

1 **RANDOMIZED ALGORITHMS FOR ROUNDING IN THE**
2 **TENSOR-TRAIN FORMAT**

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

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

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

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

23 **1. Introduction.** An increasing number of applications in science and technol-
24 ogy involve the manipulation of multi-dimensional data, or tensors that are higher
25 order equivalents of vectors (first-order) and matrices (second-order). The number
26 of elements of a tensor as well as the storage consumption grow exponentially with
27 the number of the dimensions, a phenomenon known as the *curse of dimensionality*.
28 When problems of high dimensions are concerned, beating the curse of dimensionality
29 and finding a solution efficiently remains a challenge. Nevertheless, different tensor
30 formats and methods based on tensor products [32, 36, 25, 26, 45] have shown poten-
31 tial for mitigating the curse of dimensionality and tackling high-dimensional problems
32 that could not be addressed with conventional methods. Initially, the concept of ten-
33 sor decompositions was introduced in 1927 by expressing a tensor as the sum of a
34 finite number of rank-one tensors [29] — also known as the *canonical format*. The
35 canonical format’s memory requirements are not high, though it can suffer from nu-
36 merical stability issues [15, 30]. Tensors in Tucker form [7] are well known in quantum
37 chemistry [15, 30] since they yield robust algorithms due to the ability to form an em-
38 bedded manifold [35], but one of the disadvantages of the Tucker format is its storage
39 consumption that still depends exponentially on the number of dimensions.

40 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, Virginia Tech, Blacksburg, VA 24061-1026, USA. (cazeaux@vt.edu,
amiedlar@vt.edu)

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

¶Department of Mathematics, Tufts University, USA. (mirjeta.pasha@tufts.edu)

41 tensor product format that was initially proposed in quantum physics, also known
42 as *matrix product states* (MPS) [21], and was reinvented in numerical linear algebra
43 [46, 44]. It combines both the advantages of the canonical and Tucker formats, i.e., 1)
44 the storage consumption of a tensor depends linearly on the number of dimensions and
45 2) there exist robust algorithms for the computation of best approximations. Applica-
46 tions of the TT format arise from various applications such as high-dimensional PDEs
47 like the Fokker Planck equations [18, 49], quantum physics [51, 40], high-dimensional
48 data analysis [34, 33], machine learning [8, 20, 12, 42], and uncertainty quantification
49 [53, 39] to mention just a few. Typically, those applications require an approximate
50 solution of linear systems of equations, eigenvalue problems, or completion problems
51 [23, 4, 48]. The TT format is a low-rank representation that, for TT-tensors with
52 small rank, offers a tremendous reduction in the computational complexity and often
53 exposes the structure of the problem. The use of low-rank structures such as the
54 TT format [44] to represent high-dimensional objects allows the solution of linear
55 high-dimensional problems by generalizing standard numerical linear algebra tech-
56 niques to multi-index arrays of coefficients (tensors) and the multivariate functions
57 they approximate.

58 In this paper, we focus on the problem of rounding a tensor in TT format; that is,
59 assuming that we are given a TT-tensor, we want to find a compressed representation
60 that is nearly as accurate as the original representation. There are several techniques
61 for computing the initial TT-tensors which do not require forming the entire ten-
62 sor explicitly [16, 37, 43, 50]. One such technique that is popularly used is called
63 the TT-cross approximation. The standard TT approach to rounding, proposed by
64 Oseledets [44], has two phases [44, Algorithm 2]: orthogonalization followed by com-
65 pression (typically using the SVD). Here, by orthogonalization, we mean a sweep of
66 orthogonalization steps across every tensor core. Analysis shows that the orthogonal-
67 ization step dominates the computational cost of this approach. Motivated by this
68 observation, the goal of this work is to develop randomized algorithms for rounding
69 TT-tensors that avoid expensive orthogonalization. In the following, we present the
70 main contributions of this paper.

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

- 78 1. In Algorithm 3.1, Randomize-then-Orthogonalize, we propose to form ran-
79 domized sketches of each core by nested contractions with a TT-tensor with
80 random cores in a first step, before performing the orthogonalization sweep
81 on these much smaller matrices. Our analysis and experiments show that
82 this approach allowed for the best speedup compared to the deterministic
83 algorithm while retaining excellent accuracy.
- 84 2. In Algorithm 3.2, Two-Sided-Randomization, we completely eliminate the
85 need for separate orthogonalization and compression sweeps. Instead, we
86 work with a two-sided randomized approach which computes products with
87 two random tensors followed by a compression step (which involves orthogo-
88 nalization of much smaller matrices). Although this approach is slightly more
89 expensive in terms of flops count and less accurate than techniques mentioned
90 before, it eliminates the need of extensive orthogonalization and allows for

91 the truncation phase to be more independent and highly parallelizable.

- 92 3. We extend the Randomize-then-Orthogonalize approach for compressing a
93 TT-tensor that is presented as a sum of TT-tensors (Algorithm 3.3). This
94 special case is of importance in many applications such as solving parametric
95 linear systems in the TT format. The use of randomization enables significant
96 performance improvements by exploiting the structure of the sum tensor in
97 a way that a deterministic algorithm cannot.
- 98 4. Additionally, in Algorithm SM2.1, Orthogonalize-then-Randomize, we replace
99 the SVD step in the standard TT-rounding algorithm with a randomized
100 SVD assuming that the truncated ranks are known *a priori*. This method,
101 while not competitive to the other proposed approaches, serves as a point of
102 comparison.

103 We provide an analysis of the computational cost of the proposed algorithms in Sub-
104 section 3.4 and show that they are computationally more efficient than existing algo-
105 rithms. We justify our analysis through numerical experiments in Section 4 on both
106 synthetic data and tensors generated while solving parametric partial differential equa-
107 tions (PDEs). Some conclusions and future outlook are presented in Section 5. The
108 MATLAB code for the implementation and numerical experiments is publicly available
109 at <https://github.com/SAMSI-RandTensors/randomizedTT>.

110 *Related work.* There have been several recent developments in obtaining low-rank
111 compression of tensors. We limit our literature review to the publications dealing
112 with TT-tensors described in [44], which is closest to our work, and refer the reader
113 to review papers for other developments in tensor decompositions [1, 10, 11, 25].
114 Oseledets [44] proposes a method for rounding TT-tensors. A parallel version of this
115 method is introduced and developed in [13]. Our newly proposed approaches are
116 more computationally efficient compared to existing deterministic algorithms. Other
117 works [31, 9, 2] discuss randomized algorithms for compressing tensors in the TT
118 format. These approaches differ from ours in that they require access to the entries of
119 the tensor, i.e., they do not assume that the tensor is already in TT format. A recent
120 paper [3] also uses randomization to produce a TT approximation of a full tensor but
121 relies on tensor actions (i.e., applications of the tensor on $N - 1$ vectors, where N
122 is the order of the tensor). Other methods for constructing a low-rank compression
123 in the TT format involve alternating least squares [19]. The use of tensor random
124 projections in which the random tensors are taken to be in TT format have also been
125 considered in [6, 47, 22]. While these papers use randomization in the context of TT-
126 tensors, none of them directly address the problem of rounding which is the central
127 focus of our paper.

128 **2. Background.** Here, we review the notation and necessary operations involv-
129 ing tensors in a modest amount of detail. For a more comprehensive exposition we
130 refer the reader to [44, 36, 13].

131 **2.1. Notation.** We denote tensors by boldface script letters (e.g., \mathcal{X}) and ma-
132 trices by boldface Roman letters (e.g., \mathbf{A}). We follow MATLAB-like convention and
133 denote the entries of a 3-way tensor \mathcal{X} as $\mathcal{X}(i, j, k)$. A colon denotes the entire range
134 of indices in that dimension. We denote the column fibers as $\mathcal{X}(:, j, k)$, row fibers as
135 $\mathcal{X}(j, :, k)$ and tube fibers as $\mathcal{X}(j, k, :)$. The *mode- n unfolding* (or *matricization*) of the
136 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
137 mode- n unfolding are composed of the appropriate mode- n fibers, e.g., the columns
138 of mode-1 unfolding are column fibers and the columns of mode-3 unfolding are tube
139 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

140 mode- n unfolding $\mathbf{Y}_{(n)} = \mathbf{A}\mathbf{X}_{(n)}$. The norm of a tensor is equivalent to the Frobenius
 141 norm of any of its unfoldings: $\|\mathbf{X}\| = \|\mathbf{X}_{(n)}\|_F$.

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

$$145 \quad \mathbf{X}(i_1, \dots, i_N) = \mathcal{J}_{\mathbf{X},1}(i_1, :) \cdot \dots \cdot \mathcal{J}_{\mathbf{X},n}(:, i_n, :) \cdot \dots \cdot \mathcal{J}_{\mathbf{X},N}(:, i_N),$$

146 where $1 \leq i_n \leq I_n$. Note that because $R_0 = R_N = 1$, the first and last TT-cores are
 147 (order-2) matrices so $\mathcal{J}_{\mathbf{X},1}(i_1, :) \in \mathbb{R}^{R_1}$ and $\mathcal{J}_{\mathbf{X},N}(:, i_N) \in \mathbb{R}^{R_{N-1}}$. The $R_{n-1} \times R_n$
 148 matrix $\mathcal{J}_{\mathbf{X},n}(:, i_n, :)$ is referred to as the i_n th slice of the n th TT-core of \mathbf{X} . It is worth
 149 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 $\mathcal{J}_{\mathbf{X},n}$.

150

151 In order to express the arithmetic operations on TT-cores using linear algebra,
 152 we will often use two specific matrix unfoldings of the order-3 tensors. The *horizontal*
 153 *unfolding* of a TT-core $\mathcal{J}_{\mathbf{X},n}$ corresponds to the concatenation of the slices $\mathcal{J}_{\mathbf{X},n}(:, i_n, :)$
 154 for $i_n = 1, \dots, I_n$ horizontally. We denote the corresponding operator by \mathcal{H} , so
 155 that $\mathcal{H}(\mathcal{J}_{\mathbf{X},n})$ is an $R_{n-1} \times I_n R_n$ matrix. The *vertical unfolding* of a TT-core $\mathcal{J}_{\mathbf{X},n}$
 156 corresponds to the concatenation of the slices $\mathcal{J}_{\mathbf{X},n}(:, i_n, :)$ for $i_n = 1, \dots, I_n$ vertically.
 157 We denote the corresponding operator by \mathcal{V} , so that $\mathcal{V}(\mathcal{J}_{\mathbf{X},n})$ is an $R_{n-1} I_n \times R_n$
 158 matrix, see Figure 2.1. Moreover, we will often make use of a *tensor network diagram*,
 159 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.

