

LOWER MEMORY OBLIVIOUS (TENSOR) SUBSPACE EMBEDDINGS WITH FEWER RANDOM BITS: MODEWISE METHODS FOR LEAST SQUARES

M. A. IWEN, D. NEEDELL, E. REBROVA, AND A. ZARE

ABSTRACT. In this paper new general modewise Johnson-Lindenstrauss (JL) subspace embeddings are proposed that are both considerably faster to generate and easier to store than traditional JL embeddings when working with extremely large vectors and/or tensors.

Corresponding embedding results are then proven for two different types of low-dimensional (tensor) subspaces. The first of these new subspace embedding results produces improved space complexity bounds for embeddings of rank- r tensors whose CP decompositions are contained in the span of a fixed (but unknown) set of r rank-one basis tensors. In the traditional vector setting this first result yields new and very general near-optimal oblivious subspace embedding constructions that require fewer random bits to generate than standard JL embeddings when embedding subspaces of \mathbb{C}^N spanned by basis vectors with special Kronecker structure. The second result proven herein provides new fast JL embeddings of arbitrary r -dimensional subspaces $\mathcal{S} \subset \mathbb{C}^N$ which also require fewer random bits (and so are easier to store – i.e., require less space) than standard fast JL embedding methods in order to achieve small ε -distortions. These new oblivious subspace embedding results work by (i) effectively folding any given vector in \mathcal{S} into a (not necessarily low-rank) tensor, and then (ii) embedding the resulting tensor into \mathbb{C}^m for $m \leq Cr \log^c(N)/\varepsilon^2$.

Applications related to compression and fast compressed least squares solution methods are also considered, including those used for fitting low-rank CP decompositions, and the proposed JL embedding results are shown to work well numerically in both settings.

1. MOTIVATION AND APPLICATIONS

Due to the recent explosion of massively large-scale data, the need for geometry preserving dimension reduction has become important in a wide array of applications in signal processing (see e.g. [21, 20, 3, 55, 25, 12]) and data science (see e.g. [6, 13]). This reduction is possible even on large dimensional objects when the class of such objects possesses some sort of lower dimensional intrinsic structure. For example, in classical compressed sensing [21, 20] and its related streaming applications [16, 17, 24, 30], the signals of interest are *sparse* vectors – vectors whose entries are mostly zero. In matrix recovery [13, 44], one often analogously assumes that the underlying matrix is low-rank. Under such models, tools like the Johnson-Lindenstrauss lemma [32, 2, 18, 36, 37] and the related restricted isometry property [14, 5] ask that the geometry of the signals be preserved after projection into a lower dimensional space. Typically, such projections are obtained via random linear maps that map into a dimension much smaller than the ambient dimension of the domain; s -sparse n -dimensional vectors can be projected into a dimension that scales like $s \log(n)$ and $n \times n$ rank- r matrices can be recovered from $O(rn)$ linear measurements [21, 20, 13]. Then, inference tasks or reconstruction can be performed from those lower dimensional representations.

Here, our focus is on dimension reduction of *tensors*, multi-way arrays that appear in an abundance of large-scale applications ranging from video and longitudinal imaging [39, 9] to machine learning [45, 51] and differential equations [8, 40]. Although a natural extension beyond matrices, their complicated structure leads to challenges both in defining low dimensional structure as well as dimension reduction projections. In particular, there are many notions of tensor rank, and various techniques exist to compute the corresponding decompositions [35, 54]. In this paper, we focus on tensors with low CP-rank, tensors that can be written as a sum of a few rank-1 tensors written as

outer products of basis vectors. The CP-rank and CP-decompositions are natural extensions of matrix rank and SVD, and are well motivated by applications such as topic modeling, psychometrics, signal processing, linguistics and many others [15, 26, 4]. Although there are now some nice results for low-rank tensor dimension reduction (see e.g. [43, 38, 48]), these give theoretical guarantees for dimensional reducing projections that act on tensors via their matricizations or vectorizations. Here, our goal is to provide similar guarantees but for projections that act directly on the tensors themselves without the need for unfolding. In particular, this means the projections can be defined *modewise* using the CP-decomposition, and that the low dimensional representations are also tensors, not vectors. This extends the application for such embeddings to those that cannot afford to perform unfoldings or for which it is not natural to do so. In particular, for tensors in \mathbb{C}^{n^d} for large n and d , this avoids having to store an often impossibly large $m \times n^d$ linear map. We elaborate on our main contributions next.

1.1. Contributions and Related Work. In this paper we analyze modewise tensor embedding strategies for general d -mode tensors similar to those introduced and analyzed for 2-mode tensors in [47]. In particular, herein we focus on obliviously embedding an apriori unknown r -dimensional subspace of a given tensor product space $\mathbb{C}^{n_1 \times \dots \times n_d}$ into a similarly low-dimensional vector space $\mathbb{C}^{\tilde{\mathcal{O}}(r)}$ with high probability. In contrast to the standard approach of effectively vectorizing the tensor product space and then embedding the resulting transformed subspace using standard JL methods involving a single massive $\tilde{\mathcal{O}}(r) \times \prod_{j=1}^d n_j$ matrix \mathbf{M} (see, e.g., [38]), the approaches considered herein instead result in the need to generate and store $d+1$ significantly smaller matrices $\mathbf{A} \in \mathbb{C}^{\tilde{\mathcal{O}}(r) \times \prod_{\ell=1}^d m_\ell}$, $\mathbf{A}_1 \in \mathbb{C}^{m_1 \times n_1}, \dots, \mathbf{A}_d \in \mathbb{C}^{m_d \times n_d}$ which are then combined to form a linear embedding operator $L : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{\tilde{\mathcal{O}}(r)}$ via

$$(1) \quad L(\mathcal{Z}) := \mathbf{A} (\text{vect} (\mathcal{Z} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d)),$$

where each \times_j is a j -mode product (reviewed below in §2.1), and $\text{vect} : \mathbb{C}^{m_1 \times \dots \times m_d} \rightarrow \mathbb{C}^{\prod_{\ell=1}^d m_\ell}$ is a trivial vectorization operator.

Let $m' = \tilde{\mathcal{O}}(r)$ be the number of rows one must use for both \mathbf{M} and \mathbf{A} above (as we shall see, the number of rows required for both matrices will indeed be essentially equivalent). The collective sizes of the matrices needed to define L above will be much smaller (and therefore easier to store, transmit, and generate) than \mathbf{M} whenever $\prod_{\ell=1}^d m_\ell + \sum_{\ell=1}^d n_\ell \left(\frac{m_\ell}{m'} \right) \ll \prod_{j=1}^d n_j$ holds. As a result, much of our discussion below will revolve around bounding the dominant $\prod_{\ell=1}^d m_\ell$ term on the left hand side above, which will also occasionally be referred to as the *intermediate embedding dimension* below. We are now prepared to discuss our two main results.

1.1.1. General Oblivious Subspace Embedding Results for Low Rank Tensor Subspaces Satisfying an Incoherence Condition. The first of our results provides new oblivious subspace embeddings for tensor subspaces spanned by bases of rank one tensors, as well as establishes related least squares embedding results of value in, e.g., the fitting of a general tensor with an accurate low rank CPD approximation. One of its main contributions is the generality with which it allows one to select the matrices $\mathbf{A}, \mathbf{A}_1, \dots, \mathbf{A}_d$ used to construct the JL embedding L in (1). In particular, it allows each of these matrices to be drawn independently from any desired nearly-optimal family of JL embeddings (as defined immediately below) that the user likes.

Definition 1 (ε -JL embedding). *Let $\varepsilon \in (0, 1)$. We will call a matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ an ε -JL embedding of a set $S \subset \mathbb{C}^n$ into \mathbb{C}^m if*

$$\|\mathbf{A}\mathbf{x}\|_2^2 = (1 + \varepsilon_{\mathbf{x}})\|\mathbf{x}\|_2^2$$

holds for some $\varepsilon_{\mathbf{x}} \in (-\varepsilon, \varepsilon)$ for all $\mathbf{x} \in S$.

Definition 2. Fix $\eta \in (0, 1/2)$ and let $\{\mathcal{D}_{(m,n)}\}_{(m,n) \in \mathbb{N} \times \mathbb{N}}$ be a family of probability distributions where each $\mathcal{D}_{(m,n)}$ is a distribution over $m \times n$ matrices. We will refer to any such family of distributions as being an **η -optimal family of JL embedding distributions** if there exists an absolute constant $C \in \mathbb{R}^+$ such that, for any given $\varepsilon \in (0, 1)$, $m, n \in \mathbb{N}$ with $m < n$, and nonempty set $\mathcal{S} \subset \mathbb{C}^n$ of cardinality

$$|S| \leq \eta \exp\left(\frac{\varepsilon^2 m}{C}\right),$$

a matrix $\mathbf{A} \sim \mathcal{D}_{(m,n)}$ will be an ε -JL embedding of \mathcal{S} into \mathbb{C}^m with probability at least $1 - \eta$.

In fact many η -optimal families of JL embedding distributions exist for any given $\eta \in (0, 1/2)$ including, e.g., those associated with random matrices having i.i.d. subgaussian entries (see Lemma 9.35 in [22]) as well as those associated with sparse JLT constructions [33]. The next theorem proves that any desired combination of such matrices can be used to construct a JL embedding L as per (1) for any tensor subspace spanned by a basis of rank one tensors satisfying an easily testable (and relatively mild¹) coherence condition. We utilize the notations set forth below in Section 2.

Theorem 1. Fix $\varepsilon, \eta \in (0, 1/2)$ and $d \geq 3$. Let $\mathcal{X} \in \mathbb{C}^{n_1 \times \dots \times n_d}$, $n := \max_j n_j \geq 4r + 1$, and \mathcal{L} be an r -dimensional subspace of $\mathbb{C}^{n_1 \times \dots \times n_d}$ spanned by a basis of rank one tensors $\mathcal{B} := \left\{ \bigcup_{\ell=1}^d \mathbf{y}_k^{(\ell)} \mid k \in [r] \right\}$ with modewise coherence satisfying

$$\mu_{\mathcal{B}}^{d-1} := \left(\max_{\ell \in [d]} \max_{k, h \in [r], k \neq h} \left| \langle \mathbf{y}_k^{(\ell)}, \mathbf{y}_h^{(\ell)} \rangle \right| \right)^{d-1} < 1/2r.$$

Then, one can construct a linear operator $L : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{m'}$ as per (1) with $m' \leq C'r \cdot \varepsilon^{-2} \cdot \ln\left(\frac{47}{\varepsilon \sqrt{\eta}}\right)$ for an absolute constant $C' \in \mathbb{R}^+$ so that with probability at least $1 - \eta$

$$(2) \quad \left| \|L(\mathcal{X} - \mathcal{Y})\|_2^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \leq \varepsilon \|\mathcal{X} - \mathcal{Y}\|^2$$

will hold for all $\mathcal{Y} \in \mathcal{L}$.

If $\mathcal{X} \notin \mathcal{L}$ the intermediate embedding dimension can be bounded above by

$$(3) \quad \prod_{\ell=1}^d m_{\ell} \leq C^d \cdot r^d d^{3d} / \varepsilon^{2d} \cdot \ln^d(n / \sqrt[d]{\eta})$$

for an absolute constant $C \in \mathbb{R}^+$. If, however, $\mathcal{X} \in \mathcal{L}$ then (2) holds for all $r < 1/2\mu_{\mathcal{B}}^{d-1}$ and

$$(4) \quad \prod_{\ell=1}^d m_{\ell} \leq \tilde{C}^d \cdot r^2 (d/\varepsilon)^{2d} \cdot \ln^d(2r^2 d/\eta)$$

can be achieved, where $\tilde{C} \in \mathbb{R}^+$ is another absolute constant.

Proof. This is a largely restatement of Theorem 6. When defining $L : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{m'}$ as per (1) following Theorem 6 one should draw $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ with $m_j \geq C_j \cdot r d^3 / \varepsilon^2 \cdot \ln(n / \sqrt[d]{\eta})$ from an $(\eta/4d)$ -optimal family of JL embedding distributions for each $j \in [d]$, where each $C_j \in \mathbb{R}^+$ is an absolute constant. Furthermore, $\mathbf{A} \in \mathbb{C}^{m' \times \prod_{\ell=1}^d m_{\ell}}$ should be drawn from an $(\eta/2)$ -optimal family of JL embedding distributions with m' as above. The probability bound together with (3) both then follow. The achievable intermediate embedding dimension when $\mathcal{X} \in \mathcal{L}$ in (4) can be obtained from Corollary 2 since the bound $\prod_{\ell=1}^d m_{\ell} \leq \prod_{\ell=1}^d \tilde{C}_{\ell} \cdot r^{2/d} d^2 / \varepsilon^2 \cdot \ln(2r^2 d/\eta)$ can then be utilized in that case. \square

¹In fact the coherence condition required by Theorem 1 will be satisfied by a generic basis of rank 1 tensors with high probability (see §3.2). Similar coherence results to those presented in §3.2 have also recently been considered for random tensors in more general parameter regimes by Vershynin [53].

One can vectorize the tensors and tensor spaces considered in Theorem 1 using variants of (14) to achieve subspace embedding results for subspaces spanned by basis vectors with special Kronecker structure as considered in, e.g., two other recent papers that appeared during the preparation of this manuscript [31, 41]. The most recent of these papers also produces bounds on what amounts to the intermediate embedding dimension of a JL subspace embedding along the lines of (1) when $\mathcal{X} \in \mathcal{L}$ (see Theorem 4.1 in [41]). Comparing (4) to that result we can see that Theorem 1 has reduced the r dependence of the effective intermediate embedding dimension achieved therein from r^{d+1} to r^2 (now independent of d) for a much more general set of modewise embeddings. However, Theorem 1 incurs a worse dependence on epsilon and needs the stated coherence assumption concerning $\mu_{\mathcal{B}}$ to hold. As a result, Theorem 1 provides a large new class of modewise subspace embeddings that will also have fewer rows than those in [41] for a large range of ranks r provided that $\mu_{\mathcal{B}}$ is sufficiently small and ε is sufficiently large.

Note further that the form of (2) also makes Theorem 1 useful for solving least squares problems of the type encountered while computing approximate CP decompositions for an arbitrary tensor $\mathcal{X} \notin \mathcal{L}$ using alternating least squares methods (see, e.g., §4 for a related discussion as well as [7] where modewise strategies were shown to work well for solving such problems in practice). Comparing Theorem 1 to the recent least squares result of the same kind proven in [31] (see Corollary 2.4) we can see that Theorem 1 has reduced the r dependence of the effective intermediate embedding dimension achievable in [31] from r^{2d} therein to r^d in (3) for a much more general set of modewise embeddings. In exchange, Theorem 1 again incurs a worse dependence on epsilon and needs the stated coherence assumption concerning $\mu_{\mathcal{B}}$ to hold, however. As a result, Theorem 1 guarantees that a larger class of modewise JL embeddings can be used in least squares applications, and that they will also have smaller intermediate embedding dimensions as long as $\mu_{\mathcal{B}}$ is sufficiently small and ε sufficiently large.

1.1.2. Fast Oblivious Subspace Embedding Results for Arbitrary Tensor Subspaces. Our second main result builds on Theorem 2.1 of Jin, Kolda, and Ward in [31] to provide improved fast subspace embedding results for arbitrary tensor subspaces (i.e., for low dimensional tensor subspaces whose basis tensors have arbitrary rank and coherence). Let $N := \prod_{j=1}^d n_j$. By combining elements the proof of Theorem 1 with the optimal ε -dependence of Theorem 2.1 in [31] we are able to provide a fast modewise oblivious subspace embedding L as per (1) that will simultaneously satisfy (2) for all \mathcal{Y} in an entirely arbitrary r -dimensional tensor subspace \mathcal{L} with probability at least $1 - \eta$ while also achieving an intermediate embedding dimension bounded above by

$$(5) \quad C^d \left(\frac{r}{\varepsilon} \right)^2 \cdot \log^{2d-1} \left(\frac{N}{\eta} \right) \cdot \log^4 \left(\frac{\log \left(\frac{N}{\eta} \right)}{\varepsilon} \right) \cdot \log N.$$

Above $C > 0$ is an absolute constant. Note that neither r nor ε in (5) are raised to a power of d which marks a tremendous improvement over all of the previously discussed results when d is large. See Theorem 8 for details.

As alluded to above the results herein can also be used to create new JL subspace embeddings in the traditional vector space setting. Our next and final main result does this explicitly for arbitrary vector subspaces by restating a variant of Theorem 8 in that context. We expect that this result may be of independent interest outside of the tensor setting.

