Randomized algorithms for rounding in the tensor-train format

Hussam Al Daas, Grey Ballard, Paul Cazeaux, Eric Hallman, Agnieszka Miedlar, Mirjeta Pasha, Tim W. Reid and Arvind K. Saibaba

RANDOMIZED ALGORITHMS FOR ROUNDING IN THE TENSOR-TRAIN FORMAT

HUSSAM AL DAAS[*], GREY BALLARD[†], PAUL CAZEAUX[‡], ERIC HALLMAN[§], AGNIESZKA MIĘDLAR[‡], MIRJETA PASHA[¶], TIM W. REID[§], AND ARVIND K. SAIBABA[§]

Abstract. The Tensor-Train (TT) format is a highly compact low-rank representation for high-dimensional tensors. TT is particularly useful when representing approximations to the solutions of certain types of parametrized partial differential equations. For many of these problems, computing the solution explicitly would require an infeasible amount of memory and computational time. While the TT format makes these problems tractable, iterative techniques for solving the PDEs must be adapted to perform arithmetic while maintaining the implicit structure. The fundamental operation used to maintain feasible memory and computational time is called *rounding*, which truncates the internal ranks of a tensor already in TT format. We propose several randomized algorithms for this task that are generalizations of randomized low-rank matrix approximation algorithms and provide significant reduction in computation compared to deterministic TT-rounding algorithms. Randomization is particularly effective in the case of rounding a sum of TT-tensors (where we observe 20× speedup), which is the bottleneck computation in the adaptation of GMRES to vectors in TT format. We present the randomized algorithms and compare their empirical accuracy and computational time with deterministic alternatives.

Key words. high-dimensional problems, randomized algorithms, tensor decompositions, tensor-train format

AMS subject classifications. 15A69, 65F55, 65F99, 65Y20, 68W20.

1. Introduction. An increasing number of applications in science and technology involve the manipulation of multi-dimensional data, or tensors that are higher order equivalents of vectors (first-order) and matrices (second-order). The number of elements of a tensor as well as the storage consumption grow exponentially with the number of the dimensions, a phenomenon known as the *curse of dimensionality*. When problems of high dimensions are concerned, beating the curse of dimensionality and finding a solution efficiently remains a challenge. Nevertheless, different tensor formats and methods based on tensor products [31, 35, 24, 25, 45] have shown potential for mitigating the curse of dimensionality and tackling high-dimensional problems that could not be addressed with conventional methods. Initially, the concept of tensor decompositions was introduced in 1927 by expressing a tensor as the sum of a finite number of rank-one tensors [28] — also known as the *canonical format*. The canonical format's memory requirements are not high, though it can suffer from numerical stability issues [14, 29]. Tensors in Tucker form [6] are well known in quantum chemistry [14, 29] since they yield robust algorithms due to the ability to form an embedded manifold [34], but one of the disadvantages of the Tucker format is its storage consumption that still depends exponentially on the number of dimensions.

One of the most promising tensor formats is the Tensor-Train (TT) format, a

---

[*]Computational Mathematics Group, Rutherford Appleton Laboratory, UK. (hussam.al-daas@stfc.ac.uk)

[†]Department of Computer Science, Wake Forest University, USA. (ballard@wfu.edu)

[‡]Department of Mathematics, University of Kansas, USA. (pcazeaux@ku.edu, amiedlar@ku.edu)

[§]Department of Mathematics, North Carolina State University, USA. (erhallma@ncsu.edu, twreid@ncsu.edu, asaibab@ncsu.edu)

[¶]School of Mathematical and Statistical Sciences, Arizona State University, USA. (mpasha3@asu.edu)

tensor product format that was initially proposed in quantum physics, also known as *matrix product states* (MPS) [20], and was reinvented in numerical linear algebra [46, 44]. It combines both the advantages of the canonical and Tucker formats, i.e., 1) the storage consumption of a tensor depends linearly on the number of dimensions and 2) there exist robust algorithms for the computation of best approximations. Applications of the TT format arise from various applications such as high-dimensional PDEs like the Fokker Planck equations [17, 49], quantum physics [51, 40], high-dimensional data analysis [33, 32], machine learning [7, 19, 11, 42], and uncertainty quantification [54, 38] to mention just a few. Typically, those applications require an approximate solution of linear systems of equations, eigenvalue problems, or completion problems [22, 4, 48]. The TT format is a low-rank representation that, for TT-tensors with small rank, offers a tremendous reduction in the computational complexity and often exposes the structure of the problem. The use of low-rank structures such as the TT format [44] to represent high-dimensional objects allows the solution of linear high-dimensional problems by generalizing standard numerical linear algebra techniques to multi-index arrays of coefficients (tensors) and the multivariate functions they approximate.

In this paper, we focus on the problem of rounding a tensor in TT format; that is, assuming that we are given a TT-tensor, we want to find a compressed representation that is nearly as accurate as the original representation. There are several techniques for computing the initial TT-tensors which do not require forming the entire tensor explicitly [15, 36, 43, 50]. One such technique that is popularly used is called the TT-cross approximation. The standard TT approach to rounding, proposed by Oseledets [44], has two phases [44, Algorithm 2]: orthogonalization followed by compression (typically using the SVD). Here, by orthogonalization, we mean a sweep of orthogonalization steps across every tensor core. Analysis shows that the orthogonalization step dominates the computational cost of this approach. Motivated by this observation, the goal of this work is to develop randomized algorithms for rounding TT-tensors that avoid expensive orthogonalization. In the following, we present the main contributions of this paper.

Overview of the paper and main contributions. This paper develops several new randomized algorithms for rounding tensors in the TT format and is organized as follows. In Section 2, we set some notation as well as review some basic material on randomized matrix algorithms and standard TT operations along with a detailed analysis of their computational costs. In Section 3, we propose various new randomized algorithms for TT-rounding with the focus on Randomize-then-Orthogonalize, Two-Sided-Randomization, and rounding of a sum of TT-tensors.

1. In Algorithm 3.1, Orthogonalize-then-Randomize, we replace the SVD step in the standard TT-rounding algorithm with a randomized SVD assuming that the truncated ranks are known *a priori*.

2. In Algorithm 3.2, Randomize-then-Orthogonalize, we propose to form randomized sketches of each core by nested contractions with a TT-tensor with random cores in a first step, before performing the orthogonalization sweep on these much smaller matrices. Our analysis and experiments show that this approach allowed for the best speedup compared to the deterministic algorithm while retaining excellent accuracy.

3. In Algorithm 3.3, Two-Sided-Randomization, we completely eliminate the need for separate orthogonalization and compression sweeps. Instead, we work with a two-sided randomized approach which computes products with two random tensors followed by a compression step (which involves orthogo-

2

nalization of much smaller matrices). Although this approach is slightly more expensive in terms of flops count and less accurate than techniques mentioned before, it eliminates the need of extensive orthogonalization and allows for the truncation phase to be more independent and highly parallelizable.

4. We extend the Randomize-then-Orthogonalize approach for compressing a TT-tensor that is presented as a sum of TT-tensors (Algorithm 3.4). This special case is of importance in many applications such as solving parametric linear systems in the TT format. The use of randomization enables significant performance improvements by exploiting the structure of the sum tensor in a way that a deterministic algorithm cannot.

We provide an analysis of the computational cost of the proposed algorithms in Subsection 3.5 and show that they are computationally more efficient than existing algorithms. We justify our analysis through numerical experiments in Section 4 on both synthetic data and tensors generated while solving parametric partial differential equations (PDEs). Some conclusions and future outlook are presented in Section 5. The MATLAB code for the implementation and numerical experiments is publicly available at https://github.com/SAMSI-RandTensors/randomizedTT.

Related work. There have been several recent developments in obtaining low-rank compression of tensors. We limit our literature review to the publications dealing with TT-tensors described in [44], which is closest to our work, and refer the reader to review papers for other developments in tensor decompositions [1, 9, 10, 24]. Oseledets [44] proposes a method for rounding TT-tensors. A parallel version of this method is introduced and developed in [12]. Our newly proposed approaches are more computationally efficient compared to existing deterministic algorithms. Other works [30, 8, 2] discuss randomized algorithms for compressing tensors in the TT format. These approaches differ from ours in that they require access to the entries of the tensor, i.e., they do not assume that the tensor is already in TT format. A recent paper [3] also uses randomization to produce a TT approximation of a full tensor but relies on tensor actions (i.e., applications of the tensor on $N-1$ vectors, where N is the order of the tensor). Other methods for constructing a low-rank compression in the TT format involve alternating least squares [18]. The use of tensor random projections in which the random tensors are taken to be in TT format have also been considered in [5, 47, 21]. While these papers use randomization in the context of TT-tensors, none of them directly address the problem of rounding which is the central focus of our paper.

2. Background. Here, we review the notation and necessary operations involving tensors in a modest amount of detail. For a more comprehensive exposition we refer the reader to [44, 35, 12].

2.1. Notation. We denote tensors by boldface script letters (e.g., $\mathcal{X}$) and matrices by boldface Roman letters (e.g., $\mathbf{A}$). We follow MATLAB-like convention and denote the entries of a 3-way tensor $\mathcal{X}$ as $\mathcal{X}(i,j,k)$. A colon denotes the entire range of indices in that dimension. We denote the column fibers as $\mathcal{X}(:,j,k)$, row fibers as $\mathcal{X}(j,:,k)$ and tube fibers as $\mathcal{X}(j,k,:)$. The *mode-n unfolding* (or matricization) of the tensor $\mathcal{X}$ is denoted as $\mathbf{X}_{(n)} \in \mathbb{R}^{I_n \times (I/I_n)}$, where $I = I_1 I_2 \cdots I_n$. The columns of the mode-n unfolding are composed of the appropriate mode-n fibers, e.g., the columns of mode-1 unfolding are column fibers and the columns of mode-3 unfolding are tube fibers. Given a matrix $\mathbf{A} \in \mathbb{R}^{M \times I_n}$, the mode-$n$ product $\mathcal{Y} = \mathcal{X} \times_n \mathbf{A}$ is defined by its mode-n unfolding $\mathbf{Y}_{(n)} = \mathbf{A}\mathbf{X}_{(n)}$. The norm of a tensor is equivalent to the Frobenius norm of any of its unfoldings: $\|\mathcal{X}\| = \|\mathbf{X}_{(n)}\|_F$.

An *order-N tensor* $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_1 \times \cdots \times I_N}$ is in the TT format if there exist positive integers $R_0, \ldots, R_N$ with $R_0 = R_N = 1$ and order-3 tensors $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1}, \ldots, \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N}$, called *TT-cores*, with $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n} \in \mathbb{R}^{R_{n-1} \times I_n \times R_n}$ for $1 \leq n \leq N$, such that

$$\boldsymbol{\mathcal{X}}(i_1, \ldots, i_N) = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1}(i_1,:) \cdot \ldots \cdot \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}(:, i_n, :) \cdot \ldots \cdot \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N}(:, i_N),$$

where $1 \leq i_n \leq I_n$. Note that because $R_0 = R_N = 1$, the first and last TT-cores are (order-2) matrices so $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1}(i_1,:) \in \mathbb{R}^{R_1}$ and $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N}(:, i_N) \in \mathbb{R}^{R_{N-1}}$. The $R_{n-1} \times R_n$ matrix $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}(:, i_n, :)$ is referred to as the i_nth slice of the nth TT-core of $\boldsymbol{\mathcal{X}}$. It is worth mentioning that the TT decomposition is not unique due to the multiplicative nature of the format.


FIG. 2.1. *Horizontal and vertical unfoldings of a TT-core* $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}$.

In order to express the arithmetic operations on TT-cores using linear algebra, we will often use two specific matrix unfoldings of the order-3 tensors. The *horizontal unfolding* of a TT-core $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}$ corresponds to the concatenation of the slices $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}(:, i_n, :)$ for $i_n = 1, \ldots, I_n$ horizontally. We denote the corresponding operator by $\mathcal{H}$, so that $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})$ is an $R_{n-1} \times I_n R_n$ matrix. The *vertical unfolding* of a TT-core $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}$ corresponds to the concatenation of the slices $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}(:, i_n, :)$ for $i_n = 1, \ldots, I_n$ vertically. We denote the corresponding operator by $\mathcal{V}$, so that $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})$ is an $R_{n-1}I_n \times R_n$ matrix, see Figure 2.1. Moreover, we will often make use of a *tensor network diagram*, see Figure 2.2, to graphically illustrate TT-tensor operations. Here nodes represent tensors and edges represent modes so that connected nodes can be contracted.

