

RANDOMIZED ALTERNATING LEAST SQUARES FOR CANONICAL TENSOR DECOMPOSITIONS: APPLICATION TO A PDE WITH RANDOM DATA*

MATTHEW J. REYNOLDS[†], ALIREZA DOOSTAN[†], AND GREGORY BEYLKIN[‡]

Abstract. This paper introduces a randomized variation of the alternating least squares (ALS) algorithm for rank reduction of canonical tensor formats. The aim is to address the potential numerical ill-conditioning of least squares matrices at each ALS iteration. The proposed algorithm, dubbed *randomized ALS*, mitigates large condition numbers via projections onto random tensors, a technique inspired by well-established randomized projection methods for solving overdetermined least squares problems in a matrix setting. A probabilistic bound on the condition numbers of the randomized ALS matrices is provided, demonstrating reductions relative to their standard counterparts. Additionally, results are provided that guarantee comparable accuracy of the randomized ALS solution at each iteration. The performance of the randomized algorithm is studied with three examples, including manufactured tensors and an elliptic PDE with random inputs. In particular, for the latter, tests illustrate not only improvements in condition numbers, but also improved accuracy of the iterative solver for the PDE solution represented in a canonical tensor format.

Key words. Separated representations, Tensor Decomposition, Randomized Projection, Alternating Least Squares, Canonical Tensors, Stochastic PDE

AMS subject classifications. 65C30, 15A69, 65D99

1. Introduction. The approximation of multivariate functions is an essential tool for numerous applications including computational chemistry [1, 24], data mining [2, 9, 24], and recently uncertainty quantification [11, 23]. For seemingly reasonable numbers of variables d , e.g. $\mathcal{O}(10)$, reconstructing a function, for instance, using its sampled values, requires computational costs that may be prohibitive. This is related to the so-called “curse of dimensionality.” To mitigate this phenomenon, we require such functions to have special *structures* that can be exploited by carefully crafted algorithms. One such structure is that the function of interest $u(z_1, z_2, \dots, z_d)$, depending on variables z_1, z_2, \dots, z_d , admits a separated representation, [3, 4, 24], of the form

$$(1) \quad u(z_1, z_2, \dots, z_d) = \sum_{l=1}^r \sigma_l u_1^l(z_1) u_2^l(z_2) \cdots u_d^l(z_d).$$

The number of terms, r , is called the separation rank of u and is assumed to be *small*. Any discretization of the univariate functions $u_j^l(z_j)$ in (1) with $u_{i_j}^l = u_j^l(z_{i_j})$, $i_j = 1, \dots, M_j$ and $j = 1, \dots, d$, leads to a Canonical Tensor Decomposition, or CTD,

$$(2) \quad \mathbf{U} = U(i_1 \dots i_d) = \sum_{l=1}^r \sigma_l u_{i_1}^l u_{i_2}^l \cdots u_{i_d}^l.$$

*This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0006402, and NSF grants DMS-1228359 and CMMI-1454601.

[†]Department of Aerospace Engineering Sciences, University of Colorado at Boulder, Boulder, CO (matthew.reynolds@colorado.edu, alireza.doostan@colorado.edu).

[‡]Department of Applied Mathematics, University of Colorado at Boulder, Boulder, CO (gregory.beylkin@colorado.edu).

The functions $u_j^l(z_j)$ in (1) and the corresponding vectors $u_{i_j}^l$ in (2) are normalized to unit norm so that the magnitude of the terms is carried by their positive s -values, σ_l . It is well understood that when the separation rank r is independent of d , the computation costs and storage requirements of standard algebraic operations in separated representations scale linearly in d , [4]. For this reason, such representations are widely used for approximating high-dimensional functions. To keep the computation of CTDs manageable, it is crucial to maintain as small as possible separation rank. Common operations involving CTDs, e.g. summations, lead to CTDs with separation ranks that may be larger than necessary. Therefore, a standard practice is to reduce the separation rank of a given CTD without sacrificing much accuracy, for which the workhorse algorithm is Alternating Least Squares (ALS) (see e.g., [3, 4, 6, 7, 21, 24, 31]). This algorithm optimizes the separated representation (in Frobenius norm) one direction at a time by solving least squares problems for each direction. The linear systems for each direction are obtained as normal equations by contracting over all tensor indices, $i = 1, \dots, d$, except those in the direction of optimization k .

It is well known that forming normal equations increases the condition number of the least squares problem, see e.g. [16]. In this paper we investigate the behavior of the condition numbers of linear systems that arise in the ALS algorithm, and propose an alternative formulation in order to avoid potential ill-conditioning. As we shall see later, the normal equations in the ALS algorithm are formed via the Hadamard (entry-wise) product of matrices for individual directions. We show that in order for the resulting matrix to be ill-conditioned, the matrices for all directions have to be ill-conditioned and obtain estimates of these condition numbers. To improve the conditioning of the linear systems, we propose a randomized version of ALS, called *randomized ALS*, where instead of contracting a tensor with itself (in all directions but one), we contract it with a tensor composed of random entries. We show that this random projection improves the conditioning of the linear systems. However, its straightforward use does not ensure monotonicity in error reduction, unlike in standard ALS. In order to restore monotonicity, we simply accept only random projections that do not increase the error.

Our interest here in using CTDs stems from the efficiency of such representations in tackling the issue of the curse of dimensionality arising from the solution of PDEs with random data, as studied in the context of Uncertainty Quantification (UQ). In the probabilistic framework, uncertainties are represented via a finite number of random variables z_j specified using, for example, available experimental data or expert opinion. An important task is to then quantify the dependence of quantities of interest $u(z_1, \dots, z_d)$ on these random inputs. For this purpose, approximation techniques based on separated representations have been recently studied in [12, 25, 26, 11, 27, 13, 18, 23, 10, 17, 19] and proven effective in reducing the issue of curse of dimensionality.

The paper is organized as follows. In [section 2](#), we introduce our notation and provide background information on tensors, the standard ALS algorithm, and the random matrix theory used in this paper. In [section 3](#), we introduce randomized ALS and provide analysis of the algorithm's convergence and the conditioning of matrices used. [Section 4](#) contains demonstrations of randomized ALS and comparisons with standard ALS on three examples. The most important of these examples provides background on uncertainty quantification and demonstrates the application of randomized ALS-based reduction as a step in finding the fixed point solution of a stochastic PDE. We conclude with a discussion on our new algorithm and future work in [section 5](#).

2. Notation and Background.

2.1. Notation. Our notation for tensors, i.e. d -directional arrays of numbers, is boldfaced uppercase letters, e.g. $\mathbf{F} \in \mathbb{R}^{M_1 \times \dots \times M_d}$. These tensors are assumed to be in the CTD format,

$$\mathbf{F} = \sum_{l=1}^{r_{\mathbf{F}}} s_l^{\mathbf{F}} \mathbf{F}_1^l \circ \dots \circ \mathbf{F}_d^l,$$

where the factors $\mathbf{F}_i^l \in \mathbb{R}^{M_i}$ are vectors with a subscript denoting the directional index and a superscript the rank index, and \circ denotes the standard vector outer product. We write operators in dimension d as $\mathbb{A} = A(j_1, j'_1; \dots; j_d, j'_d)$, while for standard matrices we use uppercase letters, e.g. $A \in \mathbb{R}^{N \times M}$. Vectors are represented using boldfaced lowercase letters, e.g. $\mathbf{c} \in \mathbb{R}^N$, while scalars are represented by lowercase letters. We perform three operations on CTDs: addition, inner product, and the application of a d -dimensional operator.

- When two CTDs are added together, all terms are joined into a single list and simply re-indexed. In such a case the separation rank is the sum of the ranks of the components, i.e. if the CTDs have ranks \tilde{r} and \hat{r} , the output CTD has rank $\tilde{r} + \hat{r}$.
- The inner product of two tensors in CTD format, $\tilde{\mathbf{F}}$ and $\hat{\mathbf{F}}$, is defined as

$$\langle \tilde{\mathbf{F}}, \hat{\mathbf{F}} \rangle = \sum_{\tilde{l}=1}^{\tilde{r}} \sum_{\hat{l}=1}^{\hat{r}} \tilde{s}_{\tilde{l}} \hat{s}_{\hat{l}} \langle \tilde{\mathbf{F}}_1^{\tilde{l}}, \hat{\mathbf{F}}_1^{\hat{l}} \rangle \dots \langle \tilde{\mathbf{F}}_d^{\tilde{l}}, \hat{\mathbf{F}}_d^{\hat{l}} \rangle,$$

where the inner product $\langle \cdot, \cdot \rangle$ operating on vectors is the standard vector dot product.

- When applying a d -dimensional operator to a tensor in CTD format, we have

$$\mathbb{A}\mathbf{F} = \sum_{\hat{l}=1}^{r_{\mathbb{A}}} \sum_{\tilde{l}=1}^{r_{\mathbf{F}}} s_{\hat{l}}^{\mathbb{A}} s_{\tilde{l}}^{\mathbf{F}} \left(\mathbb{A}_1^{\hat{l}} \mathbf{F}_1^{\tilde{l}} \right) \circ \dots \circ \left(\mathbb{A}_d^{\hat{l}} \mathbf{F}_d^{\tilde{l}} \right).$$

We use the symbol $\|\cdot\|$ to denote the standard spectral norm for matrices, as well as the Frobenius norm for tensors,

$$\|\mathbf{F}\| = \langle \mathbf{F}, \mathbf{F} \rangle^{\frac{1}{2}},$$

and $\|\cdot\|_1$ and $\|\cdot\|_2$ to denote the standard Euclidean ℓ_1 and ℓ_2 vector norms.

For analysis involving matrices we use three different types of multiplication in addition to the standard matrix multiplication. The Hadamard, or entry-wise, product of two matrices A and B is denoted by $A * B$. The Kronecker product of two matrices $A \in \mathbb{R}^{N_A \times M_A}$ and $B \in \mathbb{R}^{N_B \times M_B}$, is denoted as $A \otimes B$,

$$A \otimes B = \begin{bmatrix} A(1,1)B & \dots & A(1,M_A)B \\ \vdots & \ddots & \vdots \\ A(N_A,1)B & \dots & A(N_A,M_A)B \end{bmatrix}.$$

The final type of matrix product we use, the Khatri-Rao product of two matrices $A \in \mathbb{R}^{N_A \times M}$ and $B \in \mathbb{R}^{N_B \times M}$, is denoted by $A \odot B$,

$$A \odot B = \left[A(:,1) \otimes B(:,1) \quad A(:,2) \otimes B(:,2) \quad \dots \quad A(:,M) \otimes B(:,M) \right].$$

We also frequently use the maximal and minimal (non-zero) singular values of a matrix, denoted as σ_{\max} and σ_{\min} , respectively.

2.2. ALS algorithm. Operations on tensors in CTD format lead to an increase of the separation rank. However, this separation rank is not necessarily the smallest possible rank to represent the resulting tensor for a given accuracy. ALS is the most commonly used algorithm for finding low (near-minimal) separation rank approximations of CTDs given a user-supplied tolerance. Specifically, given a tensor \mathbf{G} in CTD format with separation rank $r_{\mathbf{G}}$,

$$\mathbf{G} = \sum_{l=1}^{r_{\mathbf{G}}} s_l^{\mathbf{G}} \mathbf{G}_1^l \circ \cdots \circ \mathbf{G}_d^l,$$

and an acceptable error ϵ , we attempt to find a representation

$$\mathbf{F} = \sum_{\tilde{l}=1}^{r_{\mathbf{F}}} s_{\tilde{l}}^{\mathbf{F}} \mathbf{F}_1^{\tilde{l}} \circ \cdots \circ \mathbf{F}_d^{\tilde{l}}$$

with lower separation rank, $r_{\mathbf{F}} < r_{\mathbf{G}}$, such that $\|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\| < \epsilon$.

The standard ALS algorithm starts from an initial guess, \mathbf{F} , with a small separation rank, e.g., $r_{\mathbf{F}} = 1$. A sequence of least squares problems in each direction is then constructed and solved to update the representation. Given a direction k , we freeze the factors in all other directions to produce a least squares problem for the factors in direction k . This process is then repeated for all directions k . One cycle through all k is called an ALS sweep. These ALS sweeps continue until the improvement in the residual $\|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\|$ either drops below a certain threshold or reaches the desired accuracy, i.e. $\|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\| < \epsilon$. If the residual is still above the target accuracy ϵ , the separation rank $r_{\mathbf{F}}$ is increased and we repeat the previous steps for constructing the representation with the new separation rank.

Specifically, as discussed in [4], the construction of the normal equations for direction k can be thought of as taking the derivatives of the Frobenius norm of $\|\mathbf{F} - \mathbf{G}\|^2$ with respect to the factors $\mathbf{F}_k^{\tilde{l}}$, $\tilde{l} = 1, \dots, r_{\mathbf{F}}$, and setting these derivatives to zero. This yields the normal equations

$$(3) \quad B_k \mathbf{c}_{j_k} = \mathbf{b}_{j_k},$$

where j_k corresponds to the j -th entry of $\mathbf{F}_k^{\tilde{l}}$ and $\mathbf{c}_{j_k} = c_{j_k}(\tilde{l})$ is a vector indexed by \tilde{l} . Alternatively, the normal system (3) can be obtained by contracting all directions except the optimization direction k , so that the matrix B_k is the Hadamard product of Gram matrices,

$$(4) \quad B_k(\hat{l}, \tilde{l}) = \prod_{i \neq k} \langle \mathbf{F}_i^{\tilde{l}}, \mathbf{F}_i^{\hat{l}} \rangle,$$

and, accordingly, the right-hand side is

$$\mathbf{b}_{j_k}(\hat{l}) = \sum_{l=1}^{r_{\mathbf{G}}} s_l^{\mathbf{G}} G_k^l(j_k) \prod_{i \neq k} \langle \mathbf{G}_i^l, \mathbf{F}_i^{\hat{l}} \rangle.$$

We solve (3) for \mathbf{c}_{j_k} and use the solution to update $\mathbf{F}_k^{\tilde{l}}$. Pseudocode for the ALS algorithm is provided in Algorithm 1, where *max_rank* and *max_iter* denote the maximum separation rank and the limit on the number of iterations (i.e. ALS sweeps). The threshold δ is used to decide if the separation rank needs to be increased.