Theorem 2. *Fix $\varepsilon, \eta \in (0, 1/2)$ and $d \geq 2$. Let $\mathbf{x} \in \mathbb{C}^N$ such that $\sqrt[d]{N} \in \mathbb{N}$ and $N \geq 4C'/\eta > 1$ for an absolute constant $C' > 0$, and let \mathcal{L} be an r -dimensional subspace of \mathbb{C}^N for $\max(2r^2 - r, 4r) \leq$*

N . Then, one can construct a random matrix $\mathbf{A} \in \mathbb{C}^{m \times N}$ with

$$(6) \quad m \leq C \left[r \cdot \varepsilon^{-2} \cdot \log \left(\frac{47}{\varepsilon \sqrt[4]{\eta}} \right) \cdot \log^4 \left(\frac{r \log \left(\frac{47}{\varepsilon \sqrt[4]{\eta}} \right)}{\varepsilon} \right) \cdot \log N \right],$$

for an absolute constant $C > 0$ such that with probability at least $1 - \eta$ it will be the case that

$$\left| \|\mathbf{A}(\mathbf{x} - \mathbf{y})\|_2^2 - \|\mathbf{x} - \mathbf{y}\|_2^2 \right| \leq \varepsilon \|\mathbf{x} - \mathbf{y}\|_2^2$$

holds for all $\mathbf{y} \in \mathcal{L}$. Furthermore, \mathbf{A} requires only

$$(7) \quad \mathcal{O} \left(C_1^d \left(\frac{r}{\varepsilon} \right)^2 \cdot \log^{2d-1} \left(\frac{N}{\eta} \right) \cdot \log^4 \left(\frac{\log \left(\frac{N}{\eta} \right)}{\varepsilon} \right) \cdot \log^2 N + d \sqrt[d]{N} \right)$$

random bits and memory for storage for an absolute constant $C_1 > 0$, and can be multiplied against any vector in just $\mathcal{O}(N \log N)$ -time.

Note that choosing $\mathbf{x} = \mathbf{0}$ produces an oblivious subspace embedding result for \mathcal{L} , and that choosing \mathcal{L} to be the column space of a rank r matrix produces a result useful for least squares sketching.

Proof. This follows from Theorem 8 after identifying \mathbb{C}^N with $\mathbb{C}^{\sqrt[d]{N} \times \dots \times \sqrt[d]{N}}$ (i.e., after effectively reshaping any given vectors \mathbf{x}, \mathbf{y} under consideration into d -mode tensors \mathcal{X}, \mathcal{Y} .) Note further that if $\sqrt[d]{N} \notin \mathbb{N}$ then one can implicitly pad the vectors of interest with zeros until it is (i.e., effectively trivially embedding \mathbb{C}^N into $\mathbb{C}^{\lceil \sqrt[d]{N} \rceil^d}$) before preceding. \square

1.2. Organization. The remainder of the paper is organized as follows. Section 2 provides background and notation for tensors (Subsections 2 and 2.1), as well as for Johnson-Lindenstrauss embeddings (Subsection 2.2).

We start Section 3 with the definitions of the rank of the tensor (and low-rank tensor subspaces) and the maximal modewise coherence of tensor subspace bases. Then we work our way to Corollary 2, that gives our first main result on oblivious tensor subspace embeddings via modewise tensor products (for any fixed subspace having low enough modewise coherence). This result is very general in terms of JL-embedding maps one can use as building blocks in each mode. Finally, in Subsection 3.2 we discuss the assumption of modewise incoherence and provide several natural examples of incoherent tensor subspaces.

In Section 4 we describe the fitting problem for the approximately low rank tensors and explain how modewise dimension reduction (as presented in Section 3) reduces the complexity of the problem. Then we build the machinery to show that the solution of the reduced problem will be a good solution for the original problem (in Theorem 6). We conclude Section 4 by introducing a two-step embedding procedure that allows one to further reduce the final embedding dimension (this our second main embedding result, Theorem 8). This improved procedure relies on a specific form of JL-embedding of each mode. Both embedding results can be applied to the fitting problem.

In Section 5 we present some simple experiments confirming our theoretical guarantees, and then we conclude in Section 6.

2. NOTATION, TENSOR BASICS, & LINEAR JOHNSON-LINDENSTRAUSS EMBEDDINGS

Tensors, matrices, vectors and scalars are denoted in different typeface for clarity below. Calligraphic boldface capital letters are always used for tensors, boldface capital letters for matrices, boldface lower-case letters for vectors, and regular (lower-case or capital) letters for scalars. The matrix \mathbf{I} will always represent the identity matrix. The set of the the first d natural numbers will be denoted by $[d] := \{1, \dots, d\}$ for all $d \in \mathbb{N}$.

Throughout the paper, \otimes denotes the Kronecker product of vectors or matrices, and \bigcirc denotes the tensor outer product of vectors or tensors.² The symbol \circ on the other hand represents the composition of functions (see e.g. Section 4). Numbers in parentheses used as a subscript or superscript on a tensor either denote *unfoldings* (introduced in Section 2.1) when appearing in a subscript, or else an element in a sequence when appearing in a superscript. The notation $\otimes_{\ell \neq j} \mathbf{v}^{(\ell)}$ for a given set of vectors $\{\mathbf{v}^{(\ell)}\}_{\ell=1}^d$ will always denote the vector $\mathbf{v}^{(d)} \otimes \dots \mathbf{v}^{(j+1)} \otimes \mathbf{v}^{(j-1)} \dots \otimes \mathbf{v}^{(1)}$. Additional tensor definitions and operations are reviewed below (see, e.g., [35, 19, 50, 54] for additional details and discussion).

2.1. Tensor Basics. The set of all d -mode tensors $\mathcal{X} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ forms a vector space over the complex numbers when equipped with component-wise addition and scalar multiplication. The inner product of $\mathcal{X}, \mathcal{Y} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ will be given by

$$(8) \quad \langle \mathcal{X}, \mathcal{Y} \rangle := \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \dots \sum_{i_d=1}^{n_d} \mathcal{X}_{i_1, i_2, \dots, i_d} \overline{\mathcal{Y}_{i_1, i_2, \dots, i_d}}.$$

This inner product then gives rise to the standard Euclidean norm

$$(9) \quad \|\mathcal{X}\| := \sqrt{\langle \mathcal{X}, \mathcal{X} \rangle} = \sqrt{\sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \dots \sum_{i_d=1}^{n_d} |\mathcal{X}_{i_1, i_2, \dots, i_d}|^2}.$$

If $\langle \mathcal{X}, \mathcal{Y} \rangle = 0$ we say that \mathcal{X} and \mathcal{Y} are *orthogonal*. If \mathcal{X} and \mathcal{Y} are orthogonal and also have unit norm (i.e., have $\|\mathcal{X}\| = \|\mathcal{Y}\| = 1$) we say that they are *orthonormal*.

Tensor outer products: The *tensor outer product* of two tensors $\mathcal{X} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ and $\mathcal{Y} \in \mathbb{C}^{n'_1 \times n'_2 \times \dots \times n'_{d'}}$, $\mathcal{X} \bigcirc \mathcal{Y} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d \times n'_1 \times n'_2 \times \dots \times n'_{d'}}$, is a $(d + d')$ -mode tensor whose entries are given by

$$(10) \quad (\mathcal{X} \bigcirc \mathcal{Y})_{i_1, \dots, i_d, i'_1, \dots, i'_{d'}} = \mathcal{X}_{i_1, \dots, i_d} \mathcal{Y}_{i'_1, \dots, i'_{d'}}.$$

Note that when \mathcal{X} and \mathcal{Y} are both vectors, the tensor outer product will reduce to the standard outer product.

Lemma 1. Let $\alpha, \beta \in \mathbb{C}$, $\mathcal{A}, \mathcal{B} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ and $\mathcal{C}, \mathcal{D} \in \mathbb{C}^{n'_1 \times n'_2 \times \dots \times n'_{d'}}$. Then,

$$(\dagger) \quad (\alpha \mathcal{A} + \beta \mathcal{B}) \bigcirc \mathcal{C} = \alpha \mathcal{A} \bigcirc \mathcal{C} + \beta \mathcal{B} \bigcirc \mathcal{C} = \mathcal{A} \bigcirc \alpha \mathcal{C} + \mathcal{B} \bigcirc \beta \mathcal{C}.$$

$$(\dagger\dagger) \quad \langle \mathcal{A} \bigcirc \mathcal{C}, \mathcal{B} \bigcirc \mathcal{D} \rangle = \langle \mathcal{A}, \mathcal{B} \rangle \langle \mathcal{C}, \mathcal{D} \rangle.$$

Proof. The first property follows from the fact that

$$((\alpha \mathcal{A} + \beta \mathcal{B}) \bigcirc \mathcal{C})_{i_1, \dots, i_d, i'_1, \dots, i'_{d'}} = (\alpha \mathcal{A} + \beta \mathcal{B})_{i_1, \dots, i_d} \mathcal{C}_{i'_1, \dots, i'_{d'}} = (\alpha \mathcal{A}_{i_1, \dots, i_d} + \beta \mathcal{B}_{i_1, \dots, i_d}) \mathcal{C}_{i'_1, \dots, i'_{d'}}.$$

²As (10) suggests, it can be applied to tensors with arbitrary number of modes.

To establish the second property we note that

$$\begin{aligned}
\langle \mathcal{A} \bigcirc \mathcal{C}, \mathcal{B} \bigcirc \mathcal{D} \rangle &= \sum_{i_1=1}^{n_1} \cdots \sum_{i_d=1}^{n_d} \sum_{i'_1=1}^{n'_1} \cdots \sum_{i'_{d'}=1}^{n'_{d'}} \mathcal{A}_{i_1, i_2, \dots, i_d} \mathcal{C}_{i'_1, \dots, i'_{d'}} \overline{\mathcal{B}_{i_1, i_2, \dots, i_d}} \overline{\mathcal{D}_{i'_1, \dots, i'_{d'}}} \\
&= \left(\sum_{i_1=1}^{n_1} \cdots \sum_{i_d=1}^{n_d} \mathcal{A}_{i_1, i_2, \dots, i_d} \overline{\mathcal{B}_{i_1, i_2, \dots, i_d}} \right) \left(\sum_{i'_1=1}^{n'_1} \cdots \sum_{i'_{d'}=1}^{n'_{d'}} \mathcal{C}_{i'_1, \dots, i'_{d'}} \overline{\mathcal{D}_{i'_1, \dots, i'_{d'}}} \right) \\
&= \langle \mathcal{A}, \mathcal{B} \rangle \langle \mathcal{C}, \mathcal{D} \rangle.
\end{aligned}$$

□

Fibers: Let tensor $\mathcal{X} \in \mathbb{C}^{n_1 \times \cdots \times n_{j-1} \times n_j \times n_{j+1} \times \cdots \times n_d}$. The vectors in \mathbb{C}^{n_j} obtained by fixing all of the indices of \mathcal{X} except for the one that corresponds to its j^{th} mode are called its *mode- j fibers*. Note that any such \mathcal{X} will have $\prod_{\ell \neq j} n_\ell$ mode- j fibers denoted by $\mathcal{X}_{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_d} \in \mathbb{C}^{n_j}$.

Tensor matricization (unfolding): The process of reordering the elements of the tensor into a matrix is known as matricization or unfolding. The mode- j matricization of a tensor $\mathcal{X} \in \mathbb{C}^{n_1 \times n_2 \times \cdots \times n_d}$ is denoted as $\mathbf{X}_{(j)} \in \mathbb{C}^{n_j \times \prod_{m \neq j} n_m}$ and is obtained by arranging \mathcal{X} 's mode- j fibers to be the columns of the resulting matrix.

j -mode products: The j -mode product of a d -mode tensor $\mathcal{X} \in \mathbb{C}^{n_1 \times \cdots \times n_{j-1} \times n_j \times n_{j+1} \times \cdots \times n_d}$ with a matrix $\mathbf{U} \in \mathbb{C}^{m_j \times n_j}$ is another d -mode tensor $\mathcal{X} \times_j \mathbf{U} \in \mathbb{C}^{n_1 \times \cdots \times n_{j-1} \times m_j \times n_{j+1} \times \cdots \times n_d}$. Its entries are given by

$$(11) \quad (\mathcal{X} \times_j \mathbf{U})_{i_1, \dots, i_{j-1}, \ell, i_{j+1}, \dots, i_d} = \sum_{i_j=1}^{n_j} \mathcal{X}_{i_1, \dots, i_j, \dots, i_d} \mathbf{U}_{\ell, i_j}$$

for all $(i_1, \dots, i_{j-1}, \ell, i_{j+1}, \dots, i_d) \in [n_1] \times \cdots \times [n_{j-1}] \times [m_j] \times [n_{j+1}] \times \cdots \times [n_d]$. Looking at the mode- j unfoldings of $\mathcal{X} \times_j \mathbf{U}$ and \mathcal{X} one can easily see that their j -mode matricization can be computed as a regular matrix product

$$(12) \quad (\mathcal{X} \times_j \mathbf{U})_{(j)} = \mathbf{U} \mathbf{X}_{(j)}$$

for all $j \in [d]$. The following simple lemma formally lists several important properties of mode-wise products.

Lemma 2. Let $\mathcal{X}, \mathcal{Y} \in \mathbb{C}^{n_1 \times n_2 \times \cdots \times n_d}$, $\alpha, \beta \in \mathbb{C}$, and $\mathbf{U}_\ell, \mathbf{V}_\ell \in \mathbb{C}^{m_\ell \times n_\ell}$ for all $\ell \in [d]$. The following four properties hold:

$$(\dagger) \quad (\alpha \mathcal{X} + \beta \mathcal{Y}) \times_j \mathbf{U}_j = \alpha (\mathcal{X} \times_j \mathbf{U}_j) + \beta (\mathcal{Y} \times_j \mathbf{U}_j).$$

$$(\dagger\dagger) \quad \mathcal{X} \times_j (\alpha \mathbf{U}_j + \beta \mathbf{V}_j) = \alpha (\mathcal{X} \times_j \mathbf{U}_j) + \beta (\mathcal{X} \times_j \mathbf{V}_j).$$

$$(\dagger\dagger\dagger) \quad \text{If } j \neq \ell \text{ then } \mathcal{X} \times_j \mathbf{U}_j \times_\ell \mathbf{V}_\ell = (\mathcal{X} \times_j \mathbf{U}_j) \times_\ell \mathbf{V}_\ell = (\mathcal{X} \times_\ell \mathbf{V}_\ell) \times_j \mathbf{U}_j = \mathcal{X} \times_\ell \mathbf{V}_\ell \times_j \mathbf{U}_j.$$

$$(\dagger\dagger\dagger\dagger) \quad \text{If } W \in \mathbb{C}^{p \times m_j} \text{ then } \mathcal{X} \times_j \mathbf{U}_j \times_j \mathbf{W} = (\mathcal{X} \times_j \mathbf{U}_j) \times_j \mathbf{W} = \mathcal{X} \times_j (\mathbf{W} \mathbf{U}_j) = \mathcal{X} \times_j \mathbf{W} \mathbf{U}_j.$$

Proof. The first, second, and fourth facts above are easily established using mode- j unfoldings. To establish (\dagger) above, we note that

$$\begin{aligned}
((\alpha \mathcal{X} + \beta \mathcal{Y}) \times_j \mathbf{U}_j)_{(j)} &= \mathbf{U}_j (\alpha \mathcal{X} + \beta \mathcal{Y})_{(j)} = \mathbf{U}_j (\alpha \mathbf{X}_{(j)} + \beta \mathbf{Y}_{(j)}) \\
&= \alpha \mathbf{U}_j \mathbf{X}_{(j)} + \beta \mathbf{U}_j \mathbf{Y}_{(j)} = \alpha (\mathcal{X} \times_j \mathbf{U}_j)_{(j)} + \beta (\mathcal{Y} \times_j \mathbf{U}_j)_{(j)}.
\end{aligned}$$

Reshaping both sides of the derived equality back into their original tensor forms now completes the proof.³ The proof of (††) using unfoldings is nearly identical.

To prove (††††) we may again use mode- j unfoldings to see that

$$(\mathcal{X} \times_j \mathbf{U}_j \times_j \mathbf{W})_{(j)} = \mathbf{W} (\mathcal{X} \times_j \mathbf{U}_j)_{(j)} = \mathbf{W} \mathbf{U}_j \mathbf{X}_{(j)} = (\mathcal{X} \times_j \mathbf{W} \mathbf{U}_j)_{(j)}.$$

Reshaping these expressions back into their original tensor forms again completes the proof.