FIG. 2.2. *Tensor network diagram for an order-5 TT-tensor.*

Let $\mathbf{X}_{(1:n)} \in \mathbb{R}^{(I_1 I_2 \cdots I_n) \times (I_{n+1} \cdots I_N)}$ denote an unfolding of the first n modes of a TT-tensor $\boldsymbol{\mathcal{X}}$. It has the rank R_n representation

$$\mathbf{X}_{(1:n)} = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1:n}) \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1:N}),$$

where in an extension of their earlier definitions, $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1:n}) \in \mathbb{R}^{(I_1 I_2 \cdots I_n) \times R_n}$ represents the mode-$(n+1)$ unfolding of the product of the first n TT-cores and $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1:N}) \in \mathbb{R}^{R_n \times (I_{n+1} \cdots I_N)})$ represents the mode-1 unfolding of the product of the final $N - n$

TT-cores. Likewise, we can write the same unfolding as a product of four matrices, see [12, Eq. (2.3)], i.e.,

$$(2.1) \qquad \mathbf{X}_{(1:n)} = (\mathbf{I}_{I_n} \otimes \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1:n-1}))\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1})(\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+2:N}) \otimes \mathbf{I}_{I_{n+1}}).$$

Suppose we have two tensors $\boldsymbol{\mathcal{Y}}$ and $\boldsymbol{\mathcal{Z}}$ of the same dimension, and consider their sum $\boldsymbol{\mathcal{X}}$. The cores of the tensor $\boldsymbol{\mathcal{X}}$ can be expressed as

$$\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}(:,i_n,:) = \begin{bmatrix} \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n}(:,i_n,:) & \\ & \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},n}(:,i_n,:) \end{bmatrix} \qquad 2 \le n \le N-1,$$

and for the first and the last core, we have

$$\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1}(i_1,:) = \begin{bmatrix} \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1}(i_1,:) & \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},1}(i_1,:) \end{bmatrix} \quad \text{and} \quad \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N}(:,i_N) = \begin{bmatrix} \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},N}(:,i_N) \\ \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},N}(:,i_N) \end{bmatrix}.$$

Let $\mathbf{X} \in \mathbb{R}^{m \times n}$ with $m \ge n$. We denote the *thin QR* factorization of $\mathbf{X}$ as $\mathbf{X} = \mathbf{QR}$, where $\mathbf{Q} \in \mathbb{R}^{m \times n}$ has orthonormal columns and $\mathbf{R} \in \mathbb{R}^{n \times n}$ is upper triangular; we also write $[\mathbf{Q}, \mathbf{R}] = \mathrm{QR}(\mathbf{X})$ for use in algorithms. The SVD of $\mathbf{X}$ is denoted by $\mathbf{U\Sigma V}^\top$, where the matrix $\mathbf{U} \in \mathbb{R}^{m \times n}$ has orthonormal columns containing the left singular vectors, $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with the singular values on the diagonal and $\mathbf{V} \in \mathbb{R}^{n \times n}$ is an orthogonal matrix, whose columns contain the right singular vectors. Assuming that the Householder QR algorithm is used and $\mathbf{Q}$ is formed explicitly, the computational cost of the QR factorization is $4mn^2 - \frac{4n^3}{3} + \mathcal{O}(n^2)$ flops. Given a threshold $\varepsilon > 0$, we truncate the singular values of $\mathbf{X}$ to obtain a rank-k approximation $\mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^\top$ of matrix $\mathbf{X}$, which satisfies $\|\mathbf{X} - \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^\top\|_F \le \varepsilon \|\mathbf{X}\|_F$. This is denoted as $[\mathbf{U}_k, \mathbf{\Sigma}_k, \mathbf{V}_k] = \mathrm{SVD}(\mathbf{X}, \varepsilon)$. The computational cost of computing the SVD is $\mathcal{O}(mn^2)$ flops.

2.2. Randomized matrix algorithms. An important component of our approach is the use of randomized matrix methods for low-rank matrix approximation. In this subsection, we briefly review a few well-established randomized algorithms.

The first algorithm is the basic version of the randomized SVD proposed in [26]. Suppose we want to compute a low-rank approximation of a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$; let the target rank be denoted by r and we pick an oversampling parameter p such that $r + p \le \min\{m, n\}$. We generate a random matrix $\mathbf{\Omega} \in \mathbb{R}^{n \times (r+p)}$; in practice, we take the entries of this matrix to be i.i.d. standard Gaussian random variables. Then, we compute the product $\mathbf{Y} = \mathbf{X\Omega}$ and obtain its thin QR factorization $\mathbf{Y} = \mathbf{QR}$. The main insight exploited by randomized SVD is that if the rank of $\mathbf{X}$ is close to r, or the singular values of $\mathbf{X}$ decay rapidly beyond r, then the range of $\mathbf{Q}$ approximates well the range of $\mathbf{X}$ in the sense that $\mathbf{X} \approx \mathbf{QQ}^\top\mathbf{X}$; we then use $\mathbf{QQ}^\top\mathbf{X}$ as a low-rank approximation to $\mathbf{X}$. The computational cost of this approach is

$$C_{\mathrm{randSVD}} = 2(r+p)mn + \mathcal{O}(r^2(m+n)) \quad \text{flops.}$$

Additional postprocessing can be performed to convert the low-rank approximation in the SVD format, or to truncate the low-rank approximation to rank$-r$; see [26] for additional details.

There is one variant of this algorithm that is of particular importance to our newly proposed methods, the generalized Nyström method [41]. The generalized Nyström method avoids the orthogonalization step when computing a low-rank approximation

by using a two-sided randomized approach. Let us define two Gaussian random matrices $\mathbf{\Omega} \in \mathbb{R}^{n \times s}$ and $\mathbf{\Psi} \in \mathbb{R}^{t \times m}$, where $r \leq s \leq \min\{m, n\}$ (note that t also satisfies a similar inequality). A low-rank approximation to $\mathbf{X}$ is computed as

$$(2.2) \qquad\qquad \mathbf{X} \approx \mathbf{Y}(\mathbf{\Psi}\mathbf{X}\mathbf{\Omega})^{\dagger}\mathbf{Z},$$

where $\mathbf{Y} = \mathbf{X}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{\Psi}\mathbf{X}$. To implement the pseudoinverse, [41] suggests computing the QR factorization $\mathbf{\Psi}\mathbf{X}\mathbf{\Omega} = \mathbf{Q}\mathbf{R}$ and then obtaining the low-rank approximation $(\mathbf{Y}\mathbf{R}^{-1})(\mathbf{Q}^{\top}\mathbf{Z})$. If the low-rank approximation is desired in the SVD format, this can be done by additional post-processing. In [41], the author recommends setting the sketch parameters as $s = r$ and $t = \lceil 1.5r \rceil$. The associated computational cost is

$$C_{\mathrm{genNys}} = 2mn(s + t) + \mathcal{O}(t^2(m + n) + ts^2) \quad \text{flops.}$$

2.3. Standard TT arithmetic. In this subsection, we review the standard approach to TT-rounding, first proposed in [44], using the notation of [12]. We also review the concepts of tensor contractions.

To explain the rounding procedure for TT-tensors, we consider the following analogy from matrices. Let $\mathbf{Y} = \mathbf{A}\mathbf{B}$ be an outer product matrix where $\mathbf{A}$ is $m \times r$ and $\mathbf{B}$ is $r \times n$ and $r \leq \min\{m, n\}$. To obtain an approximation of $\mathbf{Y}$ with rank $\ell < r$, we employ an orthogonalization step followed by a compression step. In the orthogonalization step, we want to make $\mathbf{Y}$ right orthogonal. That is, we compute the thin QR factorization $\mathbf{B}^{\top} = \mathbf{Q}\mathbf{R}$, and then compute $\mathbf{Z} = \mathbf{A}\mathbf{R}^{\top}$. This gives $\mathbf{Y} = \mathbf{A}\mathbf{B} = \mathbf{Z}\mathbf{Q}^{\top}$, where $\mathbf{Q}^{\top}$ has orthonormal rows. In the second step, we compress $\mathbf{Z}$ by computing the rank-ℓ truncated SVD $\mathbf{Z} \approx \mathbf{U}\mathbf{\Sigma}\mathbf{V}_Z^{\top}$. To obtain an overall low-rank approximation to $\mathbf{Y}$, we compute $\mathbf{V} = \mathbf{Q}\mathbf{V}_Z$, so that $\mathbf{Y} \approx \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\top}$.

Following [12], we say a tensor is *right orthogonal* if its horizontal unfoldings $\mathcal{H}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},n})$ have orthonormal rows for $n = 2, \ldots, N$ (all except the first core). Similarly, we say that a tensor is *left orthogonal* if its vertical unfoldings $\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},n})$ have orthonormal columns for $n = 1, \ldots, N-1$ (all except the last core).

Right-to-Left Orthogonalization. Suppose we are given a TT-tensor $\mathbf{\mathcal{Y}}$. To obtain a right orthogonal TT-tensor $\mathbf{\mathcal{X}}$ equivalent to $\mathbf{\mathcal{Y}}$, we first compute the thin QR factorization $\mathbf{Q}\mathbf{R} = \mathcal{H}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{Y}},N})^{\top}$ and set the core tensors $\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},N-1}$ and $\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},N}$ as

$$\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{Y}},N-1})\mathcal{H}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{Y}},N}) = \mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{Y}},N-1})(\mathbf{Q}\mathbf{R})^{\top} = \underbrace{(\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{Y}},N-1})\mathbf{R}^{\top})}_{\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},N-1})} \underbrace{(\mathbf{Q}^{\top})}_{\mathcal{H}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},N})} .$$

This procedure is continued through cores $N-1, \ldots, 2$ but we do not orthogonalize the first core. The details of right-to-left orthogonalization are given in Algorithm 2.1 which will form the foundation for many of the subsequent algorithms. We can similarly obtain a left orthogonal tensor by processing the modes starting from mode-1, but we omit the details here.

TT-Rounding. Suppose, now, that we want to round the tensor $\mathbf{\mathcal{Y}}$ in the TT format, i.e., compress the TT format of a tensor by decreasing the TT-ranks $\{R_n\}$. In the first step of the TT-rounding approach, we first obtain a tensor $\mathbf{\mathcal{X}}$ that is right-orthogonal by applying Algorithm 2.1. Starting with mode-1, for each mode, we compute a low-rank approximation of the vertical unfolding $\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},n})$; rather than computing an SVD directly, we first compute the thin QR factorization of $\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},n})$; followed by an SVD of the upper triangular factor $\mathbf{R}$. We then obtain a low-rank approximation $\mathcal{V}(\mathbf{\mathcal{T}}_{\mathbf{\mathcal{X}},n}) \approx \widehat{\mathbf{U}}\widehat{\mathbf{\Sigma}}\widehat{\mathbf{V}}^{\top}$. The number of singular values and vectors retained

Algorithm 2.1 Right-to-Left Orthogonalization

Require: A tensor $\boldsymbol{\mathcal{Y}}$ in TT format
Ensure: $\boldsymbol{\mathcal{X}}$ is right-orthogonal tensor equivalent to $\boldsymbol{\mathcal{Y}}$
1: **function** $\boldsymbol{\mathcal{X}} = \text{ORTHOGONALIZERL}(\boldsymbol{\mathcal{Y}})$
2: $\quad \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N} = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},N}$
3: $\quad$ **for** $n = N$ down to 2 **do**
4: $\quad\quad [\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})^{\top}, \mathbf{R}] = \text{QR}(\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})^{\top})$ $\qquad\qquad$ ▷ thin QR factorization
5: $\quad\quad \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n-1}) = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n-1}) \cdot \mathbf{R}^{\top}$ $\qquad$ ▷ $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n-1} = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n-1} \times_3 \mathbf{R}^{\top}$
6: $\quad$ **end for**
7: **end function**

---

in the low-rank approximation, depend on the threshold $\varepsilon = \frac{\|\boldsymbol{\mathcal{Y}}\|}{\sqrt{N-1}}\varepsilon_0$, where ε_0 is a user-defined threshold that controls the overall accuracy. We then rewrite $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})$ by combining it with the low-rank factor as $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}) = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})\widehat{\mathbf{U}}$. The other two factors $\widehat{\boldsymbol{\Sigma}}\widehat{\mathbf{V}}^{\top}$ are combined with the horizontal unfolding $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1})$ for processing at the next step. This process is terminated after $N-1$ steps and the resulting tensor $\boldsymbol{\mathcal{X}}$ satisfies $\|\boldsymbol{\mathcal{X}} - \boldsymbol{\mathcal{Y}}\| \leq \varepsilon_0\|\boldsymbol{\mathcal{Y}}\|$. The details are given in Algorithm 2.2.