160

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

$$163 \quad \mathbf{X}_{(1:n)} = \mathcal{V}(\mathcal{J}_{\mathbf{X},1:n})\mathcal{H}(\mathcal{J}_{\mathbf{X},n+1:N}),$$

164 where in an extension of their earlier definitions, $\mathcal{V}(\mathcal{J}_{\mathbf{X},1:n}) \in \mathbb{R}^{(I_1 I_2 \dots I_n) \times R_n}$ represents
 165 the mode- $(n+1)$ unfolding of the product of the first n TT-cores and $\mathcal{H}(\mathcal{J}_{\mathbf{X},n+1:N}) \in$

166 $\mathbb{R}^{R_n \times (I_{n+1} \cdots I_N)}$ represents the mode-1 unfolding of the product of the final $N - n$
 167 TT-cores. Likewise, we can write the same unfolding as a product of four matrices,
 168 see [13, Eq. (2.3)], i.e.,

$$169 \quad (2.1) \quad \mathbf{X}_{(1:n)} = (\mathbf{I}_{I_n} \otimes \mathcal{V}(\mathcal{J}_{\mathbf{x},1:n-1})) \mathcal{V}(\mathcal{J}_{\mathbf{x},n}) \mathcal{H}(\mathcal{J}_{\mathbf{x},n+1}) (\mathcal{H}(\mathcal{J}_{\mathbf{x},n+2:N}) \otimes \mathbf{I}_{I_{n+1}}).$$

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

$$172 \quad \mathcal{J}_{\mathbf{x},n}(:, i_n, :) = \begin{bmatrix} \mathcal{J}_{\mathcal{Y},n}(:, i_n, :) \\ \mathcal{J}_{\mathcal{Z},n}(:, i_n, :) \end{bmatrix} \quad 2 \leq n \leq N - 1,$$

173 and for the first and the last core, we have

$$174 \quad \mathcal{J}_{\mathbf{x},1}(i_1, :) = [\mathcal{J}_{\mathcal{Y},1}(i_1, :) \quad \mathcal{J}_{\mathcal{Z},1}(i_1, :)] \quad \text{and} \quad \mathcal{J}_{\mathbf{x},N}(:, i_N) = \begin{bmatrix} \mathcal{J}_{\mathcal{Y},N}(:, i_N) \\ \mathcal{J}_{\mathcal{Z},N}(:, i_N) \end{bmatrix}.$$

175 Let $\mathbf{X} \in \mathbb{R}^{m \times n}$ with $m \geq n$. We denote the *thin QR* factorization of \mathbf{X} as
 176 $\mathbf{X} = \mathbf{Q}\mathbf{R}$, where $\mathbf{Q} \in \mathbb{R}^{m \times n}$ has orthonormal columns and $\mathbf{R} \in \mathbb{R}^{n \times n}$ is upper
 177 triangular; we also write $[\mathbf{Q}, \mathbf{R}] = \text{QR}(\mathbf{X})$ for use in algorithms. The SVD of \mathbf{X}
 178 is denoted by $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where the matrix $\mathbf{U} \in \mathbb{R}^{m \times n}$ has orthonormal columns containing
 179 the left singular vectors, $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with the singular values on
 180 the diagonal and $\mathbf{V} \in \mathbb{R}^{n \times n}$ is an orthogonal matrix, whose columns contain the
 181 right singular vectors. Assuming that the Householder QR algorithm is used and \mathbf{Q} is
 182 formed explicitly, the computational cost of the QR factorization is $4mn^2 - \frac{4n^3}{3} + \mathcal{O}(n^2)$
 183 flops. Given a threshold $\varepsilon > 0$, we truncate the singular values of \mathbf{X} to obtain a
 184 rank- k approximation $\mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^\top$ of matrix \mathbf{X} , which satisfies $\|\mathbf{X} - \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^\top\|_F \leq$
 185 $\varepsilon_{\text{SVD}} \|\mathbf{X}\|_F$. This is denoted as $[\mathbf{U}_k, \mathbf{\Sigma}_k, \mathbf{V}_k] = \text{SVD}(\mathbf{X}, \varepsilon_{\text{SVD}})$. The computational
 186 cost of computing the SVD is $\mathcal{O}(mn^2)$ flops.

187 **2.2. Randomized matrix algorithms.** An important component of our ap-
 188 proach is the use of randomized matrix methods for low-rank matrix approximation.
 189 In this subsection, we briefly review a few well-established randomized algorithms.

190 The first algorithm is the basic version of the randomized SVD proposed in [27].
 191 Suppose we want to compute a low-rank approximation of a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$; let
 192 the ℓ denote the number of samples which is a sum of the target rank and a small
 193 oversampling parameter, such that $\ell \leq \min\{m, n\}$. We generate a random matrix
 194 $\mathbf{\Omega} \in \mathbb{R}^{n \times \ell}$; in practice, we take the entries of this matrix to be i.i.d. standard Gaussian
 195 random variables. Then, we compute the product $\mathbf{Y} = \mathbf{X}\mathbf{\Omega}$ and obtain its thin QR
 196 factorization $\mathbf{Y} = \mathbf{Q}\mathbf{R}$. The main insight exploited by randomized SVD is that if
 197 the rank of \mathbf{X} is close to r , or the singular values of \mathbf{X} decay rapidly beyond r , then
 198 the range of \mathbf{Q} approximates well the range of \mathbf{X} in the sense that $\mathbf{X} \approx \mathbf{Q}\mathbf{Q}^\top \mathbf{X}$; we
 199 then use $\mathbf{Q}\mathbf{Q}^\top \mathbf{X}$ as a low-rank approximation to \mathbf{X} . The computational cost of this
 200 approach is

$$201 \quad C_{\text{randSVD}} = 2\ell mn + \mathcal{O}(\ell^2(m+n)) \quad \text{flops.}$$

202 Additional postprocessing can be performed to convert the low-rank approximation
 203 in the SVD format, or to truncate the low-rank approximation to the desired target
 204 rank; see [27] for additional details.

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

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

$$211 \quad (2.2) \quad \mathbf{X} \approx \mathbf{Y}(\mathbf{\Psi}\mathbf{X}\mathbf{\Omega})^\dagger\mathbf{Z},$$

212 where $\mathbf{Y} = \mathbf{X}\mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{\Psi}\mathbf{X}$. To implement the pseudoinverse, [41] suggests comput-
 213 ing the QR factorization $\mathbf{\Psi}\mathbf{X}\mathbf{\Omega} = \mathbf{Q}\mathbf{R}$ and then obtaining the low-rank approximation
 214 $(\mathbf{Y}\mathbf{R}^{-1})(\mathbf{Q}^\top\mathbf{Z})$. If the low-rank approximation is desired in the SVD format, this can
 215 be done by additional post-processing. In [41], the author recommends setting the
 216 sketch parameters as $\rho = r$ and $t = \lceil 1.5r \rceil$. The associated computational cost is

$$217 \quad C_{\text{genNys}} = 2mn(\rho + t) + \mathcal{O}(t^2(m + n) + t\rho^2) \quad \text{flops.}$$

218 **2.3. Standard TT arithmetic.** In this subsection, we review the standard
 219 approach to TT-rounding, first proposed in [44], using the notation of [13]. We also
 220 review the concepts of tensor contractions.

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

230 Following [13], we say a tensor is *right orthogonal* if its horizontal unfoldings
 231 $\mathcal{H}(\mathcal{J}_{\mathbf{x},n})$ have orthonormal rows for $n = 2, \dots, N$ (all except the first core). Simi-
 232 larly, we say that a tensor is *left orthogonal* if its vertical unfoldings $\mathcal{V}(\mathcal{J}_{\mathbf{x},n})$ have
 233 orthonormal columns for $n = 1, \dots, N - 1$ (all except the last core).

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

$$237 \quad \mathcal{V}(\mathcal{J}_{\mathbf{y},N-1})\mathcal{H}(\mathcal{J}_{\mathbf{y},N}) = \mathcal{V}(\mathcal{J}_{\mathbf{y},N-1})(\mathbf{Q}\mathbf{R})^\top = \underbrace{(\mathcal{V}(\mathcal{J}_{\mathbf{y},N-1})\mathbf{R}^\top)}_{\mathcal{V}(\mathcal{J}_{\mathbf{x},N-1})} \underbrace{(\mathbf{Q}^\top)}_{\mathcal{H}(\mathcal{J}_{\mathbf{x},N})}.$$

238 This procedure is continued through cores $N - 1, \dots, 2$ but we do not orthogonalize
 239 the first core. The details of right-to-left orthogonalization are given in [Algorithm 2.1](#)
 240 which will form the foundation for many of the subsequent algorithms. We can simi-
 241 larly obtain a left orthogonal tensor by processing the modes starting from mode-1,
 242 but we omit the details here.