To prove (†††) it is perhaps easiest to appeal directly to the component-wise definition of the mode- j product given in equation (11). Suppose that $\ell > j$ (the case $\ell < j$ is nearly identical). Set $\mathbf{U} := \mathbf{U}_j$ and $\mathbf{V} := \mathbf{V}_\ell$ to simplify subscript notation. We have for all $k \in [m_j]$, $l \in [m_\ell]$, and $i_q \in [n_q]$ with $q \notin \{j, \ell\}$ that

$$\begin{aligned} ((\mathcal{X} \times_j \mathbf{U}) \times_\ell \mathbf{V})_{i_1, \dots, i_{j-1}, k, i_{j+1}, \dots, i_{\ell-1}, l, i_{\ell+1}, \dots, i_d} &= \sum_{i_\ell=1}^{n_\ell} (\mathcal{X} \times_j \mathbf{U})_{i_1, \dots, i_{j-1}, k, i_{j+1}, \dots, i_\ell, \dots, i_d} \mathbf{V}_{l, i_\ell} \\ &= \sum_{i_\ell=1}^{n_\ell} \left(\sum_{i_j=1}^{n_j} \mathcal{X}_{i_1, \dots, i_j, \dots, i_\ell, \dots, i_d} \mathbf{U}_{k, i_j} \right) \mathbf{V}_{l, i_\ell} \\ &= \sum_{i_j=1}^{n_j} \left(\sum_{i_\ell=1}^{n_\ell} \mathcal{X}_{i_1, \dots, i_j, \dots, i_\ell, \dots, i_d} \mathbf{V}_{l, i_\ell} \right) \mathbf{U}_{k, i_j} \\ &= \sum_{i_j=1}^{n_j} (\mathcal{X} \times_\ell \mathbf{V})_{i_1, \dots, i_j, \dots, i_{\ell-1}, l, i_{\ell+1}, \dots, i_d} \mathbf{U}_{k, i_j} \\ &= ((\mathcal{X} \times_\ell \mathbf{U}) \times_j \mathbf{U})_{i_1, \dots, i_{j-1}, k, i_{j+1}, \dots, i_{\ell-1}, l, i_{\ell+1}, \dots, i_d}. \end{aligned}$$

□

A generalization of the observation (12) is available: unfolding the tensor

$$(13) \quad \mathcal{Y} = \mathcal{X} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_d \mathbf{U}^{(d)} =: \mathcal{X} \bigtimes_{j=1}^d \mathbf{U}^{(j)},$$

along the j^{th} mode is equivalent to

$$(14) \quad \mathbf{Y}_{(j)} = \mathbf{U}^{(j)} \mathbf{X}_{(j)} \left(\mathbf{U}^{(d)} \otimes \dots \otimes \mathbf{U}^{(j+1)} \otimes \mathbf{U}^{(j-1)} \dots \otimes \mathbf{U}^{(1)} \right)^\top,$$

where \otimes is the matrix Kronecker product (see [35]). In particular, (14) implies that the matricization $(\mathcal{X} \times_j \mathbf{U}^{(j)})_{(j)} = \mathbf{U}^{(j)} \mathbf{X}_{(j)}$.⁴ On a related note, one can also express the relation between the vectorized forms of \mathcal{X} and \mathcal{Y} in (13) as

$$(15) \quad \text{vect}(\mathcal{Y}) = \left(\mathbf{U}^{(d)} \otimes \dots \otimes \mathbf{U}^{(1)} \right) \text{vect}(\mathcal{X}),$$

where $\text{vect}(\cdot)$ is the vectorization operator.

It is worth noting that trivial inner product preserving isomorphisms exist between a tensor space $\mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ and any of its matricized versions (i.e., mode- j matricization can be viewed as an isomorphism between the original tensor vector space $\mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ and its mode- j matricized target vector space $\mathbb{C}^{n_j \times \prod_{m \neq j} n_m}$). In particular, the process of matricizing tensors is linear. If,

³Here we are implicitly using that mode- j unfolding provides a vector space isomorphism between $\mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ and $\mathbb{C}^{n_j \times \prod_{\ell \in [d] \setminus \{j\}} n_\ell}$ for all $j \in [d]$.

⁴Simply set $\mathbf{U}^{(m)} = \mathbf{I}$ (the identity) for all $m \neq n$ in (14). This fact also easily follows directly from the definition of the j -mode product.

for example, $\mathcal{X}, \mathcal{Y} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ then one can see that the mode- j matricization of $\mathcal{X} + \mathcal{Y} \in \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ is $(\mathcal{X} + \mathcal{Y})_{(j)} = \mathbf{X}_{(j)} + \mathbf{Y}_{(j)}$ for all modes $j \in [d]$.

2.2. Linear Johnson-Lindenstrauss Embeddings. Many linear ε -JL embedding matrices exist [32, 2, 18, 36, 37] with the best achievable $m = \mathcal{O}(\log(|S|)/\varepsilon^2)$ for arbitrary S (see [37] for results concerning the optimality of this embedding dimension). Of course, one can define JL embedding on tensors in a similar way, namely, as linear maps approximately preserving tensor norm:

Definition 3 (Tensor ε -JL embedding). *A linear operator $L : \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d} \rightarrow \mathbb{C}^{m_1 \times \dots \times m_{d'}}$ is an ε -JL embedding of a set $S \subset \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ into $\mathbb{C}^{m_1 \times \dots \times m_{d'}}$ if*

$$\|L(\mathcal{X})\|^2 = (1 + \varepsilon_{\mathcal{X}}) \|\mathcal{X}\|^2$$

holds for some $\varepsilon_{\mathcal{X}} \in (-\varepsilon, \varepsilon)$ for all $\mathcal{X} \in S$.

It is easy to check that JL embeddings can preserve pairwise inner products.

Lemma 3. *Let $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$ and suppose that $\mathbf{A} \in \mathbb{C}^{m \times n}$ is an ε -JL embedding of the vectors*

$$\{\mathbf{x} - \mathbf{y}, \mathbf{x} + \mathbf{y}, \mathbf{x} - i\mathbf{y}, \mathbf{x} + i\mathbf{y}\} \subset \mathbb{C}^n$$

into \mathbb{C}^m . Then,

$$|\langle \mathbf{A}\mathbf{x}, \mathbf{A}\mathbf{y} \rangle - \langle \mathbf{x}, \mathbf{y} \rangle| \leq 2\varepsilon (\|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2) \leq 4\varepsilon \cdot \max \{\|\mathbf{x}\|_2^2, \|\mathbf{y}\|_2^2\}.$$

Proof. This well known result is an easy consequence of the polarization identity for inner products. We have that

$$\begin{aligned} |\langle \mathbf{A}\mathbf{x}, \mathbf{A}\mathbf{y} \rangle - \langle \mathbf{x}, \mathbf{y} \rangle| &= \left| \frac{1}{4} \sum_{\ell=0}^3 i^\ell \left(\|\mathbf{A}\mathbf{x} + i^\ell \mathbf{A}\mathbf{y}\|_2^2 - \|\mathbf{x} + i^\ell \mathbf{y}\|_2^2 \right) \right| = \left| \frac{1}{4} \sum_{\ell=0}^3 i^\ell \varepsilon_\ell \|\mathbf{x} + i^\ell \mathbf{y}\|_2^2 \right| \\ &\leq \frac{1}{4} \sum_{\ell=0}^3 \varepsilon (\|\mathbf{x}\|_2 + \|\mathbf{y}\|_2)^2 = \varepsilon (\|\mathbf{x}\|_2 + \|\mathbf{y}\|_2)^2 = \varepsilon (\|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 + 2\|\mathbf{x}\|_2 \|\mathbf{y}\|_2) \\ &\leq 2\varepsilon (\|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2) \leq 4\varepsilon \cdot \max \{\|\mathbf{x}\|_2^2, \|\mathbf{y}\|_2^2\}, \end{aligned}$$

where the second to last inequality follows from Young's inequality for products. \square

The fact that $\mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ is an inner product space means that the following trivial generalization of Lemma 3 to the tensor JL embeddings also holds.

Lemma 4. *Let $\mathcal{X}, \mathcal{Y} \in \mathbb{C}^n$ and suppose that L is an ε -JL embedding of the tensors*

$$\{\mathcal{X} - \mathcal{Y}, \mathcal{X} + \mathcal{Y}, \mathcal{X} - i\mathcal{Y}, \mathcal{X} + i\mathcal{Y}\} \subset \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$$

into $\mathbb{C}^{m_1 \times \dots \times m_{d'}}$. Then,

$$|\langle L(\mathcal{X}), L(\mathcal{Y}) \rangle - \langle \mathcal{X}, \mathcal{Y} \rangle| \leq 2\varepsilon (\|\mathcal{X}\|^2 + \|\mathcal{Y}\|^2) \leq 4\varepsilon \cdot \max \{\|\mathcal{X}\|^2, \|\mathcal{Y}\|^2\}.$$

Proof. The proof is similar to that of Lemma 3, with $L(\mathcal{X})$ replacing $\mathbf{A}\mathbf{x}$, and making use of the linearity of L . \square

In the case where a more general set S is embedded using JL embeddings, for example, a low-rank subspace of tensors, in order to pass to a smaller finite set, a discretization technique can be used. Due to linearity, it actually suffices to discretize the unit ball of the space in question. In the next lemma we present a simple subspace embedding result based on a standard covering argument (see, e.g., [5, 22]). We include its proof for the sake of completeness.

Lemma 5. Fix $\varepsilon \in (0, 1)$. Let \mathcal{L} be an r -dimensional subspace of \mathbb{C}^n , and let $\mathcal{C} \subset \mathcal{L}$ be an $(\varepsilon/16)$ -net of the $(r-1)$ -dimensional Euclidean unit sphere $\mathcal{S}_{\ell^2} \subset \mathcal{L}$. Then, if $\mathbf{A} \in \mathbb{C}^{m \times n}$ is an $(\varepsilon/2)$ -JL embedding of \mathcal{C} it will also satisfy

$$(16) \quad (1 - \varepsilon)\|\mathbf{x}\|_2^2 \leq \|\mathbf{Ax}\|_2^2 \leq (1 + \varepsilon)\|\mathbf{x}\|_2^2$$

for all $\mathbf{x} \in \mathcal{L}$. Furthermore, we note that there exists an $(\varepsilon/16)$ -net such that $|\mathcal{C}| \leq \left(\frac{47}{\varepsilon}\right)^r$.

Proof. The cardinality bound on \mathcal{C} can be obtained from the covering results in Appendix C of [22].⁵ It is enough to establish (16) for an arbitrary $\mathbf{x} \in \mathcal{S}_{\ell^2}$ due to the linearity of \mathbf{A} and \mathcal{L} . Let $\Delta := \|\mathbf{A}\|_{2 \rightarrow 2} \geq 0$, and choose an element $\mathbf{y} \in \mathcal{C}$ with $\|\mathbf{x} - \mathbf{y}\| \leq \varepsilon/16$. We have that

$$\begin{aligned} \|\mathbf{Ax}\|_2 - \|\mathbf{x}\|_2 &\leq \|\mathbf{Ay}\|_2 + \|\mathbf{A}(\mathbf{x} - \mathbf{y})\|_2 - 1 \leq \sqrt{1 + \varepsilon/2} - 1 + \|\mathbf{A}(\mathbf{x} - \mathbf{y})\|_2 \\ &\leq (1 + \varepsilon/4) - 1 + \Delta\varepsilon/16 = (\varepsilon/4)(1 + \Delta/4) \end{aligned}$$

holds for all $\mathbf{x} \in \mathcal{S}_{\ell^2}$. This, in turn, means that the upper bound above will hold for a vector \mathbf{x} realizing $\|\mathbf{Ax}\| = \|\mathbf{A}\|_{2 \rightarrow 2}$ so that $\Delta - 1 \leq (\varepsilon/4)(1 + \Delta/4)$ must also hold. As a consequence, $\Delta \leq 1 + \varepsilon/4 + \Delta\varepsilon/16 \implies \Delta \leq \frac{1 + \varepsilon/4}{1 - \varepsilon/16} \leq 1 + \varepsilon/3$. The upper bound now follows.

To establish the lower bound we define $\delta := \inf_{\mathbf{z} \in \mathcal{S}_{\ell^2}} \|\mathbf{Az}\| \geq 0$ and note that this quantity will also be realized by some element of the compact set \mathcal{S}_{ℓ^2} . As above we consider this minimizing vector $\mathbf{x} \in \mathcal{S}_{\ell^2}$ and choose an element $\mathbf{y} \in \mathcal{C}$ with $\|\mathbf{x} - \mathbf{y}\| \leq \varepsilon/16$ in order to see that

$$\begin{aligned} \delta - 1 &= \|\mathbf{Ax}\|_2 - \|\mathbf{x}\|_2 \geq \|\mathbf{Ay}\|_2 - \|\mathbf{A}(\mathbf{x} - \mathbf{y})\|_2 - 1 \geq \sqrt{1 - \varepsilon/2} - 1 - \|\mathbf{A}(\mathbf{x} - \mathbf{y})\|_2 \\ &\geq (1 - \varepsilon/3) - 1 - \Delta\varepsilon/16 \geq -(\varepsilon/3 + \varepsilon/16(1 + \varepsilon/3)) \\ &\geq -(\varepsilon/3 + \varepsilon/16 + \varepsilon/48) = -5\varepsilon/12. \end{aligned}$$

As a consequence, $\delta \geq 1 - 5\varepsilon/12$. The lower bound now follows. \square

Remark 1. We will see later in the text that the cardinality $(47/\varepsilon)^r$ (exponential in r) can be too big to produce tensor JL embeddings with optimal embedding dimensions. In this case one can use a much coarser “discretization” to improve the dependence on r based on, e.g., the next lemma.

With Lemma 4 in hand we are now able to prove a secondary subspace embedding result which, though it leads to suboptimal results in the vector setting, will be valuable for higher mode tensors.

Lemma 6. Fix $\varepsilon \in (0, 1)$ and let \mathcal{L} be an r -dimensional subspace of $\mathbb{C}^{n_1 \times \dots \times n_d}$ spanned by a set of r orthonormal basis tensors $\{\mathcal{T}_k\}_{k \in [r]}$. If L is an $(\varepsilon/4r)$ -JL embedding of the $4\binom{r}{2} + r = 2r^2 - r$ tensors

$$\left(\bigcup_{1 \leq h < k \leq r} \{\mathcal{T}_k - \mathcal{T}_h, \mathcal{T}_k + \mathcal{T}_h, \mathcal{T}_k - i\mathcal{T}_h, \mathcal{T}_k + i\mathcal{T}_h\} \right) \bigcup \{\mathcal{T}_k\}_{k \in [r]} \subset \mathcal{L}$$

into $\mathbb{C}^{m_1 \times \dots \times m_{d'}}$, then

$$\left| \|L(\mathcal{X})\|^2 - \|\mathcal{X}\|^2 \right| \leq \varepsilon \|\mathcal{X}\|^2$$

holds for all $\mathcal{X} \in \mathcal{L}$.

⁵A quick calculation reveals that projecting an ε -cover of the r -dimensional unit ball onto the $(r-1)$ -dimensional unit sphere produces an $\sqrt{2}\varepsilon$ -cover of \mathcal{S}_{ℓ^2} .

Proof. Appealing to Lemma 4 we can see that $|\varepsilon_{k,h}| := |\langle L(\mathcal{T}_k), L(\mathcal{T}_h) \rangle - \langle \mathcal{T}_k, \mathcal{T}_h \rangle| \leq \varepsilon/r$ for all $h, k \in [r]$. As a consequence, we have for any $\mathcal{X} = \sum_{k=1}^r \alpha_k \mathcal{T}_k \in \mathcal{L}$ that

$$\begin{aligned} \left| \|L(\mathcal{X})\|^2 - \|\mathcal{X}\|^2 \right| &= \left| \sum_{k=1}^r \sum_{h=1}^r \alpha_k \overline{\alpha_h} (\langle L(\mathcal{T}_k), L(\mathcal{T}_h) \rangle - \langle \mathcal{T}_k, \mathcal{T}_h \rangle) \right| = \left| \sum_{k=1}^r \sum_{h=1}^r \alpha_k \overline{\alpha_h} \varepsilon_{k,h} \right| \\ &\leq \sum_{k=1}^r |\alpha_k| \sum_{h=1}^r |\alpha_h| |\varepsilon_{k,h}| \leq \sum_{k=1}^r |\alpha_k| \|\boldsymbol{\alpha}\|_2 \left(\frac{\varepsilon}{\sqrt{r}} \right) \leq \varepsilon \|\boldsymbol{\alpha}\|_2^2. \end{aligned}$$

To finish we now note that $\|\mathcal{X}\|^2 = \|\boldsymbol{\alpha}\|_2^2$ due to the orthonormality of the basis tensors $\{\mathcal{T}_k\}_{k \in [r]}$. \square

3. MODEWISE LINEAR JOHNSON-LINDENSTRAUSS EMBEDDINGS OF LOW-RANK TENSORS

In this section, we consider low-rank tensor subspace embeddings for tensors with low-rank expansions in terms of rank-one tensors (i.e., for tensors with low-rank CP Decompositions). Our general approach will be to utilize subspace embeddings along the lines of Lemmas 5 and 6 in this setting. However, the fact that our basis tensors are rank-one will cause us some difficulties. Principally, among those difficulties will be our inability to guarantee that we can find an orthonormal, or even fairly incoherent, basis of rank-one tensors that span any particular r -dimensional tensor subspace \mathcal{L} we may be interested in below.

Going forward we will consider the *standard form* of a given rank- r d -mode tensor defined by

$$(17) \quad \mathcal{Y} := \sum_{k=1}^r \alpha_k \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \in \mathbb{C}^{n_1 \times \dots \times n_d},$$

where the vectors making up the rank-one basis tensors are normalized so that $\|\mathbf{y}_k^{(\ell)}\|_2 = 1$ for all $\ell \in [d]$ and $k \in [r]$. Given a set of rank-one tensors spanning a tensor subspace, one can define the coherence of the basis.

Definition 4 (Modewise coherence of a basis of rank-one tensors). *If a tensor subspace is spanned by a basis of rank-one tensors $\mathcal{B} := \{\bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \mid k \in [r]\}$ with $\|\mathbf{y}_k^{(\ell)}\|_2 = 1$ for all $\ell \in [d]$ and $k \in [r]$, we denote the maximum modewise coherence of the basis and the basis coherence by*

$$(18) \quad \mu_{\mathcal{B}} := \max_{\ell \in [d]} \mu_{\mathcal{B}, \ell} \quad \text{and} \quad \mu'_{\mathcal{B}} := \max_{\substack{k, h \in [r] \\ k \neq h}} \prod_{\ell=1}^d \left| \langle \mathbf{y}_k^{(\ell)}, \mathbf{y}_h^{(\ell)} \rangle \right|,$$

respectively, where $\mu_{\mathcal{B}, \ell} := \max_{\substack{k, h \in [r] \\ k \neq h}} \left| \langle \mathbf{y}_k^{(\ell)}, \mathbf{y}_h^{(\ell)} \rangle \right|$ is the modewise coherence of the basis for $\ell \in [d]$.