Algorithm 2.2 TT-Rounding

Require: A tensor $\boldsymbol{\mathcal{Y}}$ in TT format, user-defined threshold $\varepsilon_0 > 0$
Ensure: A tensor $\boldsymbol{\mathcal{X}}$ in TT format with reduced ranks such that $\|\boldsymbol{\mathcal{X}} - \boldsymbol{\mathcal{Y}}\| \leq \varepsilon_0\|\boldsymbol{\mathcal{Y}}\|$
1: **function** $\boldsymbol{\mathcal{X}} = \text{TT-ROUNDING}(\boldsymbol{\mathcal{Y}}, \varepsilon_0)$
2: $\quad \boldsymbol{\mathcal{X}} = \text{ORTHOGONALIZERL}(\boldsymbol{\mathcal{Y}})$
3: $\quad$ Compute $\|\boldsymbol{\mathcal{Y}}\|_F$ and the truncation threshold $\varepsilon = \frac{\|\boldsymbol{\mathcal{Y}}\|}{\sqrt{N-1}}\varepsilon_0$
4: $\quad$ Set $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1} = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1}$.
5: $\quad$ **for** $n = 1$ to $N-1$ **do**
6: $\quad\quad [\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}), \mathbf{R}] = \text{QR}(\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}))$ $\qquad\qquad$ ▷ thin QR factorization
7: $\quad\quad [\widehat{\mathbf{U}}, \widehat{\boldsymbol{\Sigma}}, \widehat{\mathbf{V}}] = \text{SVD}(\mathbf{R}, \varepsilon)$ $\qquad\qquad$ ▷ ε-truncated SVD factorization
8: $\quad\quad \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}) = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})\widehat{\mathbf{U}}$ $\qquad\qquad$ ▷ $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n} = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n} \times_3 \widehat{\mathbf{U}}$
9: $\quad\quad \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1}) = \widehat{\boldsymbol{\Sigma}}\widehat{\mathbf{V}}^{\top}\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1})$ $\quad$ ▷ $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1} = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1} \times_1 (\widehat{\boldsymbol{\Sigma}}\widehat{\mathbf{V}}^{\top})$
10: $\quad$ **end for**
11: **end function**

Right-to-Left Partial Contraction. We consider two TT-tensors $\boldsymbol{\mathcal{X}}$ and $\boldsymbol{\mathcal{Y}}$ with ranks $\{R_j^{\boldsymbol{\mathcal{X}}}\}$ and $\{R_j^{\boldsymbol{\mathcal{Y}}}\}$, respectively. For $n = 2, \ldots, N$ we define the partial contraction matrices

$$(2.3) \qquad \mathbf{W}_{n-1} = \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n:N})\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n:N})^{\top} \in \mathbb{R}^{R_{n-1}^{\boldsymbol{\mathcal{X}}} \times R_{n-1}^{\boldsymbol{\mathcal{Y}}}}.$$

These partial contractions can be computed sequentially as

$$(2.4) \qquad \begin{aligned} \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},n}) &= \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})\mathbf{W}_n, \\ \mathbf{W}_{n-1} &= \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},n})\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n})^{\top} \end{aligned}$$

for $n = 2, \ldots, N-1$, with $\mathbf{W}_{N-1} = \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N})\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},N})^{\top}$. Here $\boldsymbol{\mathcal{Z}}$ is a temporary TT-tensor with compatible dimensions and ranks.

The process of computing the matrices $\{\mathbf{W}_{n-1}\}_{n=2}^{N}$ according to (2.4) is called a *right-to-left partial contraction* of tensors $\boldsymbol{\mathcal{X}}$ and $\boldsymbol{\mathcal{Y}}$, and is illustrated in Figure 2.3. The corresponding algorithm is presented in Algorithm 2.3.

(a) Initial structure   (b) $\mathbf{W}_3$ (c) $\mathbf{W}_2$ (d) $\mathbf{W}_1$

FIG. 2.3. *Right-to-Left partial contraction steps for $N = 4$.*

Algorithm 2.3 Right-to-Left Contraction of Tensors $\mathcal{X}$ and $\mathcal{Y}$.

Require: Tensors $\mathcal{X}, \mathcal{Y}$ with consistent dimensions in TT format and ranks $\{R_n^{\mathcal{X}}\}$ and $\{R_n^{\mathcal{Y}}\}$, respectively.

Ensure: Matrices $\{\mathbf{W}_n\}$ satisfy $\mathbf{W}_n = \mathcal{H}(\mathcal{T}_{\mathcal{X},n+1:N})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N})^{\top}$ for $1 \le n < N$

 1: **function** $[\{\mathbf{W}_n\}] = \text{PARTIALCONTRACTIONSRL}(\mathcal{X}, \mathcal{Y})$
 2:      $\mathbf{W}_{N-1} = \mathcal{H}(\mathcal{T}_{\mathcal{X},N})\mathcal{H}(\mathcal{T}_{\mathcal{Y},N})^{\top}$
 3:      **for** $n = N - 1$ down to 2 **do**
 4:          $\mathcal{V}(\mathcal{T}_{\mathcal{Z},n}) = \mathcal{V}(\mathcal{T}_{\mathcal{X},n})\mathbf{W}_n$ $\triangleright$ $\mathcal{T}_{\mathcal{Z},n} = \mathcal{T}_{\mathcal{X},n} \times_3 \mathbf{W}_n$, for temporary $\mathcal{T}_{\mathcal{Z},n}$
 5:          $\mathbf{W}_{n-1} = \mathcal{H}(\mathcal{T}_{\mathcal{Z},n})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n})^{\top}$ $\triangleright$ matrix multiplication, $\mathbf{W}_{n-1}$ is $R_{n-1}^{\mathcal{X}} \times R_{n-1}^{\mathcal{Y}}$
 6:      **end for**
 7: **end function**

Detailed analysis of the overall computational costs of *Right-to-Left Orthogonalization* (Algorithm 2.1), *TT-Rounding* (Algorithm 2.2) and *Right-to-Left Contraction* (Algorithm 2.3) is presented in Appendix A.

3. Randomized Algorithms for TT Rounding. In this section, we propose three new randomized algorithms to perform rounding of a tensor in the TT format, i.e., given an original TT-tensor $\mathcal{Y}$ with TT-ranks $\{R_n\}$ we seek a compressed TT-tensor representation $\mathcal{X}$ with *a priori* known target ranks $\{\ell_n\}$. In randomized SVD, see Subsection 2.2, it is common to include an oversampling term; that is, if we seek a rank-r decomposition of a matrix $\mathbf{X}$, we use the number of samples (alternatively, columns of $\mathbf{\Omega}$) as $\ell = r + p$, where r is the target rank and p is the oversampling parameter. The resulting low-rank approximation $\mathbf{Q}\mathbf{Q}^{\top}\mathbf{X}$ is of rank ℓ. However, in the TT case, to save on notation, when we say target TT-ranks $\{\ell_n\}$, we assume that this rank automatically includes the necessary oversampling parameter.

3.1. Orthogonalize-then-Randomize. The first algorithm we propose is very similar to the standard TT-rounding algorithm; the main difference is that we replace the truncated SVD step in Algorithm 2.2 with the basic version of the randomized SVD reviewed in Subsection 2.2. The nomenclature of this algorithm is clear from the fact that there are two phases in this approach: an (already discussed) orthogonalization phase followed by a compression phase which utilizes randomized SVD.

In the first version, we assume that given a tensor $\mathcal{Y}$ with TT-ranks $\{R_n\}$ we want to obtain a compressed representation with ranks $\{\ell_n\}$; that is, we assume that the target TT-ranks are known in advance. The first phase, i.e., the orthogonalization phase is accomplished using Algorithm 2.1 to obtain the tensor $\mathcal{Y}$ which is right-orthogonal and equivalent to $\mathcal{Y}$. In the second phase, we loop over the cores of the

---

Algorithm 3.1 TT-Rounding: Orthogonalize-then-Randomize

Require: A tensor $\boldsymbol{\mathcal{Y}}$ in TT format with ranks $\{R_n\}$, target TT-ranks $\{\ell_n\}$
Ensure: A tensor $\boldsymbol{\mathcal{X}}$ in TT format with ranks $\{\ell_n\}$
 1: **function** $\boldsymbol{\mathcal{X}} = \text{TT-ROUNDING-ORTHRAND}(\boldsymbol{\mathcal{Y}}, \{\ell_n\})$
 2:     $\boldsymbol{\mathcal{Y}} = \text{ORTHOGONALIZERL}(\boldsymbol{\mathcal{Y}})$
 3:     **for** $n = 1$ to $N - 1$ **do**
 4:         $\mathbf{Z}_n = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n})$ $\triangleright$ $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n}$ is $\ell_{n-1} \times I_n \times R_n$
 5:         $\mathbf{Y}_n = \mathbf{Z}_n \boldsymbol{\Omega}_n$ $\triangleright$ form the sketched matrix
 6:         $[\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}), \sim] = \text{QR}(\mathbf{Y}_n)$ $\triangleright$ thin QR to compute an orthonormal basis
 7:         $\mathbf{M}_n = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})^\top \mathbf{Z}_n$ $\triangleright$ form $\ell_n \times R_n$ matrix
 8:         $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1}) = \mathbf{M}_n \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1})$ $\triangleright$ $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1} = \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1} \times_1 \mathbf{M}_n$
 9:     **end for**
10: **end function**

---

tensor $\boldsymbol{\mathcal{Y}}$ (excluding the last core): for each core, we apply randomized SVD to the vertical unfolding $\mathbf{Z}_n = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n})$. That is, we compute $\mathbf{Q}_n$ such that

$$\mathbf{Z}_n \approx \mathbf{Q}_n \mathbf{Q}_n^\top \mathbf{Z}_n,$$

by first generating a random Gaussian matrix $\boldsymbol{\Omega}_n \in \mathbb{R}^{R_n \times \ell_n}$. We then compute $\mathbf{Y}_n = \mathbf{Z}_n \boldsymbol{\Omega}_n$ and its thin QR factorization to obtain an orthonormal basis $\mathbf{Q}_n$ for the range of $\mathbf{Y}_n$. Finally, we set $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}) = \mathbf{Q}_n$ and $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1}) = (\mathbf{Q}_n^\top \mathbf{Z}_n)\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1})$ to obtain the compressed tensor $\boldsymbol{\mathcal{X}}$. The details are given in Algorithm 3.1.

We also investigated a version of Algorithm 3.1 that does not require the ranks of the rounded tensors to be known in advance. In this approach, we replace Lines 5 and 6 of Algorithm 3.1 with a randomized range finder algorithm [39, 53]. This adaptive method would produce a tensor $\boldsymbol{\mathcal{X}}$ that satisfies the desired tolerance $\|\boldsymbol{\mathcal{X}} - \boldsymbol{\mathcal{Y}}\| \leq \varepsilon_0 \|\boldsymbol{\mathcal{Y}}\|$, where ε_0 is a user defined threshold. While this approach produced TT-tensors with the desired tolerance, in numerical experiments it did not yield significant speedup over deterministic TT-rounding (Algorithm 2.2). Some insight into this behavior is given in Subsection 3.5 which shows that the orthogonalization step is the dominant computational cost of deterministic and Orthogonalize-then-Randomize algorithms. Therefore, we did not pursue this approach further.

3.2. Randomize-then-Orthogonalize. To motivate the next algorithm, we consider the overall computational cost of Algorithm 3.1, see Subsection 3.5, which is dominated by the first, orthogonalization, phase of the algorithm. First, we consider a new Randomize-then-Orthogonalize algorithm that uses randomization to reduce the overall computation cost of the TT-rounding procedure. It works by avoiding an expensive orthogonalization of the original TT-tensor $\boldsymbol{\mathcal{Y}}$ with TT-ranks $\{R_n\}$ and instead uses randomization to reduce the computational cost. In contrast to the next approach in Subsection 3.3 (Two-Sided-Randomization), here we use randomization only on one side.

We first offer a way to construct random Gaussian TT-tensors whose cores are composed of independent random Gaussian entries.

DEFINITION 3.1 (Random Gaussian TT-Tensor). *Given a set of target TT-ranks $\{\ell_n\}$, we generate a random Gaussian TT-tensor $\boldsymbol{\mathcal{R}} \in \mathbb{R}^{I_1 \times \cdots I_N}$ such that each core tensor $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{R}},n} \in \mathbb{R}^{\ell_{n-1} \times I_n \times \ell_n}$ is filled with random, independent, normally distributed entries with mean 0 and variance $1/(\ell_{n-1} I_n \ell_n)$ for $1 \leq n \leq N$.*

By this definition, while the cores of $\mathcal{R}$ have independent entries, the entries of the full tensor themselves are not independent. This normalization is chosen such that $\mathbb{E}\|\mathcal{T}_{\mathcal{R},n}\|_F^2 = 1$ and is sometimes necessary to ensure that no overflow occurs during the rounding computations. Note that constructing this random tensor requires only generating and storing $\sum_{n=1}^{N} \ell_{n-1}\ell_n I_n$ random entries. A related but distinct definition for a Gaussian TT-tensor is given in [47], but the approach taken here differs considerably in how we use the randomized tensor.