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

Algorithm 2.1 Right-to-Left Orthogonalization

Require: A tensor \mathbf{Y} in TT format**Ensure:** \mathbf{X} is right-orthogonal tensor equivalent to \mathbf{Y}

```
1: function  $\mathbf{X} = \text{ORTHOGONALIZERL}(\mathbf{Y})$ 
2:    $\mathcal{J}_{\mathbf{X},N} = \mathcal{J}_{\mathbf{Y},N}$ 
3:   for  $n = N$  down to 2 do
4:      $[\mathcal{H}(\mathcal{J}_{\mathbf{X},n})^\top, \mathbf{R}] = \text{QR}(\mathcal{H}(\mathcal{J}_{\mathbf{X},n})^\top)$  ▷ thin QR factorization
5:      $\mathcal{V}(\mathcal{J}_{\mathbf{X},n-1}) = \mathcal{V}(\mathcal{J}_{\mathbf{Y},n-1}) \cdot \mathbf{R}^\top$  ▷  $\mathcal{J}_{\mathbf{X},n-1} = \mathcal{J}_{\mathbf{Y},n-1} \times_3 \mathbf{R}^\top$ 
6:   end for
7: end function
```

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

Algorithm 2.2 TT-Rounding

Require: A tensor \mathbf{Y} in TT format, user-defined threshold $\varepsilon_0 > 0$ **Ensure:** A tensor \mathbf{X} in TT format with reduced ranks such that $\|\mathbf{X} - \mathbf{Y}\| \leq \varepsilon_0\|\mathbf{Y}\|$

```
1: function  $\mathbf{X} = \text{TT-ROUNDING}(\mathbf{Y}, \varepsilon_0)$ 
2:    $\mathbf{X} = \text{ORTHOGONALIZERL}(\mathbf{Y})$ 
3:   Compute  $\|\mathbf{Y}\|_F$  and the truncation threshold  $\varepsilon_{\text{TT}} = \frac{\|\mathbf{Y}\|}{\sqrt{N-1}}\varepsilon_0$ 
4:   Set  $\mathcal{J}_{\mathbf{X},1} = \mathcal{J}_{\mathbf{Y},1}$ .
5:   for  $n = 1$  to  $N - 1$  do
6:      $[\mathcal{V}(\mathcal{J}_{\mathbf{X},n}), \mathbf{R}] = \text{QR}(\mathcal{V}(\mathcal{J}_{\mathbf{X},n}))$  ▷ thin QR factorization
7:      $[\widehat{\mathbf{U}}, \widehat{\mathbf{\Sigma}}, \widehat{\mathbf{V}}] = \text{SVD}(\mathbf{R}, \varepsilon_{\text{TT}})$  ▷  $\varepsilon_{\text{TT}}$ -truncated SVD factorization
8:      $\mathcal{V}(\mathcal{J}_{\mathbf{X},n}) = \mathcal{V}(\mathcal{J}_{\mathbf{X},n})\widehat{\mathbf{U}}$  ▷  $\mathcal{J}_{\mathbf{X},n} = \mathcal{J}_{\mathbf{X},n} \times_3 \widehat{\mathbf{U}}$ 
9:      $\mathcal{H}(\mathcal{J}_{\mathbf{X},n+1}) = \widehat{\mathbf{\Sigma}}\widehat{\mathbf{V}}^\top \mathcal{H}(\mathcal{J}_{\mathbf{X},n+1})$  ▷  $\mathcal{J}_{\mathbf{X},n+1} = \mathcal{J}_{\mathbf{X},n+1} \times_1 (\widehat{\mathbf{\Sigma}}\widehat{\mathbf{V}}^\top)$ 
10:  end for
11: end function
```

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

$$260 \quad (2.3) \quad \mathbf{W}_{n-1} = \mathcal{H}(\mathcal{J}_{\mathbf{X},n:N})\mathcal{H}(\mathcal{J}_{\mathbf{Y},n:N})^\top \in \mathbb{R}^{R_{n-1}^{\mathbf{X}} \times R_{n-1}^{\mathbf{Y}}}.$$

261 These partial contractions can be computed sequentially as

$$262 \quad (2.4) \quad \begin{aligned} \mathcal{V}(\mathcal{J}_{\mathbf{Z},n}) &= \mathcal{V}(\mathcal{J}_{\mathbf{X},n})\mathbf{W}_n, \\ \mathbf{W}_{n-1} &= \mathcal{H}(\mathcal{J}_{\mathbf{Z},n})\mathcal{H}(\mathcal{J}_{\mathbf{Y},n})^\top \end{aligned}$$

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

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




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 $\mathbf{W}_n = \mathcal{H}(\mathcal{J}_{\mathcal{X},n+1:N})\mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N})^\top$ for $1 \leq n < N$

- 1: **function** $\{\mathbf{W}_n\} = \text{PARTIALCONTRACTIONSRL}(\mathcal{X}, \mathcal{Y})$
 - 2: $\mathbf{W}_{N-1} = \mathcal{H}(\mathcal{J}_{\mathcal{X},N})\mathcal{H}(\mathcal{J}_{\mathcal{Y},N})^\top$
 - 3: **for** $n = N - 1$ down to 2 **do**
 - 4: $\mathcal{V}(\mathcal{J}_{\mathcal{Z},n}) = \mathcal{V}(\mathcal{J}_{\mathcal{X},n})\mathbf{W}_n$ $\triangleright \mathcal{J}_{\mathcal{Z},n} = \mathcal{J}_{\mathcal{X},n} \times_3 \mathbf{W}_n$, for temporary $\mathcal{J}_{\mathcal{Z},n}$
 - 5: $\mathbf{W}_{n-1} = \mathcal{H}(\mathcal{J}_{\mathcal{Z},n})\mathcal{H}(\mathcal{J}_{\mathcal{Y},n})^\top$ \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**
-

268 Detailed analysis of the overall computational costs of *Right-to-Left Orthogonal-*
269 *ization* (Algorithm 2.1), *TT-Rounding* (Algorithm 2.2) and *Right-to-Left Contraction*
270 (Algorithm 2.3) is presented in Section SM1.

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

281 Before we present the main algorithms, we describe a naive application of ran-
282 domized SVD to rounding. The first algorithm we propose, *Orthogonalize-then-*
283 *Randomize*, is very similar to the standard TT-rounding algorithm; the main dif-
284 ference is that we replace the truncated SVD step in Algorithm 2.2 with the basic
285 version of the randomized SVD reviewed in Subsection 2.2. The nomenclature of
286 this algorithm is clear from the fact that there are two phases in this approach: an
287 (already discussed) orthogonalization phase followed by a compression phase which
288 utilizes randomized SVD. As we show in the analysis of the computational costs Sub-
289 section 3.4 and the numerical experiments, this algorithm is expensive and the costs
290 are dominated by the first, i.e., orthogonalization, phase of the algorithm. Because
291 of this, the details of this approach are relegated to the supplementary materials,
292 section SM2.

293 **3.1. Randomize-then-Orthogonalize.** First, we consider a new Randomize-
 294 then-Orthogonalize algorithm that uses randomization to reduce the overall compu-
 295 tation cost of the TT-rounding procedure. It works by avoiding an expensive or-
 296 thogonalization of the original TT-tensor \mathcal{Y} with TT-ranks $\{R_n\}$ and instead uses
 297 randomization to reduce the computational cost. In contrast to the next approach in
 298 [Subsection 3.2](#) (Two-Sided-Randomization), here we use randomization only on one
 299 side.

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

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

306 By this definition, while the cores of \mathcal{R} have independent entries, the entries of the
 307 full tensor themselves are not independent. This normalization is chosen such that
 308 $\mathbb{E}\|\mathcal{J}_{\mathcal{R},n}\|_F^2 = 1$ and is sometimes necessary to ensure that no overflow occurs during
 309 the rounding computations. Note that constructing this random tensor requires only
 310 generating and storing $\sum_{n=1}^N \ell_{n-1} \ell_n I_n$ random entries. A related but distinct defi-
 311 nition for a Gaussian TT-tensor is given in [\[47\]](#), but the approach taken here differs
 312 considerably in how we use the randomized tensor.

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

$$320 \quad (3.1) \quad \mathcal{V}(\mathcal{J}_{\mathcal{X},n})\mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N})\mathcal{H}(\mathcal{J}_{\mathcal{R},n+1:N})^\top = \mathbf{Q}_n \mathbf{R}_n.$$

321 The equation above is analogous to sketching a factored matrix $\mathbf{Y} = \mathbf{A}\mathbf{B}$ by comput-
 322 ing the QR factorization of $\mathbf{A}\mathbf{B}\mathbf{\Omega}$. We emphasize that $\mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N})$ and $\mathcal{H}(\mathcal{J}_{\mathcal{R},n+1:N})$
 323 should not be formed explicitly; instead, we compute the contraction

$$324 \quad \mathbf{W}_n = \mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N})\mathcal{H}(\mathcal{J}_{\mathcal{R},n+1:N})^\top$$

325 efficiently using the process outlined in [\(2.4\)](#), then find the QR factorization of the
 326 small matrix

$$327 \quad \mathcal{V}(\mathcal{J}_{\mathcal{X},n})\mathbf{W}_n = \mathbf{Q}_n \mathbf{R}_n.$$

328 It can be seen via [\(2.4\)](#) that to compute \mathbf{W}_n we must first find $\mathbf{W}_{n+1}, \dots, \mathbf{W}_{N-1}$, so
 329 in order to avoid redundant computation we use [Algorithm 2.3](#) to obtain the partial
 330 contractions $\{\mathbf{W}_n\}_{n=1}^{N-1}$ once, then store and reuse them.

331 Since at the n th step the first $n - 1$ cores of \mathcal{X} are already orthogonalized, they
 332 do not need to be considered explicitly in the [factorization \(3.1\)](#). By projecting
 333 $\mathcal{V}(\mathcal{J}_{\mathcal{X},n})\mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N})$ onto the column space of \mathbf{Q}_n , we approximate the product of
 334 the final $N - n + 1$ cores as

$$335 \quad \mathcal{V}(\mathcal{J}_{\mathcal{X},n})\mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N}) \approx \mathbf{Q}_n \mathbf{Q}_n^\top \mathcal{V}(\mathcal{J}_{\mathcal{X},n})\mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N}) = \mathbf{Q}_n \mathbf{M}_n \mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N}).$$

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

337 It is important to mention that the Randomize-then-Orthogonalize approach produces
 338 a left-orthogonal tensor \mathcal{X} . We can use this observation to compress the tensor fur-
 339 ther. Therefore, if the ranks are not known a priori, then we choose the ranks to be
 340 sufficiently large and truncate them further by using [Algorithm 2.2](#). In particular,
 341 since the output tensor of [Algorithm 3.1](#) is left orthogonal, we can skip the orthogo-
 342 nalization phase ([Line 2](#) of [Algorithm 2.2](#)), and execute [Lines 3](#) to [10](#). This is what
 343 we do in our numerical experiments when the rank is not known *a priori*.