Note that $\mu_{\mathcal{B}}, \mu'_{\mathcal{B}} \in [0, 1]$ and that $\mu'_{\mathcal{B}} \leq \prod_{\ell=1}^d \mu_{\mathcal{B}, \ell} \leq \mu_{\mathcal{B}}^d$ always hold. Given any tensor \mathcal{Y} in the span of a basis \mathcal{B} of rank-1 tensors we will also refer (with some abuse of notation) to its modewise coherence and maximum modewise coherence as being equal to the modewise coherence and maximum modewise coherence of the given basis \mathcal{B} defined in Definition 4. That is, we will say that

$$(19) \quad \mu_{\mathcal{Y}, \ell} = \mu_{\mathcal{B}, \ell} \quad \text{for } \ell \in [d], \quad \text{and} \quad \mu_{\mathcal{Y}} = \mu_{\mathcal{B}}$$

for all $\mathcal{Y} \in \mathcal{B}$. Similarly, the basis coherence of any such $\mathcal{Y} \in \mathcal{B}$ will be said to equal the basis coherence also defined in Definition 4, i.e., $\mu'_{\mathcal{Y}} = \mu'_{\mathcal{B}}$. It should be remembered below, however, that the quantities $\mu_{\mathcal{Y}, \ell}$, $\mu_{\mathcal{Y}}$, $\mu'_{\mathcal{Y}}$ always depend on the particular basis \mathcal{B} under consideration.

The next lemma deals with how j -mode products can change the standard form and modewise coherence of a given tensor that lies in a tensor subspace spanned by r rank-1 tensors.

Lemma 7. Let $j \in [d]$, $\mathbf{B} \in \mathbb{C}^{m \times n_j}$, and $\mathcal{Y} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ be a rank- r tensor as per (17) such that $\min_{k \in [r]} \|\mathbf{B}\mathbf{y}_k^{(j)}\|_2 > 0$. Then $\mathcal{Y}' := \mathcal{Y} \times_j \mathbf{B}$ can be written in standard form as

$$\mathcal{Y}' = \sum_{k=1}^r \alpha_k \|\mathbf{B}\mathbf{y}_k^{(j)}\|_2 \left(\left(\bigcirc_{\ell < j} \mathbf{y}_k^{(\ell)} \right) \bigcirc \frac{\mathbf{B}\mathbf{y}_k^{(j)}}{\|\mathbf{B}\mathbf{y}_k^{(j)}\|_2} \bigcirc \left(\bigcirc_{\ell > j}^d \mathbf{y}_k^{(\ell)} \right) \right).$$

Furthermore, the j -mode coherence of \mathcal{Y}' as above will satisfy

$$\mu_{\mathcal{Y}',j} = \max_{\substack{k,h \in [r] \\ k \neq h}} \frac{|\langle \mathbf{B}\mathbf{y}_k^{(j)}, \mathbf{B}\mathbf{y}_h^{(j)} \rangle|}{\|\mathbf{B}\mathbf{y}_k^{(j)}\|_2 \|\mathbf{B}\mathbf{y}_h^{(j)}\|_2}$$

so that

$$\mu_{\mathcal{Y}'} = \max \left(\mu_{\mathcal{Y}',j}, \max_{\ell \in [d] \setminus \{j\}} \max_{\substack{k,h \in [r] \\ k \neq h}} |\langle \mathbf{y}_k^{(\ell)}, \mathbf{y}_h^{(\ell)} \rangle| \right).$$

Proof. Using Lemma 2, the linearity of tensor matricization, and (14) we can see that the mode- j unfolding of \mathcal{Y}' satisfies

$$\begin{aligned} \mathbf{Y}'_{(j)} &= \mathbf{B}\mathbf{Y}_{(j)} = \mathbf{B} \sum_{k=1}^r \alpha_k \left(\bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \right)_{(j)} = \sum_{k=1}^r \alpha_k \mathbf{B}\mathbf{y}_k^{(j)} \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top \\ &= \sum_{k=1}^r \left(\alpha_k \|\mathbf{B}\mathbf{y}_k^{(j)}\|_2 \right) \frac{\mathbf{B}\mathbf{y}_k^{(j)}}{\|\mathbf{B}\mathbf{y}_k^{(j)}\|_2} \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top. \end{aligned}$$

Refolding $\mathbf{Y}'_{(j)}$ back into a d -mode tensor then gives us our first equality. The second two equalities now follow directly from the definitions of modewise coherence. \square

The next lemma gives us a useful expression for the norm of a tensor after a j -mode product in terms of vector inner products.

Lemma 8. Let $j \in [d]$, $\mathbf{B} \in \mathbb{C}^{m \times n_j}$, and $\mathcal{Y} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ be a rank- r tensor in standard form as per (17). Then,

$$\|\mathcal{Y} \times_j \mathbf{B}\|^2 = \sum_{k,h=1}^r \sum_{a=1}^{\prod_{\ell \neq j} n_\ell} \alpha_k \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)_a \overline{\alpha_h \left(\bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right)_a} \langle \mathbf{B}\mathbf{y}_k^{(j)}, \mathbf{B}\mathbf{y}_h^{(j)} \rangle.$$

Here $(\mathbf{u})_a$ denotes the a^{th} coordinate of a vector \mathbf{u} .

Proof. Using Lemma 2, the linearity of tensor matricization, and (14) once again we can see that

$$\begin{aligned} \|\mathcal{Y} \times_j \mathbf{B}\|^2 &= \left\| \sum_{k=1}^r \alpha_k \left(\bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \times_j \mathbf{B} \right) \right\|^2 = \left\| \sum_{k=1}^r \alpha_k \mathbf{B}\mathbf{y}_k^{(j)} \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top \right\|_{\text{F}}^2 \\ &= \sum_{k,h=1}^r \left\langle \alpha_k \mathbf{B}\mathbf{y}_k^{(j)} \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top, \alpha_h \mathbf{B}\mathbf{y}_h^{(j)} \left(\bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right)^\top \right\rangle_{\text{F}} \end{aligned}$$

where $\|\cdot\|_{\text{F}}$ and $\langle \cdot, \cdot \rangle_{\text{F}}$ denote the Frobenius matrix norm and inner product, respectively. Computing the Frobenius inner products above columnwise by expressing each $\mathbf{B}\mathbf{y}_k^{(j)} \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top$ as a sum of

its individual columns (each represented as a matrix with only one nonzero column) we can further see that

$$\|\mathcal{Y} \times_j \mathbf{B}\|^2 = \sum_{k,h=1}^r \sum_{a=1}^{\prod_{\ell \neq j} n_\ell} \alpha_k \left(\otimes_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)_a \overline{\alpha_h \left(\otimes_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right)_a} \left\langle \mathbf{B} \mathbf{y}_k^{(j)}, \mathbf{B} \mathbf{y}_h^{(j)} \right\rangle.$$

as we wished to show. \square

The following theorem demonstrates that a single modewise Johnson-Lindenstrauss embedding of any low-rank tensor \mathcal{Y} of the form (17) will preserve its norm up to an error depending on the overall ℓ^2 -norm of its coefficients $\boldsymbol{\alpha} \in \mathbb{C}^r$. Subsequent results will then consider when $\|\boldsymbol{\alpha}\|_2 \approx \|\mathcal{Y}\|$.

Theorem 3. *Let $j \in [d]$ and $\mathcal{Y} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ be a rank- r tensor as per (17). Suppose that $\mathbf{A} \in \mathbb{C}^{m \times n_j}$ is an $(\varepsilon/4)$ -JL embedding of the $4 \binom{r}{2} + r = 2r^2 - r$ vectors*

$$\left(\bigcup_{1 \leq h < k \leq r} \left\{ \mathbf{y}_k^{(j)} - \mathbf{y}_h^{(j)}, \mathbf{y}_k^{(j)} + \mathbf{y}_h^{(j)}, \mathbf{y}_k^{(j)} - i\mathbf{y}_h^{(j)}, \mathbf{y}_k^{(j)} + i\mathbf{y}_h^{(j)} \right\} \right) \bigcup \left\{ \mathbf{y}_k^{(j)} \right\}_{k \in [r]} \subset \mathbb{C}^{n_j}$$

into \mathbb{C}^m . Let $\mathcal{Y}' := \mathcal{Y} \times_j \mathbf{A}$ and rewrite it in standard form so that

$$\mathcal{Y}' = \sum_{k=1}^r \alpha'_k \left(\left(\bigcirc_{\ell < j} \mathbf{y}_k^{(\ell)} \right) \bigcirc \frac{\mathbf{A} \mathbf{y}_k^{(j)}}{\|\mathbf{A} \mathbf{y}_k^{(j)}\|_2} \bigcirc \left(\bigcirc_{\ell > j} \mathbf{y}_k^{(\ell)} \right) \right).$$

Then all of the following hold:

$$(\dagger) \quad |\alpha'_k - \alpha_k| \leq \varepsilon |\alpha_k|/4 \text{ for all } k \in [r] \text{ so that } \|\boldsymbol{\alpha}'\|_\infty \leq (1 + \varepsilon/4) \|\boldsymbol{\alpha}\|_\infty$$

$$(\dagger\dagger) \quad \mu_{\mathcal{Y}',j} \leq \frac{\mu_{\mathcal{Y},j} + \varepsilon}{1 - \varepsilon/4}, \text{ and } \mu_{\mathcal{Y}',\ell} = \mu_{\mathcal{Y},\ell} \text{ for all } \ell \in [d] \setminus \{j\}$$

$$(\dagger\dagger\dagger) \quad |\|\mathcal{Y}'\|^2 - \|\mathcal{Y}\|^2| \leq \varepsilon \left(1 + \sqrt{r(r-1)} \prod_{\ell \neq j} \mu_{\mathcal{Y},\ell} \right) \|\boldsymbol{\alpha}\|_2^2 \leq \varepsilon \left(1 + r \mu_{\mathcal{Y}}^{d-1} \right) \|\boldsymbol{\alpha}\|_2^2 \leq \varepsilon(r+1) \|\boldsymbol{\alpha}\|_2^2$$

Proof. We prove each property in order below.

Proof of (\dagger) : By Lemma 7 we have for all $k \in [r]$ that

$$|\alpha'_k - \alpha_k| = \left| \alpha_k \left\| \mathbf{A} \mathbf{y}_k^{(j)} \right\|_2 - \alpha_k \right| = \left| \left\| \mathbf{A} \mathbf{y}_k^{(j)} \right\|_2 - 1 \right| |\alpha_k| \leq \varepsilon |\alpha_k|/4$$

as we wished to prove.

Proof of $(\dagger\dagger)$: Appealing to Lemma 7 and the definition of j -mode coherence we have that

$$\mu_{\mathcal{Y}',j} = \max_{\substack{k,h \in [r] \\ k \neq h}} \frac{\left| \left\langle \mathbf{A} \mathbf{y}_k^{(j)}, \mathbf{A} \mathbf{y}_h^{(j)} \right\rangle \right|}{\left\| \mathbf{A} \mathbf{y}_k^{(j)} \right\|_2 \left\| \mathbf{A} \mathbf{y}_h^{(j)} \right\|_2} \leq \max_{\substack{k,h \in [r] \\ k \neq h}} \frac{\left| \left\langle \mathbf{y}_k^{(j)}, \mathbf{y}_h^{(j)} \right\rangle \right| + \varepsilon}{1 - \frac{\varepsilon}{4}} = \frac{\mu_{\mathcal{Y},j} + \varepsilon}{1 - \frac{\varepsilon}{4}},$$

where the inequality follows from Lemma 3 combined with \mathbf{A} being an $(\varepsilon/4)$ -JL embedding.

Proof of (†††): Applying Lemma 8 with $\mathbf{B} = \mathbf{A}$ and $\mathbf{B} = \mathbf{I}$, respectively, we can see that

$$(20) \quad \|\mathcal{Y}'\|^2 - \|\mathcal{Y}\|^2 = \sum_{k,h=1}^r \sum_{a=1}^{\prod_{\ell \neq j} n_\ell} \alpha_k \left(\otimes_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)_a \overline{\alpha_h \left(\otimes_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right)_a} \left(\left\langle \mathbf{A} \mathbf{y}_k^{(j)}, \mathbf{A} \mathbf{y}_h^{(j)} \right\rangle - \left\langle \mathbf{y}_k^{(j)}, \mathbf{y}_h^{(j)} \right\rangle \right).$$

Applying Lemma 3 to each inner product in (20) we can now see that

$$\left\langle \mathbf{A} \mathbf{y}_k^{(j)}, \mathbf{A} \mathbf{y}_h^{(j)} \right\rangle = \left\langle \mathbf{y}_k^{(j)}, \mathbf{y}_h^{(j)} \right\rangle + \varepsilon_{k,h}$$

for some $\varepsilon_{k,h} \in \mathbb{C}$ with $|\varepsilon_{k,h}| \leq \varepsilon$. As a result we have that

$$\begin{aligned} \left| \|\mathcal{Y} \times_j \mathbf{A}\|^2 - \|\mathcal{Y}\|^2 \right| &= \left| \sum_{k,h=1}^r \sum_{a=1}^{\prod_{\ell \neq j} n_\ell} \alpha_k \left(\otimes_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)_a \overline{\alpha_h \left(\otimes_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right)_a} \varepsilon_{k,h} \right|. \\ &= \left| \sum_{k,h=1}^r \alpha_k \overline{\alpha_h} \varepsilon_{k,h} \sum_{a=1}^{\prod_{\ell \neq j} n_\ell} \left(\otimes_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)_a \overline{\left(\otimes_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right)_a} \right| \\ &= \left| \sum_{k,h=1}^r \alpha_k \overline{\alpha_h} \varepsilon_{k,h} \left\langle \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right\rangle \right| \\ &\leq \left| \sum_{k=1}^r |\alpha_k|^2 \varepsilon_{k,k} \left\| \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right\|^2 \right| + \left| \sum_{k \neq h} \alpha_k \overline{\alpha_h} \varepsilon_{k,h} \left\langle \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right\rangle \right|. \end{aligned}$$

Noting that $\left\| \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right\|^2 = 1$ by Lemma 1 since $\left\| \mathbf{y}_k^{(\ell)} \right\|_2 = 1$ for all $\ell \in [d]$ and $k \in [r]$, we now have that

$$\begin{aligned} \left| \|\mathcal{Y} \times_j \mathbf{A}\|^2 - \|\mathcal{Y}\|^2 \right| &\leq \varepsilon \left| \sum_{k=1}^r |\alpha_k|^2 \right| + \left| \sum_{k \neq h} \alpha_k \overline{\alpha_h} \varepsilon_{k,h} \left\langle \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right\rangle \right| \\ &= \varepsilon \|\boldsymbol{\alpha}\|_2^2 + |\langle \mathbf{E}^\top \boldsymbol{\alpha}, \boldsymbol{\alpha} \rangle|, \end{aligned}$$

where $\mathbf{E} \in \mathbb{C}^{r \times r}$ is zero on its diagonal, and $E_{k,h} = \varepsilon_{k,h} \left\langle \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right\rangle$ for $k \neq h$. As a result, $\left| \|\mathcal{Y} \times_j \mathbf{A}\|^2 - \|\mathcal{Y}\|^2 \right| \leq (\varepsilon + \|\mathbf{E}^\top\|_{2 \rightarrow 2}) \|\boldsymbol{\alpha}\|_2^2$, where the operator norm $\|\mathbf{E}^\top\|_{2 \rightarrow 2}$ satisfies

$$\|\mathbf{E}^\top\|_{2 \rightarrow 2} \leq \|\mathbf{E}\|_F \leq \sqrt{\sum_{k \neq h} \left| \left\langle \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right\rangle \right|^2 \varepsilon^2} = \varepsilon \sqrt{\sum_{k \neq h} \left| \left\langle \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell \neq j} \mathbf{y}_h^{(\ell)} \right\rangle \right|^2}.$$

Finally, Lemma 1 and the definition of $\mu_{\mathcal{Y}}$ implies that

$$\|\mathbf{E}\|_{2 \rightarrow 2} \leq \varepsilon \sqrt{r(r-1)} \prod_{\ell \neq j} \mu_{\mathcal{Y},\ell} \leq \varepsilon r \mu_{\mathcal{Y}}^{d-1}.$$

Thus, we obtain the desired bound

$$\left| \|\mathcal{Y} \times_j \mathbf{A}\|^2 - \|\mathcal{Y}\|^2 \right| \leq \varepsilon \left(1 + \sqrt{r(r-1)} \prod_{\ell \neq j} \mu_{\mathcal{Y},\ell} \right) \|\boldsymbol{\alpha}\|_2^2 \leq \varepsilon \left(1 + r \mu_{\mathcal{Y}}^{d-1} \right) \|\boldsymbol{\alpha}\|_2^2.$$

□

Note that part (†††) of Theorem 3 bounds $\left| \|\mathcal{Y}'\|^2 - \|\mathcal{Y}\|^2 \right|$ with respect to $\|\boldsymbol{\alpha}\|_2^2$. Traditional JL-type error guarantees typically want to prove error bounds of the form $\left| \|\mathcal{Y}'\|^2 - \|\mathcal{Y}\|^2 \right| \leq C_\varepsilon \|\mathcal{Y}\|^2$,

however. The next lemma bounds $\|\boldsymbol{\alpha}\|_2^2$ by $\|\mathcal{Y}\|^2$ so that the reader who desires such bounds can obtain them easily for any tensor with sufficiently small modewise coherence.