Algorithm 1 Alternating least squares algorithm for rank reduction

input : $\epsilon > 0$, $\delta > 0$, \mathbf{G} with rank $r_{\mathbf{G}}$, max_rank , max_iter
initialize $r_{\mathbf{F}} = 1$ tensor $\mathbf{F} = \mathbf{F}_1^1 \circ \dots \circ \mathbf{F}_d^1$ with randomly generated \mathbf{F}_k^1
while $r_{\mathbf{F}} \leq max_rank$ **do**
 $iter = 1$
 if $r_{\mathbf{F}} > 1$ **then**
 add a random rank 1 contribution to \mathbf{F} : $\mathbf{F} = \mathbf{F} + \mathbf{F}_1^{r_{\mathbf{F}}} \circ \dots \circ \mathbf{F}_d^{r_{\mathbf{F}}}$
 end if
 $res = \|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\|$
 while $iter \leq max_iter$ **do**
 $res_old = res$
 for $k = 1, \dots, d$ **do**
 solve $B_k \mathbf{c}_{j_k} = \mathbf{b}_{j_k}$ for every j_k in direction k
 define $\mathbf{v}_{\tilde{l}} = (c_1(\tilde{l}), \dots, c_{M_k}(\tilde{l}))$ for $\tilde{l} = 1, \dots, r_{\mathbf{F}}$
 $s_{\tilde{l}}^{\mathbf{F}} = \|\mathbf{v}_{\tilde{l}}\|_2$ for $\tilde{l} = 1, \dots, r_{\mathbf{F}}$
 $F_k^{\tilde{l}}(j_k) = c_{j_k}(\tilde{l}) / s_{\tilde{l}}^{\mathbf{F}}$ for $\tilde{l} = 1, \dots, r_{\mathbf{F}}$
 end for
 $res = \|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\|$
 if $res < \epsilon$ **then**
 return \mathbf{F}
 else if $|res - res_old| < \delta$ **then**
 break
 else
 $iter = iter + 1$
 end if
 end while
 $r_{\mathbf{F}} = r_{\mathbf{F}} + 1$
end while
return \mathbf{F}

A potential pitfall of the ALS algorithm is poor-conditioning of the matrix B_k since the construction of normal equations squares the condition number as is well known in matrix problems. An alternative that avoids the normal equations is mentioned in the review paper [24], but it is not feasible for problems with even moderately large dimension (e.g. $d = 5$).

2.3. Estimate of condition numbers of least squares matrices. It is an empirical observation that the condition number of the matrices B_k is sometimes significantly better than the condition numbers of some of the Gram matrices comprising the Hadamard product in (4). In fact we have

LEMMA 1. *Let A and B be Gram matrices with all diagonal entries equal to 1. Then we have*

$$\sigma_{\min}(B) \leq \sigma_{\min}(A * B) \leq \sigma_{\max}(A * B) \leq \sigma_{\max}(B).$$

If the matrix B is positive definite, then

$$\kappa(A * B) \leq \kappa(B).$$

Since Gram matrices are symmetric non-negative definite, the proof of [Lemma 1](#) follows directly from [[22](#), Theorem 5.3.4]. This estimate implies that it is sufficient for only one of the matrices to be well conditioned to assure that the Hadamard product is also well conditioned. In other words, it is necessary for all directional Gram matrices to be ill-conditioned to cause the ill-conditioning of the Hadamard product. Clearly, this situation can occur and we address it in the paper.

2.4. Modification of normal equations: motivation for randomized methods. We motivate our approach by first considering an alternative to forming normal equations for ordinary matrices (excluding the QR factorization that can be easily used for matrices). Given a matrix $A \in \mathbb{R}^{N \times n}$, $N \geq n$, we can multiply $A\mathbf{x} = \mathbf{b}$ by a matrix $R \in \mathbb{R}^{n' \times N}$ with independent random entries and then solve

$$(5) \quad R A \mathbf{x} = R \mathbf{b},$$

instead (see, e.g. [[20](#), [29](#), [30](#), [33](#)]). The solution of this system, given that R is of appropriate size (i.e., n' is large enough), will be close to the least squares solution [[29](#), Lemma 2]. In [[29](#)], (5) is used to form a preconditioner and an initial guess for solving $\min \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ via a preconditioned conjugate gradient method. However, for our application we are interested in using equations of the form (5) in the Hadamard product in (4). We observe that RA typically has a smaller condition number than $A^T A$. To see why, recall that for full-rank, square matrices A and B , a bound on the condition number is

$$\kappa(AB) \leq \kappa(A)\kappa(B).$$

However, for rectangular full-rank matrices $A \in \mathbb{R}^{r' \times N}$ and $B \in \mathbb{R}^{N \times r}$, $r \leq r' \leq N$, this inequality does not necessarily hold. Instead, we have the inequality

$$(6) \quad \kappa(AB) \leq \kappa(A) \frac{\sigma_1(B)}{\sigma_{\min}(P_{A^T}(B))},$$

where $P_{A^T}(B)$ is the projection of B onto the row space of A (for a proof of this inequality, see [Appendix A](#)). If A has a small condition number (for example, when A is a Gaussian random matrix, see [[8](#), [14](#), [15](#)]) and we were to assume $\sigma_{\min}(P_{A^T}(B))$ is close to $\sigma_{\min}(B)$, we obtain condition numbers smaller than $\kappa^2(B)$. The assumption that $\sigma_{\min}(P_{A^T}(B))$ is close to $\sigma_{\min}(B)$ is the same as assuming the columns of B lie within the subspace spanned by the row of A . This is achieved by choosing r' to be larger than r when A is a randomized matrix.

2.5. Definitions and random matrix theory. The main advantage of our approach is an improved condition number for the linear system solved at every step of the ALS algorithm. We use a particular type of random matrices to derive bounds on the condition number: the rows are independently distributed random vectors, but the columns are not (instead of the standard case where all entries are i.i.d). Such matrices were studied extensively by Vershynin [[32](#)] and we rely heavily on this work for our estimates. To proceed, we need the following definitions from [[32](#)].

Remark 2. Definitions involving random variables, and vectors composed of random variables, are not consistent with the notation of the rest of the paper, outlined in [subsection 2.1](#).

DEFINITION 3. [[32](#), Definition 5.7] Let $\mathbb{P}\{\cdot\}$ denote the probability of a set and \mathbb{E} the mathematical expectation operator. Also, let X be a random variable that satisfies

one of the three following equivalent properties,

1. $\mathbb{P}\{|X| > t\} \leq \exp(1 - t^2/K_1^2)$ for all $t \geq 0$
2. $(\mathbb{E}|X|^p)^{1/p} \leq K_2\sqrt{p}$ for all $p \geq 1$
3. $\mathbb{E} \exp(X^2/K_3^2) \leq e$,

where the constants K_i , $i = 1, 2, 3$, differ from each other by at most an absolute constant factor (see [32, Lemma 5.5] for a proof of the equivalence of these properties). Then X is called a sub-Gaussian random variable. The sub-Gaussian norm of X is defined as the smallest K_2 in property 2, i.e.,

$$\|X\|_{\psi_2} = \sup_{p \geq 1} \frac{(\mathbb{E}|X|^p)^{1/p}}{\sqrt{p}}.$$

Examples of sub-Gaussian random variables include Gaussian and Bernoulli random variables. We also present definitions for sub-Gaussian random vectors and their norm.

DEFINITION 4. [32, Definition 5.7] A random vector $X \in \mathbb{R}^n$ is called a sub-Gaussian random vector if $\langle X, \mathbf{x} \rangle$ is a sub-Gaussian random variable for all $\mathbf{x} \in \mathbb{R}^n$. The sub-Gaussian norm of X is subsequently defined as

$$\|X\|_{\psi_2} = \sup_{\mathbf{x} \in \mathcal{S}^{n-1}} \|\langle X, \mathbf{x} \rangle\|_{\psi_2},$$

where \mathcal{S}^{n-1} is the unit Euclidean sphere.

DEFINITION 5. [32, Definition 5.19] A random vector $X \in \mathbb{R}^n$ is called isotropic if its second moment matrix, $\Sigma = \Sigma(X) = \mathbb{E}[XX^T]$, is equal to identity, i.e. $\Sigma(X) = I$. This definition is equivalent to

$$\mathbb{E} \langle X, \mathbf{x} \rangle^2 = \|\mathbf{x}\|_2^2 \quad \text{for all } \mathbf{x} \in \mathbb{R}^n.$$

The following theorem from [32] provides bounds on the condition numbers of matrices whose rows are independent sub-Gaussian isotropic random variables.

THEOREM 6. [32, Theorem 5.38] Let A be an $N \times n$ matrix whose rows $A(i, :)$ are independent, sub-Gaussian isotropic random vectors in \mathbb{R}^n . Then for every $t \geq 0$, with probability at least $1 - 2 \exp(-ct^2)$, one has

$$(7) \quad \sqrt{N} - C\sqrt{n} - t \leq \sigma_{\min}(A) \leq \sigma_{\max}(A) \leq \sqrt{N} + C\sqrt{n} + t.$$

Here $C = C_K$, $c = c_K > 0$, depend only on the sub-Gaussian norm $K = \max_i \|A(i, :)\|_{\psi_2}$. ■

An outline of the proof of [Theorem 6](#) will be useful for deriving our own results, so we provide a sketch in [Appendix A](#). The following lemma is used to prove [Theorem 6](#), and will also be useful later on in the paper. We later modify it to prove a version of [Theorem 6](#) that works for sub-Gaussian, non-isotropic random vectors.

LEMMA 7. [32, Lemma 5.36] Consider a matrix B that satisfies

$$\|B^T B - I\| < \max(\delta, \delta^2)$$

for some $\delta > 0$. Then

$$(8) \quad 1 - \delta \leq \sigma_{\min}(B) \leq \sigma_{\max}(B) \leq 1 + \delta.$$

Conversely, if B satisfies (8) for some $\delta > 0$, then $\|B^T B - I\| < 3 \max(\delta, \delta^2)$.

3. Randomized ALS algorithm.

3.1. Alternating least squares algorithm using random matrices. We propose the following alternative to using the normal equations in ALS algorithms: instead of (4), define the entries of B_k via randomized projections,

$$(9) \quad B_k(\hat{l}, \tilde{l}) = \prod_{i \neq k} \langle \mathbf{F}_i^{\tilde{l}}, \mathbf{R}_i^{\hat{l}} \rangle,$$

where $\mathbf{R}_i^{\hat{l}}$ is the \hat{l} -th column of a matrix $R_i \in \mathbb{R}^{M_i \times r'}$, $r' > r$, with random entries corresponding to direction i . The choice of $r' > r$ is made to reduce the condition number of B_k . As shown in subsection 3.3, as $r/r' \rightarrow 0$ the bound on $\kappa(B_k)$ goes to $\kappa(B_k) \leq \kappa\left((B_k^{\text{ALS}})^{\frac{1}{2}}\right)$, where B_k^{ALS} is the B_k matrix for standard ALS, i.e. (4). In this paper we consider independent signed Bernoulli random variables, i.e., $R_i(j_i, \hat{l})$ is either -1 or 1 each with probability $1/2$. The proposed change also alters the right-hand side of the normal equations (3),

$$(10) \quad \mathbf{b}_{j_k}(\hat{l}) = \sum_{l=1}^{r_G} s_l^G G_k^l(j_k) \prod_{i \neq k} \langle \mathbf{G}_i^l, \mathbf{R}_i^{\hat{l}} \rangle.$$

Equivalently, B_k may be written as

$$B_k = \prod_{i \neq k} R_i^T F_i.$$

Looking ahead, we choose random matrices R_i such that B_k is a tall, rectangular matrix. Solving the linear system (3) with rectangular B_k will require a pseudo-inverse, computed via either the singular value decomposition (SVD) or a QR algorithm.

To further contrast the randomized ALS algorithm with the standard ALS algorithm, we highlight two differences: firstly, the randomized ALS trades the monotonic reduction of approximation error (a property of the standard ALS algorithm) for better conditioning. To adjust we use a simple tactic: if a randomized ALS sweep (over all directions) decreases the error, we keep the resulting approximation. Otherwise, we discard the sweep, generate independent random matrices R_i , and rerun the sweep. Secondly, the randomized ALS algorithm can be more computationally expensive than the standard one. This is due to the rejection scheme outlined above and the fact that B_k in the randomized algorithm has a larger number of rows than its standard counterpart, i.e., $r' > r$. Pseudocode of our new algorithm is presented in Algorithm 2.

Remark 8. We have explored an alternative approach using projections onto random tensors, different from Algorithm 2. Instead of using B_k in (9) to solve for \mathbf{c}_{j_k} , we use the QR factorization of B_k to form a preconditioner matrix, similar to the approach of [29] for solving overdetermined least squares problems in a matrix setting. This preconditioner is used to improve the condition number of B_k in (4). The approach is different from Algorithm 2: we solve the same equations as the standard ALS algorithm, but in a better conditioned manner. Solving the same equations preserves the monotone error reduction property of standard ALS. With Algorithm 2 the equations we solve are different, but, as shown in subsection 3.2, the solutions of each least squares problem are close to the those obtained by the standard ALS algorithm.