Algorithm 3.1 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  $\mathcal{X} = \text{TT-ROUNDING-RANDORTH}(\mathcal{Y}, \{\ell_n\})$ 
2:   Select a random Gaussian TT-tensor  $\mathcal{R}$  with target TT-ranks  $\{\ell_n\}$ 
3:    $\{\mathbf{W}_n\} = \text{PARTIALCONTRACTIONSRL}(\mathcal{Y}, \mathcal{R})$   $\triangleright$  compute partial random
   contractions
4:    $\mathcal{J}_{\mathcal{X},1} = \mathcal{J}_{\mathcal{Y},1}$ 
5:   for  $n = 1$  to  $N - 1$  do
6:      $\mathbf{Z}_n = \mathcal{V}(\mathcal{J}_{\mathcal{X},n})$   $\triangleright \mathcal{J}_{\mathcal{X},n}$  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:      $[\mathcal{V}(\mathcal{J}_{\mathcal{X},n}), \sim] = \text{QR}(\mathbf{Y}_n)$   $\triangleright$  thin QR to compute an orthonormal basis
9:      $\mathbf{M}_n = \mathcal{V}(\mathcal{J}_{\mathcal{X},n})^\top \mathbf{Z}_n$   $\triangleright$  form  $\ell_n \times R_n$  matrix
10:     $\mathcal{H}(\mathcal{J}_{\mathcal{X},n+1}) = \mathbf{M}_n \mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1})$   $\triangleright \mathcal{J}_{\mathcal{X},n+1} = \mathcal{J}_{\mathcal{Y},n+1} \times_1 \mathbf{M}_n$ 
11:  end for
12: end function

```

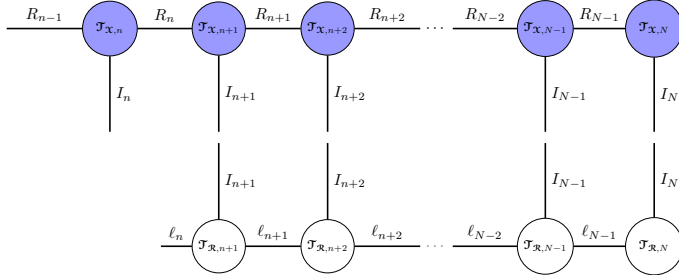


FIG. 3.1. Random projection for the Randomize-then-Orthogonalize [Algorithm 3.1](#).

344 **3.2. Two-Sided-Randomization.** Analogous to the one-sided [Algorithm 3.1](#),
 345 we start with generating two TT random Gaussian tensors \mathcal{L} and \mathcal{R} with given target
 346 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
 347 \mathcal{Y} from the left and right, respectively (*randomization phase*, see [Figure SM2](#)). Next,
 348 for each $n = 1, \dots, N - 1$ we compute the SVD of a product of partial contractions
 349 $\mathbf{W}_n^L \mathbf{W}_n^R$, 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

$$350 \quad (3.2) \quad \mathbf{L}_n = \mathbf{W}_n^R \mathbf{V}_n (\mathbf{\Sigma}_n^\dagger)^{1/2} \quad \text{and} \quad \mathbf{R}_n = (\mathbf{\Sigma}_n^\dagger)^{1/2} \mathbf{U}_n^\top \mathbf{W}_n^L.$$

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

$$353 \quad \mathbf{Y}_{(1:n)} = \mathcal{V}(\mathcal{J}_{\mathcal{Y},1:n}) \mathcal{H}(\mathcal{J}_{\mathcal{Y},n+1:N}),$$

Algorithm 3.2 TT-Rounding: Two-Sided-Randomization (Generalized Nyström)

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

Ensure: A tensor \mathbf{X} in TT format with ranks $\{\ell_n\}$

```

1: function  $\mathbf{X} = \text{TT-ROUNDING-RANDORTH}(\mathbf{Y}, \{\ell_n\})$ 
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:    $\{\mathbf{W}_n^L\} = \text{PARTIALCONTRACTIONSRL}(\mathbf{Y}, \mathcal{L})$ 
                                     ▷ Precompute sketches from the left
5:    $\{\mathbf{W}_n^R\} = \text{PARTIALCONTRACTIONSRL}(\mathbf{Y}, \mathcal{R})$ 
                                     ▷ 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:      $\mathbf{R}_n = (\mathbf{\Sigma}_n^\dagger)^{1/2} \mathbf{U}_n^L \mathbf{W}_n^L$        ▷ Determine internal  $\ell_n \times R_n$  right factor  $\mathbf{R}_n$ 
10:  end for
11:   $\mathcal{V}(\mathcal{J}_{\mathbf{X},1}) = \mathcal{V}(\mathcal{J}_{\mathbf{Y},1}) \mathbf{L}_1$ 
12:  for  $n = 2$  to  $N - 1$  do
13:     $\mathcal{H}(\mathcal{J}_{\mathbf{X},n}) = \mathbf{R}_{n-1} \mathcal{H}(\mathcal{V}(\mathcal{J}_{\mathbf{Y},n}) \mathbf{L}_n)$ 
                                                         ▷  $\mathcal{J}_{\mathbf{X},n} = \mathcal{J}_{\mathbf{Y},n} \times_1 \mathbf{R}_{n-1} \times_3 \mathbf{L}_n$ 
                                                         ▷ hence  $\mathcal{J}_{\mathbf{X},n}$  is  $\ell_{n-1} \times I_n \times \ell_n$ 
14:  end for
15:   $\mathcal{H}(\mathcal{J}_{\mathbf{X},N}) = \mathbf{R}_{N-1} \mathcal{H}(\mathcal{J}_{\mathbf{Y},N})$ 
16: end function

```

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

$$\begin{aligned}
 356 \quad \Psi_n &:= \mathcal{V}(\mathcal{J}_{\mathcal{L},1:n})^\top \in \mathbb{R}^{\ell_n \times (I_1 \cdots I_n)}, \\
 357 \quad \Omega_n &:= \mathcal{H}(\mathcal{J}_{\mathcal{R},n+1:N})^\top \in \mathbb{R}^{(I_{n+1} \cdots I_N) \times \rho_n}
 \end{aligned}$$

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

$$361 \quad \Psi_n \mathbf{Y}_{(1:n)} \Omega_n = (\Psi_n \mathcal{V}(\mathcal{J}_{\mathbf{Y},1:n})) (\mathcal{H}(\mathcal{J}_{\mathbf{Y},n+1:N}) \Omega_n) = \mathbf{W}_n^L \mathbf{W}_n^R.$$

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

$$\begin{aligned}
 364 \quad \mathbf{Y}_{(1:n)} &\approx (\mathbf{Y}_{(1:n)} \Omega_n) (\Psi_n \mathbf{Y}_{(1:n)} \Omega_n)^\dagger (\Psi_n \mathbf{Y}_{(1:n)}) \\
 &= \mathcal{V}(\mathcal{J}_{\mathbf{Y},1:n}) \mathbf{W}_n^R (\mathbf{W}_n^L \mathbf{W}_n^R)^\dagger \mathbf{W}_n^L \mathcal{H}(\mathcal{J}_{\mathbf{Y},n+1:N}) \\
 &= \mathcal{V}(\mathcal{J}_{\mathbf{Y},1:n}) \mathbf{L}_n \mathbf{R}_n \mathcal{H}(\mathcal{J}_{\mathbf{Y},n+1:N}),
 \end{aligned}$$

365 see Figure SM3. Having all left and right factors at hand, for each core of the tensor
 366 \mathbf{Y} (treating the first and last core separately), we distribute them according to the
 367 formula

$$368 \quad \mathcal{H}(\mathcal{J}_{\mathbf{X},n}) = \mathbf{R}_{n-1} \mathcal{H}(\mathcal{V}(\mathcal{J}_{\mathbf{Y},n}) \mathbf{L}_n)$$

369 forming the cores of the resulting tensor \mathbf{X} , see Figure SM4.

370 In contrast to Algorithm 3.1, the Two-Sided-Randomization approach does not
 371 produce an orthogonal tensor. However, with a little restructuring, it can be adapted

372 to produce an orthogonal tensor (we do not discuss that here). This variation may
 373 be useful for the case when the target TT-ranks are not known in advance, and
 374 producing an orthogonal tensor can be used in conjunction with [Algorithm 2.2](#) to
 375 further compress the tensor.

376 **3.3. Rounding of TT-sums.** One of the most common arithmetic operations
 377 that depends on TT-rounding is TT-summation, i.e., we want to compress a tensor
 378 \mathbf{Y} that is available as the sum of s TT-tensors: $\mathbf{Y} = \mathbf{y}^{(1)} + \dots + \mathbf{y}^{(s)}$. To reuse
 379 existing algorithms, there are two options available to us. For example, we can form
 380 the TT-tensor \mathbf{Y} explicitly and then apply one of the compression algorithms pro-
 381 posed previously. As we will argue in [Subsection 3.4](#), the computational costs of this
 382 approach has cubic scaling with respect to the number of summands s using the TT-
 383 rounding approach, and a quadratic scaling with respect to s using the randomized
 384 approaches (Randomize-then-Orthogonalize and Two-Sided-Randomization). This is
 385 computationally infeasible as s becomes large. Alternatively, we can form the partial
 386 sum $\mathbf{y}^{(1)} + \mathbf{y}^{(2)}$, compress this partial sum, add the resulting truncated term to the
 387 summand $\mathbf{y}^{(3)}$ and proceed in the same way with remaining terms; see e.g., [5]. Vari-
 388 ations of this approach can be performed using ideas from the summation methods
 389 described in [28, Chapter 4.1]. [These approaches scale linearly with \$s\$, but assuming](#)
 390 [that one prescribes a certain tolerance at each pairwise summation step, this could](#)
 391 [lead to large intermediate TT-ranks even if the rank of the overall sum is relatively](#)
 392 [small.](#)

393 In this subsection, we show how to combine the addition and randomized rounding
 394 operations to reduce further the computational costs, which is particularly effective
 395 when the number of summands s is large. The basic idea is to exploit the nonzero
 396 structure of the TT-cores of the sum of TT-tensors to avoid computing with zeros.
 397 Applying the orthogonalization phase, as required to perform the deterministic trun-
 398 cation phase, requires assembling the TT representation of the sum. Furthermore,
 399 the orthogonalization and multiplication by the triangular factor destroys the struc-
 400 ture in the middle cores. By using randomization, we can avoid this explicit TT
 401 assembly of the sum and avoid unnecessary computations on zeros. [Algorithm 3.3](#)
 402 provides the pseudocode for rounding the sum of s input TT-tensors. To simplify the
 403 notation, we derive the efficient computations considering the case $s = 2$, as the gen-
 404 eralization will be clear. Let \mathbf{Y} and \mathbf{Z} be two TT-tensors and consider the TT-tensor
 405 $\mathbf{X} = \mathbf{Y} + \mathbf{Z}$. Let \mathcal{R} be a given random Gaussian TT-tensor and let $\{\mathbf{W}_n^{\mathbf{y}}\}$ and $\{\mathbf{W}_n^{\mathbf{z}}\}$,
 406 for $n = 1, \dots, N - 1$ be the right-to-left partial contractions of \mathbf{Y} and \mathbf{Z} with \mathcal{R} . We
 407 have

$$408 \quad \mathbf{X}_{(1:n)} = \mathcal{V}(\mathcal{J}_{\mathbf{x},1:n})\mathcal{H}(\mathcal{J}_{\mathbf{x},n+1:N}) = \begin{bmatrix} \mathcal{V}(\mathcal{J}_{\mathbf{y},1:n}) & \mathcal{V}(\mathcal{J}_{\mathbf{z},1:n}) \end{bmatrix} \begin{bmatrix} \mathcal{H}(\mathcal{J}_{\mathbf{y},n+1:N}) \\ \mathcal{H}(\mathcal{J}_{\mathbf{z},n+1:N}) \end{bmatrix},$$

409 and by using [\(2.1\)](#),

$$410 \quad \mathcal{V}(\mathcal{J}_{\mathbf{x},1:n}) = \begin{bmatrix} (\mathbf{I}_{I_n} \otimes \mathcal{V}(\mathcal{J}_{\mathbf{y},1:n-1})) & (\mathbf{I}_{I_n} \otimes \mathcal{V}(\mathcal{J}_{\mathbf{z},1:n-1})) \end{bmatrix} \begin{bmatrix} \mathcal{V}(\mathcal{J}_{\mathbf{y},n}) \\ \mathcal{V}(\mathcal{J}_{\mathbf{z},n}) \end{bmatrix},$$

$$411 \quad \mathcal{H}(\mathcal{J}_{\mathbf{x},n+1:N}) = \begin{bmatrix} \mathcal{H}(\mathcal{J}_{\mathbf{y},n+1}) & \\ & \mathcal{H}(\mathcal{J}_{\mathbf{z},n+1}) \end{bmatrix} \begin{bmatrix} \mathcal{H}(\mathcal{J}_{\mathbf{y},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \\ \mathcal{H}(\mathcal{J}_{\mathbf{z},n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \end{bmatrix}.$$

413 The matrix $\mathbf{W}_n^{\mathbf{x}}$ can be expressed as

$$414 \quad \mathbf{W}_n^{\mathbf{x}} = \mathcal{H}(\mathcal{J}_{\mathbf{x},n+1:N})\mathcal{H}(\mathcal{J}_{\mathcal{R},n+1:N})^\top = \begin{bmatrix} \mathcal{H}(\mathcal{J}_{\mathbf{y},n+1:N}) \\ \mathcal{H}(\mathcal{J}_{\mathbf{z},n+1:N}) \end{bmatrix} \mathcal{H}(\mathcal{J}_{\mathcal{R},n+1:N})^\top = \begin{bmatrix} \mathbf{W}_n^{\mathbf{y}} \\ \mathbf{W}_n^{\mathbf{z}} \end{bmatrix}.$$

Algorithm 3.3 TT-Rounding of a Sum: Randomize-then-Orthogonalize

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

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