Lemma 9. *Let $\mathcal{Y} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ be a rank- r tensor as per (17) with the basis coherence $\mu'_{\mathcal{Y}} < (r-1)^{-1}$. Then,*

$$\|\boldsymbol{\alpha}\|_2^2 \leq \left(\frac{1}{1 - (r-1)\mu'_{\mathcal{Y}}} \right) \|\mathcal{Y}\|^2 \leq \left(\frac{1}{1 - (r-1)\prod_{\ell=1}^d \mu_{\mathcal{Y},\ell}^d} \right) \|\mathcal{Y}\|^2 \leq \left(\frac{1}{1 - (r-1)\mu_{\mathcal{Y}}^d} \right) \|\mathcal{Y}\|^2.$$

Proof. Utilizing Lemma 1 and the standard form of \mathcal{Y} we can see that

$$\begin{aligned} |\|\mathcal{Y}\|^2 - \|\boldsymbol{\alpha}\|_2^2| &= \left| \sum_{k,h=1}^r \alpha_k \overline{\alpha_h} \left\langle \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)}, \bigcirc_{\ell=1}^d \mathbf{y}_h^{(\ell)} \right\rangle - \sum_{k=1}^r |\alpha_k|^2 \right| \\ &= \left| \sum_{k \neq h}^r \alpha_k \overline{\alpha_h} \prod_{\ell=1}^d \left\langle \mathbf{y}_k^{(\ell)}, \mathbf{y}_h^{(\ell)} \right\rangle \right| \leq \mu'_{\mathcal{Y}} \sum_{k \neq h}^r |\alpha_k \overline{\alpha_h}| \\ &= \mu'_{\mathcal{Y}} \left(\left(\sum_{k=1}^r |\alpha_k| \right)^2 - \sum_{k=1}^r |\alpha_k|^2 \right) \leq \mu'_{\mathcal{Y}} \left((\sqrt{r} \|\boldsymbol{\alpha}\|_2)^2 - \|\boldsymbol{\alpha}\|_2^2 \right) \end{aligned}$$

where the last inequality follows from Cauchy-Schwarz. As a result we have that

$$|\|\mathcal{Y}\|^2 - \|\boldsymbol{\alpha}\|_2^2| \leq \mu'_{\mathcal{Y}} (r-1) \|\boldsymbol{\alpha}\|_2^2$$

which in turn implies that

$$\|\mathcal{Y}\|^2 \geq (1 - (r-1)\mu'_{\mathcal{Y}}) \|\boldsymbol{\alpha}\|_2^2.$$

□

The following simple technical lemma will be used repeatedly in our next theorem.

Lemma 10. *Let $c, d \in \mathbb{R}^+$. Then, $\mathbb{e}^c \geq \left(1 + \frac{c}{d}\right)^d$.*

We are now prepared to prove our main theorem for this section. Recall that combining it with Lemma 9 provides traditional JL-embedding error bounds.

Theorem 4. *Let $\varepsilon \in (0, 3/4]$, $\mathcal{Y} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ be a rank- r tensor expressed in standard form as per (17), and $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ be an $(\varepsilon/4d)$ -JL embedding of the $2r^2 - r$ vectors*

$$\mathcal{S}'_j := \left(\bigcup_{1 \leq h < k \leq r} \left\{ \mathbf{y}_k^{(j)} - \mathbf{y}_h^{(j)}, \mathbf{y}_k^{(j)} + \mathbf{y}_h^{(j)}, \mathbf{y}_k^{(j)} - i\mathbf{y}_h^{(j)}, \mathbf{y}_k^{(j)} + i\mathbf{y}_h^{(j)} \right\} \right) \bigcup \left\{ \mathbf{y}_k^{(j)} \right\}_{k \in [r]} \subset \mathbb{C}^{n_j}$$

into \mathbb{C}^{m_j} for each $j \in [d]$. Then,

$$\begin{aligned} (21) \quad |\|\mathcal{Y}\|^2 - \|\mathcal{Y} \times_1 \mathbf{A}_1 \dots \times_d \mathbf{A}_d\|^2| &\leq \varepsilon \left(\mathbb{e} + \mathbb{e}^2 \sqrt{r(r-1)} \cdot \max \left(\varepsilon^{d-1}, \mu_{\mathcal{Y}}^{d-1} \right) \right) \|\boldsymbol{\alpha}\|_2^2 \\ &\leq \varepsilon \mathbb{e}^2 (r+1) \|\boldsymbol{\alpha}\|_2^2 \end{aligned}$$

always holds. Here, $\mu_{\mathcal{Y}}$ is maximum modewise coherence of the tensor defined by (19). Furthermore, if $\mu_{\mathcal{Y}} = 0$ then

$$|\|\mathcal{Y}\|^2 - \|\mathcal{Y} \times_1 \mathbf{A}_1 \dots \times_d \mathbf{A}_d\|^2| \leq \left(\varepsilon + \mathbb{e} \sqrt{r(r-1)} \varepsilon^d \right) \mathbb{e} \|\boldsymbol{\alpha}\|_2^2.$$

Proof. Let $\mathcal{Y}^{(0)} := \mathcal{Y}$, and for each $j \in [d]$ define the tensor

$$\mathcal{Y}^{(j)} := \mathcal{Y} \times_1 \mathbf{A}_1 \cdots \times_j \mathbf{A}_j = \sum_{k=1}^r \alpha_{j,k} \bigcirc_{\ell=1}^d \mathbf{y}_{j,k}^{(\ell)}$$

expressed in standard form via j applications of Lemma 7. Note that parts (\dagger) and $(\ddagger\ddagger)$ of Theorem 3 imply that

- (i) $|\alpha_{j,k} - \alpha_{j-1,k}| \leq \varepsilon |\alpha_{j-1,k}|/4d$ so that $|\alpha_{j,k}| \leq (1 + \varepsilon/4d)|\alpha_{j-1,k}|$ holds for all $k \in [r]$, and
- (ii) $\mu_{\mathcal{Y}^{(j)},j} \leq (\mu_{\mathcal{Y}^{(j-1)},j} + \varepsilon/d)/(1 - \varepsilon/4d)$, and $\mu_{\mathcal{Y}^{(j)},\ell} = \mu_{\mathcal{Y}^{(j-1)},\ell}$ for all $\ell \in [d] \setminus \{j\}$,

both hold for all $j \in [d]$. Using these facts it is not too difficult to inductively establish that both

$$(22) \quad |\alpha_{j,k}| \leq (1 + \varepsilon/4d)^j |\alpha_k|,$$

and

$$(23) \quad \prod_{\ell \neq j} \mu_{\mathcal{Y}^{(j-1)},\ell} \leq \left(\prod_{\ell < j} \frac{\mu_{\mathcal{Y},\ell} + \varepsilon/d}{1 - \varepsilon/4d} \right) \prod_{\ell > j} \mu_{\mathcal{Y},\ell} \leq \left(\frac{\mu_{\mathcal{Y}} + \varepsilon/d}{1 - \varepsilon/4d} \right)^{j-1} \mu_{\mathcal{Y}}^{d-j},$$

also hold for all $k \in [r]$ and $j \in [d]$. Note that in (23) we will let $\mu_{\mathcal{Y}}^0 = 1$ even if $\mu_{\mathcal{Y}} = 0$ since this still yields the correct bound in the $j = d$ and $\mu_{\mathcal{Y}} = 0$ case.

Preceding with the desired error bound we can now see that

$$\begin{aligned} \left| \|\mathcal{Y}\|^2 - \|\mathcal{Y} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d\|^2 \right| &= \left| \sum_{j=0}^{d-1} \left\| \mathcal{Y}^{(j)} \right\|^2 - \left\| \mathcal{Y}^{(j+1)} \right\|^2 \right| \\ &\leq \frac{\varepsilon}{d} \sum_{j=0}^{d-1} \left(1 + \sqrt{r(r-1)} \prod_{\ell \neq j+1} \mu_{\mathcal{Y}^{(j)},\ell} \right) \|\boldsymbol{\alpha}_j\|_2^2 \\ &\leq \frac{\varepsilon}{d} \sum_{j=0}^{d-1} \left(1 + \sqrt{r(r-1)} \left(\frac{\mu_{\mathcal{Y}} + \varepsilon/d}{1 - \varepsilon/4d} \right)^j \mu_{\mathcal{Y}}^{d-1-j} \right) (1 + \varepsilon/4d)^{2j} \|\boldsymbol{\alpha}\|_2^2 \\ &\leq \frac{\varepsilon}{d} \sum_{j=0}^{d-1} \left(1 + \sqrt{r(r-1)} \left(\frac{\mu_{\mathcal{Y}} + \varepsilon/d}{1 - \varepsilon/4d} \right)^j \mu_{\mathcal{Y}}^{d-1-j} \right) (1 + 9\varepsilon/16d)^j \|\boldsymbol{\alpha}\|_2^2 \end{aligned}$$

where we have used part $(\ddagger\ddagger\ddagger)$ of Theorem 3, (22), and (23). Considering each term in the upper bound above separately, we have that

$$\left| \|\mathcal{Y}\|^2 - \|\mathcal{Y} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d\|^2 \right| \leq \frac{\varepsilon}{d} \|\boldsymbol{\alpha}\|_2^2 \left(T_1 + \sqrt{r(r-1)} T_2 \right)$$

where

$$T_1 := \sum_{j=0}^{d-1} (1 + 9\varepsilon/16d)^j = \frac{(1 + 9\varepsilon/16d)^d - 1}{9\varepsilon/16d} \leq \varepsilon d$$

using Lemma 10 and that $9\varepsilon/16 < 1$, and where

$$T_2 := \sum_{j=0}^{d-1} \left(\frac{\mu_{\mathcal{Y}} + \varepsilon/d}{1 - \varepsilon/4d} \right)^j \mu_{\mathcal{Y}}^{d-1-j} (1 + 9\varepsilon/16d)^j \leq \sum_{j=0}^{d-1} (\mu_{\mathcal{Y}} + \varepsilon/d)^j \mu_{\mathcal{Y}}^{d-1-j} (1 + \varepsilon/d)^j$$

for $\varepsilon \leq 3/4$.

Continuing to bound the second term we will consider three cases. First, if $\mu_{\mathcal{Y}} = 0$ then

$$T_2 \leq (\varepsilon/d)^{d-1} (1 + \varepsilon/d)^{d-1} \leq \varepsilon (\varepsilon/d)^{d-1},$$

using Lemma 10 and that $\varepsilon < 1$. Second, if $0 < \mu_{\mathcal{Y}} \leq \varepsilon$ then

$$\begin{aligned} T_2 &\leq \sum_{j=0}^{d-1} (\varepsilon + \varepsilon/d)^j \varepsilon^{d-1-j} (1 + \varepsilon/d)^j = \varepsilon^{d-1} \sum_{j=0}^{d-1} (1 + 1/d)^j (1 + \varepsilon/d)^j \\ &\leq \varepsilon^{d-1} d (1 + 1/d)^d (1 + \varepsilon/d)^d \leq d \mathbb{e}^2 \varepsilon^{d-1}, \end{aligned}$$

using Lemma 10 and that $\varepsilon < 1$ once more. If, however, $\mu_{\mathcal{Y}} > \varepsilon$ then we can see that

$$\begin{aligned} T_2 &\leq \mu_{\mathcal{Y}}^{d-1} \sum_{j=0}^{d-1} (1 + \varepsilon/\mu_{\mathcal{Y}} d)^j (1 + \varepsilon/d)^j \leq \mu_{\mathcal{Y}}^{d-1} \sum_{j=0}^{d-1} (1 + 1/d)^j (1 + \varepsilon/d)^j \\ &\leq \mu_{\mathcal{Y}}^{d-1} \cdot d (1 + 1/d)^d (1 + \varepsilon/d)^d \leq \mu_{\mathcal{Y}}^{d-1} d \mathbb{e}^{1+\varepsilon} \leq d \mathbb{e}^2 \mu_{\mathcal{Y}}^{d-1}, \end{aligned}$$

where we have again utilized Lemma 10. The desired result now follows. \square

3.1. Extension of Theorem 4 to Oblivious Tensor Subspace Embedding. First, Theorem 4 can be extended to show that the modewise compression preserves scalar products between two tensors \mathcal{X} and \mathcal{Y} spanned by the same rank one tensors. We have the following corollary of Theorem 4.

Corollary 1. *Suppose that $\mathcal{X}, \mathcal{Y} \in \mathcal{L} \subset \mathbb{C}^{n_1 \times \dots \times n_d}$ have standard forms given by*

$$\mathcal{X} = \sum_{k=1}^r \beta_k \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)}, \text{ and } \mathcal{Y} = \sum_{k=1}^r \alpha_k \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)}.$$

Let $\varepsilon \in (0, 3/4]$, and $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ be a $(\varepsilon/4d)$ -JL embedding of the set S'_j defined as in the statement of Theorem 4 for each $j \in [d]$. Then,

$$\begin{aligned} \left| \left\langle \mathcal{X} \times_{j=1}^d \mathbf{A}_j, \mathcal{Y} \times_{j=1}^d \mathbf{A}_j \right\rangle - \langle \mathcal{X}, \mathcal{Y} \rangle \right| &\leq 2\varepsilon' (\|\boldsymbol{\beta}\|_2^2 + \|\boldsymbol{\alpha}\|_2^2) \leq 4\varepsilon' \cdot \max \{ \|\boldsymbol{\beta}\|_2^2, \|\boldsymbol{\alpha}\|_2^2 \} \\ &\leq 4\varepsilon' \cdot \frac{\max \{ \|\mathcal{X}\|^2, \|\mathcal{Y}\|^2 \}}{1 - (r-1)\mu'_{\mathcal{Y}}}, \end{aligned}$$

where

$$(24) \quad \varepsilon' := \begin{cases} \left(\varepsilon + \mathbb{e} \sqrt{r(r-1)} \varepsilon^d \right) \mathbb{e} & \text{if } \mu_{\mathcal{Y}} = 0, \\ \varepsilon \left(\mathbb{e} + \mathbb{e}^2 \sqrt{r(r-1)} \cdot \max \left(\varepsilon^{d-1}, \mu_{\mathcal{Y}}^{d-1} \right) \right) & \text{otherwise.} \end{cases}$$

Proof. Using the polarization identity in combination with Lemma 2 and Theorem 4 we can see that

$$\begin{aligned} \left| \left\langle \mathcal{X} \times_{j=1}^d \mathbf{A}_j, \mathcal{Y} \times_{j=1}^d \mathbf{A}_j \right\rangle - \langle \mathcal{X}, \mathcal{Y} \rangle \right| &= \left| \frac{1}{4} \sum_{\ell=0}^3 \mathbb{i}^\ell \left(\left\| \mathcal{X} \times_{j=1}^d \mathbf{A}_j + \mathbb{i}^\ell \mathcal{Y} \times_{j=1}^d \mathbf{A}_j \right\|_2^2 - \left\| \mathcal{X} + \mathbb{i}^\ell \mathcal{Y} \right\|_2^2 \right) \right| \\ &\leq \frac{1}{4} \sum_{\ell=0}^3 \varepsilon' \left\| \boldsymbol{\beta} + \mathbb{i}^\ell \boldsymbol{\alpha} \right\|_2^2 \leq \varepsilon' (\|\boldsymbol{\beta}\|_2 + \|\boldsymbol{\alpha}\|_2)^2 \\ &\leq 2\varepsilon' (\|\boldsymbol{\beta}\|_2^2 + \|\boldsymbol{\alpha}\|_2^2) \leq 4\varepsilon' \cdot \max \{ \|\boldsymbol{\beta}\|_2^2, \|\boldsymbol{\alpha}\|_2^2 \}, \end{aligned}$$

where the second to last inequality follows from Young's inequality for products. An application of Lemma 9 yields the final inequality. \square

Theorem 4 and Corollary 1 guarantee that modewise JL-embeddings approximately preserve the norms and the scalar products between all tensors in the span of the set

$$\mathcal{B} := \left\{ \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \mid k \in [r] \right\} \subset \mathbb{C}^{n_1 \times \cdots \times n_d}.$$

Let

$$\mathcal{L} := \text{span} \left(\left\{ \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \mid k \in [r] \right\} \right).$$

Then, employing η -optimal JL embeddings (as per Definition 2), we can get a subspace oblivious version of Theorem 4.

Corollary 2. *Fix $\delta, \eta \in (0, 1/2)$ and $d \geq 2$. Let \mathcal{L} be an r -dimensional subspace of $\mathbb{C}^{n_1 \times \cdots \times n_d}$ spanned by a basis of rank-1 tensors $\mathcal{B} := \left\{ \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \mid k \in [r] \right\}$ with modewise coherence (as per (18)) satisfying $\mu_{\mathcal{B}}^{d-1} < 1/2r$. For each $j \in [d]$ draw $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ with*

$$(25) \quad m_j \geq \tilde{C} \cdot r^{2/d} d^2 / \varepsilon^2 \cdot \ln(2r^2 d / \eta)$$

from an (η/d) -optimal family of JL embedding distributions, where $\tilde{C} \in \mathbb{R}^+$ is an absolute constant. Then with probability at least $1 - \eta$ we have

$$\left| \|\mathcal{Y} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d\|^2 - \|\mathcal{Y}\|^2 \right| \leq \varepsilon \|\mathcal{Y}\|^2,$$

for all $\mathcal{Y} \in \mathcal{L}$.