In Algorithm 3.2, we first generate a random Gaussian tensor $\mathcal{R}$ with given target TT-ranks $\{\ell_n\}$ following Definition 3.1. Next, we use the efficient multiplication of tensor $\mathcal{R}$ with a given tensor $\mathcal{Y}$, see Algorithm 2.3, to obtain the sketches (sometimes also referred to as partial random projections) $\{\mathbf{W}_n\}$ of $\mathcal{Y}$ (*randomization phase*). A visualization of this process is provided in Figure 3.1. Finally, we construct a left-orthogonal compressed TT-tensor $\mathcal{X}$. Starting with $n = 1$ and $\mathcal{T}_{\mathcal{X},1} = \mathcal{T}_{\mathcal{Y},1}$, we compute the QR factorization of the sketched matrix, i.e.,

$$[\mathcal{V}(\mathcal{T}_{\mathcal{X},n})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N})]\,\mathcal{H}(\mathcal{T}_{\mathcal{R},n+1:N})^\top = \mathcal{V}(\mathcal{T}_{\mathcal{X},n})\mathbf{W}_n = \mathbf{Q}_n\mathbf{R}_n.$$

Since at the nth step the first $n-1$ cores of $\mathcal{X}$ are already orthogonalized, they do not need to be considered explicitly in the above factorization. By projecting $\mathcal{V}(\mathcal{T}_{\mathcal{X},n})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N})$ onto the column space of $\mathbf{Q}_n$, we approximate the product of the final $N - n + 1$ cores as

$$\mathcal{V}(\mathcal{T}_{\mathcal{X},n})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N}) \approx \mathbf{Q}_n\mathbf{Q}_n^\top\mathcal{V}(\mathcal{T}_{\mathcal{X},n})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N}) = \mathbf{Q}_n\mathbf{M}_n\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N}).$$

Then, the cores are updated, i.e., $\mathcal{V}(\mathcal{T}_{\mathcal{X},n}) = \mathbf{Q}_n$ and $\mathcal{H}(\mathcal{T}_{\mathcal{X},n+1}) = \mathbf{M}_n\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1})$.

It is important to mention that the Randomize-then-Orthogonalize approach produces a left-orthogonal tensor $\mathcal{X}$. We can use this observation to compress the tensor further. Therefore, if the ranks are not known a priori, then we choose the ranks to be sufficiently large and truncate them further by using Algorithm 2.2. In particular, since the output tensor of Algorithm 3.2 is left orthogonal, we can skip the orthogonalization phase (Line 2 of Algorithm 2.2), and execute Lines 3 to 10. This is what we do in our numerical experiments when the rank is not known *a priori*.

Algorithm 3.2 TT-Rounding: Randomize-then-Orthogonalize

Require: A tensor $\mathcal{Y}$ in TT format with ranks $\{R_n\}$, target TT-ranks $\{\ell_n\}$
Ensure: A tensor $\mathcal{X}$ in TT format with ranks $\{\ell_n\}$
1: **function** $\mathcal{X} = $ TT-ROUNDING-RANDORTH($\mathcal{Y}, \{\ell_n\}$)
2:    Select a random Gaussian TT-tensor $\mathcal{R}$ with target TT-ranks $\{\ell_n\}$
3:    $\{\mathbf{W}_n\} = $ PARTIALCONTRACTIONSRL($\mathcal{Y}, \mathcal{R}$) $\triangleright$ compute partial random contractions
4:    $\mathcal{T}_{\mathcal{X},1} = \mathcal{T}_{\mathcal{Y},1}$
5:    **for** $n = 1$ to $N - 1$ **do**
6:        $\mathbf{Z}_n = \mathcal{V}(\mathcal{T}_{\mathcal{X},n})$ $\triangleright$ $\mathcal{T}_{\mathcal{X},n}$ is $\ell_{n-1} \times I_n \times R_n$
7:        $\mathbf{Y}_n = \mathbf{Z}_n\mathbf{W}_n$ $\triangleright$ form the sketched matrix
8:        $[\mathcal{V}(\mathcal{T}_{\mathcal{X},n}), \sim] = $ QR($\mathbf{Y}_n$) $\triangleright$ thin QR to compute an orthonormal basis
9:        $\mathbf{M}_n = \mathcal{V}(\mathcal{T}_{\mathcal{X},n})^\top\mathbf{Z}_n$ $\triangleright$ form $\ell_n \times R_n$ matrix
10:        $\mathcal{H}(\mathcal{T}_{\mathcal{X},n+1}) = \mathbf{M}_n\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1})$ $\triangleright$ $\mathcal{T}_{\mathcal{X},n+1} = \mathcal{T}_{\mathcal{Y},n+1} \times_1 \mathbf{M}_n$
11:    **end for**
12: **end function**

FIG. 3.1. *Random projection for the Randomize-then-Orthogonalize Algorithm 3.2.*

Algorithm 3.3 TT-Rounding: Two-Sided-Randomization (Generalized Nyström)

Require: A tensor $\mathcal{Y}$ in TT format with ranks $\{R_n\}$, target TT-ranks $\{\ell_n\}$ and $\{\rho_n\}$
Ensure: A tensor $\mathcal{X}$ in TT format with ranks $\{\ell_n\}$

1: **function** $\mathcal{X} = \text{TT-Rounding-RandOrth}(\mathcal{Y}, \{\ell_n\})$
2:     Generate random Gaussian TT-tensor $\mathcal{L}$ with ranks $\{\ell_n\}$
3:     Generate random Gaussian TT-tensor $\mathcal{R}$ with ranks $\{\rho_n\}$ ▷ choose $\rho_n > \ell_n$
4:     $\{\mathbf{W}_n^L\} = \text{PartialContractionsLR}(\mathcal{Y}, \mathcal{L})$
                                    ▷ Precompute sketches from the left
5:     $\{\mathbf{W}_n^R\} = \text{PartialContractionsRL}(\mathcal{Y}, \mathcal{R})$
                                    ▷ Precompute sketches from the right
6:     **for** $n = 1$ to $N - 1$ **do**
7:         $[\mathbf{U}_n, \mathbf{\Sigma}_n, \mathbf{V}_n] = \text{SVD}(\mathbf{W}_n^L \mathbf{W}_n^R)$ ▷ Compute SVD of $\mathbf{W}_n^L \mathbf{W}_n^R$
                                    ▷ $\mathbf{V}_n$ is $\rho_n \times \ell_n$
8:         $\mathbf{L}_n = \mathbf{W}_n^R \mathbf{V}_n (\mathbf{\Sigma}_n^{\dagger})^{1/2}$ ▷ Determine internal $R_n \times \ell_n$ left factor $\mathbf{L}_n$
9:         $\mathbf{R}_n = (\mathbf{\Sigma}_n^{\dagger})^{1/2} \mathbf{U}_n^{\top} \mathbf{W}_n^L$ ▷ Determine internal $\ell_n \times R_n$ right factor $\mathbf{R}_n$
10:     **end for**
11:     $\mathcal{V}(\mathcal{T}_{\mathcal{X},1}) = \mathcal{V}(\mathcal{T}_{\mathcal{Y},1})\mathbf{L}_1$
12:     **for** $n = 2$ to $N - 1$ **do**
13:         $\mathcal{H}(\mathcal{T}_{\mathcal{X},n}) = \mathbf{R}_{n-1}\mathcal{H}(\mathcal{V}(\mathcal{T}_{\mathcal{Y},n})\mathbf{L}_n)$ ▷ $\mathcal{T}_{\mathcal{X},n} = \mathcal{T}_{\mathcal{Y},n} \times_1 \mathbf{R}_{n-1} \times_3 \mathbf{L}_n$
                                    ▷ hence $\mathcal{T}_{\mathcal{X},n}$ is $\ell_{n-1} \times I_n \times \ell_n$
14:     **end for**
15:     $\mathcal{H}(\mathcal{T}_{\mathcal{X},N}) = \mathbf{R}_{N-1}\mathcal{H}(\mathcal{T}_{\mathcal{Y},N})$
16: **end function**

3.3. Two-Sided-Randomization. Analogous to the one-sided Algorithm 3.2, we start with generating two TT random Gaussian tensors $\mathcal{L}$ and $\mathcal{R}$ with given target TT-ranks $\{\ell_n\}$ and $\{\rho_n\}$ (with $\rho_n > \ell_n$) and computing the sketches $\{\mathbf{W}_n^L\}$, $\{\mathbf{W}_n^R\}$ of $\mathcal{Y}$ from the left and right, respectively (*randomization phase*, see Figure B.2). Next, for each $n = 1, \ldots, N - 1$ we compute the SVD of a product of partial contractions $\mathbf{W}_n^L \mathbf{W}_n^R$, i.e., $\mathbf{W}_n^L \mathbf{W}_n^R = \mathbf{U}_n \mathbf{\Sigma}_n \mathbf{V}_n^{\top}$, and form left and right factor matrices

$$(3.1) \qquad \mathbf{L}_n = \mathbf{W}_n^R \mathbf{V}_n (\mathbf{\Sigma}_n^{\dagger})^{1/2} \quad \text{and} \quad \mathbf{R}_n = (\mathbf{\Sigma}_n^{\dagger})^{1/2} \mathbf{U}_n^{\top} \mathbf{W}_n^L.$$

In order to highlight the significance of matrices $\mathbf{L}_n$ and $\mathbf{R}_n$, we consider the following unfolding of the TT-tensor $\mathcal{Y}$

$$\mathbf{Y}_{(1:n)} = \mathcal{V}(\mathcal{T}_{\mathcal{Y},1:n})\mathcal{H}(\mathcal{T}_{\mathcal{Y},n+1:N}),$$

11

with factors $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1:n}) \in \mathbb{R}^{(I_1\cdots I_n)\times R_n}$ and $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1:N}) \in \mathbb{R}^{R_n\times(I_{n+1}\cdots I_N)}$. Similarly, we define matrices

$$\boldsymbol{\Psi}_n := \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{L}},1:n})^\top \in \mathbb{R}^{\ell_n\times(I_1\cdots I_n)},$$
$$\boldsymbol{\Omega}_n := \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{R}},n+1:N})^\top \in \mathbb{R}^{(I_{n+1}\cdots I_N)\times\rho_n}$$

as the partial unfoldings of random Gaussian TT-tensors $\boldsymbol{\mathcal{L}}$ and $\boldsymbol{\mathcal{R}}$, respectively. Then multiplying matrix $\mathbf{Y}_{(1:n)}$ on the left by $\boldsymbol{\Psi}_n$ and on the right by $\boldsymbol{\Omega}_n$ yields

$$\boldsymbol{\Psi}_n \mathbf{Y}_{(1:n)}\boldsymbol{\Omega}_n = (\boldsymbol{\Psi}_n\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1:n}))(\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1:N})\boldsymbol{\Omega}_n) = \mathbf{W}_n^L \mathbf{W}_n^R.$$

Following identity (2.2) illustrating the main idea of the Generalized Nyström method for matrices discussed in Subsection 2.2, we have

$$\begin{aligned}
\mathbf{Y}_{(1:n)} &\approx \big(\mathbf{Y}_{(1:n)}\boldsymbol{\Omega}_n\big)\big(\boldsymbol{\Psi}_n\mathbf{Y}_{(1:n)}\boldsymbol{\Omega}_n\big)^\dagger\big(\boldsymbol{\Psi}_n\mathbf{Y}_{(1:n)}\big)\\
&= \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1:n})\mathbf{W}_n^R\big(\mathbf{W}_n^L\mathbf{W}_n^R\big)^\dagger\mathbf{W}_n^L\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1:N})\\
&= \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1:n})\mathbf{L}_n\mathbf{R}_n\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1:N}),
\end{aligned}$$

see Figure B.3.

Having all left and right factors at hand, for each core of the tensor $\boldsymbol{\mathcal{Y}}$ (treating the first and last core separately), we distribute them according to the formula

$$\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}) = \mathbf{R}_{n-1}\mathcal{H}(\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n})\mathbf{L}_n)$$

forming the cores of the resulting tensor $\boldsymbol{\mathcal{X}}$, see Figure B.4.