Algorithm 2 Randomized alternating least squares algorithm for rank reduction

input : $\epsilon > 0$, \mathbf{G} with rank $r_{\mathbf{G}}$, max_tries , max_rank , max_iter
initialize $r_{\mathbf{F}} = 1$ tensor $\mathbf{F} = \mathbf{F}_1^1 \circ \dots \circ \mathbf{F}_d^1$ with randomly generated \mathbf{F}_k^1
while $r_{\mathbf{F}} \leq max_rank$ **do**
 $tries = 1$
 $iter = 1$
 construct randomized tensor \mathbf{R}
 if $r_{\mathbf{F}} > 1$ **then**
 add a random rank 1 contribution to \mathbf{F} : $\mathbf{F} = \mathbf{F} + \mathbf{F}_1^{r_{\mathbf{F}}} \circ \dots \circ \mathbf{F}_d^{r_{\mathbf{F}}}$
 end if
 while $iter \leq max_iter$ and $tries \leq max_tries$ **do**
 $\mathbf{F}_{old} = \mathbf{F}$
 for $k = 1, \dots, d$ **do**
 construct B_k , using (9)
 solve $B_k \mathbf{c}_{j_k} = \mathbf{b}_{j_k}$ for every j_k in direction k
 define $\mathbf{v}_{\tilde{l}} = (c_1(\tilde{l}), \dots, c_{M_k}(\tilde{l}))$ for $\tilde{l} = 1, \dots, r_{\mathbf{F}}$
 $s_{\tilde{l}}^{\mathbf{F}} = \|\mathbf{v}_{\tilde{l}}\|_2$ for $\tilde{l} = 1, \dots, r_{\mathbf{F}}$
 $F_k^{\tilde{l}}(j_k) = c_{j_k}(\tilde{l})/s_{\tilde{l}}^{\mathbf{F}}$ for $\tilde{l} = 1, \dots, r_{\mathbf{F}}$
 end for
 if $\|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\| < \epsilon$ **then**
 return \mathbf{F}
 else if $\|\mathbf{F}_{old} - \mathbf{G}\| / \|\mathbf{G}\| < \|\mathbf{F} - \mathbf{G}\| / \|\mathbf{G}\|$ **then**
 $\mathbf{F} = \mathbf{F}_{old}$
 $tries = tries + 1$
 $iter = iter + 1$
 else
 $tries = 1$
 $iter = iter + 1$
 end if
 end while
 $r_{\mathbf{F}} = r_{\mathbf{F}} + 1$
end while
return \mathbf{F}

Remark 9. A possible application of [Algorithm 2](#) is to use it in concert with [Algorithm 1](#). If (4) becomes poorly conditioned during standard ALS iterations, future iterations can be performed using randomized ALS sweeps from [Algorithm 2](#).

We provide convergence results and theoretical bounds on the condition number of B_k with entries (9) in [subsections 3.2](#) and [3.3](#), respectively. Additionally, in [section 4](#), we empirically demonstrate the superior conditioning properties of B_k defined in (9) relative to those given by the standard ALS in (4).

3.2. Convergence of the randomized ALS algorithm. Before deriving bounds on the condition number of (9), it is important to discuss the convergence properties of our algorithm. To do so for our tensor algorithm, we derive a convergence result similar to [33, Lemma 4.8]. In this analysis, we flatten our tensors into large matrices and use results from random matrix theory to show convergence. First, we construct the large matrices used in this section from (9). Writing the inner product as a sum

allows us to group all the summations together,

We have

$$\begin{aligned} B_k(\hat{l}, \tilde{l}) &= \prod_{i \neq k} \sum_{j_i=1}^{M_i} F_i^{\tilde{l}}(j_i) R_i^{\hat{l}}(j_i) \\ &= \sum_{j_1=1}^{M_1} \cdots \sum_{j_{k-1}=1}^{M_{k-1}} \sum_{j_{k+1}=1}^{M_{k+1}} \cdots \sum_{j_d=1}^{M_d} \left(F_1^{\tilde{l}}(j_1) \cdots \right) \left(R_1^{\hat{l}}(j_1) \cdots \right), \end{aligned}$$

where we have expanded the product to get the sum of the products of individual entries. Introducing a multi-index $j = (j_1, \dots, j_{k-1}, j_{k+1}, \dots, j_d)$, we define two matrices, $A_k \in \mathbb{R}^{M \times r}$ and $R_k \in \mathbb{R}^{M \times r'}$, where $M = \prod_{i \neq k} M_i$ is large, i.e., we write

$$\begin{aligned} A_k(j, \tilde{l}) &= F_1^{\tilde{l}}(j_1) \cdots F_{k-1}^{\tilde{l}}(j_{k-1}) F_{k+1}^{\tilde{l}}(j_{k+1}) \cdots F_d^{\tilde{l}}(j_d) \\ S_k(j, \hat{l}) &= R_1^{\hat{l}}(j_1) \cdots R_{k-1}^{\hat{l}}(j_{k-1}) R_{k+1}^{\hat{l}}(j_{k+1}) \cdots R_d^{\hat{l}}(j_d). \end{aligned}$$

We note that these matrices can also be written as Khatri-Rao products,

$$\begin{aligned} (11) \quad A_k &= F_1 \odot \cdots \odot F_{k-1} \odot F_{k+1} \odot \cdots \odot F_d \\ S_k &= R_1 \odot \cdots \odot R_{k-1} \odot R_{k+1} \odot \cdots \odot R_d. \end{aligned}$$

Since M is large, $M \gg r' > r$, both A_k and S_k are rectangular matrices. Similarly, we rewrite a vector \mathbf{b} in (10),

$$\mathbf{b}_{j_k}(\hat{l}) = \sum_{l=1}^{r_{\mathbf{G}}} s_l^{\mathbf{G}} G_k^l(j_k) \sum_{j_1=1}^{M_1} \cdots \sum_{j_{k-1}=1}^{M_{k-1}} \sum_{j_{k+1}=1}^{M_{k+1}} \cdots \sum_{j_d=1}^{M_d} \left(G_1^l(j_1) \cdots \right) \left(R_1^{\hat{l}}(j_1) \cdots \right),$$

using the multi-index j as

$$\mathbf{b}_k(j) = \sum_{l=1}^{r_{\mathbf{G}}} s_l^{\mathbf{G}} G_k^l(j_k) \left(G_1^l(j_1) \cdots G_{k-1}^l(j_{k-1}) G_{k+1}^l(j_{k+1}) \cdots G_d^l(j_d) \right).$$

Using the introduced notation, A_k , S_k , and \mathbf{b}_k , we rewrite the normal equations (3) for direction k and coordinate j_k as

$$(12) \quad A_k^T A_k \mathbf{c}_k = A_k^T \mathbf{b}_k,$$

and the randomized version of those equations as

$$(13) \quad S_k^T A_k \mathbf{c}_k = S_k^T \mathbf{b}_k.$$

We highlight the notable difference between the random matrix S_k above and those found in the usual matrix settings, for instance, in randomized least squares regression [20, 29, 30, 33]. Specifically, in the former, the entries of S_k are not statistically independent and are products of random variables, whereas in the latter the entries are often i.i.d realizations of single random variables.

Next, we present a convergence result showing that the solution to the least squares problem at each iteration of randomized ALS is close to the solution we would get using standard ALS.

LEMMA 10. *Given arbitrary A , S , and \mathbf{b} such that $A \in \mathbb{R}^{M \times r}$, $S \in \mathbb{R}^{M \times r'}$, $\mathbf{b} \in \mathbb{R}^M$, and $r < r' \leq M$, and assuming that $\mathbf{x} \in \mathbb{R}^r$ is the solution that minimizes $\|S^T A \hat{\mathbf{x}} - S^T \mathbf{b}\|_2$ and $\mathbf{y} \in \mathbb{R}^{r'}$ is the solution that minimizes $\|A \hat{\mathbf{y}} - \mathbf{b}\|_2$, then*

$$(14) \quad \|A\mathbf{x} - \mathbf{b}\|_2 \leq \kappa(S^T Q) \|A\mathbf{y} - \mathbf{b}\|_2,$$

where $Q \in \mathbb{R}^{M \times r_Q}$, $r_Q \leq r + 1$, is a matrix with orthonormal columns from the QR factorization of the augmented matrix $[A \mid \mathbf{b}]$ and where $S^T Q$ is assumed to have full rank.

Proof. We form the augmented matrix $[A \mid \mathbf{b}]$ and find its QR decomposition, $[A \mid \mathbf{b}] = QT$, where $T = [T_A \mid T_{\mathbf{b}}]$, $T_A \in \mathbb{R}^{r_Q \times r}$ and $T_{\mathbf{b}} \in \mathbb{R}^{r_Q}$, and $Q \in \mathbb{R}^{M \times r_Q}$ has orthonormal columns. Therefore, we have

$$\begin{aligned} A &= QT_A \\ \mathbf{b} &= QT_{\mathbf{b}}. \end{aligned}$$

Using these decompositions of A and \mathbf{b} , we define a matrix Θ such that

$$\begin{aligned} \Theta S^T A &= A \\ \Theta S^T \mathbf{b} &= \mathbf{b}, \end{aligned}$$

and arrive at

$$\Theta = Q \left((S^T Q)^T (S^T Q) \right)^{-1} (S^T Q)^T, \quad \square$$

where $(S^T Q)^T (S^T Q)$ has full rank because $\text{rank}(S^T Q) = r_Q$.

Starting from the left-hand side of (14), we have

$$\begin{aligned} \|A\mathbf{x} - \mathbf{b}\|_2 &= \|\Theta S^T A \mathbf{x} - \Theta S^T \mathbf{b}\|_2 \\ &\leq \|\Theta\| \|S^T A \mathbf{x} - S^T \mathbf{b}\|_2 \\ &\leq \|\Theta\| \|S^T A \mathbf{y} - S^T \mathbf{b}\|_2. \end{aligned}$$

Since multiplication by A maps a vector to the column space of A , there exists $T_{\mathbf{y}} \in \mathbb{R}^{r_Q}$ such that $A\mathbf{y} = QT_{\mathbf{y}}$. Hence, we obtain

$$\begin{aligned} \|A\mathbf{x} - \mathbf{b}\|_2 &\leq \|\Theta\| \|S^T QT_{\mathbf{y}} - S^T QT_{\mathbf{b}}\|_2 \\ &\leq \|\Theta\| \|S^T Q\| \|T_{\mathbf{y}} - T_{\mathbf{b}}\|_2 \\ &\leq \|\Theta\| \|S^T Q\| \|A\mathbf{y} - \mathbf{b}\|_2, \end{aligned}$$

where in the last step we used the orthonormality of the columns of Q .

Next we estimate norms, $\|\Theta\|$ and $\|S^T Q\|$. First, we decompose $S^T Q$ using the singular value decomposition, $S^T Q = U \Sigma V^T$. From the definition of the spectral norm we know $\|S^T Q\| = \sigma_{\max}(S^T Q)$. Using the SVD of $S^T Q$ and the definition of Θ we write

$$\begin{aligned} \|\Theta\| &= \left\| Q \left((S^T Q)^T (S^T Q) \right)^{-1} (S^T Q)^T \right\| \\ &= \left\| (V \Sigma^2 V^T)^{-1} V \Sigma U^T \right\| \\ &= \|V \Sigma^{-2} V^T V \Sigma U^T\| \\ &= \|V \Sigma^{-1} U^T\|. \end{aligned}$$

Hence $\|\Theta\| = 1/\sigma_{\min}(S^T Q)$, and the bound is

$$\begin{aligned} \|\mathbf{Ax} - \mathbf{b}\|_2 &\leq \sigma_{\max}(S^T Q) / \sigma_{\min}(S^T Q) \|\mathbf{Ay} - \mathbf{b}\|_2 \\ &\leq \kappa(S^T Q) \|\mathbf{Ay} - \mathbf{b}\|_2. \end{aligned}$$

We later substitute $S = S_k$ in [Lemma 10](#) and use results from [\[32\]](#) to bound $\kappa(S_k^T Q)$, since $S_k^T Q$ is a random matrix whose rows are independent from one another but whose columns are not. To use this machinery, specifically [Theorem 6](#), we require the following lemma.

LEMMA 11. *$S_k^T Q$, where $Q \in \mathbb{R}^{M \times r_Q}$ has orthonormal columns, $S_k \in \mathbb{R}^{M \times r'}$ is defined in [\(11\)](#), and $r_Q < r' \leq M$, is a random matrix with isotropic rows.*

Proof. Using the second moment matrix, we show that the rows of $S_k^T Q$ are isotropic. Given a row of $S_k^T Q$ written in column form, $\left[S_k(:, \hat{l})^T Q \right]^T = Q^T S_k(:, \hat{l})$, we form the second moment matrix,

$$\mathbb{E} \left[Q^T S_k(:, \hat{l}) S_k(\hat{l}, :)^T Q \right] = Q^T \mathbb{E} \left[S_k(:, \hat{l}) S_k(\hat{l}, :)^T \right] Q.$$

and show

$$\mathbb{E} \left[S_k(:, \hat{l}) S_k(\hat{l}, :)^T \right] = I_{M \times M}.$$

Hence $\mathbb{E} \left[Q^T S_k(:, \hat{l}) S_k(\hat{l}, :)^T Q \right] = Q^T Q = I_{r_Q \times r_Q}$ and $S_k^T Q$ is isotropic.

From the Khatri-Rao product definition of the matrix S_k [\(11\)](#), we write a column of S_k as

$$S_k(:, \hat{l}) = \bigotimes_{\substack{i=1:d \\ i \neq k}} R_i(:, \hat{l}).$$

Therefore, using properties of the Kronecker product (see, e.g. [\[24, equation \(2.2\)\]](#)) we can switch the order of the regular matrix product and the Kronecker products,

$$S_k(:, \hat{l}) S_k(\hat{l}, :)^T = \bigotimes_{\substack{i=1:d \\ i \neq k}} R_i(:, \hat{l}) R_i(\hat{l}, :)^T.$$

□

Taking the expectation and moving it inside the Kronecker product gives us

$$\begin{aligned} \mathbb{E} \left[S_k(:, \hat{l}) S_k(\hat{l}, :)^T \right] &= \bigotimes_{\substack{i=1:d \\ i \neq k}} \mathbb{E} \left[R_i(:, \hat{l}) R_i(\hat{l}, :)^T \right] \\ &= \bigotimes_{\substack{i=1:d \\ i \neq k}} I_{M_i \times M_i} \\ &= I_{M \times M}. \end{aligned}$$

Since $S_k^T Q$ is a tall rectangular matrix ($S_k^T Q \in \mathbb{R}^{r' \times r_Q}$) with independent sub-Gaussian isotropic rows, we may use Theorem 5.39 from Vershynin to bound the extreme singular values.

LEMMA 12. For $S_k^T Q$ defined in Lemma 11, and every $t \geq 0$, with probability at least $1 - 2 \exp(-ct^2)$ we have

$$(15) \quad \kappa(S_k^T Q) \leq \frac{1 + C\sqrt{(r+1)/r'} + t/\sqrt{r'}}{1 - C\sqrt{(r+1)/r'} - t/\sqrt{r'}},$$

where $C = C_K$ and $c = c_K > 0$ depend only on the sub-Gaussian norm $K = \max_i \|S_k(:, i)^T Q\|_{\psi_2}$ of the rows of $S_k^T Q$.