Proof. Let $\mathcal{Y} \in \mathcal{L}$. By Corollary 1, the linear operator L defined as $L(\mathcal{Z}) = \mathcal{Z} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d$ is an ε -JL embedding of \mathcal{Y} if

- $4 / (1 - (r - 1)\mu'_{\mathcal{B}}) \leq 8$ and
- each \mathbf{A}_j is an $(\delta/4)$ d -JL embedding of the set S'_j of cardinality $|S'_j| \leq 2r^2 - r$, where the dependence $\varepsilon'(\delta)$ is defined by (24), and $\varepsilon \geq 8\varepsilon'$.

The first condition is satisfied since basis incoherence condition implies

$$\mu'_{\mathcal{B}} \leq \mu_{\mathcal{B}}^d \leq 1/2(r - 1).$$

Hence, $8(1 - (r - 1)\mu'_{\mathcal{B}}) \geq 4$. To check the second condition, note that due to (24), it is enough to take ε such that

$$\varepsilon \geq 8\delta e + 8\delta e^2 r \max \left(\delta^{d-1}, \mu_{\mathcal{B}}^{d-1} \right),$$

and $\delta := \varepsilon/16e \cdot (1/r)^{1/d}$ satisfies that. Then, the matrix \mathbf{A}_j taken from an (η/d) -optimal family of JL distributions will be an $(\delta/4d)$ -JL embedding of S'_j to \mathbb{C}^{m_j} with probability $1 - \eta/d$ as long as

$$|S'_j| = 2r^2 - r \leq \frac{\eta}{d} \exp \left(\frac{\delta^2 m_j}{16d^2 C} \right),$$

which is satisfied for each m_j defined by (25). Taking union bound over d modes, we conclude the proof of Corollary 2. \square

Remark 2 (JL-type embedding for low-rank matrices). *Corollary 2 (as well as the above results, including Theorem 4) can be applied in the special case where $\mathcal{X} = \mathbf{X}$ is a matrix in $\mathbb{C}^{n_1 \times n_2}$. In this case, the CP-rank is the usual matrix rank, and the CP decomposition becomes the regular SVD decomposition of the matrix which can be computed efficiently in parallel (see, e.g., [27]). In particular, the basis vectors are orthogonal to each other in this case. The result of Corollary 2 implies that taking A and B as matrices belonging to the $(\eta/2)$ -JL embedding family and of sizes $n_1 \times m_1$ and $n_2 \times m_2$, respectively, such that $m_j \gtrsim r \ln(r/\sqrt{\eta})/\varepsilon^2$ (for $j = 1, 2$), we get the following JL-type result for the Frobenius matrix norm: with probability $1 - \eta$,*

$$\|A^T \mathbf{X} B\|_F^2 = (1 + \tilde{\varepsilon}) \|\mathbf{X}\|_F^2 \quad \text{for some } |\tilde{\varepsilon}| \leq \varepsilon.$$

3.2. Naturally incoherent tensor bases. Again, we remind the reader that Lemma 9 can be used in combination with the theorems and corollaries above/below in order to provide JL-embedding results of the usual type. In order for Lemma 9 to apply, however, we need the coherence $\mu'_{\mathcal{B}}$ of the basis \mathcal{B} to satisfy $\mu'_{\mathcal{B}} < (r-1)^{-1}$. One popular set of bases with this property are those that result from considering tensors whose Tucker decompositions [49, 34, 27] have core tensors with a small number of nonzero entries. More specifically, let $\mathcal{C} \in \mathbb{C}^{n_1 \times \dots \times n_d}$, $\mathbf{U}^{(j)} \in \mathbb{C}^{n_j \times n_j}$ be unitary for all $j \in [d]$, and $\mathcal{S} \subset [n_1] \times \dots \times [n_d]$ be a set of r indices in \mathcal{C} . Now consider the r -dimensional tensor subspace

$$\mathcal{L}_{\text{Tucker}} := \left\{ \mathcal{X} \mid \mathcal{X} = \mathcal{C} \times_{j=1}^d \mathbf{U}^{(j)} \text{ with } \mathcal{C}_{\mathbf{i}} = 0 \text{ for all } \mathbf{i} \notin \mathcal{S} \right\}.$$

One can see that any tensor $\mathcal{Y} \in \mathcal{L}_{\text{Tucker}}$ can be written in standard form as per (17) with, for all $\ell \in [d]$, $\mathbf{y}_k^{(\ell)} = \mathbf{U}_{k'}^{(\ell)}$ for some column $k' \in [n_{\ell}]$. As a result, $\mu'_{\mathcal{Y}} = \mu'_{\mathcal{B}} = 0$ will hold due to the orthogonality of the columns of each $\mathbf{U}^{(\ell)}$ matrix. We therefore have the following special case of Theorem 4 in this setting.

Corollary 3. *Suppose that $\mathcal{Y} \in \mathcal{L}_{\text{Tucker}} \subset \mathbb{C}^{n_1 \times \dots \times n_d}$. Let $\varepsilon \in (0, 3/4]$, and $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ be defined as per Theorem 4 for each $j \in [d]$. Then,*

$$\left| \|\mathcal{Y}\|^2 - \|\mathcal{Y} \times_1 \mathbf{A}_1 \times_2 \mathbf{A}_2 \times_3 \mathbf{A}_3 \times_4 \mathbf{A}_4\|^2 \right| \leq \varepsilon' \|\mathcal{Y}\|^2,$$

where

$$\varepsilon' := \begin{cases} \left(\varepsilon + \mathbb{E} \sqrt{r(r-1)} \varepsilon^d \right) \mathbb{E} & \text{if } \mu_{\mathcal{B}} = 0, \\ \varepsilon \left(\mathbb{E} + \mathbb{E}^2 \sqrt{r(r-1)} \cdot \max \left(\varepsilon^{d-1}, \mu_{\mathcal{B}}^{d-1} \right) \right) & \text{otherwise.} \end{cases}$$

Proof. This follows from Theorem 4 combined with Lemma 9 after noting that $\mu'_{\mathcal{B}} = 0$ holds. \square

Another natural set of bases on which the property $\mu'_{\mathcal{B}} < (r-1)^{-1}$ is satisfied is random family of sub-gaussian tensors. The following Lemma 11 shows that if all the components of all vectors $\mathbf{y}_k^{(j)}$ (for $j \in [d], k \in [r]$) are normalized independent K -subgaussian random variables (see Definition 5 below), the coherence is actually low with high probability.

Definition 5. *A random variable ξ is called K -subgaussian, if for all $t \geq 0$*

$$\mathbb{P} \{ |\xi| > t \} \leq 2 \exp \left(-t^2/K^2 \right).$$

Informally, all normal random variables (with any mean and variance), and also those with lighter tails are K -subgaussian with some proper constant K . All bounded random variables are subgaussian.

Lemma 11. *Let $\mu > 0$. Let $j \in [d]$ and $\mathcal{Y} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ be a rank- r tensor as per (17). Let $n = \min_{i \in [d]} n_i$. If all components of all vectors $\mathbf{y}_k^{(j)}$ are normalized independent mean zero K -subgaussian random variables, with probability at least $1 - 2r^2 d \exp(-c\mu^2 n)$ maximum modewise coherence parameter of the tensor \mathcal{Y} is at most μ . Here, c is a positive constant depending only on K .*

Proof. For any $k \in [r]$ and $j \in [d]$ denote $\tilde{\mathbf{y}}_k^{(j)} := \mathbf{y}_k^{(j)} / \|\mathbf{y}_k^{(j)}\|$. By definition, $\tilde{\mathbf{y}}_k^{(j)}$ are independent K -subgaussian random variables for all $k \in [r]$ and $j \in [d]$. Therefore, their norms are of order \sqrt{n} with high probability: for any fixed k, j ,

$$\mathbb{P} \left\{ n/2 \leq \|\tilde{\mathbf{y}}_k^{(j)}\|_2^2 \leq 2n \right\} \geq 1 - 2 \exp \left(-c_1 n/K^4 \right)$$

(see, e.g. [[52], Section 3.1]). Taking union bound, we can conclude that with probability at least $1 - 2rd \exp(-c_1 n/K^4)$, all vectors $\tilde{\mathbf{y}}_k^{(j)}$ have their norms between $[\sqrt{n/2}, \sqrt{2n}]$.

For any mean zero independent K -subgaussian vectors \mathbf{x} and \mathbf{y} ,

$$(26) \quad \begin{aligned} & \mathbb{P}\{|\langle \mathbf{x}, \mathbf{y} \rangle| \geq \mu \|\mathbf{x}\| \|\mathbf{y}\|\} \\ & \leq \mathbb{P}\{|\langle \mathbf{x}, \mathbf{y} \rangle| \geq \mu \|\mathbf{y}\| \sqrt{n/2}\} + \mathbb{P}\{\|\mathbf{x}\| < \sqrt{n/2}\}. \end{aligned}$$

To bound the first term, let us use Hoeffding's inequality (see, e.g. [[52], Theorem 2.6.3]). Conditioning on \mathbf{y} , we have

$$\mathbb{P}_{\mathbf{x}}\left\{\left|\sum_i x_i y_i\right| \geq \mu \|\mathbf{y}\| \sqrt{n/2}\right\} \leq 2 \exp\left(-\frac{c_2 \mu^2 n}{2K^2}\right).$$

Now, let $\tilde{\mathbf{y}}_k^{(j)} = \mathbf{x}$ and $\tilde{\mathbf{y}}_l^{(j)} = \mathbf{y}$. Integrating over $\tilde{\mathbf{y}}_l^{(j)}$ and then taking union bound over all choices of k, l and j , we get $|\langle \mathbf{y}_k^{(j)}, \mathbf{y}_l^{(j)} \rangle| \leq \mu$ for all component vectors in the tensor \mathcal{Y} with probability at least

$$1 - 2r^2 d \exp\left(-\frac{c_2 \mu^2 n}{2K^2}\right) - 2rd \exp\left(-\frac{c_1 n}{K^4}\right) \geq 1 - 2r^2 d \exp(-c \mu^2 n).$$

Lemma 11 is proved. \square

The following two elementary corollaries illustrate the applicability of our theory to independent subgaussian tensors. In these corollaries, the term *subgaussian tensor* always refers to a tensor defined as per Lemma 11, and should not be confused with a tensor with subgaussian elements.

Corollary 4. *Let $\varepsilon \in (0, 3/4]$. Let \mathcal{Y} be a subgaussian tensor defined as in Lemma 11. For low-rank tensors in high-dimensional spaces, such that*

$$n := \min_{i \in [d]} \geq \frac{\log(r^2 d)}{\varepsilon^2 c}$$

(the small constant c is the same as in Lemma 11), with probability at least $1 - \exp(-c' \varepsilon^2 n)$, Theorem 4 holds with better dependence on ε , namely,

$$\left| \|\mathcal{Y}\|^2 - \|\mathcal{Y} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d\|^2 \right| \leq \left(\varepsilon^d r + \varepsilon \right) e^2 \|\boldsymbol{\alpha}\|_2^2.$$

Here, $c' > 0$ is an absolute constant.

Proof. Apply Lemma 11 with $\mu = \varepsilon$. \square

Corollary 5. *Let \mathcal{Y} be a subgaussian tensor defined as in Lemma 11. If*

$$n := \min_{i=1, \dots, d} n_i \geq C r^{2/d} \log(\max(r, d)),$$

with probability at least $1 - \exp(-c' n / r^{2/d})$, Lemma 9 gives a non-trivial lower bound $\|\mathcal{Y}\| \geq 0.99 \|\boldsymbol{\alpha}\|$. Here, $c' > 0$ is an absolute constant.

In particular, the claim holds when $r \leq C_1^d$ and $n \geq C_2 \max\{r, d\}$.

Proof. Apply Lemma 11 with $\mu = \left(\frac{0.01}{r-1}\right)^{d-1}$. \square

Remark 3. Note that in the general case, when r can be as large as $O(n^d)$, the $\mu_{\mathcal{Y}}$ estimate given in Lemma 11 is not strong enough. Indeed, to have a non-trivial probability estimate, one must take $\mu > \sqrt{2d \log n / n}$. However, $\mu_{\mathcal{Y}} \sim \sqrt{d \log n / n}$ together with $r \sim n^d$ do not satisfy the condition of Lemma 9, since $(r-1) \mu_{\mathcal{Y}}^d = (d \log n)^d \gg 1$.

One could use alternative more sophisticated anti-concentration results instead of Lemma 9. For example, it was shown recently in [53] that for any $r \leq 0.99n^d$ and under some mild conditions, $\|\mathcal{Y}\| \geq cn^{-d/2}\|\boldsymbol{\alpha}\|_2$ (in the independent subgaussian setting as discussed above). Note that this result contains additional non-favorable dependence on n . To the best of our knowledge, it is an open question whether general systems of independent (sub)gaussian vectors form tensors that satisfy norm anti-concentration like the one in Lemma 9. See also the discussion in [53].

4. APPLICATIONS TO LEAST SQUARES PROBLEMS AND FITTING CP MODELS

Now, let us consider the following *fitting* problem. Given tensor \mathcal{X} , which is suspected to have (approximately) low CP-rank r , we would like to find the rank- r tensor \mathcal{Y} in the standard form, as per (17), being closest to \mathcal{X} in the tensor Euclidean norm. Although the r -dimensional basis (subspace) of \mathcal{Y} is naturally unknown, a common way to tackle the fitting problem is to start with a randomly generated basis, and then update the basis tensors mode by mode improving the least square error. This brings us to a framework considered in the previous section: a tensor \mathcal{Y} being in some fixed low-dimensional subspace at each step. Since this subspace is changing throughout the fitting process, the oblivious subspace dimension reduction technique is desirable. The fitting problem can be considered as a generalization of the embedding problem introduced in the previous section (with the addition of a potentially full rank tensor \mathcal{X} that is being approximated).

In this section, we formalize the fitting problem and explain how we propose to use modewise dimension reduction for it. Then, we develop the machinery generalizing our methods from Section 2 to incorporate an unknown tensor \mathcal{X} . Finally, we propose a more-sophisticated two-step dimension reduction process that further improves the resulting dimension for both embedding and fitting problems to almost log-optimal order $\mathcal{O}(r\varepsilon^{-2})$.

As explained above, the common alternating least squares approach for fitting a low-rank CP decomposition along the lines of (17) to an arbitrary tensor $\mathcal{X} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ involves solving a sequence of least squares problems

$$(27) \quad \arg \min_{\tilde{\mathbf{y}}_1^{(j)}, \dots, \tilde{\mathbf{y}}_r^{(j)} \in \mathbb{C}^{n_j}} \left\| \mathcal{X} - \sum_{k=1}^r \alpha_k \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \right\|$$

for each $j \in [d]$ after fixing $\{\mathbf{y}_k^{(\ell)}\}_{k \in [r], \ell \in [d] \setminus \{j\}}$. Here, $\mathbf{y}_k^{(j)} = \tilde{\mathbf{y}}_k^{(j)} / \|\tilde{\mathbf{y}}_k^{(j)}\|_2 \forall j, k$ and $\alpha_k = \prod_{\ell=1}^d \|\tilde{\mathbf{y}}_k^{(\ell)}\|_2$.

One then varies j through all values in $[d]$ computing (27) for each j in order to update $\mathbf{y}_k^{(j)} \forall j, k$ (potentially cycling through all d modes many times). This makes it particularly important to solve each least squares problem (27) efficiently.

Fix $j \in [d]$ and let $\mathbf{e}_h \in \mathbb{C}^{n_j}$ be the h^{th} column of the $n_j \times n_j$ identity matrix. To see how our modewise tensor subspace embeddings can be of value for solving (27), one can begin by noting that

$$\begin{aligned} \left\| \mathcal{X} - \sum_{k=1}^r \alpha_k \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \right\|^2 &= \left\| \mathbf{X}_{(j)} - \sum_{k=1}^r \alpha_k \mathbf{y}_k^{(j)} \left(\otimes_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top \right\|_{\text{F}}^2 \\ &= \left\| \sum_{h=1}^{n_j} \left(\mathbf{X}_{(j)}^{(h)} - \sum_{k=1}^r \alpha_k y_{k,h}^{(j)} \mathbf{e}_h \left(\otimes_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right)^\top \right) \right\|_{\text{F}}^2 \end{aligned}$$

where $\mathbf{X}_{(j)}$ denotes mode- j matricization of \mathcal{X} , and all the rows of $\mathbf{X}_{(j)}^{(h)} \in \mathbb{C}^{n_j \times \prod_{\ell \neq j} n_\ell}$ are zero except for its h^{th} -row which matches that of $\mathbf{X}_{(j)}$. We may now compute the squared Frobenius

norm directly above row-wise and get that

$$\begin{aligned} \left\| \mathcal{X} - \sum_{k=1}^r \alpha_k \bigcirc_{\ell=1}^d \mathbf{y}_k^{(\ell)} \right\|^2 &= \sum_{h=1}^{n_j} \left\| \mathbf{x}_{j,h} - \sum_{k=1}^r \alpha_k y_{k,h}^{(j)} \left(\bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right) \right\|_{\text{F}}^2 \\ &= \sum_{h=1}^{n_j} \left\| \mathcal{X}^{(j,h)} - \sum_{k=1}^r \alpha_k y_{k,h}^{(j)} \bigcirc_{\ell \neq j} \mathbf{y}_k^{(\ell)} \right\|^2 \end{aligned}$$

where $\mathbf{x}_{j,h} \in \mathbb{C}^{\prod_{\ell \neq j} n_{\ell}}$ denotes the h^{th} -row of $\mathbf{X}_{(j)}$, and $\mathcal{X}^{(j,h)}$ its tensorized version. As a consequence, (27) can be decoupled into n_j separate least squares problems of the form

$$(28) \quad \arg \min_{\alpha'_{j,h} \in \mathbb{C}^r} \left\| \mathcal{X}^{(j,h)} - \sum_{k=1}^r \alpha'_{j,h,k} \bigcirc_{\ell \neq j}^d \mathbf{y}_k^{(\ell)} \right\|$$

each involving one $(d-1)$ -mode mode- j slice, $\mathcal{X}^{(j,h)}$, of the original tensor \mathcal{X} . Here $\alpha'_{j,h,k} := \alpha_k y_{k,h}^{(j)}$ where α_k is known $\forall k \in [r]$ from (27). Note also that these n_j separate least squares problems can, if desired, be solved in parallel for each different $h \in [n_j]$.