```

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

```

415 This justifies the procedure in [Algorithm 3.3](#) of computing the partial contractions
 416 separately for each summand ([Line 4](#)) and concatenating them ([Line 9](#)).

417 After the QR factorization of the projected matrix produces the truncated core,
 418 we compute the contraction between the new and old cores ([Line 11](#)) and store the
 419 result in a matrix $\mathbf{M}_n = [\mathbf{M}_n^{\mathbf{y}} \ \mathbf{M}_n^{\mathbf{z}}]$ of size $\ell_n \times (R_n^{\mathbf{y}} + R_n^{\mathbf{z}})$. This matrix is now
 420 multiplied from the right by $\mathcal{H}(\mathcal{J}\mathbf{x}_{,n+1:N})$ to compute the updated right factor of
 421 $\mathbf{X}_{(1:n)}$. This multiplication can be absorbed by the $(n+1)$ th core as follows:

$$\begin{aligned}
 422 \quad \mathbf{M}_n \mathcal{H}(\mathcal{J}\mathbf{x}_{,n+1:N}) &= [\mathbf{M}_n^{\mathbf{y}} \ \mathbf{M}_n^{\mathbf{z}}] \begin{bmatrix} \mathcal{H}(\mathcal{J}\mathbf{y}_{,n+1}) & & \\ & \mathcal{H}(\mathcal{J}\mathbf{z}_{,n+1}) & \\ & & \begin{bmatrix} \mathcal{H}(\mathcal{J}\mathbf{y}_{,n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \\ \mathcal{H}(\mathcal{J}\mathbf{z}_{,n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \end{bmatrix} \end{bmatrix}, \\
 423 \quad &= [\mathbf{M}_n^{\mathbf{y}} \mathcal{H}(\mathcal{J}\mathbf{y}_{,n+1}) \ \mathbf{M}_n^{\mathbf{z}} \mathcal{H}(\mathcal{J}\mathbf{z}_{,n+1})] \begin{bmatrix} \mathcal{H}(\mathcal{J}\mathbf{y}_{,n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \\ \mathcal{H}(\mathcal{J}\mathbf{z}_{,n+2:N}) \otimes \mathbf{I}_{I_{n+1}} \end{bmatrix}.
 \end{aligned}$$

424
 425

426 Hence we update $\mathcal{H}(\mathcal{J}\mathbf{x}_{,n+1})$ as in [Line 13](#). For $n = N - 1$ we have

$$427 \quad \mathcal{J}\mathbf{x}_{,N} = [\mathbf{M}_{N-1}^{\mathbf{y}} \ \mathbf{M}_{N-1}^{\mathbf{z}}] \begin{bmatrix} \mathcal{J}\mathbf{y}_{,N} \\ \mathcal{J}\mathbf{z}_{,N} \end{bmatrix} = \mathbf{M}_{N-1}^{\mathbf{y}} \mathcal{J}\mathbf{y}_{,N} + \mathbf{M}_{N-1}^{\mathbf{z}} \mathcal{J}\mathbf{z}_{,N}.$$

428 Generalizing this expression to s terms yields [Line 15](#).

429 **3.4. Computational costs.** To analyze the computational cost, we make the
 430 following assumptions that will simplify the analysis. Let $\mathbf{Y} \in \mathbb{R}^{I \times \dots \times I}$ be a tensor of
 431 order N with ranks $(1, R, \dots, R, 1)$ in TT format. We want to compress \mathbf{Y} to obtain
 432 a TT-tensor \mathbf{X} with ranks $(1, \ell, \dots, \ell, 1)$. Here and in [Section SM1](#), we assume that
 433 $\ell = \Theta(R)$.

434 **3.4.1. Randomized compression algorithms.** We now analyze the compu-
 435 tational cost of [Algorithm 3.1](#), and [Algorithm 3.2](#). The computational cost of [Algo-](#)
 436 [rithm SM2.1](#) is given in [Section SM2](#).

437 *Randomize-then-Orthogonalize* ([Algorithm 3.1](#)). Denoting the total cost of [Algo-](#)
 438 [rithm 3.1](#) with C_{RtO} , we analyze the main components that contribute to the total
 439 computational cost. [Line 3](#) invokes [Algorithm 2.3](#) with the corresponding cost de-
 440 noted by C_{Contr} , see [Section SM1](#). [Lines 7](#) and [9](#) contribute by a factor of $2IR\ell^2$ that
 441 is the cost of performing the multiplication $\mathcal{V}(\mathcal{J}_{\mathbf{x},n})\mathbf{W}_n$ of sizes $\ell I \times R$ with $R \times \ell$
 442 and $\mathcal{V}(\mathcal{J}_{\mathbf{x},n})^T\mathbf{Z}_n$ of sizes $\ell \times I\ell$ and $I\ell \times R$ that prepare the matrices \mathbf{Y}_n and \mathbf{M}_n ,
 443 respectively, for the next steps. The term $4I\ell^3 + \mathcal{O}(\ell^3)$ represents the cost of the thin
 444 QR factorization, in [Line 8](#), of the matrix \mathbf{Y}_n of size $I\ell \times \ell$. [Line 10](#) that involves
 445 multiplication of matrices of size $\ell \times R$ and $R \times IR$ which costs $2IR^2\ell$ flops. The total
 446 cost of [Algorithm 3.1](#) is

$$447 \quad C_{\text{RtO}} = C_{\text{Contr}} + (N-2)(2IR^2\ell + 4IR\ell^2 + 4I\ell^3) + \mathcal{O}(IR^2 + NR^3)$$

$$448 \quad = I(N-2) \cdot (4R^2\ell + 6R\ell^2 + 4\ell^3) + \mathcal{O}(IR^2 + NR^3) \quad \text{flops.}$$

449 *Two-Sided-Randomization* ([Algorithm 3.2](#)). Here we analyze the total cost of
 450 [Algorithm 3.2](#) and denote it by $C_{2\text{SR}}$. The cost of [Lines 4](#) and [5](#) is $2C_{\text{Contr}}$. Here,
 451 we recall that the cost of the partial contractions from the left and from the right
 452 is the same since we assume that ranks are the same, i.e., $\rho_n = \ell_n = \ell$. However,
 453 in practice, we choose $\rho_n > \ell_n$ to be different for numerical stability. There are two
 454 other contributions to the cost that come from two different matrix multiplication:
 455 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
 456 the for loop starting with [Line 12](#) is $\mathcal{O}(NR^2)$. Hence, the total computational cost of
 457 [Algorithm 3.2](#) is

$$458 \quad C_{2\text{SR}} = 2 \cdot C_{\text{Contr}} + (N-2)(2IR^2\ell + 2IR\ell^2) + \mathcal{O}(NR^2)$$

$$459 \quad = I(N-2) \cdot (6R^2\ell + 6R\ell^2) + \mathcal{O}(IR\ell + NR^2) \quad \text{flops.}$$

460 Nakatsukasa [\[41\]](#) suggests oversampling from the left, i.e., taking $\rho_n = \lceil 1.5\ell_n \rceil$. We
 461 follow this suggestion in our numerical experiments. With this assumption, the cost
 462 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
 463 increased cost of the contraction ([Algorithm 2.3](#)) with a larger random tensor.

464 *Comparison of different algorithms.* To enable the comparison of the different
 465 algorithms, we set a target rank as $\ell = \beta R$, where $\beta \in (0, 1]$ is the ratio between
 466 the target rank ℓ and the current rank R . This allows us to compare the different
 467 algorithms more clearly. A summary of the dominant costs of the algorithms is pro-
 468 vided in [Table 3.1](#). For simplicity, we also provide a simplified representation of the
 469 computational costs with $\beta = \ell/R$. In [Figure 3.2](#), we plot the speedup of the ran-
 470 domized algorithms compared to the TT-Rounding algorithm; we used the simplified
 471 representation of the costs while generating the figure. It can be easily seen that
 472 all the proposed methods are faster than the TT-Rounding algorithm. However, the
 473 speedup using the Orthogonalize-then-Randomize algorithm is very incremental. In
 474 contrast, the other two algorithms, Randomize-then-Orthogonalize and Two-Sided-
 475 Randomization, have very similar costs (Randomize-then-Orthogonalize is slightly

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	$(N-2)I(6R^2\ell + 6R\ell^2)$	$(N-2)IR^3(6\beta + 6\beta^2)$




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.

476 more efficient for smaller β) and have much higher speedups especially if $\beta \ll 1$. If
 477 $\beta \approx 1$, both algorithms are very close to the TT-Rounding algorithm. Therefore, the
 478 proposed methods are most efficient if $\beta \ll 1$, i.e., the target rank ℓ is much smaller
 479 compared to the original rank R .

480 **3.4.2. Rounding of TT-sums.** To explain the benefits of the algorithm for
 481 rounding TT-sums in Subsection 3.3, consider the summation of s tensors of order
 482 N (size I in each dimension) each with TT-ranks $(1, R, \dots, R, 1)$. Suppose we form
 483 the TT-tensor \mathbf{Y} , which represents the summation $\mathbf{Y} = \sum_{j=1}^s \mathbf{Y}^{(j)}$, explicitly. The
 484 intermediate cores have size $sR \times I \times sR$; the first core is of size $I \times sR$ and the last
 485 core is of size $sR \times I$. Suppose the target compression rank in each mode is ℓ . To
 486 leading order, the cost of executing TT-Rounding and Orthogonalize-then-Randomize