Proof. Using Lemma 11 and Theorem 6, we have the following bound on the extreme condition numbers of $S_k^T Q \in \mathbb{R}^{r' \times r_Q}$ for every $t \geq 0$, with probability at least $1 - 2 \exp(-ct^2)$,

$$\sqrt{r'} - C\sqrt{r_Q} - t \leq \sigma_{\min}(S_k^T Q) \leq \sigma_{\max}(S_k^T Q) \leq \sqrt{r'} + C\sqrt{r_Q} + t,$$

where $C = C_K$ and $c = c_K > 0$ depend only on the sub-Gaussian norm $K = \max_i \|(S_k^T Q)_i\|_{\psi_2}$ of the rows of $S_k^T Q$. Since $r_Q \leq r + 1$, we have

$$\sqrt{r'} - C\sqrt{r+1} - t \leq \sigma_{\min}(S_k^T Q) \leq \sigma_{\max}(S_k^T Q) \leq \sqrt{r'} + C\sqrt{r+1} + t,$$

with the same probability. \square

We now state the convergence result.

THEOREM 13. Given $A \in \mathbb{R}^{M \times r}$ and $S_k \in \mathbb{R}^{M \times r'}$ defined in (11), where $r < r' \leq M$, and assuming that $\mathbf{x} \in \mathbb{R}^r$ is the solution that minimizes $\|S_k^T A \hat{\mathbf{x}} - S_k^T \mathbf{b}\|_2$ and $\mathbf{y} \in \mathbb{R}^r$ is the solution that minimizes $\|A \hat{\mathbf{y}} - \mathbf{b}\|_2$, then for every $t \geq 0$, with probability at least $1 - 2 \exp(-ct^2)$ we have

$$\|A \mathbf{x} - \mathbf{b}\| \leq \frac{1 + C\sqrt{(r+1)/r'} + t/\sqrt{r'}}{1 - C\sqrt{(r+1)/r'} - t/\sqrt{r'}} \|A \mathbf{y} - \mathbf{b}\|,$$

where $C = C_K$ and $c = c_K > 0$ depend only on the sub-Gaussian norm $K = \max_i \|S_k(:, i)^T Q\|_{\psi_2}$. The matrix $Q \in \mathbb{R}^{M \times r_Q}$, $r_Q \leq r + 1$, is composed of orthonormal columns from the QR factorization of the augmented matrix $[A | \mathbf{b}]$ and $S_k^T Q$ is assumed to have full rank.

Proof. The proof consists of substituting $S = S_k$ in Lemma 10 and then bounding $\kappa(S_k^T Q)$ with Lemma 12 via Lemma 11. \square

Remark 14. The convergence properties communicated by Lemma 10 and Theorem 13 rely on the residual $\|A \mathbf{y} - \mathbf{b}\|$ being small. If it is not small, i.e., the original tensor does not admit a low-rank separated representation, then Algorithm 2 will converge slowly or not at all.

3.3. Bounding the condition number of B_k . To bound the condition number of B_k , we use a modified version of Theorem 6. If the rows of B_k were isotropic then we could use Theorem 6 directly. However, unlike $S_k^T Q$ in Lemma 11, this is not the case for B_k . While the second moment matrix Σ is not the identity, it

does play a special role in the bound of the condition number since it is the matrix B_k^{ALS} from the standard ALS algorithm. To see this, we take the \hat{l} -th row of B_k , $B_k(\hat{l}, \cdot) = \left\{ \prod_{i \neq k} \langle \mathbf{F}_i^{\hat{l}}, \mathbf{R}_i^{\hat{l}} \rangle \right\}_{\hat{l}=1, \dots, n}$, and form the second moment matrix of X ,

$$\begin{aligned} \Sigma(l, l') &= \mathbb{E} \left[\prod_{i \neq k} \langle \mathbf{F}_i^l, \mathbf{R}_i^{\hat{l}} \rangle \langle \mathbf{F}_i^{l'}, \mathbf{R}_i^{\hat{l}} \rangle \right] \\ &= \prod_{i \neq k} \mathbb{E} \left[(\mathbf{F}_i^l)^T \mathbf{R}_i^{\hat{l}} (\mathbf{R}_i^{\hat{l}})^T \mathbf{F}_i^{l'} \right] \\ &= \prod_{i \neq k} (\mathbf{F}_i^l)^T \mathbb{E} \left[\mathbf{R}_i^{\hat{l}} (\mathbf{R}_i^{\hat{l}})^T \right] \mathbf{F}_i^{l'}. \end{aligned}$$

Since $\mathbf{R}_i^{\hat{l}}$ is a vector composed of either Bernoulli or standard Gaussian random variables, $\mathbb{E} \left[\mathbf{R}_i^{\hat{l}} (\mathbf{R}_i^{\hat{l}})^T \right] = I$. Therefore, we are left with $\Sigma(l, l') = B_k^{\text{ALS}}(l, l') = \prod_{i \neq k} (\mathbf{F}_i^l)^T \mathbf{F}_i^{l'}$.

We need to modify [Theorem 6](#) for matrices that have independent, non-isotropic rows. In [\[32, Remark 5.40\]](#) it is noted in the case of a random matrix $A \in \mathbb{R}^{N \times n}$ with non-isotropic rows that we can apply [Theorem 6](#) to $A\Sigma^{-\frac{1}{2}}$ instead of A , where Σ is the second moment matrix of A . The matrix $A\Sigma^{-\frac{1}{2}}$ has isotropic rows and, thus, we obtain the following inequality that holds with probability at least $1 - 2 \exp(-ct^2)$,

$$(16) \quad \left\| \frac{1}{N} A^T A - \Sigma \right\| \leq \max(\delta, \delta^2) \|\Sigma\|, \quad \text{where } \delta = C \sqrt{\frac{n}{N}} + \frac{t}{\sqrt{N}},$$

and $C = C_K$, $c = c_K > 0$.

To clarify how [\(16\)](#) changes the bounds on the singular values $\sigma_{\min}(A)$ and $\sigma_{\max}(A)$ of a matrix A with non-isotropic rows, we modify [Lemma 7](#).

LEMMA 15. *Consider matrices $B \in \mathbb{R}^{N \times n}$ and $\Sigma^{-\frac{1}{2}} \in \mathbb{R}^{n \times n}$ (non-singular) that satisfy*

$$(17) \quad 1 - \delta \leq \sigma_{\min}(B\Sigma^{-\frac{1}{2}}) \leq \sigma_{\max}(B\Sigma^{-\frac{1}{2}}) \leq 1 + \delta,$$

for $\delta > 0$. Then we have the following bounds on the extreme singular values of B :

$$\sigma_{\min}\left(\Sigma^{\frac{1}{2}}\right) \cdot (1 - \delta) \leq \sigma_{\min}(B) \leq \sigma_{\max}(B) \leq \sigma_{\max}\left(\Sigma^{\frac{1}{2}}\right) \cdot (1 + \delta).$$

The proof of [Lemma 15](#) is included in [Appendix A](#). Using [Lemma 15](#), we observe that the bound on the condition number of a matrix B satisfying [\(17\)](#) has the following form:

$$\kappa(B) \leq \frac{(1 + \delta)}{(1 - \delta)} \kappa\left(\Sigma^{\frac{1}{2}}\right).$$

Using [Lemma 15](#), we prove an extension of [Theorem 6](#) for matrices with non-isotropic rows.

THEOREM 16. *Let A be an $N \times n$ matrix whose rows, $A(i, \cdot)$, are independent, sub-Gaussian random vectors in \mathbb{R}^n . Then for every $t \geq 0$, with probability at least $1 - 2 \exp(-ct^2)$ one has*

$$\sigma_{\min}\left(\Sigma^{\frac{1}{2}}\right) \cdot (\sqrt{N} - C\sqrt{n} - t) \leq \sigma_{\min}(A) \leq \sigma_{\max}(A) \leq \sigma_{\max}\left(\Sigma^{\frac{1}{2}}\right) \cdot (\sqrt{N} + C\sqrt{n} + t). \blacksquare$$

Here $C = C_K$, $c = c_K > 0$, depend only on the sub-Gaussian norm $K = \max_i \|A(i, :)\|_{\psi_2}$ and the norm of $\Sigma^{-\frac{1}{2}}$.

Proof. We form the second moment matrix Σ using rows $A(i, :)$ and apply [Theorem 6](#) to the matrix $A\Sigma^{-\frac{1}{2}}$, which has isotropic rows. Therefore, for every $t \geq 0$, with probability at least $1 - 2 \exp(-ct^2)$, we have

$$(18) \quad \sqrt{N} - C\sqrt{n} - t \leq \sigma_{\min} \left(A\Sigma^{-\frac{1}{2}} \right) \leq \sigma_{\max} \left(A\Sigma^{-\frac{1}{2}} \right) \leq \sqrt{N} + C\sqrt{n} + t,$$

where $C = C_{\tilde{K}}$, $c = c_{\tilde{K}} > 0$, depend only on the sub-Gaussian norm $\tilde{K} = \max_i \left\| \Sigma^{-\frac{1}{2}} A(i, :)^T \right\|_{\psi_2}$.

Applying [Lemma 15](#) to (18) with $B = A/\sqrt{N}$ and $\delta = C\sqrt{n/N} + t/\sqrt{N}$, results in the bound

$$\sigma_{\min} \left(\Sigma^{\frac{1}{2}} \right) \cdot (\sqrt{N} - C\sqrt{n} - t) \leq \sigma_{\min}(A) \leq \sigma_{\max}(A) \leq \sigma_{\max} \left(\Sigma^{\frac{1}{2}} \right) \cdot (\sqrt{N} + C\sqrt{n} + t),$$

with the same probability as (18).

To move $\Sigma^{-\frac{1}{2}}$ outside the sub-Gaussian norm, we bound \tilde{K} from above using the sub-Gaussian norm of A , $K = \max_i \|A(i, :)\|_{\psi_2}$,

$$\begin{aligned} \left\| \Sigma^{-\frac{1}{2}} A(i, :)^T \right\|_{\psi_2} &= \sup_{x \in \mathcal{S}^{n-1}} \left\| \left\langle \Sigma^{-\frac{1}{2}} A(i, :)^T, x \right\rangle \right\|_{\psi_2} \\ &= \sup_{x \in \mathcal{S}^{n-1}} \frac{\left\| \left\langle A(i, :)^T, \Sigma^{-\frac{1}{2}} x \right\rangle \right\|_{\psi_2}}{\left\| \Sigma^{-\frac{1}{2}} x \right\|_2} \left\| \Sigma^{-\frac{1}{2}} x \right\|_2 \\ &\leq \sup_{y \in \mathcal{S}^{n-1}} \left\| \left\langle A(i, :)^T, y \right\rangle \right\|_{\psi_2} \sup_{x \in \mathcal{S}^{n-1}} \left\| \Sigma^{-\frac{1}{2}} x \right\|_2 \\ &= \left\| A(i, :)^T \right\|_{\psi_2} \left\| \Sigma^{-\frac{1}{2}} \right\|, \end{aligned}$$

hence $\tilde{K} \leq K \left\| \Sigma^{-\frac{1}{2}} \right\|$. Using this inequality, we bound the probability in (31) for the case of [Theorem 6](#) applied to $A\Sigma^{-\frac{1}{2}}$.

$$\begin{aligned} \mathbb{P} \left\{ \max_{x \in \mathcal{N}} \left| \frac{1}{N} \left\| A\Sigma^{-\frac{1}{2}} x \right\|_2^2 - 1 \right| \geq \frac{\epsilon}{2} \right\} &\leq 9^n \cdot 2 \exp \left[-\frac{c_1}{\tilde{K}^4} (C^2 n + t^2) \right] \\ &\leq 9^n \cdot 2 \exp \left[-\frac{c_1}{K^4 \left\| \Sigma^{-\frac{1}{2}} \right\|^4} (C^2 n + t^2) \right] \\ &\leq 2 \exp \left(-\frac{c_1 t^2}{K^4 \left\| \Sigma^{-\frac{1}{2}} \right\|^4} \right). \end{aligned}$$

The last step, similar to the proof of [Theorem 6](#), comes from choosing C large enough, for example $C = K^2 \left\| \Sigma^{-\frac{1}{2}} \right\|^2 \sqrt{\ln(9)/c_1}$. \square

The combination of [Lemma 15](#), the fact that Σ for (9) equals B_k^{ALS} , and [Theorem 16](#), leads to our bound on the condition number of B_k in (9): for every $t \geq 0$ and with

probability at least $1 - 2 \exp(-ct^2)$,

$$(19) \quad \kappa(B_k) \leq \frac{1 + C\sqrt{r/r'} + t/\sqrt{r'}}{1 - C\sqrt{r/r'} - t/\sqrt{r'}} \kappa\left((B_k^{\text{ALS}})^{\frac{1}{2}}\right),$$

where the definitions of C and c are the same as in [Theorem 16](#).

Remark 17. In both (15) and (19) the ratios r/r' and $t/\sqrt{r'}$ are present. As both ratios go to zero, our bound on the condition number of B_k goes to $\kappa(B_k) \leq \kappa\left((B_k^{\text{ALS}})^{\frac{1}{2}}\right)$, and the bound on the condition number of $S_k^T Q$ goes to $\kappa(S_k^T Q) \leq 1$. These properties explain our choice to set r' as a constant multiple of r in the randomized ALS algorithm. As with similar bounds for randomized matrix algorithms, these bounds are pessimistic. Hence r' does not have to be very large with respect to r in order to get acceptable results.

Remark 18. [Algorithm 2](#) and the proofs in the present work use products of random variables extensively. For such problems the choice of signed Bernoulli random variables is a natural one, since their products are also signed Bernoulli random variables. While we had some success using Gaussian random variables in experiments, we have not included their use in this paper as they theoretically result in slower concentration of $B_k^T B_k$ around its expectation $\Sigma = B_k^{\text{ALS}}$ from Section 3.3. These experiments have led to an interesting question: is there an optimal choice of distribution for setting the entries of R_i ? This question requires a careful examination which is beyond the scope of this paper.

4. Examples.

4.1. Sine function. Our first test of the randomized ALS algorithm is to reduce a CTD generated from samples of the multivariate function $\sin(z_1 + \dots + z_d)$. This reduction problem was studied in [4], where the output of the standard ALS algorithm suggested a new trigonometric identity yielding a rank d separated representation of $\sin(z_1 + \dots + z_d)$. As input, we use standard trigonometric identities to produce a rank 2^{d-1} initial CTD.