In order to solve each least squares problem (28) we can now utilize modewise JL embeddings and instead solve the smaller least squares problem

$$(29) \quad \arg \min_{\alpha'_{j,h} \in \mathbb{C}^r} \left\| \mathcal{X}^{(j,h)} \bigtimes_{\ell \neq j} \mathbf{A}_{\ell} - \sum_{k=1}^r \alpha'_{j,h,k} \bigcirc_{\ell \neq j}^d \mathbf{y}_k^{(\ell)} \bigtimes_{\ell \neq j} \mathbf{A}_{\ell} \right\|$$

provided that the $\left\{ \mathbf{y}_k^{(\ell)} \right\}_{k \in [r]}$ are sufficiently incoherent for all $\ell \in [d] \setminus \{j\}$ (an easy to check condition). We can then update each entry of $\tilde{\mathbf{y}}_k^{(j)}$ by setting $\tilde{y}_{k,h}^{(j)} = \alpha'_{j,h,k} / \alpha_k$ for all $h \in [n_j]$ and $k \in [r]$.

We prove that the method described above works in Theorem 6. Namely, Theorem 6 shows that the solution to (29) will be close to that of (28) in terms of quality if the matrices \mathbf{A}_j are chosen from appropriate η -optimal JL families of distributions. In order to do that, we first establish that $\left\| \mathcal{X}^{(j,h)} \bigtimes_{\ell \neq j} \mathbf{A}_{\ell} \right\| \approx \left\| \mathcal{X}^{(j,h)} \right\|$ can also hold for all $j \in [d]$ and $h \in [n_j]$. This is proven in Lemma 12. With Lemma 12 in hand, we prove a more general result in Theorem 5 which directly applies to least squares problems as per (29) when $L(\mathcal{Z}) := \mathcal{Z} \bigtimes_{\ell \neq j} \mathbf{A}_{\ell}$ and $\mathbf{A} = \mathbf{I}$.

Lemma 12. *Let $\varepsilon \in (0, 1)$, $\mathcal{Z}^{(1)}, \dots, \mathcal{Z}^{(p)} \in \mathbb{C}^{n_1 \times \dots \times n_d}$, and $\mathbf{A}_1 \in \mathbb{C}^{m_1 \times n_1}$ be an $(\varepsilon/\varepsilon d)$ -JL embedding of the all $p \left(\prod_{\ell=2}^d n_{\ell} \right)$ mode-1 fibers of all p of these tensors,*

$$\mathcal{S}_1 := \bigcup_{t \in [p]} \left\{ \mathcal{Z}_{:,i_2,\dots,i_d}^{(t)} \mid \forall i_{\ell} \in [n_{\ell}], \ell \in [d] \setminus \{1\} \right\} \subset \mathbb{C}^{n_1},$$

into \mathbb{C}^{m_1} . Next, set $\mathcal{Z}^{(1,t)} := \mathcal{Z}^{(t)} \bigtimes_1 \mathbf{A}_1 \in \mathbb{C}^{m_1 \times n_2 \times \dots \times n_d} \forall t \in [p]$, and then let $\mathbf{A}_2 \in \mathbb{C}^{m_2 \times n_2}$ be an $(\varepsilon/\varepsilon d)$ -JL embedding of all $p \left(m_1 \prod_{\ell=3}^d n_{\ell} \right)$ mode-2 fibers

$$\mathcal{S}_2 := \bigcup_{t \in [p]} \left\{ \mathcal{Z}_{i_1,:,i_3,\dots,i_d}^{(1,t)} \mid \forall i_1 \in [m_1] \& i_{\ell} \in [n_{\ell}], \ell \in [d] \setminus \{2\} \right\} \subset \mathbb{C}^{n_2}$$

into \mathbb{C}^{m_2} . Continuing inductively, for each $j \in [d] \setminus [2]$ and $t \in [p]$ set $\mathcal{Z}^{(j-1,t)} := \mathcal{Z}^{(j-2,t)} \bigtimes_{j-1} \mathbf{A}_{j-1} \in \mathbb{C}^{m_1 \times \dots \times m_{j-1} \times n_j \times \dots \times n_d}$, and then let $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ be an $(\varepsilon/\varepsilon d)$ -JL embedding of all

$p \left(\prod_{\ell=1}^{j-1} m_\ell \right) \left(\prod_{\ell=j+1}^d n_\ell \right)$ mode- j fibers

$$\mathcal{S}_j := \bigcup_{t \in [p]} \left\{ \mathcal{Z}_{i_1, \dots, i_{j-1}, :, i_{j+1}, \dots, i_d}^{(j-1, t)} \mid \forall i_\ell \in [m_\ell], \ell \in [j-1] \text{ \& } i_\ell \in [n_\ell], \ell \in [d] \setminus [j], \right\} \subset \mathbb{C}^{n_j}$$

into \mathbb{C}^{m_j} . Then,

$$\left| \left\| \mathcal{Z}^{(t)} \right\|^2 - \left\| \mathcal{Z}^{(t)} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d \right\|^2 \right| \leq \varepsilon \left\| \mathcal{Z}^{(t)} \right\|^2$$

will hold for all $t \in [p]$.

Proof. Fix $t \in [p]$ and let $\mathcal{X}^{(0)} := \mathcal{Z}^{(t)}$, $\mathcal{X}^{(j)} := \mathcal{Z}^{(j, t)}$ for all $j \in [d-1]$, and $\mathcal{X}^{(d)} := \mathcal{Z}^{(d-1, t)} \times_d \mathbf{A}_d = \mathcal{Z}^{(t)} \times_1 \mathbf{A}_1 \cdots \times_d \mathbf{A}_d$. Choose any $j \in [d]$, and let $\mathbf{x}_{j,h} \in \mathbb{C}^{n_j}$ denote the h^{th} column of the mode- j unfolding of $\mathcal{X}^{(j-1)}$, denoted by $\mathbf{X}_{(j)}^{(j-1)}$. It is easy to see that each $\mathbf{x}_{j,h}$ is a mode- j fiber of $\mathcal{X}^{(j-1)} = \mathcal{Z}^{(j-1, t)}$ for each $1 \leq h \leq N'_j := \left(\prod_{\ell=1}^{j-1} m_\ell \right) \left(\prod_{\ell=j+1}^d n_\ell \right)$. Thus, we can see that

$$\begin{aligned} \left| \left\| \mathcal{X}^{(j-1)} \right\|^2 - \left\| \mathcal{X}^{(j)} \right\|^2 \right| &= \left| \left\| \mathcal{X}^{(j-1)} \right\|^2 - \left\| \mathcal{X}^{(j-1)} \times_j \mathbf{A}_j \right\|^2 \right| = \left| \left\| \mathbf{X}_{(j)}^{(j-1)} \right\|_{\text{F}}^2 - \left\| \mathbf{A}_j \mathbf{X}_{(j)}^{(j-1)} \right\|_{\text{F}}^2 \right| \\ &= \left| \sum_{h=1}^{N'_j} \left\| \mathbf{x}_{j,h} \right\|_2^2 - \left\| \mathbf{A}_j \mathbf{x}_{j,h} \right\|_2^2 \right| \leq \sum_{h=1}^{N'_j} \left| \left\| \mathbf{x}_{j,h} \right\|_2^2 - \left\| \mathbf{A}_j \mathbf{x}_{j,h} \right\|_2^2 \right| \\ &\leq \frac{\varepsilon}{\mathfrak{e}d} \sum_{h=1}^{N'_j} \left\| \mathbf{x}_{j,h} \right\|_2^2 = \frac{\varepsilon}{\mathfrak{e}d} \left\| \mathbf{X}_{(j)}^{(j-1)} \right\|_{\text{F}}^2 = \frac{\varepsilon}{\mathfrak{e}d} \left\| \mathcal{X}^{(j-1)} \right\|^2. \end{aligned}$$

A short induction argument now reveals that $\left\| \mathcal{X}^{(j)} \right\|^2 \leq \left(1 + \frac{\varepsilon}{\mathfrak{e}d} \right)^j \left\| \mathcal{X}^{(0)} \right\|^2$ holds for all $j \in [d]$. As a result we can now see that

$$\begin{aligned} \left| \left\| \mathcal{X}^{(0)} \right\|^2 - \left\| \mathcal{X}^{(d)} \right\|^2 \right| &= \left| \sum_{j=1}^d \left\| \mathcal{X}^{(j-1)} \right\|^2 - \left\| \mathcal{X}^{(j)} \right\|^2 \right| \leq \sum_{j=1}^d \left| \left\| \mathcal{X}^{(j-1)} \right\|^2 - \left\| \mathcal{X}^{(j)} \right\|^2 \right| \leq \frac{\varepsilon}{\mathfrak{e}d} \sum_{j=1}^d \left\| \mathcal{X}^{(j-1)} \right\|^2 \\ &\leq \frac{\varepsilon}{\mathfrak{e}d} \sum_{j=1}^d \left(1 + \frac{\varepsilon}{\mathfrak{e}d} \right)^{j-1} \left\| \mathcal{X}^{(0)} \right\|^2 \leq \frac{\varepsilon}{\mathfrak{e}} \left(1 + \frac{\varepsilon}{\mathfrak{e}d} \right)^d \left\| \mathcal{X}^{(0)} \right\|^2. \end{aligned}$$

holds. The desired result now follows from Lemma 10. \square

With Lemma 12 in hand we can now prove that the solution to (29) will be close to that of (28) in terms of quality if the matrices \mathbf{A}_j are chosen appropriately. We have the following general result which directly applies to least squares problems as per (29) when $L(\mathcal{Z}) := \mathcal{Z} \times_{\ell \neq j} \mathbf{A}_\ell$ and $\mathbf{A} = \mathbf{I}$.

Theorem 5 (Embeddings for Compressed Least Squares). *Let $\mathcal{X} \in \mathbb{C}^{n_1 \times \cdots \times n_d}$, \mathcal{L} be an r -dimensional subspace of $\mathbb{C}^{n_1 \times \cdots \times n_d}$ spanned by a set of orthonormal basis tensors $\{\mathcal{T}_k\}_{k \in [r]}$, and $\mathbb{P}_{\mathcal{L}^\perp} : \mathbb{C}^{n_1 \times \cdots \times n_d} \rightarrow \mathbb{C}^{n_1 \times \cdots \times n_d}$ be the orthogonal projection operator on the orthogonal complement of \mathcal{L} . Fix $\varepsilon \in (0, 1)$ and suppose that the linear operator $L : \mathbb{C}^{n_1 \times n_2 \times \cdots \times n_d} \rightarrow \mathbb{C}^{m_1 \times \cdots \times m_{d'}}$ has both of the following properties:*

- (i) L is an $(\varepsilon/6)$ -JL embedding of all $\mathcal{Y} \in \mathcal{L} \cup \{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\}$ into $\mathbb{C}^{m_1 \times \cdots \times m_{d'}}$, and
- (ii) L is an $(\varepsilon/24\sqrt{r})$ -JL embedding of the $4r$ tensors

$$\mathcal{S}' := \bigcup_{k \in [r]} \left\{ \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\left\| \mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X}) \right\|} - \mathcal{T}_k, \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\left\| \mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X}) \right\|} + \mathcal{T}_k, \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\left\| \mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X}) \right\|} - \mathfrak{i}\mathcal{T}_k, \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\left\| \mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X}) \right\|} + \mathfrak{i}\mathcal{T}_k \right\} \subset \mathbb{C}^{n_1 \times n_2 \times \cdots \times n_d}$$

into $\mathbb{C}^{m_1 \times \cdots \times m_{d'}}$.

Furthermore, let $\text{vect} : \mathbb{C}^{m_1 \times \cdots \times m_{d'}} \rightarrow \mathbb{C}^{\prod_{\ell=1}^{d'} m_\ell}$ be a reshaping vectorization operator, and $\mathbf{A} \in \mathbb{C}^{m \times \prod_{\ell=1}^{d'} m_\ell}$ be an $(\varepsilon/3)$ -JL embedding of the $(r+1)$ -dimensional subspace

$$\mathcal{L}' := \text{span} \{ \text{vect} \circ L(\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})), \text{vect} \circ L(\mathcal{T}_1), \dots, \text{vect} \circ L(\mathcal{T}_r) \} \subset \mathbb{C}^{\prod_{\ell=1}^{d'} m_\ell}$$

into \mathbb{C}^m . Then,

$$\left| \|\mathbf{A}(\text{vect} \circ L(\mathcal{X} - \mathcal{Y}))\|_2^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \leq \varepsilon \|\mathcal{X} - \mathcal{Y}\|^2$$

holds for all $\mathcal{Y} \in \mathcal{L}$.

Proof. Note that the theorem will be proven if L is an $(\varepsilon/3)$ -JL embedding of all tensors of the form $\{\mathcal{X} - \mathcal{Y} \mid \mathcal{Y} \in \mathcal{L}\}$ into $\mathbb{C}^{m_1 \times \cdots \times m_{d'}}$ since any such tensor $\mathcal{X} - \mathcal{Y}$ will also have $\text{vect} \circ L(\mathcal{X} - \mathcal{Y}) \in \mathcal{L}'$ so that

$$\begin{aligned} & \left| \|\mathbf{A}(\text{vect} \circ L(\mathcal{X} - \mathcal{Y}))\|_2^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \\ & \leq \left| \|\mathbf{A}(\text{vect} \circ L(\mathcal{X} - \mathcal{Y}))\|_2^2 - \|L(\mathcal{X} - \mathcal{Y})\|^2 \right| + \left| \|L(\mathcal{X} - \mathcal{Y})\|^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \\ & \leq \left| \|\mathbf{A}(\text{vect} \circ L(\mathcal{X} - \mathcal{Y}))\|_2^2 - \|\text{vect} \circ L(\mathcal{X} - \mathcal{Y})\|_2^2 \right| + \frac{\varepsilon}{3} \|\mathcal{X} - \mathcal{Y}\|^2 \\ & \leq \frac{\varepsilon}{3} \|\text{vect} \circ L(\mathcal{X} - \mathcal{Y})\|_2^2 + \frac{\varepsilon}{3} \|\mathcal{X} - \mathcal{Y}\|^2 \\ & = \frac{\varepsilon}{3} \|L(\mathcal{X} - \mathcal{Y})\|^2 + \frac{\varepsilon}{3} \|\mathcal{X} - \mathcal{Y}\|^2 \\ & \leq \frac{\varepsilon}{3} \left(1 + \frac{\varepsilon}{3} \right) \|\mathcal{X} - \mathcal{Y}\|^2 + \frac{\varepsilon}{3} \|\mathcal{X} - \mathcal{Y}\|^2 \leq \varepsilon \|\mathcal{X} - \mathcal{Y}\|^2. \end{aligned}$$

Let $\mathbb{P}_{\mathcal{L}}$ be the orthogonal projection operator onto \mathcal{L} . Our first step in establishing that L is an $(\varepsilon/3)$ -JL embedding of all tensors of the form $\{\mathcal{X} - \mathcal{Y} \mid \mathcal{Y} \in \mathcal{L}\}$ into $\mathbb{C}^{m_1 \times \cdots \times m_{d'}}$ will be to show that L preserves all the angles between $\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})$ and \mathcal{L} well enough that the Pythagorean theorem

$$\|\mathcal{X} - \mathcal{Y}\|^2 = \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X}) + \mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|^2 = \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|^2 + \|\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|^2$$

still approximately holds for all $\mathcal{Y} \in \mathcal{L}$ after L is applied. Toward that end, let $\boldsymbol{\gamma} \in \mathbb{C}^r$ be such that $\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y} = \sum_{k \in [r]} \gamma_k \mathcal{T}_k$ and note that $\|\boldsymbol{\gamma}\|_2 = \|\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|$ due to the orthonormality of $\{\mathcal{T}_k\}_{k \in [r]}$. Appealing to Lemma 4 we now have that

$$\begin{aligned} \langle L(\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}), L(\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})) \rangle &= \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\| \left| \sum_{k \in [r]} \gamma_k \left\langle L(\mathcal{T}_k), L\left(\frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|}\right) \right\rangle \right| \\ (30) \quad &\leq \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\| \left(\frac{\varepsilon}{6\sqrt{r}} \right) \sum_{k \in [r]} |\gamma_k| \leq \frac{\varepsilon}{6} \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\| \|\boldsymbol{\gamma}\|_2 \\ &\leq \frac{\varepsilon}{12} \left(\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|^2 + \|\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|^2 \right) = \frac{\varepsilon}{12} \|\mathcal{X} - \mathcal{Y}\|^2. \end{aligned}$$

Using (30) we can now see that

$$\begin{aligned}
& \left| \|L(\mathcal{X} - \mathcal{Y})\|_2^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \\
&= \left| \|L(\mathcal{X} - \mathcal{Y})\|_2^2 - \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|^2 - \|\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|^2 \right| \\
&\leq \left| \|L(\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X}))\|^2 - \|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|^2 \right| + \left| \|L(\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y})\|^2 - \|\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|^2 \right| \\
&\quad + 2 \langle L(\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}), L(\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})) \rangle \\
&\leq \frac{\varepsilon}{6} (\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|^2 + \|\mathbb{P}_{\mathcal{L}}(\mathcal{X}) - \mathcal{Y}\|^2 + \|\mathcal{X} - \mathcal{Y}\|^2) = \frac{\varepsilon}{3} \|\mathcal{X} - \mathcal{Y}\|^2.
\end{aligned}$$

Thus, L has the desired JL-embedding property required to conclude the proof. \square

Theorems 4 and 5 together with Lemma 12 can now be used to demonstrate the existence of a large range of modewise Johnson-Lindenstrauss Transforms (JLTs) for oblivious tensor subspace embeddings. The following modewise JLT result for tensors describes the compression one can achieve from Theorem 5 if the linear operator L one employs is formed using j -mode products (as considered in Theorem 4) with $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ taken from η -optimal families of JL embedding distributions (in the sense of Definition 2).