 487 is $\mathcal{O}(Ns^3R^3I)$. In contrast, the cost of using Randomize-then-Orthogonalize and the
 488 Two-Sided-Randomization approach are both $\mathcal{O}(Nls^2R^2I)$. This can be beneficial if
 489 the number of summands s is large, or the rank R is large. This simple cost analysis
 490 does not take into account any structure present in the summation.

491 Carefully exploiting the structure, as in Subsection 3.3, can reduce this cost. In
 492 particular, by using TT-Rounding of a sum with s tensors of order N with Randomize-
 493 then-Orthogonalize summarized in Algorithm 3.3 the leading order computational
 494 cost is $\mathcal{O}(NlsR^2I)$ flops. Notice that by exploiting the structure of the tensor and
 495 using randomization, the leading order of the cost is decreased to be linear in s in

496 contrast to cubic in s when no structure was taken into account and the TT-sum
 497 tensor was formed explicitly. This decrease in the computational cost is obviously
 498 more pronounced when the number of summands s is large. In what follows, we
 499 present the analysis of the computational cost of computing the sum of s TT-tensors
 500 by randomization and by exploiting the underlying tensor structure.

501 *TT-rounding of a sum: Randomize-then-Orthogonalize (Algorithm 3.3).* We an-
 502alyze the computational cost of Algorithm 3.3 which we denote by C_{RtOsum} . The
 503leading order term is sourced from two main contributions (1) Line 4 that executes
 504Algorithm 2.3 s times resulting in a total computational cost of $sI(N-2)(2R\ell^2+2R^2\ell)$
 505flops and (2) Line 13 that represents s multiplications between matrices of size $IR \times R$
 506and $R \times \ell$ resulting in a total cost of $2sIR^2\ell$ flops. Next we analyze the source of the
 507second leading order term present in the total cost. Line 9 contributes by a factor of
 508two to the second leading term with a cost $2sIR\ell^2$ flops resulting from multiplying
 509matrices of size $I\ell \times sR$ and $sR \times \ell$. Another factor of two comes from the multipli-
 510cation of two matrices of size $\ell \times I\ell$ and $I\ell \times sR$ in Line 11, resulting in a total cost
 511of $2sIR\ell^2$ flops. The last factor of two comes from the second leading order term of
 512Algorithm 2.3. Computing the thin QR factorization of the matrix of size $I\ell \times \ell$ in
 513Line 10 costs $4I\ell^3$ flops. Hence, the total cost of Algorithm 3.3 is

$$514 \quad C_{\text{RtOsum}} = s \cdot C_{\text{Contr}} + (N-2)(2sIR^2\ell + 4sIR\ell^2 + 4I\ell^3) + \mathcal{O}(NIR^2)$$

$$515 \quad = I(N-2) \cdot (4sR^2\ell + 6sR\ell^2 + 4\ell^2) + \mathcal{O}(NIR^2) \text{ flops.}$$

516 **4. Numerical results.** In this section, we illustrate numerically the perfor-
 517mance of the newly developed algorithms using tensor data in TT format. We consider
 518both synthetic as well as more realistic test examples. Additional numerical experi-
 519ments are available in Section SM4. All the numerical experiments were performed
 520on MATLAB R2021a running on a laptop computer with CPU Intel(R) Core(TM)
 521i9-9980H and 64GB of RAM, using multithreading with 4 computational threads.

522 **4.1. TT-tensor with a fixed target rank.** In our first numerical experiment,
 523we illustrate the performance of our rounding algorithms by rounding a random TT-
 524tensor with a known low-rank representation. Throughout, we choose the ranks of the
 525right side randomization in the Two-Sided-Randomization approach (Algorithm 3.2)
 526to be $\rho = \lceil 1.5\ell \rceil$ as discussed in Subsection 3.4.1.

527 The random TT-tensor \mathcal{X} is constructed by perturbing a random TT-tensor \mathcal{X}_1
 528with the random TT-tensor $\epsilon\mathcal{X}_2$ as follows: $\mathcal{X} = \mathcal{X}_1 + \epsilon\mathcal{X}_2$. TT-tensors $\mathcal{X}_1, \mathcal{X}_2 \in$
 529 $\mathbb{R}^{100 \times \dots \times 100}$ are order $N = 10$ normalized random TT-tensors of ranks $(1, 50, \dots, 50, 1)$
 530(normalized according to their dimension as described in Definition 3.1), and ϵ is
 531a perturbation scalar taking the values $\epsilon \in \{10^{-2}, 10^{-6}, 10^{-10}\}$. The ranks of the
 532perturbed tensor \mathcal{X} are $(1, 50 + 50, \dots, 50 + 50, 1)$, and the perturbation parameter ϵ
 533determines how well tensor \mathcal{X} is approximated by the lower rank tensor \mathcal{X}_1 , i.e., if ϵ
 534is small, then \mathcal{X}_1 is a good rank- $(1, 50, \dots, 50, 1)$ approximation of \mathcal{X} .

535 We round the random TT-tensor \mathcal{X} using Algorithms SM2.1, 2.2, 3.1 and 3.2 to
 536have ranks $(1, \ell, \dots, \ell, 1)$, where we vary the parameter ℓ from 35 to 80 by an increment
 537of 5. We present these results in Figure 4.1. The approximation error is the relative
 538norm error between tensor \mathcal{X} and the approximate rounded tensor $\widehat{\mathcal{X}}$, i.e., $\frac{\|\mathcal{X} - \widehat{\mathcal{X}}\|}{\|\mathcal{X}\|}$.
 539We also present the time speedup (computed with the average of 5 independent runs)
 540of the randomized algorithms compared to the deterministic algorithm.

541 For all values of perturbation ϵ , the error resulting from the deterministic al-
 542gorithm (Algorithm 2.2) decreases slightly until the ranks of the rounded tensor are

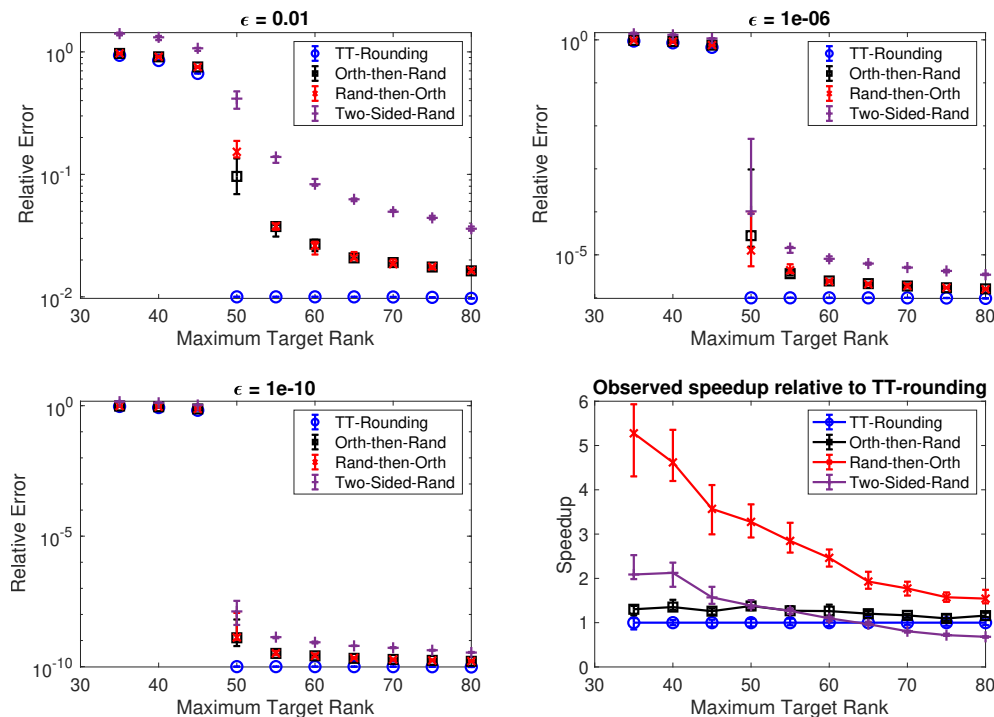


FIG. 4.1. Comparison of error between a low-rank tensor and full rank perturbed tensor, and timings using the deterministic and randomized TT-rounding algorithms for ϵ perturbed tensor for different values of perturbation ϵ . Statistics were based on 5 independent runs.

543 $\ell = 50$. When $\ell > 50$, the error appears to be very close to ϵ . The errors resulting from
 544 the randomized algorithms (Algorithms SM2.1, 3.1 and 3.2) are greater than the error
 545 resulting from the deterministic algorithm. Additionally, the randomized algorithms
 546 produce a more accurate approximation when the target rank is larger and are more
 547 accurate when ϵ is small. The Orthogonalize-then-Randomize and the Randomize-
 548 then-Orthogonalize method (Algorithms SM2.1 and 3.1) produce a rounded tensor
 549 with similar levels of accuracy while the Two-Sided-Randomization approach is the
 550 least accurate. For smaller values of ϵ , there is less difference in the accuracy between
 551 the different algorithms. The Randomize-then-Orthogonalize algorithm is the fastest
 552 compared to the deterministic algorithm followed by the Two-Sided-Randomization
 553 and Orthogonalize-then-Randomize algorithms.

554 **4.2. Solving a parametric PDE in the TT format.** As a realistic test exam-
 555 ple, we consider the parameter dependent PDE referred in the literature as the
 556 *cookie problem* [38, 52], see Section SM5 for details. Since it is known that the set
 557 of solutions of problem (SM5.1) admits a low-rank representation [24, 14], we consider
 558 a global linear system encapsulating all these linear systems, i.e.,

$$559 \quad \left(\sum_{i=1}^N \mathbf{A}_{i,1} \otimes \dots \otimes \mathbf{A}_{i,N} \right) \mathbf{x} = \mathcal{F},$$

560 where $\mathbf{A}_{1,1}$ is the discretization of the operator over the spatial domain with constant
 561 parameter values, for each $2 \leq i \leq N$, $\mathbf{A}_{i,1}$ is the discretization of the operator over the

562 domain multiplied by the characteristic function corresponding to the corresponding
563 subdomain and $\mathbf{A}_{i,i}$ is a diagonal matrix containing the parameter values for the
564 corresponding parameter, and for each $2 \leq j \neq i \leq N$, $\mathbf{A}_{i,j}$ is the identity matrix.
565 We use the TT-GMRES algorithm [17] to solve this global linear system of equations.