We ran 500 tests using both standard ALS and the new randomized algorithm to reduce the separation rank of a CTD of samples of $\sin(z_1 + \dots + z_d)$. The tests differed in that each one had a different random initial guess with separation rank $r_{\mathbf{F}} = 1$. In this example we chose $d = 5$ and sampled each variable z_i , $i = 1, \dots, d$, with $M = 64$ equispaced samples in the interval $[0, 2\pi]$. Our input CTD for both algorithms was rank 16 and was generated via a standard trigonometric identity. The reduction tolerance for both algorithms was set to $\epsilon = 10^{-5}$, and the maximum number of iterations per rank, i.e. *max_iter* in [Algorithm 1](#) and [Algorithm 2](#), was set to 1000. For tests involving the standard ALS algorithm we used a stuck tolerance of $\delta = 10^{-8}$. To test the randomized ALS algorithm we used B_k matrices of size $(25 r_{\mathbf{F}}) \times r_{\mathbf{F}}$ and set *max_tries* in [Algorithm 2](#) to 50.

According to Lemma 2.4 in [4], there exists exact rank 5 separated representations of $\sin(z_1 + \dots + z_d)$. Using $\epsilon = 10^{-5}$ for our reduction tolerance, we were able to find rank 5 approximations with both standard ALS and our randomized ALS whose relative errors were less than the requested ϵ (for a histogram of residuals of the tests, see [Figure 1](#)).

Due to the random initial guess \mathbf{F} and our choices of parameters (in particular the stuck tolerance and *max_tries*) both algorithms had a small number of runs that

did not find rank 5 approximations with the requested tolerance ϵ . The randomized ALS algorithm produced fewer of these outcomes than standard ALS.

Large differences in maximum condition number (of ALS solves) are illustrated in [Figure 2](#), where we compare tests of the standard and randomized ALS algorithms. We observe that the maximum condition numbers produced by the randomized ALS algorithm are much smaller than those from the standard ALS algorithm. This is consistent with our theory.

Furthermore, as shown in [Figure 3](#), the number of iterations required for randomized ALS to converge was smaller than the number required by standard ALS. For these experiments the ALS sweeps rejected by [Algorithm 2](#) are included in the iteration count. It is important to remember that the number of iterations required by the standard ALS algorithm to reduce a CTD can be optimized by adjusting the tolerance, stuck tolerance, and maximum number of iterations per rank. In these experiments we chose the stuck tolerance and maximum number of iterations to reduce the number of tests of the standard ALS algorithm that did not meet the requested tolerance ϵ .

To better illustrate the behavior of the condition numbers in these experiments, we display condition numbers from a single experiment in [Figure 4](#). Specifically, [Figure 4\(a\)](#) shows the condition number of the matrix B_k defined by [\(4\)](#) and used in [Algorithm 1](#), and [Figure 4\(b\)](#) shows the condition number of the non-rejected matrix B_k defined by [\(9\)](#) and used in [Algorithm 2](#). We emphasize that these plots contain the condition numbers of all matrices B_k used in this experiment corresponding to each directional update and ALS sweep. In [Figure 4\(a\)](#) we observe the condition number of B_k growing rapidly as the iteration number increases. [Figure 4\(b\)](#), however, displays a milder increase in the condition number of B_k for the randomized ALS algorithm.

4.2. A manufactured tensor. Our next test is to compare the performance of the standard and randomized ALS algorithms on a manufactured random tensor example. To construct this example we generate factors by drawing $M = 128$ random samples from the standard Gaussian distribution. We chose $d = 20$ and set the separation rank of the input tensor to $r = 50$. Then we normalized the factors and set the s-values of the tensor equal to $s_l = e^{-l}$, $l = 0, \dots, r - 1$, where r was predetermined such that s_{end} is small.

Similar to the sine example, we ran 500 experiments and requested an accuracy of $\epsilon = 10^{-4}$ from both algorithms. The maximum number of iterations for both algorithms was set to 1000, while the stuck tolerance for the standard ALS algorithm was set to 10^{-6} . We used the following parameters for the randomized ALS algorithm: the B_k matrices were of size $(25 r_{\mathbf{F}}) \times r_{\mathbf{F}}$, and the repetition parameter, *max_tries* in [Algorithm 2](#), was set to 50. We started all tests from randomized guesses with rank $r_{\mathbf{F}} = 9$. This value was chosen because in all previous runs the reduced separation rank never fell below $r_{\mathbf{F}} = 10$. Such an experiment allows us to compare how the algorithms perform when the initial approximation has rank greater than one.

We show in [Figure 5](#) the output separation ranks from 500 tests of both the randomized and standard ALS algorithms. The CTD outputs from randomized ALS had, on average, lower separation ranks than those from standard ALS. Furthermore, as seen in [Figure 5](#), some of the output CTDs from the standard ALS algorithm had separation rank of 40. In these instances, standard ALS failed to reduce the separation rank of the input CTD because simple truncation to $r_{\mathbf{F}} = 35$ would have given double precision. These failures did not occur with the randomized ALS algorithm. We can also see a contrast in performance in [Figure 6](#): all tests of the randomized ALS

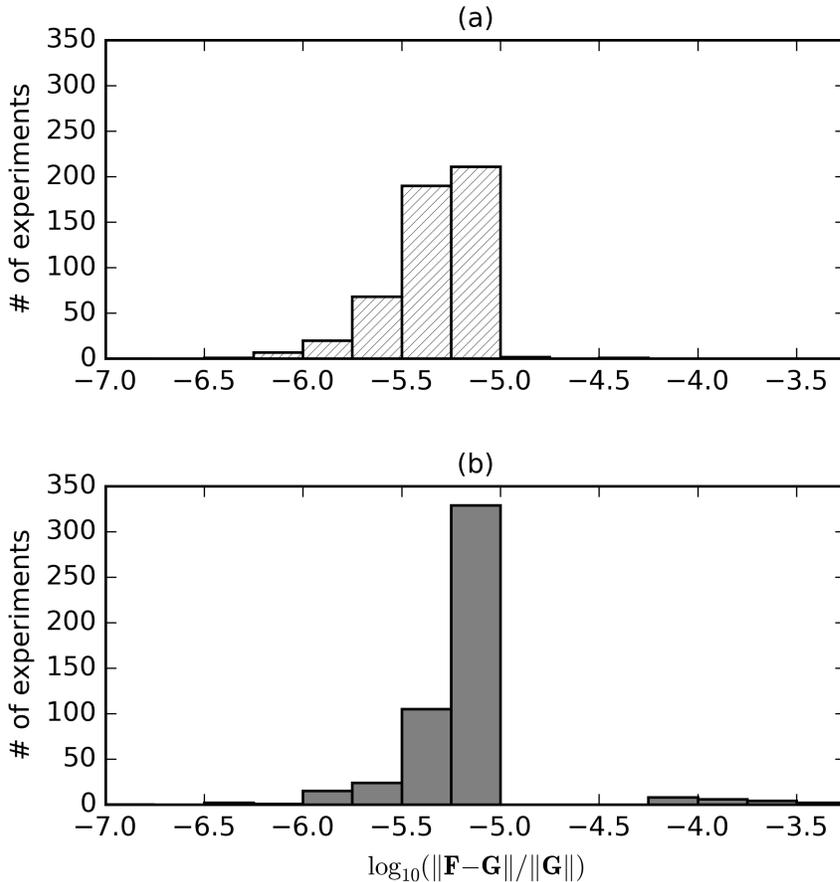


FIG. 1. Histograms displaying ALS reduction residuals, in \log_{10} scale, for reducing the length of a CTD of samples of $\sin(z_1 + \dots + z_5)$. The experiments shown in (a) used randomized ALS, whereas the experiments shown in (b) used standard ALS. We note that both algorithms produced a small number of results with approximation errors worse than the requested tolerance. However, the randomized ALS method produced fewer results that did not meet our requested tolerance.

algorithm produced CTDs with reduced separation rank whose relative reduction errors were less than the accuracy ϵ . Also, in Figure 6, we observe instances where the standard ALS algorithm failed to output a reduced separation rank CTD with relative error less than ϵ .

There was a significant difference in the maximum condition numbers of matrices used in the two algorithms. In Figure 7, we see that matrices produced by standard ALS had much larger condition numbers (by a factor of roughly 10^{10}) than their counterparts in the randomized ALS algorithm. Such large condition numbers may explain the failures of the standard ALS algorithm to output reduced separation rank CTDs with relative errors less than ϵ .

From Figure 8, we see that in most of the tests standard ALS required fewer iterations than randomized ALS to converge. However, there were a few tests where standard ALS required a larger number of iterations. These are the tests that failed

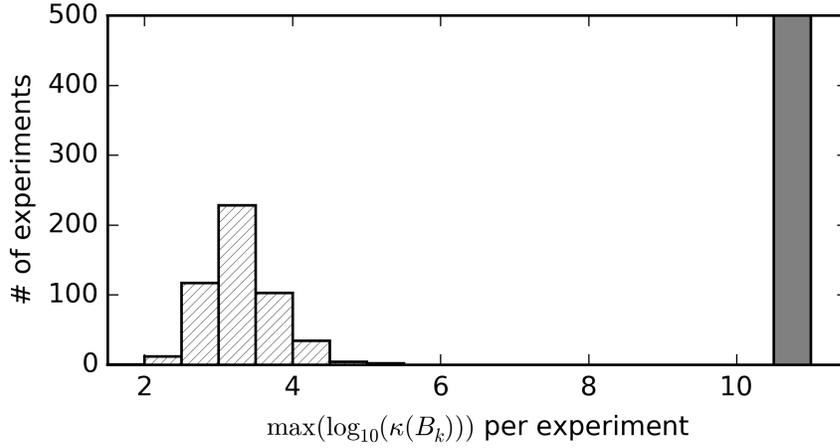


FIG. 2. Histogram showing the maximum condition numbers from our experiments reducing the length of a CTD of samples of $\sin(z_1 + \dots + z_5)$. The condition numbers are shown in \log_{10} scale; the solid gray pattern represents condition numbers from standard ALS, while the hatch pattern represents condition numbers from the randomized ALS algorithm.

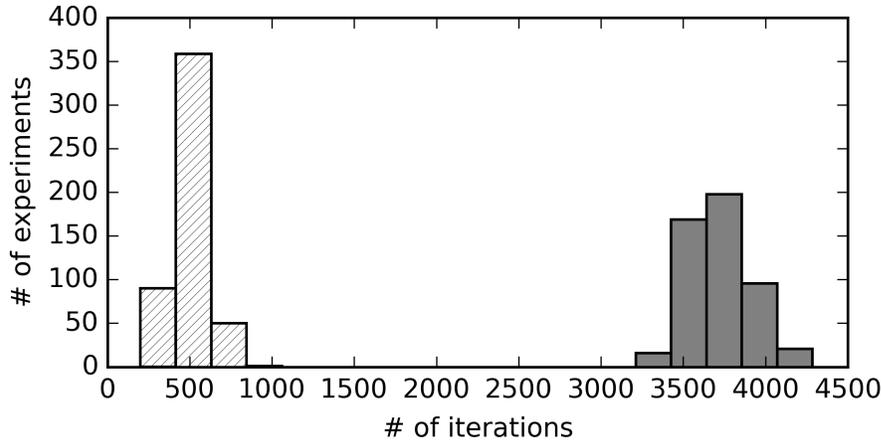


FIG. 3. Histogram showing the number of iterations required by randomized ALS (hatch pattern) and the standard ALS algorithm (gray pattern) to reduce the length of a CTD of samples of $\sin(z_1 + \dots + z_5)$.

to output a reduced separation rank CTD with relative error less than ϵ . We note that similar to the experiments in [subsection 4.1](#), the randomized ALS sweeps rejected by [Algorithm 2](#) are included in the iteration count.

4.3. Elliptic PDE with random coefficient. As the key application of the randomized ALS algorithm, we consider the separated representation of the solution $u(\mathbf{x}, \mathbf{z})$ to the linear elliptic PDE

$$(20) \quad \begin{aligned} -\nabla \cdot (a(\mathbf{x}, \mathbf{z}) \nabla u(\mathbf{x}, \mathbf{z})) &= 1, & \mathbf{x} \in \mathcal{D}, \\ u(\mathbf{x}, \mathbf{z}) &= 0, & \mathbf{x} \in \partial \mathcal{D}, \end{aligned}$$

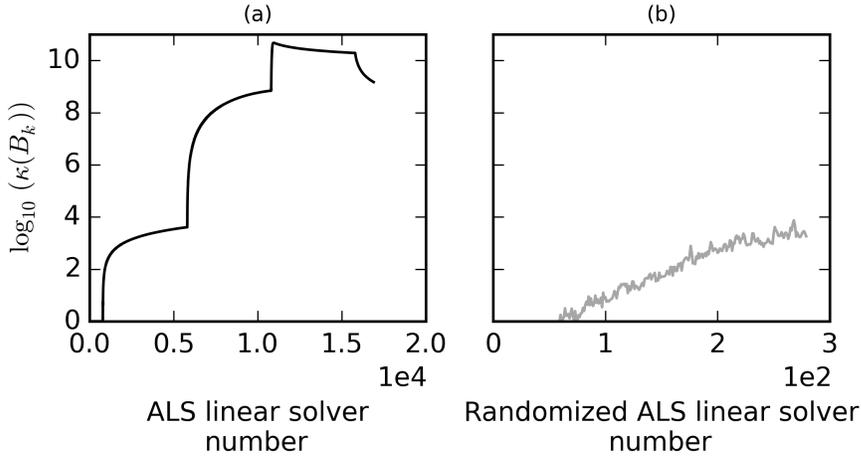


FIG. 4. Side-by-side comparison of condition numbers of matrices B_k from reducing the rank of a CTD of samples of $\sin(z_1 + \dots + z_5)$ using (a) standard ALS and (b) randomized ALS. The horizontal axis of each plot measures the total number of times the linear system (3) is solved for standard ALS (a), and (3) is solved for randomized ALS using B_k and \mathbf{b}_{j_k} specified in (9) and (10), respectively (b). Note we only display the condition numbers from experiments not rejected in Algorithm 2.

defined on the unit square $\mathcal{D} = (0, 1) \times (0, 1)$ with boundary $\partial\mathcal{D}$. The diffusion coefficient $a(\mathbf{x}, \mathbf{z})$ is considered random and is modeled by

$$(21) \quad a(\mathbf{x}, \mathbf{z}) = a_0 + \sigma_a \sum_{k=1}^d \sqrt{\zeta_k} \varphi_k(\mathbf{x}) z_k,$$

where $\mathbf{z} = (z_1, \dots, z_d)$ and the random variables z_k are independent and uniformly distributed over the interval $[-1, 1]$, and we choose $a_0 = 0.1$, $\sigma_a = 0.01$, and $d = 5$. In (21), $\{\zeta_k\}_{k=1}^d$ are the d largest eigenvalues associated with $\{\varphi_k\}_{k=1}^d$, the $L_2(\mathcal{D})$ -orthonormalized eigenfunctions of the exponential covariance function

$$(22) \quad C_{aa}(\mathbf{x}_1, \mathbf{x}_2) = \exp\left(-\frac{\|\mathbf{x}_1 - \mathbf{x}_2\|_1}{l_c}\right),$$

where l_c denotes the correlation length, here set to $l_c = 2/3$. Given the choices of parameters a_0 , σ_a , d , and l_c , the model in (21) leads to strictly positive realizations of $a(\mathbf{x}, \mathbf{z})$.