In contrast to Algorithm 3.2, the Two-Sided-Randomization approach does not produce an orthogonal tensor. However, with a little restructuring, it can be adapted to produce an orthogonal tensor (we do not discuss that here). This variation may be useful for the case when the target TT-ranks are not known in advance, and producing an orthogonal tensor can be used in conjunction with Algorithm 2.2 to further compress the tensor.

3.4. Rounding of TT-sums. One of the most common arithmetic operations that depends on TT-rounding is TT-summation, i.e., we want to compress a tensor $\boldsymbol{\mathcal{Y}}$ that is available as the sum of s TT-tensors: $\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{Y}}^{(1)} + \cdots + \boldsymbol{\mathcal{Y}}^{(s)}$. To reuse existing algorithms, there are two options available to us. For example, we can form the TT-tensor $\boldsymbol{\mathcal{Y}}$ explicitly and then apply one of the compression algorithms proposed previously. As we will argue in Subsection 3.5, the computational costs of this approach has cubic scaling with respect to the number of summands s using the TT-rounding approach, and a quadratic scaling with respect to s using the randomized approaches (Randomize-then-Orthogonalize and Two-Sided-Randomization). This is computationally infeasible as s becomes large. Alternatively, we can form the partial sum $\boldsymbol{\mathcal{Y}}^{(1)} + \boldsymbol{\mathcal{Y}}^{(2)}$, compress this partial sum, add the resulting truncated term to the summand $\boldsymbol{\mathcal{Y}}^{(3)}$ and proceed in the same way with remaining terms. Variations of this approach can be performed using ideas from the summation methods described in [27, Chapter 4.1]. These approaches scale linearly with s, but can result in large errors in the compressed representation due to the repeated truncation.

In this subsection, we show how to combine the addition and randomized rounding operations to reduce further the computational costs, which is particularly effective when the number of summands s is large. The basic idea is to exploit the nonzero structure of the TT-cores of the sum of TT-tensors to avoid computing with zeros.

Algorithm 3.4 TT-Rounding of a Sum: Randomize-then-Orthogonalize

---

Require: Tensors $\{\boldsymbol{\mathcal{Y}}^{(j)}\}_{1\leq j\leq s}$ in TT format with ranks $\{R_n^{(j)}\}_{1\leq j\leq s}$, target TT-ranks $\{\ell_n\}$

Ensure: A tensor $\boldsymbol{\mathcal{X}} \approx \sum_{j=1}^{s} \boldsymbol{\mathcal{Y}}^{(j)}$ in TT format with ranks $\{\ell_n\}$

1: **function** $\boldsymbol{\mathcal{X}} = \text{TT-ROUNDING-SUM-RANDORTH}(\{\boldsymbol{\mathcal{Y}}^{(j)}\}_{1\leq j\leq s}, \{\ell_n\})$
2:     Select random Gaussian TT-tensor $\boldsymbol{\mathcal{R}}$ with ranks $\{\ell_n\}$
3:     **for** $j = 1$ to s **do**
4:         $\left\{\mathbf{W}_n^{(j)}\right\} = \text{PARTIALCONTRACTIONSRL}(\boldsymbol{\mathcal{Y}}^{(j)}, \boldsymbol{\mathcal{R}})$ ▷ precompute sketches from the right
5:     **end for**
6:     $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1} = \begin{bmatrix} \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}}^{(1)},1} & \cdots & \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}}^{(s)},1} \end{bmatrix}$
7:     **for** $n = 1$ to $N - 1$ **do**
8:         $\mathbf{Z}_n = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})$ ▷ $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}$ is $\ell_{n-1} \times I_n \times \sum_{j=1}^{s} R_n^{(j)}$
9:         $\mathbf{Y}_n = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}) \begin{bmatrix} \mathbf{W}_n^{(1)} \\ \vdots \\ \mathbf{W}_n^{(s)} \end{bmatrix}$ ▷ complete random sketch
10:         $[\mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n}), \sim] = \text{QR}(\mathbf{Y}_n)$ ▷ thin QR factorization
11:         $\begin{bmatrix} \mathbf{M}_n^{(1)} & \cdots & \mathbf{M}_n^{(s)} \end{bmatrix} = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n})^{\top}\mathbf{Z}_n$
12:         **if** $n < N - 1$ **then**   ▷ exploit structure in next internal core
13:             $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1}) = \begin{bmatrix} \mathbf{M}_n^{(1)}\mathcal{H}\left(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}}^{(1)},n+1}\right) & \cdots & \mathbf{M}_n^{(s)}\mathcal{H}\left(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}}^{(s)},n+1}\right) \end{bmatrix}$
14:         **else**
15:             $\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},N} = \sum_{j=1}^{s} \mathbf{M}_{n-1}^{(j)} \boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}}^{(j)},N}$
16:         **end if**
17:     **end for**
18: **end function**

Applying the orthogonalization phase, as required to perform the deterministic truncation phase, requires assembling the TT representation of the sum. Furthermore, the orthogonalization and multiplication by the triangular factor destroys the structure in the middle cores. By using randomization, we can avoid this explicit TT assembly of the sum and avoid unnecessary computations on zeros. Algorithm 3.4 provides the pseudocode for rounding the sum of s input TT-tensors. To simplify the notation, we derive the efficient computations considering the case $s = 2$, as the generalization will be clear. Let $\boldsymbol{\mathcal{Y}}$ and $\boldsymbol{\mathcal{Z}}$ be two TT-tensors and consider the TT-tensor $\boldsymbol{\mathcal{X}} = \boldsymbol{\mathcal{Y}} + \boldsymbol{\mathcal{Z}}$. Let $\boldsymbol{\mathcal{R}}$ be a given random Gaussian TT-tensor and let $\{\mathbf{W}_n^{\boldsymbol{\mathcal{Y}}}\}$ and $\{\mathbf{W}_n^{\boldsymbol{\mathcal{Z}}}\}$, for $n = 1, \ldots, N - 1$ be the right-to-left partial contractions of $\boldsymbol{\mathcal{Y}}$ and $\boldsymbol{\mathcal{Z}}$ with $\boldsymbol{\mathcal{R}}$. We have

$$\mathbf{X}_{(1:n)} = \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},1:n})\mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n+1:N}) = \begin{bmatrix} \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},1:n}) & \mathcal{V}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},1:n}) \end{bmatrix} \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Y}},n+1:N}) \\ \mathcal{H}(\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{Z}},n+1:N}) \end{bmatrix},$$

and by using (2.1),

$$\mathcal{V}(\boldsymbol{\mathcal{T}}_{\mathcal{X},1:n}) = \begin{bmatrix} (\mathbf{I}_{I_n} \otimes \mathcal{V}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},1:n-1})) & (\mathbf{I}_{I_n} \otimes \mathcal{V}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},1:n-1})) \end{bmatrix} \begin{bmatrix} \mathcal{V}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n}) & \\ & \mathcal{V}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n}) \end{bmatrix},$$

$$\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{X},n+1:N}) = \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+1}) & \\ & \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+1}) \end{bmatrix} \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \\ \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \end{bmatrix}.$$

The matrix $\mathbf{W}_n^{\mathcal{X}}$ can be expressed as

$$\mathbf{W}_n^{\mathcal{X}} = \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{X},n+1:N})\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{R},n+1:N})^{\top} = \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+1:N}) \\ \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+1:N}) \end{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{R},n+1:N})^{\top} = \begin{bmatrix} \mathbf{W}_n^{\mathcal{Y}} \\ \mathbf{W}_n^{\mathcal{Z}} \end{bmatrix}.$$

This justifies the procedure in Algorithm 3.4 of computing the partial contractions separately for each summand (Line 4) and concatenating them (Line 9).

After the QR factorization of the projected matrix produces the truncated core, we compute the contraction between the new and old cores (Line 11) and store the result in a matrix $\mathbf{M}_n = \begin{bmatrix} \mathbf{M}_n^{\mathcal{Y}} & \mathbf{M}_n^{\mathcal{Z}} \end{bmatrix}$ of size $\ell_n \times (R_n^{\mathcal{Y}} + R_n^{\mathcal{Z}})$. This matrix is now multiplied from the right by $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{X},n+1:N})$ to compute the updated right factor of $\mathbf{X}_{(1:n)}$. This multiplication can be absorbed by the $(n+1)$th core as follows:

$$\mathbf{M}_n\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{X},n+1:N}) = \begin{bmatrix} \mathbf{M}_n^{\mathcal{Y}} & \mathbf{M}_n^{\mathcal{Z}} \end{bmatrix} \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+1}) & \\ & \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+1}) \end{bmatrix} \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \\ \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \end{bmatrix},$$

$$= \begin{bmatrix} \mathbf{M}_n^{\mathcal{Y}}\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+1}) & \mathbf{M}_n^{\mathcal{Z}}\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+1}) \end{bmatrix} \begin{bmatrix} \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Y},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \\ \mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{Z},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \end{bmatrix}.$$

Hence we update $\mathcal{H}(\boldsymbol{\mathcal{T}}_{\mathcal{X},n+1})$ as in Line 13. For $n = N - 1$ we have

$$\boldsymbol{\mathcal{T}}_{\mathcal{X},N} = \begin{bmatrix} \mathbf{M}_{N-1}^{\mathcal{Y}} & \mathbf{M}_{N-1}^{\mathcal{Z}} \end{bmatrix} \begin{bmatrix} \boldsymbol{\mathcal{T}}_{\mathcal{Y},N} \\ \boldsymbol{\mathcal{T}}_{\mathcal{Z},N} \end{bmatrix} = \mathbf{M}_{N-1}^{\mathcal{Y}}\boldsymbol{\mathcal{T}}_{\mathcal{Y},N} + \mathbf{M}_{N-1}^{\mathcal{Z}}\boldsymbol{\mathcal{T}}_{\mathcal{Z},N}.$$

Generalizing this expression to s terms yields Line 15.

3.5. Computational costs. To analyze the computational cost, we make the following assumptions that will simplify the analysis. Let $\boldsymbol{\mathcal{Y}} \in \mathbb{R}^{I \times \cdots \times I}$ be a tensor of order N with ranks $(1, R, \ldots, R, 1)$ in TT format. We want to compress $\boldsymbol{\mathcal{Y}}$ to obtain a TT-tensor $\boldsymbol{\mathcal{X}}$ with ranks $(1, \ell, \ldots, \ell, 1)$. Here and in Appendix A, we assume that $\ell = \Theta(R)$.

3.5.1. Randomized compression algorithms. We now analyze the computational cost of Algorithm 3.1, Algorithm 3.2, and Algorithm 3.3.

Orthogonalize-then-Randomize (Algorithm 3.1). We denote the overall computational cost of Algorithm 3.1 by C_{OtR}. Line 2 of Algorithm 3.1 invokes Algorithm 2.1, and its cost is C_{Orth}, see Appendix A. We now analyze the cost of the loop in Lines 3 to 9. In Line 5, we multiply the matrix $\mathcal{V}(\boldsymbol{\mathcal{T}}_{\mathcal{X},n})$ of size $I\ell \times R$ with the random matrix of size $R \times \ell$; this costs $2IR\ell^2$ flops. In Line 6, we compute the thin QR factorization of the matrix $\mathbf{Y}_n$ of size $I\ell \times \ell$, which costs $4I\ell^3 + \mathcal{O}(\ell^3)$ flops. Finally, in Lines 7 and 8, we have two different matrix-matrix multiplications which cost $2IR\ell^2 + 2IR^2\ell$ flops, respectively. The total computational cost of Algorithm 3.1 is

$$\begin{aligned} C_{\text{OtR}} = \ & C_{\text{Orth}} + (N-2)(2IR^2\ell + 4IR\ell^2 + 4I\ell^3) + \mathcal{O}(IR^2 + NR^3) \\ = \ & I(N-2) \cdot (5R^3 + 2R^2\ell + 4R\ell^2 + 4\ell^3) + \mathcal{O}(IR^2 + NR^3) \quad \text{flops}. \end{aligned}$$

