oseledets, I. Exponential machines. *arXiv preprint arXiv:1605.03795*, 2017.
- Oseledets, I. and Tyrtyshnikov, E. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications*, 432(1):70–88, 2010.
- Oseledets, I. V. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.
- Oseledets, I. V., Savostianov, D. V., and Tyrtyshnikov, E. E. Tucker dimensionality reduction of three-dimensional arrays in linear time. *SIAM Journal on Matrix Analysis and Applications*, 30(3):939–956, 2008.
- Papalexakis, E. E., Faloutsos, C., and Sidiropoulos, N. D. Tensors for data mining and data fusion: Models, applications, and scalable algorithms. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(2):1–44, 2016.
- Rigamonti, R., Sironi, A., Lepetit, V., and Fua, P. Learning separable filters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2754–2761, 2013.
- Stoudenmire, E. M. and Schwab, D. J. Supervised learning with quantum-inspired tensor networks. *arXiv preprint arXiv:1605.05775*, May 2017.
- Sun, Y., Guo, Y., Luo, C., Tropp, J., and Udell, M. Low-rank Tucker approximation of a tensor from streaming data. *SIAM Journal on Mathematics of Data Science*, 2(4):1123–1150, 2020.
- Tarzanagh, D. A. and Michailidis, G. Fast randomized algorithms for t-product based tensor operations and decompositions with applications to imaging data. *SIAM Journal on Imaging Sciences*, 11(4):2629–2664, 2018.
- Tsourakakis, C. E. MACH: Fast randomized tensor decompositions. In *Proceedings of the 2010 SIAM International Conference on Data Mining*, pp. 689–700. SIAM, 2010.
- Vasilescu, M. A. O. and Terzopoulos, D. Multilinear analysis of image ensembles: Tensorfaces. In *European Conference on Computer Vision*, pp. 447–460. Springer, 2002.
- Wang, W., Aggarwal, V., and Aeron, S. Efficient low rank tensor ring completion. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5697–5705, 2017.
- Wang, Y., Tung, H.-Y., Smola, A. J., and Anandkumar, A. Fast and guaranteed tensor decomposition via sketching. In *Advances in Neural Information Processing Systems*, pp. 991–999, 2015.
- Woodruff, D. P. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2):1–157, 2014.

- Yang, B., Zamzam, A., and Sidiropoulos, N. D. ParaSketch: Parallel tensor factorization via sketching. In *Proceedings of the 2018 SIAM International Conference on Data Mining*, pp. 396–404. SIAM, 2018.
- Yang, Y., Krompass, D., and Tresp, V. Tensor-train recurrent neural networks for video classification. *arXiv preprint arXiv:1707.01786*, 2017.
- Ye, J., Wang, L., Li, G., Chen, D., Zhe, S., Chu, X., and Xu, Z. Learning compact recurrent neural networks with block-term tensor decomposition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9378–9387, 2018.
- Ye, K. and Lim, L.-H. Tensor network ranks. *arXiv preprint arXiv:1801.02662*, 2018.
- Yu, R., Zheng, S., Anandkumar, A., and Yue, Y. Long-term forecasting using higher order tensor RNNs. *arXiv preprint arXiv:1711.00073*, 2017.
- Yuan, L., Li, C., Cao, J., and Zhao, Q. Randomized tensor ring decomposition and its application to large-scale data reconstruction. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2127–2131. IEEE, 2019.
- Zhang, J., Saibaba, A. K., Kilmer, M. E., and Aeron, S. A randomized tensor singular value decomposition based on the t-product. *Numerical Linear Algebra with Applications*, 25:e2179, 2018.
- Zhao, Q., Zhou, G., Xie, S., Zhang, L., and Cichocki, A. Tensor ring decomposition. *arXiv preprint arXiv:1606.05535*, 2016.