We discretize (20) in the spatial domain \mathcal{D} via triangular finite elements of size $h = 1/32$. This discretization, along with the affine representation of $a(\mathbf{x}, \mathbf{z})$ in z_k , yields the random linear system of equations

$$(23) \quad \left(K_0 + \sum_{k=1}^d K_k z_k \right) \mathbf{u}(\mathbf{z}) = \mathbf{f},$$

for the approximate vector of nodal solutions $\mathbf{u}(\mathbf{z}) \in \mathbb{R}^N$. The sparse matrices K_0 and K_k are obtained from the finite element discretization of the differential operator in (20) assuming $a(\mathbf{x}, \mathbf{z})$ is replaced by a_0 and $\sigma_a \sqrt{\zeta_k} \varphi_k(\mathbf{x})$, respectively.

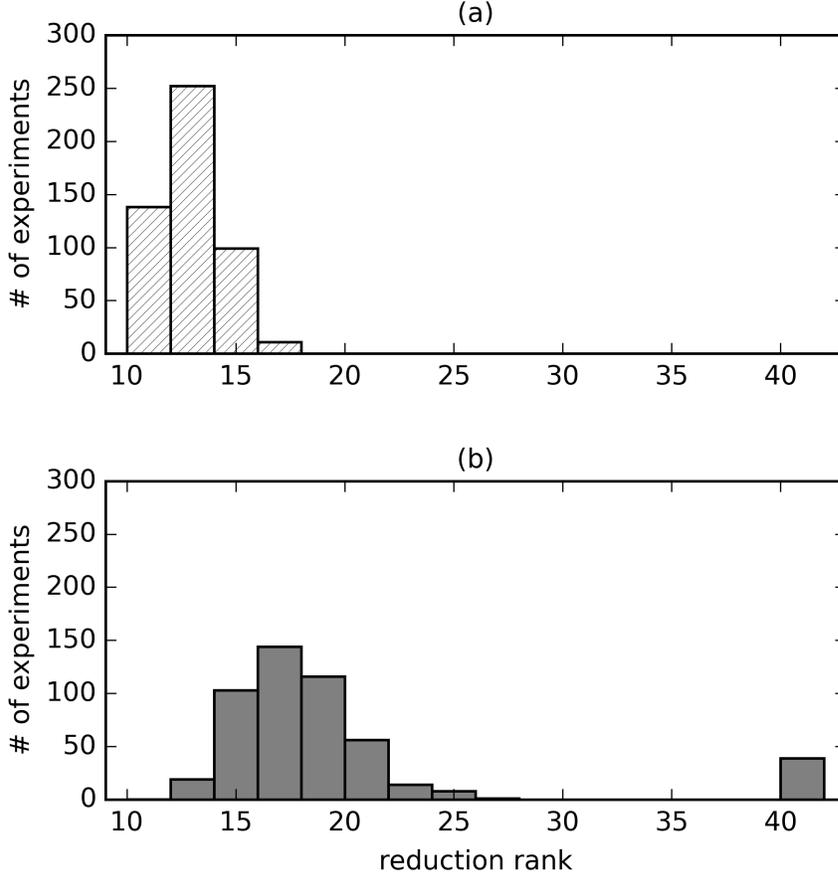


FIG. 5. Histograms showing output ranks from experiments in reducing the length of CTDs. (a) shows ranks of CTDs output by the randomized ALS algorithm. (b) shows ranks of CTDs output by the standard ALS algorithm. The CTDs output by the randomized ALS method typically have a smaller separation rank. In many examples the standard ALS algorithm required 40 terms, i.e. it failed since truncation of the input tensor to $r_{\mathbf{F}} = 35$ should give double precision.

To fully discretize (23), we consider the representation of $\mathbf{u}(\mathbf{z})$ at a tensor-product grid $\{(z_1(j_1), \dots, z_d(j_d)) : j_k = 1, \dots, M_k\}$ where, for each k , the grid points $z_k(j_k)$ are selected to be the Gauss-Legendre abscissas. In our numerical experiments, we used the same number of abscissas $M_k = M = 8$ for all $k = 1, \dots, d$. The discrete representation of (23) is then given by the tensor system of equations

$$(24) \quad \mathbb{K}\mathbf{U} = \mathbf{F},$$

where the linear operation $\mathbb{K}\mathbf{U}$ is defined as

$$\mathbb{K}\mathbf{U} = \sum_{i=0}^d \sum_{\tilde{l}=1}^{r_{\mathbf{U}}} s_{\tilde{l}}^{\mathbf{U}} \left(\mathbb{K}_0^i \mathbf{U}_0^{\tilde{l}} \right) \circ \left(\mathbb{K}_1^i \mathbf{U}_1^{\tilde{l}} \right) \circ \dots \circ \left(\mathbb{K}_d^i \mathbf{U}_d^{\tilde{l}} \right),$$

$$\mathbb{K}_0^i = K_{\tilde{l}},$$

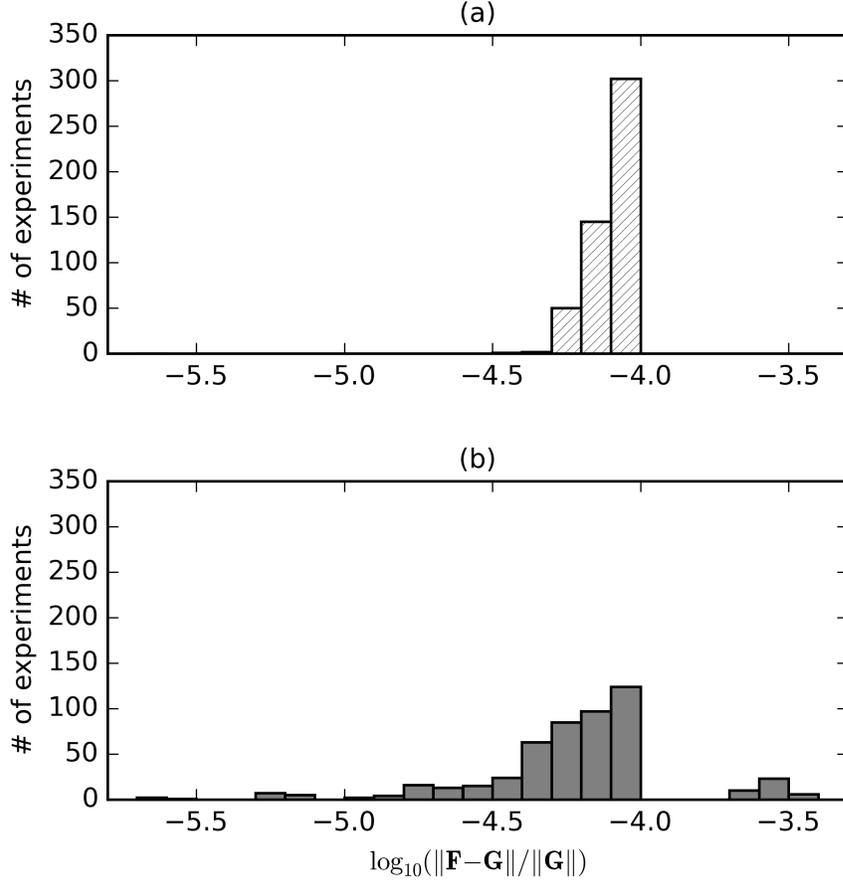


FIG. 6. Histograms displaying ALS reduction errors, in \log_{10} scale, for reduced-rank CTDs of the random tensor example. (a) shows that in our 500 tests, the randomized ALS method always produced a result that met the required tolerance. (b) shows how the standard ALS algorithm fared with the same problem. Note that the standard ALS algorithm failed to reach the requested tolerance in a small number of tests.

and for $k = 1, \dots, d$,

$$\mathbb{K}_k^{\hat{l}} = \begin{cases} D & \hat{l} = k \\ I_M & \hat{l} \neq k, \end{cases}$$

where

$$D = \begin{bmatrix} z_k(1) & & 0 \\ & \ddots & \\ 0 & & z_k(M) \end{bmatrix},$$

and I_M is the $M \times M$ identity matrix. The tensor \mathbf{F} in (24) is defined as

$$\mathbf{F} = \mathbf{f} \circ \mathbf{1}_M \circ \dots \circ \mathbf{1}_M,$$

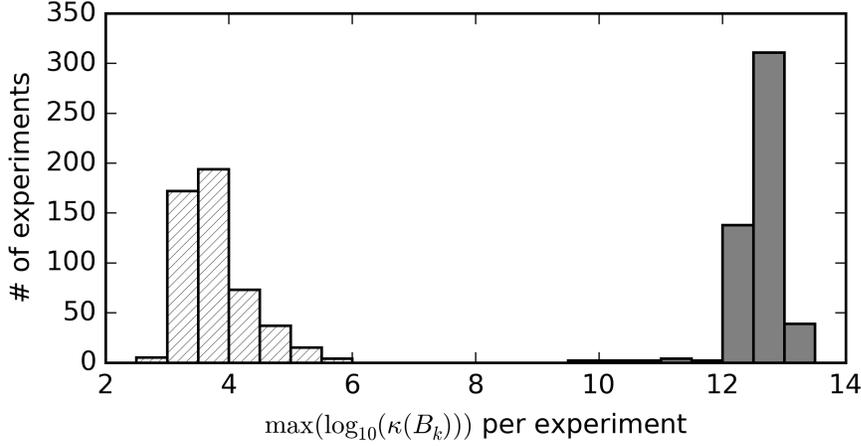


FIG. 7. Histogram showing the maximum condition numbers from experiments in reducing the length of CTDs of the random tensor example. The condition numbers of B_k are shown in \log_{10} scale; solid gray represents condition numbers from standard ALS while the hatch pattern represents condition numbers from the randomized ALS algorithm. Similar to the sine example, the condition numbers from randomized ALS are much smaller than those from the standard ALS algorithm

where $\mathbf{1}_M$ is an M -vector of ones. We seek to approximate \mathbf{U} in (24) with a CTD,

$$(25) \quad \mathbf{U} = \sum_{\tilde{i}=1}^{r_{\mathbf{U}}} s_{\tilde{i}}^{\mathbf{U}} \mathbf{U}_0^{\tilde{i}} \circ \mathbf{U}_1^{\tilde{i}} \circ \cdots \circ \mathbf{U}_d^{\tilde{i}},$$

where the separation rank $r_{\mathbf{U}}$ will be determined by a target accuracy. In (25) $\mathbf{U}_0^{\tilde{i}} \in \mathbb{R}^N$ and $\mathbf{U}_k^{\tilde{i}} \in \mathbb{R}^M$, $k = 1, \dots, d$. To solve (24), we use a fixed point iteration similar to those used for solving matrix equations and recently employed to solve tensor equations in [23]. In detail, the iteration starts with an initial tensor \mathbf{U} of the form in (25). At each iteration i , \mathbf{U} is updated according to

$$\mathbf{U}_{i+1} = (\mathbb{I} - \mathbb{K}) \mathbf{U}_i + \mathbf{F},$$

while requiring $\|\mathbb{I} - \mathbb{K}\| < 1$. To assure this requirement is satisfied we solve

$$(26) \quad \mathbf{U}_{i+1} = c(\mathbf{F} - \mathbb{K}\mathbf{U}_i) + \mathbf{U}_i,$$

where c is chosen such that $\|\mathbb{I} - c\mathbb{K}\| < 1$. We compute the operator norm $\|\mathbb{I} - \mathbb{K}\|$ via power method; see, e.g., [4, 5].