Randomize-then-Orthogonalize (Algorithm 3.2). Denoting the total cost of Algorithm 3.2 with C_{RtO}, we analyze the main components that contribute to the total computational cost. Line 3 invokes Algorithm 2.3 with the corresponding cost denoted by C_{Contr}, see Appendix A. Lines 7 and 9 contribute by a factor of $2IR\ell^2$ that is the cost of performing the multiplication $\mathcal{V}(\mathfrak{T}_{\mathfrak{X},n})\mathbf{W}_n$ of sizes $\ell I \times R$ with $R \times \ell$ and $\mathcal{V}(\mathfrak{T}_{\mathfrak{X},n})^T\mathbf{Z}_n$ of sizes $\ell \times I\ell$ and $I\ell \times R$ that prepare the matrices $\mathbf{Y}_n$ and $\mathbf{M}_n$, respectively, for the next steps. The term $4I\ell^3 + \mathcal{O}(\ell^3)$ represents the cost of the thin QR factorization, in Line 8, of the matrix $\mathbf{Y}_n$ of size $I\ell \times \ell$. Line 10 that involves multiplication of matrices of size $\ell \times R$ and $R \times IR$ which costs $2IR^2\ell$ flops. The total cost of Algorithm 3.2 is

$$
\begin{aligned}
C_{\text{RtO}} &= C_{\text{Contr}} + (N-2)(2IR^2\ell + 4IR\ell^2 + 4I\ell^3) + \mathcal{O}(IR^2 + NR^3) \\
&= I(N-2) \cdot (4R^2\ell + 6R\ell^2 + 4\ell^3) + \mathcal{O}(IR^2 + NR^3) \quad \text{flops.}
\end{aligned}
$$

Two-Sided-Randomization (Algorithm 3.3). Here we analyze the total cost of Algorithm 3.3 and denote it by C_{2SR}. The cost of Lines 4 and 5 is $2C_{\text{Contr}}$. Here, we recall that the cost of the partial contractions from the left and from the right is the same since we assume that ranks are the same, i.e., $\rho_n = \ell_n = \ell$. However, in practice, we choose $\rho_n > \ell_n$ to be different for numerical stability. There are two other contributions to the cost that come from two different matrix multiplication: first, of sizes $\ell \times R$ and $R \times IR$, and second, of sizes $\ell \times R$ and $R \times I\ell$. The cost of the for loop starting with Line 12 is $\mathcal{O}(NR^2)$. Hence, the total computational cost of Algorithm 3.3 is

$$
\begin{aligned}
C_{\text{2SR}} &= 2 \cdot C_{\text{Contr}} + (N-2)(2IR^2\ell + 2IR\ell^2) + \mathcal{O}(NR^2) \\
&= I(N-2) \cdot (6R^2\ell + 6R\ell^2) + \mathcal{O}(IR\ell + NR^2) \quad \text{flops.}
\end{aligned}
$$

Nakatsukasa [41] suggests oversampling from the left, i.e., taking $\rho_n = \lceil 1.5\ell_n \rceil$. We follow this suggestion in our numerical experiments. With this assumption, the cost is slightly higher, i.e., $I(N-2) \cdot (7R^2\ell + 8.5R\ell^2) + \mathcal{O}(IR\ell + NR^2)$ flops, due to the increased cost of the contraction (Algorithm 2.3) with a larger random tensor.

TABLE 3.1
Summary of the computational costs (discarding lower order terms) of the randomized algorithms proposed in this paper. For completeness, we also include the computational costs of the deterministic algorithms in Subsection 2.3. Orth and Contr refer to Algorithm 2.1 and Algorithm 2.3, respectively.

Algorithms	Computational cost (flops)	Simplified Cost (flops)
Orth	$(N-2)I(5R^3)$	—
Contr	$(N-2)I(2R^2\ell + 2R\ell^2)$	—
TT-Rounding	$(N-2)I(5R^3 + 6R^2\ell + 2R\ell^2)$	$(N-2)IR^3(5 + 6\beta + 2\beta^2)$
Orth-then-Rand	$(N-2)I(5R^3 + 2R^2\ell + 4R\ell^2 + 4\ell^3)$	$(N-2)IR^3(5 + 2\beta + 4\beta^2 + 4\beta^3)$
Rand-then-Orth	$(N-2)I(2R^2\ell + 4R\ell^2 + 4\ell^3)$	$(N-2)IR^3(4\beta + 6\beta^2 + 4\beta^3)$
Two-Sided-Rand	$(N-2)I(6R^2\ell + 6R\ell^2)$	$(N-2)IR^3(6\beta + 6\beta^2)$

Comparison of different algorithms. To enable the comparison of the different algorithms, we set a target rank as $\ell = \beta R$, where $\beta \in (0, 1]$ is the ratio between the target rank ℓ and the current rank R. This allows us to compare the different algorithms more clearly. A summary of the dominant costs of the algorithms is provided in Table 3.1. For simplicity, we also provide a simplified representation of the computational costs with $\beta = \ell/R$. In Figure 3.2, we plot the speedup of the randomized algorithms compared to the TT-Rounding algorithm; we used the simplified

FIG. 3.2. *Illustration of the speedups obtained by the randomized algorithms compared with the TT-rounding. Here $\beta = \ell/R$ is the ratio between the target rank and the original rank of the tensor. The speedup computations are based on the simplified cost analysis in* Table *3.1.*

representation of the costs while generating the figure. It can be easily seen that all the proposed methods are faster than the TT-Rounding algorithm. However, the speedup using the Orthogonalize-then-Randomize algorithm is very incremental. In contrast, the other two algorithms, Randomize-then-Orthogonalize and Two-Sided-Randomization, have very similar costs (Randomize-then-Orthogonalize is slightly more efficient for smaller β) and have much higher speedups especially if $\beta \ll 1$. If $\beta \approx 1$, both algorithms are very close to the TT-Rounding algorithm. Therefore, the proposed methods are most efficient if $\beta \ll 1$, i.e., the target rank ℓ is much smaller compared to the original rank R.

3.5.2. Rounding of TT-sums. To explain the benefits of the algorithm for rounding TT-sums in Subsection 3.4, consider the summation of s tensors of order N (size I in each dimension) each with TT-ranks $(1, R, \ldots, R, 1)$. Suppose we form the TT-tensor $\boldsymbol{\mathcal{Y}}$, which represents the summation $\boldsymbol{\mathcal{Y}} = \sum_{j=1}^{s} \boldsymbol{\mathcal{Y}}^{(j)}$, explicitly. The intermediate cores have size $sR \times I \times sR$; the first core is of size $I \times sR$ and the last core is of size $sR \times I$. Suppose the target compression rank in each mode is ℓ. To leading order, the cost of executing TT-Rounding and Orthogonalize-then-Randomize is $\mathcal{O}(Ns^3 R^3 I)$. In contrast, the cost of using Randomize-then-Orthogonalize and the Two-Sided-Randomization approach are both $\mathcal{O}(N\ell s^2 R^2 I)$. This can be beneficial if the number of summands s is large, or the rank R is large. This simple cost analysis does not take into account any structure present in the summation.

Carefully exploiting the structure, as in Subsection 3.4, can reduce this cost. In particular, by using TT-Rounding of a sum with s tensors of order N with Randomize-then-Orthogonalize summarized in Algorithm 3.4 the leading order computational cost is $\mathcal{O}(N\ell s R^2 I)$ flops. Notice that by exploiting the structure of the tensor and using randomization, the leading order of the cost is decreased to be linear in s in contrast to cubic in s when no structure was taken into account and the TT-sum tensor was formed explicitly. This decrease in the computational cost is obviously more pronounced when the number of summands s is large. In what follows, we present the analysis of the computational cost of computing the sum of s TT-tensors by randomization and by exploiting the underlying tensor structure.

TT-rounding of a sum: Randomize-then-Orthogonalize (Algorithm 3.4). We analyze the computational cost of Algorithm 3.4 which we denote by C_{RtOsum}. The

leading order term is sourced from two main contributions (1) Line 4 that executes Algorithm 2.3 s times resulting in a total computational cost of $sI(N-2)(2R\ell^2+2R^2\ell)$ flops and (2) Line 13 that represents s multiplications between matrices of size $IR \times R$ and $R \times \ell$ resulting in a total cost of $2sIR^2\ell$ flops. Next we analyze the source of the second leading order term present in the total cost. Line 9 contributes by a factor of two to the second leading term with a cost $2sIR\ell^2$ flops resulting from multiplying matrices of size $I\ell \times sR$ and $sR \times \ell$. Another factor of two comes from the multiplication of two matrices of size $\ell \times I\ell$ and $I\ell \times sR$ in Line 11, resulting in a total cost of $2sIR\ell^2$ flops. The last factor of two comes from the second leading order term of Algorithm 2.3. Computing the thin QR factorization of the matrix of size $I\ell \times \ell$ in Line 10 costs $4I\ell^3$ flops. Hence, the total cost of Algorithm 3.4 is

$$C_{\mathrm{RtOsum}} = s \cdot C_{\mathrm{Contr}} + (N-2)(2sIR^2\ell + 4sIR\ell^2 + 4I\ell^3) + \mathcal{O}(NIR^2)$$
$$= I(N-2) \cdot (4sR^2\ell + 6sR\ell^2 + 4\ell^2) + \mathcal{O}(NIR^2) \quad \text{flops.}$$

4. Numerical results. In this section, we illustrate numerically the performance of the newly developed algorithms using tensor data in TT format. We consider both synthetic as well as more realistic test examples. Additional numerical experiments are available in Appendix C. All the numerical experiments were performed on MATLAB R2021a running on a laptop computer with CPU Intel(R) Core(TM) i9-9980H and 64GB of RAM, using multithreading with 4 computational threads.

FIG. 4.1. *Comparison of error between a low-rank tensor and full rank perturbed tensor, and timings using the deterministic and randomized TT-rounding algorithms for ϵ perturbed tensor for different values of perturbation ϵ. Statistics were based on 5 independent runs.*

4.1. TT-tensor with a fixed target rank. In our first numerical experiment, we illustrate the performance of our rounding algorithms by rounding a random TT-tensor with a known low-rank representation. Throughout, we choose the ranks of the right side randomization in the Two-Sided-Randomization approach (Algorithm 3.3) to be $\rho = \lceil 1.5\ell \rceil$ as discussed in Subsection 3.5.1.

The random TT-tensor $\mathcal{X}$ is constructed by perturbing a random TT-tensor $\mathcal{X}_1$ with the random TT-tensor $\epsilon\mathcal{X}_2$ as follows: $\mathcal{X} = \mathcal{X}_1 + \epsilon\mathcal{X}_2$. TT-tensors $\mathcal{X}_1$, $\mathcal{X}_2 \in \mathbb{R}^{100\times\cdots\times100}$ are order $N = 10$ normalized random TT-tensors of ranks $(1, 50, ..., 50, 1)$ (normalized according to their dimension as described in Definition 3.1), and ϵ is a perturbation scalar taking the values $\epsilon \in \{10^{-2}, 10^{-6}, 10^{-10}\}$. The ranks of the perturbed tensor $\mathcal{X}$ are $(1, 50 + 50, ..., 50 + 50, 1)$, and the perturbation parameter ϵ determines how well tensor $\mathcal{X}$ is approximated by the lower rank tensor $\mathcal{X}_1$, i.e., if ϵ is small, then $\mathcal{X}_1$ is a good rank-$(1, 50, ..., 50, 1)$ approximation of $\mathcal{X}$.

We round the random TT-tensor $\mathcal{X}$ using Algorithms 2.2 and 3.1 to 3.3 to have ranks $(1, \ell, \ldots, \ell, 1)$, where we vary the parameter ℓ from 35 to 80 by an increment of 5. We present these results in Figure 4.1. The approximation error is the relative norm error between tensor $\mathcal{X}$ and the approximate rounded tensor $\widehat{\mathcal{X}}$, i.e., $\frac{\|\mathcal{X}-\widehat{\mathcal{X}}\|}{\|\mathcal{X}\|}$. We also present the time speedup (computed with the average of 5 independent runs) of the randomized algorithms compared to the deterministic algorithm.

For all values of perturbation ϵ, the error resulting from the deterministic algorithm (Algorithm 2.2) decreases slightly until the ranks of the rounded tensor are $\ell = 50$. When $\ell > 50$, the error appears to be very close to ϵ. The errors resulting from the randomized algorithms (Algorithms 3.1 to 3.3) are greater than the error resulting from the deterministic algorithm. Additionally, the randomized algorithms produce a more accurate approximation when the target rank is larger and are more accurate when ϵ is small. The Orthogonalize-then-Randomize and the Randomize-then-Orthogonalize method (Algorithms 3.1 and 3.2) produce a rounded tensor with similar levels of accuracy while the Two-Sided-Randomization approach is the least accurate. For smaller values of ϵ, there is less difference in the accuracy between the different algorithms. The Randomize-then-Orthogonalize algorithm is the fastest compared to the deterministic algorithm followed by the Two-Sided-Randomization and Orthogonalize-then-Randomize algorithms.