566 The preconditioned TT-GMRES algorithm builds the basis vectors in TT format
567 $\mathbf{V}_1, \mathbf{V}_2, \dots$ using the inexact Arnoldi procedure; since at each step, the corresponding
568 TT-tensors are rounded, this results in an inexact Krylov subspace method. The main
569 bottleneck is the computation of two linear combinations in each iteration. First, the
570 following sum of N tensors with the same ranks as \mathbf{V}_k is formed when applying the
571 operator to the k -th basis vector computed at the previous iteration, i.e.,

$$572 \quad \mathbf{W} = \sum_{i=1}^N (\mathbf{A}_{i,1} \otimes \dots \otimes \mathbf{A}_{i,N}) \mathbf{Y},$$

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

$$576 \quad \mathbf{Z} = \left(\mathbf{W} - \sum_{j=1}^k h_{jk} \mathbf{V}_j \right), \quad \mathbf{V}_{k+1} = \frac{1}{h_{k+1,k}} \mathbf{Z}, \quad \begin{cases} h_{jk} = \langle \mathbf{V}_j, \mathbf{W} \rangle, & j = 1, \dots, k, \\ h_{k+1,k} = \|\mathbf{Z}\|_F. \end{cases}$$

577 In both cases, the addition of TT-tensors is followed by a TT-rounding operation in
578 order to reduce the ranks and keep the computations tractable. Hence, these steps are
579 amenable to acceleration by using the randomized [Algorithm 3.3](#) as a rounding pro-
580 cedure in the aforementioned computations. Because the second linear combination
581 involves $k + 1$ summands, the randomized implementation reduces greatly the cost of
582 later TT-GMRES iterations in particular, as the leading order of its computational
583 cost is decreased from cubic to linear in k , as detailed in [Subsection 3.4](#).

584 We perform numerical experiments with $N = 5$ using a piecewise linear finite
585 element discretization with the mesh presented in [Figure SM7](#), for various choices of
586 the number of parameter samples, $I = I_2 = \dots = I_N$, with values of ρ_i distributed
587 linearly between 1 and 10. The relative tolerance of the TT-GMRES solver is set
588 to 10^{-8} . We compare the naive, deterministic implementation of the preconditioned
589 TT-GMRES algorithm with one using randomized summation and rounding steps.
590 Results of the comparison are reported in [Figure 4.2](#). On the right, we display all in-
591 ternal ranks of the TT representation of the Krylov vector computed at that iteration.
592 We observe that the ranks of the basis vectors and number of iterations are the same
593 using both implementations, and they depend only weakly on the dimension I of the
594 parameter modes of the tensors. The speedup achieved by the randomized approach
595 increases consistently with the number of parameter samples. Taking a closer look
596 at the timing statistics as presented in [Figure 4.3](#), we note that the computation and
597 rounding of the linear combinations identified above indeed dominates the computa-
598 tional cost in both cases, and is a clear computational bottleneck for the deterministic
599 implementation in particular, as the ranks of the sum tensor increase to as much as
600 2491 in these experiments. This explains the remarkable speedup obtained with the
601 randomized approach.

602 **5. Conclusions and outlook.** In this paper, we present randomized algorithms
603 for rounding a tensor, assuming that we have an initial representation in the TT

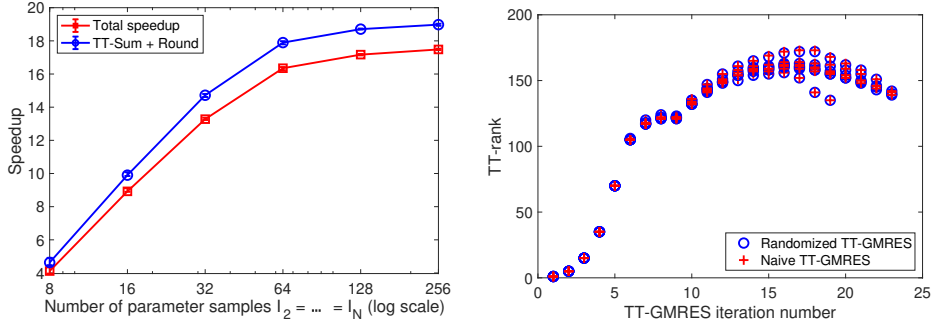


FIG. 4.2. Illustration of the speedups (left) and TT-ranks of the Krylov basis vectors (right) obtained by the deterministic and randomized summation and rounding algorithms within the TT-GMRES algorithm to solve problem (SM5.1).

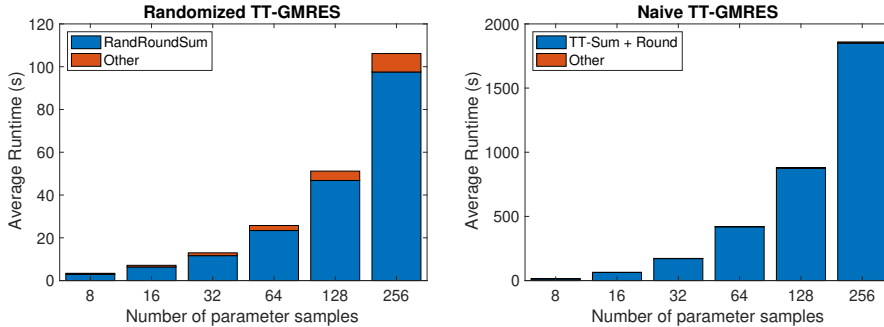


FIG. 4.3. Illustration of the timings using the deterministic and randomized summation and rounding algorithms within the TT-GMRES algorithm to solve problem (SM5.1).

604 format. This initial representation may not be optimal in terms of storage, and
605 the randomized compression techniques can be used to obtain a more efficient rep-
606 resentation. We derive three different algorithms: Orthogonalize-then-Randomize,
607 Randomize-then-Orthogonalize, and Two-Sided-Randomization. We study the com-
608 putational cost of these algorithms in some detail and show that it can be much
609 smaller than the standard TT-rounding algorithm. Additionally, we consider the spe-
610 cial case of rounding a TT-tensor that is represented as the sum of many TT-tensors.
611 While applying each of the randomized algorithms proposed here can reduce the com-
612 putational cost over standard TT-rounding, we further exploit the structure of the
613 problem to reduce the computational cost to be linear in the number of summands.
614 We perform extensive numerical experiments and achieve over $20\times$ speedups on test
615 problems compared to standard algorithms. There are many avenues for future inves-
616 tigations. First, it would be interesting to derive probabilistic bounds for the accuracy
617 of the rounding approach. Second, we could consider extending our algorithms to the
618 case where the TT-tensor is obtained as the Hadamard (or elementwise) product of
619 two tensors. Finally, another extension worth considering is developing randomized
620 rounding algorithms in the \mathcal{H} -Tucker format.

621 **6. Acknowledgements.** This work was initiated as a part of the Statistical and
622 Applied Mathematical Sciences Institute (SAMSI) Program on Numerical Analysis in
623 Data Science in Fall 2020. Any opinions, findings, and conclusions or recommenda-

624 tions expressed in this material are those of the authors and do not necessarily reflect
 625 the views of the National Science Foundation (NSF). Additionally, the authors would
 626 like to acknowledge partial support through the NSF: G.B. (CCF-1942892), P.C.
 627 (DMS-1819220), E.H. (DMS-1745654), A.M. (CCF-1812927), M.P. (DMS-1502640),
 628 T.W.R. (DMS-1745654), A.K.S. (DMS-1821149). The authors thank Dr. Jocelyn Chi
 629 and Prof. Eric de Sturler for the constructive discussions and suggestions.

630

REFERENCES

- 631 [1] S. Ahmadi-Asl, S. Abukhovich, M. G. Asante-Mensah, A. Cichocki, A. H. Phan, T. Tanaka,
 632 and I. Oseledets. Randomized algorithms for computation of Tucker decomposition and
 633 higher order SVD (HOSVD). *IEEE Access*, 9:28684–28706, 2021.
- 634 [2] S. Ahmadi-Asl, A. Cichocki, A. H. Phan, M. G. Asante-Mensah, M. M. Ghazani, T. Tanaka,
 635 and I. Oseledets. Randomized algorithms for fast computation of low rank tensor ring
 636 model. *Machine Learning: Science and Technology*, 2(1):011001, 2020.
- 637 [3] N. Alger, P. Chen, and O. Ghattas. Tensor train construction from tensor actions, with ap-
 638 plication to compression of large high order derivative tensors. *SIAM J. Sci. Comput.*,
 639 42(5):A3516–A3539, 2020.
- 640 [4] J. Ballani and L. Grasedyck. A projection method to solve linear systems in tensor format.
 641 *Numer. Linear Algebra Appl.*, 20(1):27–43, 2013.
- 642 [5] R. Ballester-Ripoll and R. Pajarola. Tensor decompositions for integral histogram compression
 643 and look-up. *IEEE Transactions on Visualization and Computer Graphics*, 25(2):1435–
 644 1446, 2018.
- 645 [6] K. Batselier, W. Yu, L. Daniel, and N. Wong. Computing low-rank approximations of large-
 646 scale matrices with the tensor network randomized SVD. *SIAM J. Matrix Anal. Appl.*,
 647 39(3):1221–1244, 2018.
- 648 [7] M. H. Beck, A. Jäckle, G. A. Worth, and H.-D. Meyer. The multiconfiguration time-dependent
 649 Hartree (MCTDH) method: a highly efficient algorithm for propagating wavepackets. *Phys.*
 650 *Rep.*, 324(1):1–105, 2000.
- 651 [8] G. Beylkin, J. Garcke, and M. J. Mohlenkamp. Multivariate regression and machine learning
 652 with sums of separable functions. *SIAM J. Sci. Comput.*, 31(3):1840–1857, 2009.