One aspect of applying such an iteration to a CTD is an increase in the output separation rank. For example, if we take a tensor \mathbf{U} of separation rank $r_{\mathbf{U}}$ and use it as input for (26), one iteration would increase the rank to $r_{\mathbf{F}} + (d + 2)r_{\mathbf{U}}$. Therefore we require a reduction algorithm to decrease the separation rank as we iterate. This is where either the standard or randomized ALS algorithm is required: to truncate the separated representation after we have run an iteration. Both ALS methods work with a user-supplied truncation accuracy ϵ , so we denote the reduction operator as τ_{ϵ} . Including this operator into our iteration, we have

$$(27) \quad \mathbf{U}_{i+1} = \tau_{\epsilon}(c(\mathbf{F} - \mathbb{K}\mathbf{U}_i) + \mathbf{U}_i).$$

Pseudocode for our fixed point algorithm is shown in [Algorithm 3](#)

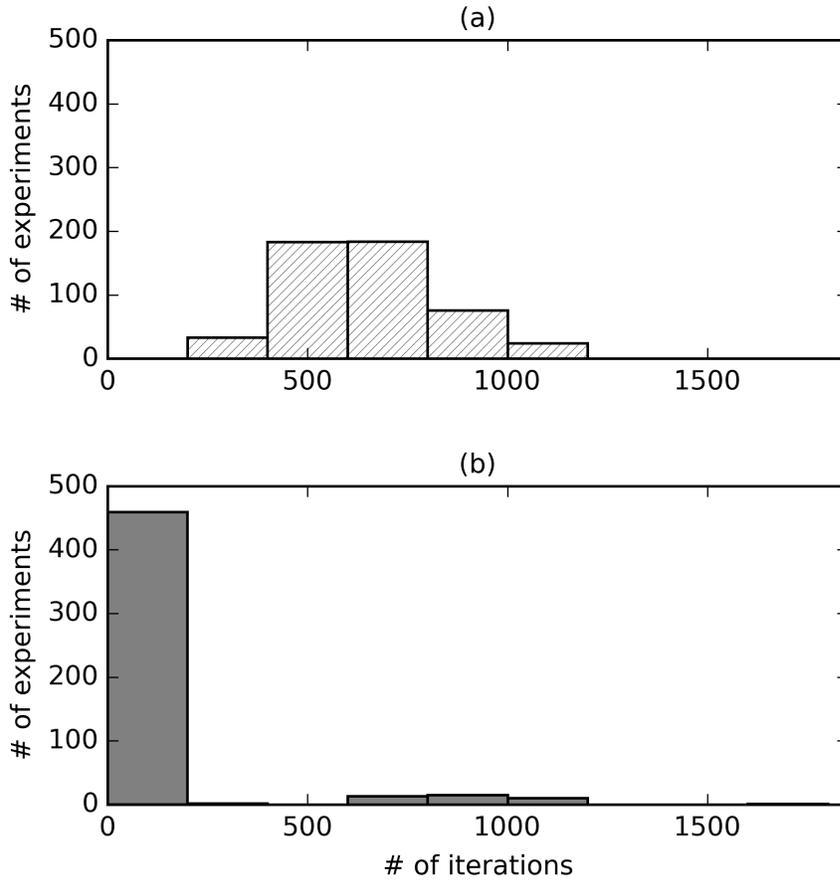


FIG. 8. Histograms showing iterations required to produce reduced-length CTDs for the random tensor example. (a) shows iterations required by randomized ALS, while (b) shows the iterations required by the standard ALS algorithm. As seen in (b), a few examples using standard ALS required large numbers of iterations to output CTDs. These examples failed to produce a reduced separation rank CTD with relative error less than ϵ . However, for most of the experiments the standard ALS algorithm required fewer iterations than the randomized ALS algorithm.

Remark 19. In this example, the separation rank of \mathbb{K} is directly related to the problem dimension d , i.e. $r_{\mathbb{K}} = d + 1$, which is a consequence of using a Karhunen-Loeve-type expansion for finite-dimensional noise representation of $a(\mathbf{x}, \mathbf{z})$. This will increase the computational cost of the algorithm to more than linear with respect to d , e.g. quadratic in d when an iterative solver is used and $N \gg M$. Alternatively, one can obtain the finite-dimensional noise representation of $a(\mathbf{x}, \mathbf{z})$ by applying the separated rank reduction technique of this study on the stochastic differential operator itself to possibly achieve $r_{\mathbb{K}} < d$. The interested reader is referred to [3, 4] for more details.

First, we examine the convergence of the iterative algorithm given a fixed ALS reduction tolerance in Figure 9. The randomized ALS method converges to more accurate solutions in all of these tests (see Table 1). However, the ranks of the randomized ALS

Algorithm 3 Fixed point iteration algorithm for solving (27)

input : $\epsilon > 0$, $\mu > 0$, operator \mathbb{K} , \mathbf{F} , c , *max_iter*, *max_rank*, $\delta > 0$
(for standard ALS), *max_tries* (for randomized ALS)
initialize $r_{\mathbf{U}} = 1$ tensor $\mathbf{U}_0 = \mathbf{U}_1^1 \circ \dots \circ \mathbf{U}_d^1$ with either randomly generated \mathbf{U}_k^1 or \mathbf{U}_k^1 generated from the solution of the deterministic version of (20), i.e., when $a(\mathbf{x}, \mathbf{z})$ is replaced by a_0 in (21). Also initialize the fixed point iteration counter to $i = 0$.
 $\mathbf{D}_0 = \mathbf{F} - \mathbb{K}\mathbf{U}_0$
 $res = \|\mathbf{D}_0\| / \|\mathbf{F}\|$
while $res > \mu$ **do**
 $i = i + 1$
 $\mathbf{U}_i = c\mathbf{D}_{i-1} + \mathbf{U}_{i-1}$
 $\mathbf{U}_i = \tau_\epsilon(\mathbf{U}_i)$
 $\mathbf{D}_i = \mathbf{F} - \mathbb{K}\mathbf{U}_i$
 $res = \|\mathbf{D}_i\| / \|\mathbf{F}\|$
end while
return \mathbf{U}_i

ALS type	ALS tol	max $\kappa(B_k)$	max rank	rank	residual
standard	1×10^{-3}	5.35×10^1	5	4	4.16×10^{-2}
	1×10^{-4}	5.29×10^5	13	11	5.72×10^{-3}
	1×10^{-5}	1.07×10^9	37	34	4.18×10^{-4}
randomized	1×10^{-3}	2.59×10^2	7	6	2.36×10^{-2}
	1×10^{-4}	3.59×10^3	22	19	2.35×10^{-3}
	1×10^{-5}	2.72×10^4	57	54	3.00×10^{-4}

TABLE 1

Table containing ranks, maximum condition numbers, and final relative residual errors of experiments with fixed ALS tolerance.

solutions are larger than the ranks required for solutions produced by the standard ALS algorithm.

In Figure 10, we observe different behavior in the relative residuals using fixed ranks instead of fixed accuracies. For these experiments the ALS-based linear solve using the standard algorithm out-performs the randomized version, except in the rank $r = 30$ case (see Table 2). In this case, the standard ALS algorithm has issues reaching the requested ALS reduction tolerance, thus leading to convergence problems in the iterative linear solve. The randomized ALS algorithm does not have the same difficulty with the rank $r = 30$ example. This difference in decay between the standard and randomized ALS residuals corresponds to a significant difference between the maximum condition numbers of B_k . For the $r = 30$ case, the maximum condition number of B_k matrices generated by randomized ALS was 3.94×10^7 , whereas the maximum condition number of B_k matrices generated by standard ALS was 3.00×10^{13} .

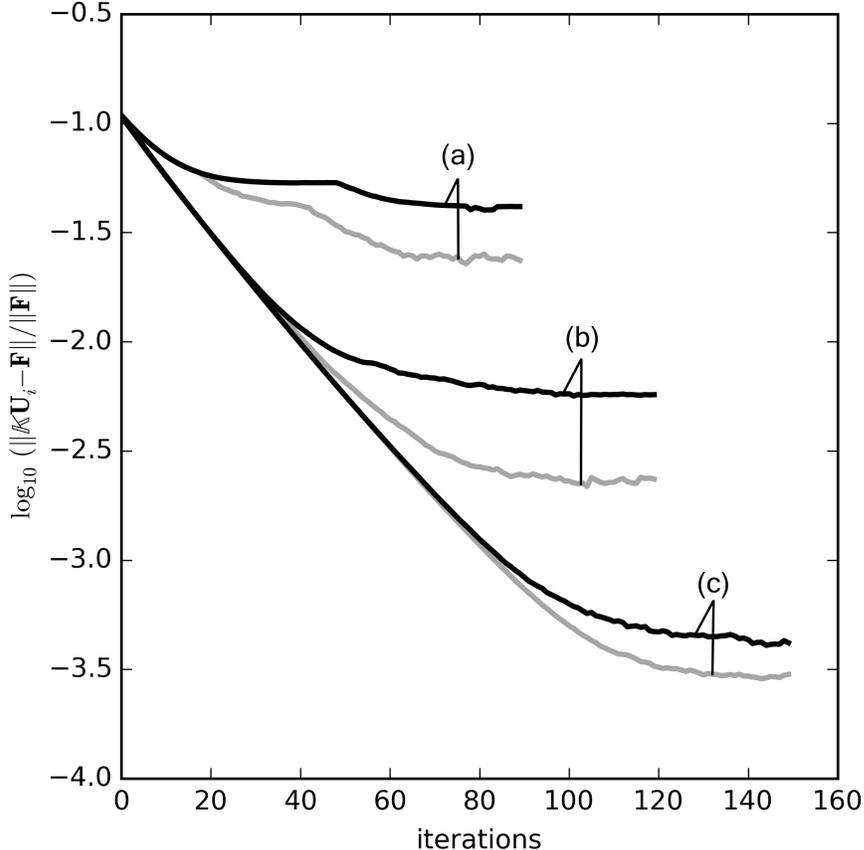


FIG. 9. Residual error versus fixed point iteration number of results from linear solvers. The black lines represent linear solve residuals where standard ALS was used for reduction, while the gray lines represent linear solve residuals where randomized ALS was used for reduction. In the three examples shown above the ALS tolerances, for both standard and randomized ALS, were set to 1×10^{-3} for curves labeled (a), 1×10^{-4} for curves labeled (b), and 1×10^{-5} for curves labeled (c).

5. Discussion and conclusions. We have proposed a new ALS algorithm for reducing the rank of tensors in canonical format that relies on projections onto random tensors. Tensor rank reduction is one of the primary operations for approximations with tensors. Additionally, we have presented a general framework for the analysis of this new algorithm. The benefit of using such random projections is the improved conditioning of matrices associated with the least squares problem at each ALS iteration. While significant reductions of condition numbers may be achieved, unlike in the standard ALS, the application of random projections results in a loss of monotonic error reduction. In order to restore monotonicity, we have employed a simple rejection approach, wherein several random tensors are applied and only those that do not increase the error are accepted. This, however, comes at the expense of additional computational cost as compared to the standard ALS algorithm. Finally, a set of numerical experiments has been studied to illustrate the efficiency of the randomized

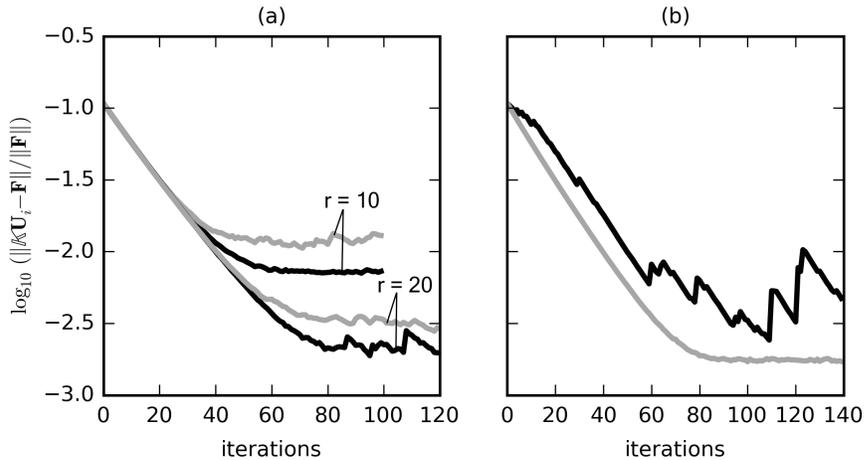


FIG. 10. Plots showing relative residuals of fixed point solutions versus fixed point iteration number. (a) two fixed point solutions are shown here: solutions for fixed ranks $r = 10$, and $r = 20$. Gray lines are residuals corresponding to reductions with randomized ALS and black lines correspond to reductions with the standard ALS algorithm. (b) One experiment with $r = 30$ and the same color scheme as in (a).

ALS type	ALS tol	max $\kappa(B_k)$	rank	residual
standard	1×10^{-5}	9.45×10^{11}	10	7.29×10^{-3}
	5×10^{-6}	1.27×10^{13}	20	1.97×10^{-3}
	1×10^{-6}	3.00×10^{13}	30	4.73×10^{-3}
randomized	1×10^{-5}	9.39×10^5	10	1.30×10^{-2}
	5×10^{-6}	4.12×10^6	20	2.93×10^{-3}
	1×10^{-6}	3.94×10^7	30	1.72×10^{-3}

TABLE 2

Table containing maximum condition numbers and final relative residual errors of experiments with fixed separation ranks.

ALS in improving numerical properties of its standard counterpart.

The optimal choice of random variables to use in the context of projecting onto random tensors is a question to be addressed in future work. In our examples we have used signed Bernoulli random variables, a choice that worked well with both our numerical experiments and analysis. On the other hand, there are limitations of such a construction of random tensors, which motivate further investigations. A topic of interest for future work is the extension of the proposed randomized framework to other tensor formats including the Tucker, [24], and tensor-train, [28].

Another area of future work involves directly solving systems such as (24) with a randomized ALS variant, instead of utilizing a fixed point algorithm (e.g. Algorithm 3). A standard ALS algorithm for this purpose was derived in [4, Section 4.2],

where the approach is to solve $\mathbb{A}\mathbf{F} = \mathbf{G}$ by minimizing $\|\mathbb{A}\mathbf{F} - \mathbf{G}\|$ with an imposed separation rank constraint on \mathbf{F} . The resulting equations are also normal equations similar to (3) and (4), thus we anticipate our randomized machinery will extend to this approach for solving $\mathbb{A}\mathbf{F} = \mathbf{G}$.

Finally we have suggested an alternative approach to using projections onto random tensors that merits further examination. This approach uses the QR factorization to construct a preconditioner for the least squares problem at each ALS iteration. Hence it solves the same equations as the standard ALS, but the matrices have better conditioning. Also, because it solves the same equations, the monotonic error reduction property is preserved. This is an important distinction from randomized ALS, which solves different linear systems, but the solutions to which are close to the solutions from standard ALS.

Appendix A. Proofs of (6), Theorem 6, and Lemma 15. First, we prove (6).