4.2. Solving a parametric PDE in the TT format. As a realistic test example, we consider the parameter dependent PDE referred in the literature as the *cookie problem* [37, 52], see Appendix D for details. Since it is known that the set of solutions of problem (D.1) admits a low-rank representation [23, 13], we consider a global linear system encapsulating all these linear systems, i.e.,

$$\left(\sum_{i=1}^{N} \mathbf{A}_{i,1} \otimes \ldots \otimes \mathbf{A}_{i,N}\right)\mathcal{X} = \mathcal{F},$$

where $\mathbf{A}_{1,1}$ is the discretization of the operator over the spatial domain with constant parameter values, for each $2 \leq i \leq N$, $\mathbf{A}_{i,1}$ is the discretization of the operator over the domain multiplied by the characteristic function corresponding to the corresponding subdomain and $\mathbf{A}_{i,i}$ is a diagonal matrix containing the parameter values for the corresponding parameter, and for each $2 \leq j \neq i \leq N$, $\mathbf{A}_{i,j}$ is the identity matrix. We use the TT-GMRES algorithm [16] to solve this global linear system of equations. The preconditioned TT-GMRES algorithm builds the basis vectors in TT format $\mathcal{V}_1, \mathcal{V}_2, \ldots$ using the inexact Arnoldi procedure; since at each step, the corresponding

TT-tensors are rounded, this results in an inexact Krylov subspace method. The main bottleneck is the computation of two linear combinations in each iteration. First, the following sum of N tensors with the same ranks as $\boldsymbol{\mathcal{V}}_k$ is formed when applying the operator to the k-th basis vector computed at the previous iteration, i.e.,

$$\boldsymbol{\mathcal{W}} = \sum_{i=1}^{N} \left(\mathbf{A}_{i,1} \otimes \ldots \otimes \mathbf{A}_{i,N} \right) \boldsymbol{\mathcal{Y}},$$

after application of the preconditioner, $\boldsymbol{\mathcal{Y}} = \left(\left(\sum_{i=1}^{N} \mathbf{A}_{i,1} \right)^{-1} \otimes I \otimes \cdots \otimes I \right) \boldsymbol{\mathcal{V}}_{\mathcal{K}}$. Second, a linear combination of $k+1$ tensors appears when using the Gram-Schmidt algorithm to orthogonalize $\boldsymbol{\mathcal{W}}$ with respect to the previous basis vectors,

$$\boldsymbol{\mathcal{Z}} = \left(\boldsymbol{\mathcal{W}} - \sum_{j=1}^{k} h_{jk} \boldsymbol{\mathcal{V}}_k \right), \qquad \boldsymbol{\mathcal{V}}_{k+1} = \frac{1}{h_{k+1,k}} \boldsymbol{\mathcal{Z}}, \qquad \begin{cases} h_{jk} = \langle \boldsymbol{\mathcal{V}}_k, \boldsymbol{\mathcal{W}} \rangle, & j = 1, \ldots, k, \\ h_{k+1,k} = \|\boldsymbol{\mathcal{Z}}\|_F. \end{cases}$$

In both cases, the addition of TT-tensors is followed by a TT-rounding operation in order to reduce the ranks and keep the computations tractable. Hence, these steps are amenable to acceleration by using the randomized Algorithm 3.4 as a rounding procedure in the aforementioned computations. Because the second linear combination involves $k+1$ summands, the randomized implementation reduces greatly the cost of later TT-GMRES iterations in particular, as the leading order of its computational cost is decreased from cubic to linear in k, as detailed in Subsection 3.5.

We perform numerical experiments with $N = 5$ using a piecewise linear finite element discretization with the mesh presented in Figure D.1, for various choices of the number of parameter samples, $I = I_2 = \ldots = I_N$, with values of ρ_i distributed linearly between 1 and 10. The relative tolerance of the TT-GMRES solver is set to 10^{-8}. We compare the naive, deterministic implementation of the preconditioned TT-GMRES algorithm with one using randomized summation and rounding steps. Results of the comparison are reported in Figure 4.2. On the right, we display all internal ranks of the TT representation of the Krylov vector computed at that iteration. We observe that the ranks of the basis vectors and number of iterations are the same using both implementations, and they depend only weakly on the dimension I of the parameter modes of the tensors. The speedup achieved by the randomized approach increases consistently with the number of parameter samples. Taking a closer look at the timing statistics as presented in Figure 4.3, we note that the computation and rounding of the linear combinations identified above indeed dominates the computational cost in both cases, and is a clear computational bottleneck for the deterministic implementation in particular, as the ranks of the sum tensor increase to as much as 2491 in these experiments. This explains the remarkable speedup obtained with the randomized approach.

5. Conclusions and outlook. In this paper, we present randomized algorithms for rounding a tensor, assuming that we have an initial representation in the TT format. This initial representation may not be optimal in terms of storage, and the randomized compression techniques can be used to obtain a more efficient representation. We derive three different algorithms: Orthogonalize-then-Randomize, Randomize-then-Orthogonalize, and Two-Sided-Randomization. We study the computational cost of these algorithms in some detail and show that it can be much smaller than the standard TT-rounding algorithm. Additionally, we consider the special case of rounding a TT-tensor that is represented as the sum of many TT-tensors.

FIG. 4.2. *Illustration of the speedups (left) and TT-ranks of the Krylov basis vectors (right) obtained by the deterministic and randomized summation and rounding algorithms within the TT-GMRES algorithm to solve problem* (D.1).


FIG. 4.3. *Illustration of the timings using the deterministic and randomized summation and rounding algorithms within the TT-GMRES algorithm to solve problem* (D.1).

While applying each of the randomized algorithms proposed here can reduce the computational cost over standard TT-rounding, we further exploit the structure of the problem to reduce the computational cost to be linear in the number of summands. We perform extensive numerical experiments and achieve over $20\times$ speedups on test problems compared to standard algorithms. There are many avenues for future investigations. First, it would be interesting to derive probabilistic bounds for the accuracy of the rounding approach. Second, we could consider extending our algorithms to the case where the TT-tensor is obtained as the Hadamard (or elementwise) product of two tensors. Finally, another extension worth considering is developing randomized rounding algorithms in the $\mathcal{H}-$Tucker format.

Appendix A. Computational cost of standard TT algorithms.

In this section, we analyze the computational cost of the standard approach to TT-rounding which is presented in Subsection 2.3.

To analyze the computational cost, we make the following assumptions that will simplify the analysis. Let $\boldsymbol{\mathcal{Y}} \in \mathbb{R}^{I \times \cdots \times I}$ be a tensor of order N with ranks $(1, R, \ldots, R, 1)$ in TT format. We want to compress $\boldsymbol{\mathcal{Y}}$ to obtain a TT-tensor $\boldsymbol{\mathcal{X}}$ with target ranks $(1, \ell, \ldots, \ell, 1)$ assuming that $\ell = \Theta(R)$. We explain in some detail the analysis of the computational costs of the algorithms described in Subsection 2.3, since it will help analyze the computational costs of the proposed algorithms.

Right-to-Left Orthogonalization (Algorithm 2.1). The two main contributions to the leading order flop cost are: (1) in Line 4, the thin QR factorization of the matrix $\mathcal{H}(\boldsymbol{\mathcal{T}_{\mathcal{X},n}})$ of size $IR \times R$, which cost $4IR^3 + \mathcal{O}(R^3)$ flops; (2) in Line 5, the multiplication of a matrix of size $IR \times R$ with another matrix of size $R \times R$ that forms the product $\mathcal{V}(\boldsymbol{\mathcal{T}_{\mathcal{Y},n-1}})\mathbf{R}^\top$. Since $\mathbf{R}$ is triangular, the computation in Line 5 requires only IR^3 flops to leading order. Therefore, the resulting total cost of Algorithm 2.1 is

$$(\text{A.1}) \qquad C_{\text{Orth}} = I(N-2) \cdot 5R^3 + \mathcal{O}(IR^2 + NR^3) \quad \text{flops.}$$

TT-Rounding (Algorithm 2.2). The main contributions to the total cost of TT-rounding come from two sources: in Line 2, we first right-orthogonalize $\boldsymbol{\mathcal{Y}}$ to obtain $\boldsymbol{\mathcal{X}}$ using Algorithm 2.1; this cost has just been analyzed and is given by C_{Orth}. Next we analyze the cost of the loop in Lines 5 to 10. In Line 6, we compute the QR factorization of matrix $\mathcal{V}(\boldsymbol{\mathcal{T}_{\mathcal{X},n}})$ of size $I\ell \times R$ using the Householder-QR algorithm, which costs $4I\ell R^2 + \mathcal{O}(R^3)$ flops. Next, in Line 7 performing the SVD requires $\mathcal{O}(R^3)$ flops. The next two steps in lines Lines 8 and 9 prepare the factors for the next iteration. The matrix-matrix multiplication in Line 8 costs $2IR\ell^2$. Notice that the size of $\boldsymbol{\mathcal{T}_{\mathcal{X},n}}$ has been reduced from the previous iteration, so that Line 9 requires $2IR^2\ell$ flops. Hence, the total cost of Algorithm 2.2 is

$$\begin{aligned} C_{\text{TTR}} &= C_{\text{Orth}} + (N-2) \cdot \left(6IR^2\ell + 2IR\ell^2\right) + \mathcal{O}(IR^2 + NR^3) \\ &= I(N-2) \cdot \left(5R^3 + 6R^2\ell + 2R\ell^2\right) + \mathcal{O}(IR^2 + NR^3) \quad \text{flops.} \end{aligned}$$

Right-to-Left Partial Contraction (Algorithm 2.3). Consider two TT-tensors $\boldsymbol{\mathcal{X}}$ and $\boldsymbol{\mathcal{Y}}$ with ranks $\{R_j^{\boldsymbol{\mathcal{X}}}\}$ and $\{R_j^{\boldsymbol{\mathcal{Y}}}\}$. For simplicity of exposition, we take $\ell = R_j^{\boldsymbol{\mathcal{X}}}$ and $R = R_j^{\boldsymbol{\mathcal{Y}}}$ for $j = 1, \ldots, N-1$; that is we take the same ranks in each core tensor. We denote the computational cost of Algorithm 2.3 by C_{Contr}. The first contribution in the computational cost (Lines 4 and 5) comes from two different matrix multiplications, i.e., matrix of size $I\ell \times \ell$ with the one of size $\ell \times R$ which costs $2I\ell^2R$ flops, and matrix of size $\ell \times IR$ with the one of size $IR \times R$ which costs $2IR^2\ell$ flops. The cost of Line 2 is of order $\mathcal{O}(IR\ell)$ flops. The overall cost of Algorithm 2.3 is, therefore,

$$(\text{A.2}) \qquad C_{\text{Contr}} = I(N-2) \cdot (2R\ell^2 + 2R^2\ell) + \mathcal{O}(IR\ell) \quad \text{flops.}$$

If $\ell = R$, then we are computing an inner product and the associated cost is $C_{\text{Contr}} = 4(N-2)IR^3 + \mathcal{O}(IR^2)$ flops.

Appendix B. Additional Tensor network diagrams.

In Figures 3.1 and B.1, we give tensor network diagrams that help illustrate the random projection steps in Algorithm 3.1 and Algorithm 3.2. In Figures B.2 to B.4, we provide the tensor network diagrams illustrating Algorithm 3.3.

FIG. B.1. *Random projection steps for Orthogonalize-then-Randomize* Algorithm *3.1.*

FIG. B.2. *Random projections for Two-Sided-Randomization (Generalized Nyström)* Algorithm *3.3.*


FIG. B.3. *Internal left and right factors in the Two-Sided-Randomization (Generalized Nyström)* Algorithm *3.3.*

FIG. B.4. *Left and right factors distributions within the Two-Sided-Randomization (Generalized Nyström)* Algorithm *3.2.*

Appendix C. Additional Numerical Experiments: Hilbert-type TT-tensor.

To supplement the numerical experiments in the main manuscript, we propose a Hilbert-type TT-tensor, i.e., a TT-tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_1 \times \cdots \times I_N}$ with the cores of the form

$$\boldsymbol{\mathcal{T}}_{\boldsymbol{\mathcal{X}},n} = \frac{1}{i_1 + i_2 + i_3 - 1}, \qquad 1 \le i_1 \le r_{n-1},\ 1 \le i_2 \le I_n,\ 1 \le i_3 \le r_n,$$

22

for $n = 1, \ldots, N$. Our definition of a Hilbert-type TT-tensor is inspired by Hilbert matrices, i.e., symmetric and positive definite matrices with rapidly decaying singular values. Since the singular values of each mode-unfolding of the cores of the Hilbert-type TT-tensor decay very quickly, it is a useful tensor for validating the accuracy of the randomized algorithms proposed in this paper.