- 653 [9] M. Che and Y. Wei. Randomized algorithms for the approximations of Tucker and the tensor
 654 train decompositions. *Adv. Comput. Math.*, 45(1):395–428, 2019.
- 655 [10] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, and D. P. Mandic. Tensor networks
 656 for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decom-
 657 positions. *Foundations and Trends® in Machine Learning*, 9(4-5):249–429, 2016.
- 658 [11] A. Cichocki, A.-H. Phan, Q. Zhao, N. Lee, I. V. Oseledets, M. Sugiyama, and D. Mandic. Tensor
 659 Networks for Dimensionality Reduction and Large-Scale Optimizations. Part 2 Applications
 660 and Future perspectives. *arXiv preprint arXiv:1708.09165*, 2017.
- 661 [12] N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor
 662 analysis. In *Conference on Learning Theory*, pages 698–728. PMLR, 2016.
- 663 [13] H. A. Daas, G. Ballard, and P. Benner. Parallel Algorithms for Tensor Train Arithmetic. *arXiv*
 664 *preprint arXiv:2011.06532*, 2020.
- 665 [14] W. Dahmen, R. DeVore, L. Grasedyck, and E. Süli. Tensor-sparsity of solutions to high-
 666 dimensional elliptic partial differential equations. *FoCM*, 16(4):813–874, 2016.
- 667 [15] V. De Silva and L.-H. Lim. Tensor rank and the ill-posedness of the best low-rank approximation
 668 problem. *SIAM J. Matrix Anal. Appl.*, 30(3):1084–1127, 2008.
- 669 [16] S. Dolgov and D. Savostyanov. Parallel cross interpolation for high-precision calculation of
 670 high-dimensional integrals. *Comput. Phys. Commun.*, 246:106869, 2020.
- 671 [17] S. V. Dolgov. TT-GMRES: solution to a linear system in the structured tensor format. *Russ.*
 672 *J. Numer. Anal. Math. Model.*, 28(2):149–172, 2013.
- 673 [18] S. V. Dolgov, B. N. Khoromskij, and I. V. Oseledets. Fast solution of parabolic problems in the
 674 tensor train/quantized tensor train format with initial application to the Fokker–Planck
 675 equation. *SIAM J. Sci. Comput.*, 34(6):A3016–A3038, 2012.
- 676 [19] S. V. Dolgov and D. V. Savostyanov. Alternating minimal energy methods for linear systems
 677 in higher dimensions. *SIAM J. Sci. Comput.*, 36(5):A2248–A2271, 2014.
- 678 [20] D. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and
 679 R. Adams. Advances in Neural Information Processing Systems 28. *Cortes C., Lawrence*
 680 *ND, Lee DD, Sugiyama M., Garnett R., Eds*, pages 2224–2232, 2015.
- 681 [21] M. Fannes, B. Nachtergaele, and R. F. Werner. Finitely correlated states on quantum spin

- 682 chains. *Commun. Math. Phys.*, 144(3):443–490, 1992.
- 683 [22] Y. Feng, K. Tang, L. He, P. Zhou, and Q. Liao. Tensor Train random projection. *arXiv preprint*
684 *arXiv:2010.10797*, 2020.
- 685 [23] P. Gelß, S. Klus, S. Matera, and C. Schütte. Nearest-neighbor interaction systems in the
686 tensor-train format. *J. Comput. Phys.*, 341:140–162, 2017.
- 687 [24] L. Grasedyck. Existence and computation of low Kronecker-rank approximations for large linear
688 systems of tensor product structure. *Computing*, 72(3-4):247–265, 2004.
- 689 [25] L. Grasedyck, D. Kressner, and C. Tobler. A literature survey of low-rank tensor approximation
690 techniques. *GAMM-Mitteilungen*, 36(1):53–78, 2013.
- 691 [26] W. Hackbusch. *Tensor spaces and numerical tensor calculus*, volume 42. Springer, 2012.
- 692 [27] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic
693 algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–
694 288, 2011.
- 695 [28] N. J. Higham. *Accuracy and stability of numerical algorithms*. Society for Industrial and
696 Applied Mathematics (SIAM), Philadelphia, PA, second edition, 2002.
- 697 [29] F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys.*,
698 6(1-4):164–189, 1927.
- 699 [30] S. Holtz, T. Rohwedder, and R. Schneider. The alternating linear scheme for tensor optimization
700 in the tensor train format. *SIAM J. Sci. Comput.*, 34(2):A683–A713, 2012.
- 701 [31] B. Huber, R. Schneider, and S. Wolf. A randomized tensor train singular value decomposition.
702 In *Compressed Sensing and its Applications*, pages 261–290. Springer, 2017.
- 703 [32] B. N. Khoromskij. Tensors-structured numerical methods in scientific computing: Survey on
704 recent advances. *Chemom. Intell. Lab. Syst.*, 110(1):1–19, 2012.
- 705 [33] S. Klus, P. Gelß, S. Peitz, and C. Schütte. Tensor-based dynamic mode decomposition. *Non-*
706 *linearity*, 31(7):3359, 2018.
- 707 [34] S. Klus, P. Koltai, and C. Schütte. On the numerical approximation of the Perron-Frobenius
708 and Koopman operator. *arXiv preprint arXiv:1512.05997*, 2015.
- 709 [35] O. Koch and C. Lubich. Dynamical tensor approximation. *SIAM J. Matrix Anal. Appl.*,
710 31(5):2360–2375, 2010.
- 711 [36] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*,
712 51(3):455–500, 2009.
- 713 [37] D. Kressner, R. Kumar, F. Nobile, and C. Tobler. Low-rank tensor approximation for high-order
714 correlation functions of gaussian random fields. *SIAM-ASA J. Uncertain.*, 3(1):393–416,
715 2021/09/21 2015.
- 716 [38] D. Kressner and C. Tobler. Low-rank tensor Krylov subspace methods for parametrized linear
717 systems. *SIAM J. Matrix Anal. Appl.*, 32(4):1288–1316, 2021/04/09 2011.
- 718 [39] D. Loukrezis, U. Römer, T. Casper, S. Schöps, and H. De Gerssem. High-dimensional uncertainty
719 quantification for an electrothermal field problem using stochastic collocation on sparse
720 grids and tensor train decompositions. *Int. J. Numer. Model.: Electron. Netw. Devices*
721 *Fields*, 31(2):e2222, 2018.
- 722 [40] H.-D. Meyer, F. Gatti, and G. A. Worth. Basic MCTDH theory. *Multidimensional Quantum*
723 *Dynamics: MCTDH Theory and Applications*, pages 17–30, 2009.
- 724 [41] Y. Nakatsukasa. Fast and stable randomized low-rank matrix approximation. *arXiv preprint*
725 *arXiv:2009.11392*, 2020.
- 726 [42] A. Obukhov, M. Rakhuba, A. Liniger, Z. Huang, S. Georgoulis, D. Dai, and L. Van Gool.
727 Spectral tensor train parameterization of deep learning layers. In *Int. Conf. Artif. Intell.*
728 *Stat., AISTATS*, pages 3547–3555. PMLR, 2021.
- 729 [43] I. Oseledets and E. Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear*
730 *Algebra Appl.*, 432(1):70–88, 2010.
- 731 [44] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33(5):2295–2317, 2011.
- 732 [45] I. V. Oseledets and S. V. Dolgov. Solution of linear systems and matrix inversion in the TT-
733 format. *SIAM J. Sci. Comput.*, 34(5):A2718–A2739, 2012.
- 734 [46] I. V. Oseledets and E. E. Tyrtyshnikov. Breaking the curse of dimensionality, or how to use
735 SVD in many dimensions. *SIAM J. Sci. Comput.*, 31(5):3744–3759, 2009.
- 736 [47] B. Rakhshan and G. Rabusseau. Tensorized random projections. In *Int. Conf. Artif. Intell.*
737 *Stat., AISTATS*, pages 3306–3316. PMLR, 2020.
- 738 [48] H. Rauhut, R. Schneider, and Ž. Stojanac. Tensor completion in hierarchical tensor represen-
739 tations. In *Compressed Sensing and its Applications*, pages 419–450. Springer, 2015.
- 740 [49] L. Richter, L. Sallandt, and N. Nüsken. Solving high-dimensional parabolic PDEs using the
741 tensor train format. *arXiv preprint arXiv:2102.11830*, 2021.
- 742 [50] D. Savostyanov and I. Oseledets. Fast adaptive interpolation of multi-dimensional arrays in ten-
743 sor train format. In *The 2011 International Workshop on Multidimensional (nD) Systems*,

- 744 pages 1–8, 2011.
- 745 [51] U. Schollwöck. The density-matrix renormalization group. *Rev. Mod. Phys.*, 77(1):259, 2005.
- 746 [52] C. Tobler. *Low-rank Tensor Methods for Linear Systems and Eigenvalue Problems*. PhD thesis,
- 747 ETH Zurich, 2012.
- 748 [53] Z. Zhang, X. Yang, I. V. Oseledets, G. E. Karniadakis, and L. Daniel. Enabling high-
- 749 dimensional hierarchical uncertainty quantification by ANOVA and tensor-train decom-
- 750 position. *EEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 34(1):63–76, 2014.