Proof. To bound the condition number of AB we bound $\sigma_{\max}(AB)$ from above and $\sigma_{\min}(AB)$ from below. The bound we use of $\sigma_{\max}(AB)$ is straightforward; it comes from the properties of the two norm,

$$\sigma_{\max}(AB) \leq \sigma_{\max}(A) \sigma_{\max}(B).$$

To bound $\sigma_{\min}(AB)$ we first note that AA^T is nonsingular, and write $\sigma_{\min}(AB)$ as follows,

$$\sigma_{\min}(AB) = \left\| AA^T (AA^T)^{-1} AB\mathbf{x}^* \right\|_2,$$

where $\|\mathbf{x}^*\| = 1$ is the value of \mathbf{x} such that the minimum of the norm is obtained (see the minimax definition of singular values, e.g. [22, Theorem 3.1.2]). If we define $\mathbf{y} = A^T (AA^T)^{-1} AB\mathbf{x}^*$, then

$$\sigma_{\min}(AB) = \frac{\|\mathbf{A}\mathbf{y}\| \cdot \|\mathbf{y}\|}{\|\mathbf{y}\|} \geq \sigma_{\min}(A) \cdot \|\mathbf{y}\|,$$

from the minimax definition of singular values. To bound $\|\mathbf{y}\|$, we observe that $A^T (AA^T)^{-1} AB$ is the projection of B onto the row space of A . Denoting this projection as $P_{A^T}(B)$ we have,

$$\|\mathbf{y}\| = \|P_{A^T}(B)\mathbf{x}^*\| \geq \sigma_{\min}(P_{A^T}(B)),$$

since $\|\mathbf{x}^*\| = 1$. Combining our bounds on the first and last singular values gives us the bound on the condition number,

$$\kappa(AB) \leq \frac{\sigma_{\max}(A) \sigma_{\max}(B)}{\sigma_{\min}(A) \sigma_{\min}(P_{A^T}(B))} = \kappa(A) \cdot \frac{\sigma_{\max}(B)}{\sigma_{\min}(P_{A^T}(B))}. \quad \square$$

The proof of Theorem 6 is broken down into three steps in order to control $\|\mathbf{A}\mathbf{x}\|$ for all \mathbf{x} on the unit sphere: an approximation step, where the unit sphere is covered using a finite epsilon-net \mathcal{N} (see [32, Section 5.2.2] for background on nets); a concentration step, where tight bounds are applied to $\|\mathbf{A}\mathbf{x}\|$ for every $\mathbf{x} \in \mathcal{N}$; and the final step where a union bound is taken over all the vectors $\mathbf{x} \in \mathcal{N}$.

Proof. (of Theorem 6)

Vershynin observes that if we set B in [Lemma 7](#) to A/\sqrt{N} , the bounds on the extreme singular values $\sigma_{\min}(A)$ and $\sigma_{\max}(A)$ in [\(7\)](#) are equivalent to

$$(28) \quad \left\| \frac{1}{N} A^T A - I \right\| < \max(\delta, \delta^2) =: \epsilon,$$

where $\delta = C\sqrt{\frac{n}{N}} + \frac{t}{\sqrt{N}}$. In the approximation step of the proof, he chooses a $\frac{1}{4}$ -net \mathcal{N} to cover the unit sphere \mathcal{S}^{n-1} . Evaluating the operator norm [\(28\)](#) on \mathcal{N} , it is sufficient to show

$$\max_{x \in \mathcal{N}} \left| \frac{1}{N} \|Ax\|_2^2 - 1 \right| < \frac{\epsilon}{2},$$

with the required probability to prove the theorem.

Starting the concentration step, [\[32\]](#) defines $Z_i = \langle A_i, \mathbf{x} \rangle$, where A_i is the i -th row of A and $\|\mathbf{x}\|_2 = 1$. Hence, the vector norm may be written as

$$(29) \quad \|Ax\|_2^2 = \sum_{i=1}^N Z_i^2.$$

Using an exponential deviation inequality to control [\(29\)](#), and that $K \geq \frac{1}{\sqrt{2}}$, the following probabilistic bound for a fixed $\mathbf{x} \in \mathcal{S}^{n-1}$ is,

$$(30) \quad \begin{aligned} \mathbb{P} \left\{ \left| \frac{1}{N} \|Ax\|_2^2 - 1 \right| \geq \frac{\epsilon}{2} \right\} &= \mathbb{P} \left\{ \left| \frac{1}{N} \sum_{i=1}^N Z_i^2 - 1 \right| \geq \frac{\epsilon}{2} \right\} \leq 2 \exp \left[-\frac{c_1}{K^4} \min(\epsilon^2, \epsilon) N \right] \\ &= 2 \exp \left[-\frac{c_1}{K^4} \delta^2 N \right] \leq 2 \exp \left[-\frac{c_1}{K^4} (C^2 n + t^2) \right], \end{aligned}$$

where c_1 is an absolute constant.

Finally, [\(30\)](#) is applied to every vector $\mathbf{x} \in \mathcal{N}$ resulting in the union bound,

$$(31) \quad \mathbb{P} \left\{ \max_{x \in \mathcal{N}} \left| \frac{1}{N} \|Ax\|_2^2 - 1 \right| \geq \frac{\epsilon}{2} \right\} \leq 9^n \cdot 2 \exp \left[-\frac{c_1}{K^4} (C^2 n + t^2) \right] \leq 2 \exp \left(-\frac{c_1 t^2}{K^4} \right),$$

where we arrive at the second inequality by choosing a sufficiently large $C = C_K$ ([\[32\]](#) gives the example $C = K^2 \sqrt{\ln(9)/c_1}$). \square

We now prove [Lemma 15](#).

Proof. To prove this lemma we use the following inequality derived from [\(17\)](#),

$$(1 - \delta)^2 \leq \sigma_{\min} \left(\Sigma^{-\frac{1}{2}} B^T B \Sigma^{-\frac{1}{2}} \right) \leq \sigma_{\max} \left(\Sigma^{-\frac{1}{2}} B^T B \Sigma^{-\frac{1}{2}} \right) \leq (1 + \delta)^2.$$

First we bound $\sigma_{\max}(B)$ from above:

$$\begin{aligned} \sigma_{\max}(B)^2 &\leq \left\| \Sigma^{\frac{1}{2}} \right\| \cdot \left\| \Sigma^{-\frac{1}{2}} B^T B \Sigma^{-\frac{1}{2}} \right\| \cdot \left\| \Sigma^{\frac{1}{2}} \right\| \\ &\leq \left\| \Sigma^{\frac{1}{2}} \right\|^2 \cdot \sigma_{\max} \left(\Sigma^{-\frac{1}{2}} B^T B \Sigma^{-\frac{1}{2}} \right) \\ &\leq \sigma_{\max} \left(\Sigma^{\frac{1}{2}} \right)^2 \cdot (1 + \delta)^2, \end{aligned}$$

implying $\sigma_{\max}(B) \leq \sigma_{\max}\left(\Sigma^{\frac{1}{2}}\right) \cdot (1 + \delta)$. Second we bound $\sigma_{\min}(B)$ from below:

$$\begin{aligned} \sigma_{\min}(B)^2 &= \sigma_{\min}\left(\Sigma^{\frac{1}{2}}\Sigma^{-\frac{1}{2}}B^TB\Sigma^{-\frac{1}{2}}\Sigma^{\frac{1}{2}}\right) \\ &\geq \sigma_{\min}\left(\Sigma^{\frac{1}{2}}\right)^2 \cdot \sigma_{\min}\left(\Sigma^{-\frac{1}{2}}B^TB\Sigma^{-\frac{1}{2}}\right) \\ (32) \quad &\geq \sigma_{\min}\left(\Sigma^{\frac{1}{2}}\right)^2 \cdot (1 - \delta)^2, \end{aligned}$$

implying $\sigma_{\min}(B) \geq \sigma_{\min}\left(\Sigma^{\frac{1}{2}}\right) \cdot (1 - \delta)$. The first inequality in (32) is from [22, prob. 3.3.12]. Finally, using properties of singular values we combine the inequalities:

$$\sigma_{\min}\left(\Sigma^{\frac{1}{2}}\right) \cdot (1 - \delta) \leq \sigma_{\min}(B) \leq \sigma_{\max}(B) \leq \sigma_{\max}\left(\Sigma^{\frac{1}{2}}\right) \cdot (1 + \delta). \quad \square$$

REFERENCES

- [1] C. APPELLOF AND E. DAVIDSON, *Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents*, Analytical Chemistry, 53 (1982), pp. 2053–2056.
- [2] B. BADER, M. BERRY, AND M. BROWNE, *Discussion tracking in enron email using parafac*, in Survey of Text Mining II, Springer London, 2008, pp. 147–163.
- [3] G. BEYLKIN AND M. J. MOHLENKAMP, *Numerical operator calculus in higher dimensions*, Proc. Natl. Acad. Sci. USA, 99 (2002), pp. 10246–10251.
- [4] G. BEYLKIN AND M. J. MOHLENKAMP, *Algorithms for numerical analysis in high dimensions*, SIAM J. Sci. Comput., 26 (2005), pp. 2133–2159.
- [5] D. J. BIAGIONI, D. BEYLKIN, AND G. BEYLKIN, *Randomized interpolative decomposition of separated representations*, Journal of Computational Physics, 281 (2015), pp. 116–134.
- [6] R. BRO, *Parafac. Tutorial & Applications.*, in Chemom. Intell. Lab. Syst., Special Issue 2nd Internet Conf. in Chemometrics (incinc’96), vol. 38, 1997, pp. 149–171.
- [7] J. D. CARROLL AND J. J. CHANG, *Analysis of individual differences in multidimensional scaling via an N-way generalization of Eckart-Young decomposition*, Psychometrika, 35 (1970), pp. 283–320.
- [8] Z. CHEN AND J. DONGARRA, *Condition number of gaussian random matrices*, SIAM J. Matrix Anal. Appl., 27 (2005), pp. 603–620.
- [9] P. CHEW, B. BADER, T. KOLDA, AND A. ABDELALI, *Cross-language information retrieval using parafac2*, in Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’07, New York, NY, USA, 2007, ACM, pp. 143–152, <http://dx.doi.org/10.1145/1281192.1281211>.
- [10] F. CHINESTA, P. LADEVEZE, AND E. CUETO, *A short review on model order reduction based on proper generalized decomposition*, Archives of Computational Methods in Engineering, 18 (2011), pp. 395–404.
- [11] A. DOOSTAN AND G. IACCARINO, *A least-squares approximation of partial differential equations with high-dimensional random inputs*, Journal of Computational Physics, 228 (2009), pp. 4332–4345.
- [12] A. DOOSTAN, G. IACCARINO, AND N. ETEMADI, *A least-squares approximation of high-dimensional uncertain systems*, Tech. Report Annual Research Brief, Center for Turbulence Research, Stanford University, 2007.
- [13] A. DOOSTAN, A. VALIDI, AND G. IACCARINO, *Non-intrusive low-rank separated approximation of high-dimensional stochastic models*, Comput. Methods Appl. Mech. Engrg., 263 (2013), pp. 42–55.
- [14] A. EDELMAN, *Eigenvalues and condition numbers of random matrices*, SIAM J. Matrix Anal. Appl., 9 (1988), pp. 543–560.
- [15] A. EDELMAN, *Eigenvalues and condition numbers of random matrices*, ph.d. thesis, Massachusetts Institute of Technology, 1989.
- [16] G. GOLUB AND C. V. LOAN, *Matrix Computations*, Johns Hopkins University Press, 3rd ed., 1996.
- [17] L. GRASEDYCK, D. KRESSNER, AND C. TOBLER, *A literature survey of low-rank tensor approximation techniques*, CoRR, abs/1302.7121 (2013).

- [18] W. HACKBUSCH, *Tensor spaces and numerical tensor calculus*, vol. 42, Springer, 2012.
- [19] M. HADIGOL, A. DOOSTAN, H. MATTHIES, AND R. NIEKAMP, *Partitioned treatment of uncertainty in coupled domain problems: A separated representation approach*, *Computer Methods in Applied Mechanics and Engineering*, 274 (2014), pp. 103–124.
- [20] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions*, *SIAM Review*, 53 (2011), pp. 217–288, <http://dx.doi.org/10.1137/090771806>.
- [21] R. A. HARSHMAN, *Foundations of the Parafac procedure: model and conditions for an “explanatory” multi-mode factor analysis*, Working Papers in Phonetics 16, UCLA, 1970.
- [22] R. A. HORN AND C. R. JOHNSON, *Topics in matrix analysis*, Cambridge Univ. Press, Cambridge, 1994.
- [23] B. KHOROMSKIJ AND C. SCHWAB, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, *SIAM Journal on Scientific Computing*, 33 (2011), pp. 364–385.
- [24] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, *SIAM Review*, 51 (2009), pp. 455–500, <http://dx.doi.org/10.1137/07070111X>.
- [25] A. NOUY, *A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations*, *Computer Methods in Applied Mechanics and Engineering*, 196 (2007), pp. 4521–4537.
- [26] A. NOUY, *Generalized spectral decomposition method for solving stochastic finite element equations: Invariant subspace problem and dedicated algorithms*, *Computer Methods in Applied Mechanics and Engineering*, 197 (2008), pp. 4718–4736.
- [27] A. NOUY, *Proper generalized decompositions and separated representations for the numerical solution of high dimensional stochastic problems*, *Archives of Computational Methods in Engineering*, 17 (2010), pp. 403–434.
- [28] I. V. OSELEDETS, *Tensor-train decomposition*, *SIAM Journal on Scientific Computing*, 33 (2011), pp. 2295–2317.
- [29] V. ROKHLIN AND M. TYGERT, *A fast randomized algorithm for overdetermined linear least-squares regression*, in *Proceedings of the National Academy of Sciences*, vol. 105(36), National Academy of Sciences, September 2008, pp. 13212–13217.
- [30] T. SARLOS, *Improved approximation algorithms for large matrices via random projections*, in *Foundations of Computer Science, 2006. FOCS '06. 47th Annual IEEE Symposium on*, 2006, pp. 143–152, <http://dx.doi.org/10.1109/FOCS.2006.37>.
- [31] G. TOMASI AND R. BRO, *A comparison of algorithms for fitting the PARAFAC model*, *Comput. Statist. Data Anal.*, 50 (2006), pp. 1700–1734.
- [32] R. VERSHYNIN, *Introduction to the non-asymptotic analysis of random matrices*, in *Compressed Sensing, Theory and Applications*, Y. Eldar and G. Kutyniok, eds., Cambridge University Press, 2012, ch. 5, pp. 210–268.
- [33] F. WOOLFE, E. LIBERTY, V. ROKHLIN, AND M. TYGERT, *A fast randomized algorithm for the approximation of matrices*, *Appl. Comput. Harmon. Anal.*, 25 (2008), pp. 335–366, <http://dx.doi.org/10.1016/j.acha.2007.12.002>.