In our test, we consider the order-10 Hilbert-type TT-tensor with dimensions $I_1 = 6, I_2 = 8, \ldots, I_{10} = 24$ and ranks $(1, 4, 5, 6, \ldots, 11, 12, 1)$ and compare numerically the deterministic approach (Algorithm 2.2), the Orthogonalize-then-Randomize algorithm and the Randomize-then-Orthogonalize algorithm (Algorithms 3.1 and 3.2, respectively). Results of our experiments are presented in Figure C.1, with all three algorithms exhibiting a very similar accuracy.

FIG. C.1. *Hilbert TT-tensor truncation. Relative Error between a low rank approximation obtained via deterministic or randomized TT-rounding, and a full rank Hilbert TT-tensor.*

Appendix D. Additional details regarding Subsection 4.2.

In this section, we present additional details regarding the *cookie problem* we examined in Subsection 4.2.

(D.1)
$$
\begin{aligned}
-\mathrm{div}(\sigma(x, y; \boldsymbol{\rho})\nabla(u(x, y; \boldsymbol{\rho}))) &= f(x, y) && \text{in } \Omega, \\
u(x, y; \boldsymbol{\rho}) &= 0 && \text{on } \partial\Omega,
\end{aligned}
$$

where Ω is $(-1, 1) \times (-1, 1)$, $\partial\Omega$ is the boundary of Ω and σ is defined as follows

$$
\sigma(x, y; \boldsymbol{\rho}) = \begin{cases} 1 + \rho_i & \text{if } (x, y) \in D_i, \\ 1 & \text{elsewhere,} \end{cases}
$$

where D_i for $i = 1, \ldots, N-1$ are disjoint disks distributed in Ω such that their centers are equidistant (see Figure D.1) and ρ_i is selected from a discrete set of I_{i+1} samples $J_i \subset \mathbb{R}$ for $i = 1, \ldots, N-1$.

Discretizing this problem in Ω yields the set of $\prod_{i=1}^{N-1} I_{i+1}$ parametrized linear systems of equations

$$\left(\mathbf{A}_{1,1} + \sum_{i=1}^{N-1} \rho_i \mathbf{A}_{i+1,1}\right) \boldsymbol{x} = \boldsymbol{f},$$

where $\mathbf{A}_{1,1} \in \mathbb{R}^{I_1 \times I_1}$ is the discretization of the operator $-\mathrm{div}(\sigma(x, y; \boldsymbol{\rho})\nabla(u(x, y; \boldsymbol{\rho})))$ in Ω and $\mathbf{A}_{i+1,1} \in \mathbb{R}^{I_1 \times I_1}$ are the discretizations of $-\mathrm{div}(\chi_{D_i}\nabla(\cdot))$ in Ω where χ_S denotes the indicator function of the set S, and $\boldsymbol{f}$ is the discretization of the function f.

REFERENCES

[1] S. Ahmadi-Asl, S. Abukhovich, M. G. Asante-Mensah, A. Cichocki, A. H. Phan, T. Tanaka, and I. Oseledets. Randomized algorithms for computation of Tucker decomposition and higher order SVD (HOSVD). *IEEE Access*, 9:28684–28706, 2021.

[2] S. Ahmadi-Asl, A. Cichocki, A. H. Phan, M. G. Asante-Mensah, M. M. Ghazani, T. Tanaka, and I. Oseledets. Randomized algorithms for fast computation of low rank tensor ring model. *Machine Learning: Science and Technology*, 2(1):011001, 2020.

[3] N. Alger, P. Chen, and O. Ghattas. Tensor train construction from tensor actions, with application to compression of large high order derivative tensors. *SIAM J. Sci. Comput.*, 42(5):A3516–A3539, 2020.

[4] J. Ballani and L. Grasedyck. A projection method to solve linear systems in tensor format. *Numer. Linear Algebra Appl.*, 20(1):27–43, 2013.

[5] K. Batselier, W. Yu, L. Daniel, and N. Wong. Computing low-rank approximations of large-scale matrices with the tensor network randomized SVD. *SIAM J. Matrix Anal. Appl.*, 39(3):1221–1244, 2018.

[6] M. H. Beck, A. Jäckle, G. A. Worth, and H.-D. Meyer. The multiconfiguration time-dependent Hartree (MCTDH) method: a highly efficient algorithm for propagating wavepackets. *Phys. Rep.*, 324(1):1–105, 2000.

[7] G. Beylkin, J. Garcke, and M. J. Mohlenkamp. Multivariate regression and machine learning with sums of separable functions. *SIAM J. Sci. Comput.*, 31(3):1840–1857, 2009.

[8] M. Che and Y. Wei. Randomized algorithms for the approximations of Tucker and the tensor train decompositions. *Adv. Comput. Math.*, 45(1):395–428, 2019.

[9] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, and D. P. Mandic. Tensor networks

Fig. D.1. *Geometry and mesh used for the cookie problem* (D.1) *with* $N = 5$.

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.

[10] A. Cichocki, A.-H. Phan, Q. Zhao, N. Lee, I. V. Oseledets, M. Sugiyama, and D. Mandic. Tensor Networks for Dimensionality Reduction and Large-Scale Optimizations. Part 2 Applications and Future perspectives. *arXiv preprint arXiv:1708.09165*, 2017.

[11] N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on Learning Theory*, pages 698–728. PMLR, 2016.

[12] H. A. Daas, G. Ballard, and P. Benner. Parallel Algorithms for Tensor Train Arithmetic. *arXiv preprint arXiv:2011.06532*, 2020.

[13] W. Dahmen, R. DeVore, L. Grasedyck, and E. Süli. Tensor-sparsity of solutions to high-dimensional elliptic partial differential equations. *FoCM*, 16(4):813–874, 2016.

[14] V. De Silva and L.-H. Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM J. Matrix Anal. Appl.*, 30(3):1084–1127, 2008.

[15] S. Dolgov and D. Savostyanov. Parallel cross interpolation for high-precision calculation of high-dimensional integrals. *Comput. Phys. Commun.*, 246:106869, 2020.

[16] S. V. Dolgov. TT-GMRES: solution to a linear system in the structured tensor format. *Russ. J. Numer. Anal. Math. Model.*, 28(2):149–172, 2013.

[17] S. V. Dolgov, B. N. Khoromskij, and I. V. Oseledets. Fast solution of parabolic problems in the tensor train/quantized tensor train format with initial application to the Fokker–Planck equation. *SIAM J. Sci. Comput.*, 34(6):A3016–A3038, 2012.

[18] S. V. Dolgov and D. V. Savostyanov. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM J. Sci. Comput.*, 36(5):A2248–A2271, 2014.

[19] D. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. Adams. Advances in Neural Information Processing Systems 28. *Cortes C., Lawrence ND, Lee DD, Sugiyama M., Garnett R., Eds*, pages 2224–2232, 2015.

[20] M. Fannes, B. Nachtergaele, and R. F. Werner. Finitely correlated states on quantum spin chains. *Commun. Math. Phys.*, 144(3):443–490, 1992.

[21] Y. Feng, K. Tang, L. He, P. Zhou, and Q. Liao. Tensor Train random projection. *arXiv preprint arXiv:2010.10797*, 2020.

[22] P. Gelß, S. Klus, S. Matera, and C. Schütte. Nearest-neighbor interaction systems in the tensor-train format. *J. Comput. Phys.*, 341:140–162, 2017.

[23] L. Grasedyck. Existence and computation of low Kronecker-rank approximations for large linear systems of tensor product structure. *Computing*, 72(3-4):247–265, 2004.

[24] L. Grasedyck, D. Kressner, and C. Tobler. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen*, 36(1):53–78, 2013.

[25] W. Hackbusch. *Tensor spaces and numerical tensor calculus*, volume 42. Springer, 2012.

[26] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.

[27] N. J. Higham. *Accuracy and stability of numerical algorithms*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, second edition, 2002.

[28] F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys*, 6(1-4):164–189, 1927.

[29] S. Holtz, T. Rohwedder, and R. Schneider. The alternating linear scheme for tensor optimization in the tensor train format. *SIAM J. Sci. Comput.*, 34(2):A683–A713, 2012.

[30] B. Huber, R. Schneider, and S. Wolf. A randomized tensor train singular value decomposition. In *Compressed Sensing and its Applications*, pages 261–290. Springer, 2017.

[31] B. N. Khoromskij. Tensors-structured numerical methods in scientific computing: Survey on recent advances. *Chemom. Intell. Lab. Syst.*, 110(1):1–19, 2012.

[32] S. Klus, P. Gelß, S. Peitz, and C. Schütte. Tensor-based dynamic mode decomposition. *Nonlinearity*, 31(7):3359, 2018.

[33] S. Klus, P. Koltai, and C. Schütte. On the numerical approximation of the Perron-Frobenius and Koopman operator. *arXiv preprint arXiv:1512.05997*, 2015.

[34] O. Koch and C. Lubich. Dynamical tensor approximation. *SIAM J. Matrix Anal. Appl.*, 31(5):2360–2375, 2010.

[35] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[36] D. Kressner, R. Kumar, F. Nobile, and C. Tobler. Low-rank tensor approximation for high-order correlation functions of gaussian random fields. *SIAM-ASA J. Uncertain.*, 3(1):393–416, 2021/09/21 2015.

[37] D. Kressner and C. Tobler. Low-rank tensor Krylov subspace methods for parametrized linear systems. *SIAM J. Matrix Anal. Appl.*, 32(4):1288–1316, 2021/04/09 2011.

[38] D. Loukrezis, U. Römer, T. Casper, S. Schöps, and H. De Gersem. High-dimensional uncertainty quantification for an electrothermal field problem using stochastic collocation on sparse grids and tensor train decompositions. *Int. J. Numer. Model.: Electron. Netw. Devices Fields*, 31(2):e2222, 2018.

[39] P.-G. Martinsson and S. Voronin. A randomized blocked algorithm for efficiently computing rank-revealing factorizations of matrices. *SIAM J. Sci. Comput.*, 38(5):S485–S507, 2016.

[40] H.-D. Meyer, F. Gatti, and G. A. Worth. Basic MCTDH theory. *Multidimensional Quantum Dynamics: MCTDH Theory and Applications*, pages 17–30, 2009.

[41] Y. Nakatsukasa. Fast and stable randomized low-rank matrix approximation. *arXiv preprint arXiv:2009.11392*, 2020.

[42] A. Obukhov, M. Rakhuba, A. Liniger, Z. Huang, S. Georgoulis, D. Dai, and L. Van Gool. Spectral tensor train parameterization of deep learning layers. In *Int. Conf. Artif. Intell. Stat., AISTATS*, pages 3547–3555. PMLR, 2021.

[43] I. Oseledets and E. Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra Appl.*, 432(1):70–88, 2010.

[44] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33(5):2295–2317, 2011.

[45] I. V. Oseledets and S. V. Dolgov. Solution of linear systems and matrix inversion in the TT-format. *SIAM J. Sci. Comput.*, 34(5):A2718–A2739, 2012.

[46] I. V. Oseledets and E. E. Tyrtyshnikov. Breaking the curse of dimensionality, or how to use SVD in many dimensions. *SIAM J. Sci. Comput.*, 31(5):3744–3759, 2009.

[47] B. Rakhshan and G. Rabusseau. Tensorized random projections. In *Int. Conf. Artif. Intell. Stat., AISTATS*, pages 3306–3316. PMLR, 2020.

[48] H. Rauhut, R. Schneider, and Ž. Stojanac. Tensor completion in hierarchical tensor representations. In *Compressed Sensing and its Applications*, pages 419–450. Springer, 2015.

[49] L. Richter, L. Sallandt, and N. Nüsken. Solving high-dimensional parabolic PDEs using the tensor train format. *arXiv preprint arXiv:2102.11830*, 2021.

[50] D. Savostyanov and I. Oseledets. Fast adaptive interpolation of multi-dimensional arrays in tensor train format. In *The 2011 International Workshop on Multidimensional (nD) Systems*, pages 1–8, 2011.

[51] U. Schollwöck. The density-matrix renormalization group. *Rev. Mod. Phys.*, 77(1):259, 2005.

[52] C. Tobler. *Low-rank Tensor Methods for Linear Systems and Eigenvalue Problems*. PhD thesis, ETH Zurich, 2012.

[53] W. Yu, Y. Gu, and Y. Li. Efficient randomized algorithms for the fixed-precision low-rank matrix approximation. *SIAM J. Matrix Anal. Appl.*, 39(3):1339–1359, 2018.

[54] Z. Zhang, X. Yang, I. V. Oseledets, G. E. Karniadakis, and L. Daniel. Enabling high-dimensional hierarchical uncertainty quantification by ANOVA and tensor-train decomposition. *EEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 34(1):63–76, 2014.