Theorem 6. *Fix $\varepsilon, \eta \in (0, 1/2)$ and $d \geq 3$. Let $\mathcal{X} \in \mathbb{C}^{n_1 \times \dots \times n_d}$, $n := \max_j n_j \geq 4r + 1$, and \mathcal{L} be an r -dimensional subspace of $\mathbb{C}^{n_1 \times \dots \times n_d}$ spanned by a basis $\mathcal{B} := \left\{ \bigcup_{\ell=1}^d \mathbf{y}_k^{(\ell)} \mid k \in [r] \right\}$ of rank-1 tensors, with modewise coherence satisfying $\mu_{\mathcal{B}}^{d-1} < 1/2r$. For each $j \in [d]$ draw $\mathbf{A}_j \in \mathbb{C}^{m_j \times n_j}$ with*

$$(31) \quad m_j \geq C_j \cdot r d^3 / \varepsilon^2 \cdot \ln(n / \sqrt[d]{\eta})$$

from an $(\eta/4d)$ -optimal family of JL embedding distributions, where $C_j \in \mathbb{R}^+$ is an absolute constant. Furthermore, let $\mathbf{A} \in \mathbb{C}^{m' \times \prod_{\ell=1}^d m_\ell}$ with

$$m' \geq C' r \cdot \varepsilon^{-2} \cdot \ln \left(\frac{47}{\varepsilon \sqrt[d]{\eta}} \right)$$

be drawn from an $(\eta/2)$ -optimal family of JL embedding distributions, where $C' \in \mathbb{R}^+$ is an absolute constant. Define $L : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{m_1 \times \dots \times m_d}$ by $L(\mathcal{Z}) = \mathcal{Z} \times_1 \mathbf{A}_1 \times_2 \dots \times_d \mathbf{A}_d$. Then with probability at least $1 - \eta$ the linear operator $\mathbf{A} \circ \text{vect} \circ L : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{m'}$ satisfies

$$\left| \|\mathbf{A}(\text{vect} \circ L(\mathcal{X} - \mathcal{Y}))\|_2^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \leq \varepsilon \|\mathcal{X} - \mathcal{Y}\|^2$$

for all $\mathcal{Y} \in \mathcal{L}$.

Proof. To begin, we note that \mathbf{A} will satisfy the conditions required by Theorem 5 with probability at least $1 - \eta/2$ as a consequence of Lemma 5. Thus, if we can also establish that the L will satisfy the conditions required by Theorem 5 with probability at least $1 - \eta/2$ we will be finished with our proof by Theorem 5 and the union bound.

To establish that L satisfies the conditions required by Theorem 5 with probability at least $1 - \eta/2$, it suffices to prove that

- (a) L will be an $(\varepsilon/6)$ -JL embedding of all $\mathcal{Y} \in \mathcal{L}$ into $\mathbb{C}^{m_1 \times \dots \times m_d}$ with probability at least $1 - \eta/4$, and that
- (b) L will be an $(\varepsilon/24\sqrt{r})$ -JL embedding of the $4r + 1$ tensors $\mathcal{S}' \cup \{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\} \subset \mathbb{C}^{n_1 \times n_2 \times \dots \times n_d}$ into $\mathbb{C}^{m_1 \times \dots \times m_d}$ with probability at least $1 - \eta/4$, where the set \mathcal{S}' is defined as in Theorem 5 and apply yet another union bound.

To show that (a) holds we will utilize Theorem 4 and Lemma 9. Since each \mathbf{A}_j matrix is an $(\eta/4d)$ -optimal JL embedding and the sets \mathcal{S}'_j (defined as in Theorem 4) are such that $|\mathcal{S}'_j| < n^d$,

we know that each \mathbf{A}_j is an $(\varepsilon/480d\sqrt{r})$ -JL embedding of \mathcal{S}'_j into \mathbb{C}^{m_j} with probability⁶ at least $1 - \eta/4d$. Thus, Theorem 4 holds with $\varepsilon \rightarrow \varepsilon/120\sqrt{r}$ with probability at least $1 - \eta/4$. Note that the modewise coherence assumption that $\mu_{\mathcal{B}}^{d-1} < 1/2r$ both allows ε^{d-1} to reduce the $\sqrt{r(r-1)}$ factor in (21) to a size less than one for any $\varepsilon \leq 1/\sqrt{r} \leq (1/r)^{1/(d-1)}$, and also allows Lemma 9 to guarantee that $\|\boldsymbol{\alpha}\|_2^2 < 2\|\mathcal{Y}\|^2$ holds for all $\mathcal{Y} \in \mathcal{L}$. Hence, applying Theorem 4 with $\varepsilon \rightarrow \varepsilon/120\sqrt{r}$ will ensure that L is an $(\varepsilon/6)$ -JL embedding of all $\mathcal{Y} \in \mathcal{L}$ into $\mathbb{C}^{m_1 \times \dots \times m_d}$.

To show that (b) holds we will utilize Lemma 12. Note that the \mathcal{S}_j sets defined in Lemma 12 all have cardinalities $|\mathcal{S}_j| \leq p \cdot n^{d-1}$, where $p = 4r + 1 \leq n$ in our current setting. As a consequence we can see that the conditions of Lemma 12 will be satisfied with $\varepsilon \rightarrow \varepsilon/24\sqrt{r}$ for all $j \in [d]$ with probability at least $1 - \eta/4$ by the union bound. Hence, both (a) and (b) hold and our proof is concluded. \square

Remark 4 (About r -dependence). *Fix ε and η . Looking at Theorem 6 we can see that it's intermediate embedding dimension is*

$$\prod_{\ell=1}^d m_\ell \leq C_{\varepsilon, \eta}^d r^d$$

which effectively determines its overall storage complexity. Hence, Theorem 6 will only result in an improved memory complexity over the straightforward single stage vectorization approach if the rank r of \mathcal{L} is relatively small. The purpose of facultative vectorization and subsequent multiplication by an additional JL transform \mathbf{A} in Theorem 6 is to reduce the resulting embedding dimension to the order $O(r/\varepsilon^2)$ from total dimension $O(r^d)$ that we have after the modewise compression.

We will now consider a final tensor subspace embedding result concerning a special case of modewise JL embeddings that is also made possible by our work above. This result will exhibit better dependence with respect to both ε and r than what is achieved by the more general modewise embedding constructions in Theorem 6.

4.1. Fast and Memory Efficient Modewise JL embeddings for Tensors. In this section we consider a fast Johnson-Lindenstrauss transform for tensors recently introduced in [31] which are effectively based on applying fast JL transforms [36] in a modewise fashion.⁷ Given a tensor $\mathcal{Z} \in \mathbb{C}^{n_1 \times \dots \times n_d}$ the transform takes the form

$$(32) \quad L_{\text{FJL}}(\mathcal{Z}) := \mathbf{R}(\text{vect}(\mathcal{Z} \times_1 \mathbf{F}_1 \mathbf{D}_1 \cdots \times_d \mathbf{F}_d \mathbf{D}_d))$$

where $\text{vect} : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^N$ for $N := \prod_{\ell=1}^d n_\ell$ is the vectorization operator, $\mathbf{R} \in \{0, 1\}^{m \times N}$ is a matrix containing m rows selected randomly from the $N \times N$ identity matrix, $\mathbf{F}_\ell \in \mathbb{C}^{n_\ell \times n_\ell}$ is a unitary discrete Fourier transform matrix for all $\ell \in [d]$, and $\mathbf{D}_\ell \in \mathbb{C}^{n_\ell \times n_\ell}$ is a diagonal matrix with n_ℓ random ± 1 entries for all $\ell \in [d]$. The following theorem is proven about this transform in [31, 36].

Theorem 7 (See Theorem 2.1 and Remark 4 in [31]). *Fix $d \geq 1$, $\varepsilon, \eta \in (0, 1)$, and $N \geq C'/\eta$ for a sufficiently large absolute constant $C' \in \mathbb{R}^+$. Consider a finite set $\mathcal{S} \subset \mathbb{C}^{n_1 \times \dots \times n_d}$ of cardinality $p = |\mathcal{S}|$, and let $L_{\text{FJL}} : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^m$ be defined as above in (32) with*

$$m \geq C \left[\varepsilon^{-2} \cdot \log^{2d-1} \left(\frac{\max(p, N)}{\eta} \right) \cdot \log^4 \left(\frac{\log \left(\frac{\max(p, N)}{\eta} \right)}{\varepsilon} \right) \cdot \log N \right],$$

⁶Here we also implicitly use the fact that $\sqrt[d]{d} \leq \sqrt[\ell]{\ell}$ holds for all $d > 0$ in order to avoid a $\sqrt[d]{d}$ term appearing inside the logarithm in (31).

⁷In fact, the fast transform described here differs cosmetically from the form in which it is presented in [31]. However, one can easily see they are equivalent using (15).

where $C > 0$ is an absolute constant. Then with probability at least $1 - \eta$ the linear operator L_{FJL} is an ε -JL embedding of \mathcal{S} into \mathbb{C}^m . If $d = 1$ then we may replace $\max(p, N)$ with p inside all of the logarithmic factors above (see [36]).

Note that the fast transform L_{FJL} requires only $\mathcal{O}(m \log N + \sum_\ell n_\ell)$ i.i.d. random bits and memory for storage. Thus, it can be used to produce fast and low memory complexity oblivious subspace embeddings. The next Theorem does so.

Theorem 8. Fix $\varepsilon, \eta \in (0, 1/2)$ and $d \geq 2$. Let $\mathcal{X} \in \mathbb{C}^{n_1 \times \dots \times n_d}$, $N = \prod_{\ell=1}^d n_\ell \geq 4C'/\eta$ for an absolute constant $C' > 0$, \mathcal{L} be an r -dimensional subspace of $\mathbb{C}^{n_1 \times \dots \times n_d}$ for $\max(2r^2 - r, 4r) \leq N$, and $L_{\text{FJL}} : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{m_1}$ be defined as above in (32) with

$$m_1 \geq C_1 \left[C_2^d \left(\frac{r}{\varepsilon} \right)^2 \cdot \log^{2d-1} \left(\frac{N}{\eta} \right) \cdot \log^4 \left(\frac{\log \left(\frac{N}{\eta} \right)}{\varepsilon} \right) \cdot \log N \right],$$

where $C_1, C_2 > 0$ are absolute constants. Furthermore, let $\mathbf{L}'_{\text{FJL}} \in \mathbb{C}^{m_2 \times m_1}$ be defined as above in (32) for $d = 1$ with

$$m_2 \geq C_3 \left[r \cdot \varepsilon^{-2} \cdot \log \left(\frac{47}{\varepsilon \sqrt[4]{\eta}} \right) \cdot \log^4 \left(\frac{r \log \left(\frac{47}{\varepsilon \sqrt[4]{\eta}} \right)}{\varepsilon} \right) \cdot \log m_1 \right],$$

where $C_3 > 0$ is an absolute constant. Then with probability at least $1 - \eta$ it will be the case that

$$\left| \left\| \mathbf{L}'_{\text{FJL}} (L_{\text{FJL}} (\mathcal{X} - \mathcal{Y})) \right\|_2^2 - \|\mathcal{X} - \mathcal{Y}\|^2 \right| \leq \varepsilon \|\mathcal{X} - \mathcal{Y}\|^2$$

holds for all $\mathcal{Y} \in \mathcal{L}$.

In addition, the $(\mathbf{L}'_{\text{FJL}}, L_{\text{FJL}})$ transform pair requires only $\mathcal{O}(m_1 \log N + \sum_\ell n_\ell)$ random bits and memory for storage (assuming w.l.o.g. that $m_2 \leq m_1$), and $\mathbf{L}'_{\text{FJL}} \circ L_{\text{FJL}} : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{m_2}$ can be applied to any tensor in just $\mathcal{O}(N \log N)$ -time.

Proof. Let $\{\mathcal{T}_k\}_{k \in [r]}$ be an orthonormal basis for \mathcal{L} (note that these basis tensors need not be low-rank), and $\mathbb{P}_{\mathcal{L}^\perp} : \mathbb{C}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{C}^{n_1 \times \dots \times n_d}$ be the orthogonal projection operator onto the orthogonal complement of \mathcal{L} . Theorem 5 combined with Lemmas 6 and 5 imply that the result will be proven if all of the following hold:

(i) L_{FJL} is an $(\varepsilon/24r)$ -JL embedding of the $2r^2 - r$ tensors

$$\left(\bigcup_{1 \leq h < k \leq r} \{\mathcal{T}_k - \mathcal{T}_h, \mathcal{T}_k + \mathcal{T}_h, \mathcal{T}_k - i\mathcal{T}_h, \mathcal{T}_k + i\mathcal{T}_h\} \right) \bigcup \{\mathcal{T}_k\}_{k \in [r]} \subset \mathcal{L}$$

into \mathbb{C}^{m_1} ,

(ii) L_{FJL} is an $(\varepsilon/6)$ -JL embedding of $\{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\}$ into \mathbb{C}^{m_1} ,
 (iii) L_{FJL} is an $(\varepsilon/24\sqrt{r})$ -JL embedding of the $4r$ tensors

$$\bigcup_{k \in [r]} \left\{ \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|} - \mathcal{T}_k, \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|} + \mathcal{T}_k, \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|} - i\mathcal{T}_k, \frac{\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})}{\|\mathbb{P}_{\mathcal{L}^\perp}(\mathcal{X})\|} + i\mathcal{T}_k \right\} \subset \mathbb{C}^{n_1 \times \dots \times n_d}$$

into \mathbb{C}^{m_1} , and

(iv) \mathbf{L}'_{FJL} is an $(\varepsilon/6)$ -JL embedding of a minimal $(\varepsilon/16)$ -cover, \mathcal{C} , of the r -dimensional Euclidean unit sphere in the subspace $\mathcal{L}' \subset \mathbb{C}^{m_1}$ from Theorem 5 with $L = L_{\text{FJL}}$ into \mathbb{C}^{m_2} . Here we note that $|\mathcal{C}| \leq \left(\frac{47}{\varepsilon}\right)^r$.

Furthermore, if m_1 and m_2 are chosen as above for sufficiently large absolute constants C_1 , C_2 , and C_3 then Theorem 7 implies that each of (i) – (iv) above will fail to hold with probability at most $\eta/4$. The desired result now follows from the union bound.

The number of random bits and storage complexity follows directly from Theorem 7 after noting that each row of \mathbf{R} in (32) is determined by $\mathcal{O}(\log N)$ bits. The fact that $\mathbf{L}'_{\text{FJL}} \circ L_{\text{FJL}}$ can be applied to any tensor \mathcal{Z} in $\mathcal{O}(N \log N)$ -time again follows from the form of (32). Note that each j -mode product with $\mathbf{F}_j \mathbf{D}_j$ involves $\prod_{\ell \neq j} n_\ell$ multiplications of $\mathbf{F}_j \mathbf{D}_j$ against all the mode- j fibers of the given tensor \mathcal{Z} , each of which can be performed in $\mathcal{O}(n_j \log(n_j))$ -time using fast Fourier transform techniques (or approximated even more quickly using sparse Fourier transform techniques if n_j is itself very large – see e.g. [23, 42, 10, 28, 29, 46]). The required vectorization and applications of \mathbf{R} can then be performed in just $\mathcal{O}(N)$ -time thereafter. Finally, Fourier transform techniques can again be used to also apply \mathbf{L}'_{FJL} in $\mathcal{O}(m_1 \log m_1)$ -time. \square

We are now prepared to consider the numerical performance of such modewise JL transforms.

5. EXPERIMENTS

In this section, it is shown that the norms of several different types of (approximately) low-rank data can be preserved using JL embeddings, and trial least squares experiments with compressed tensor data have been performed to show the effect of these embeddings on solutions to least squares problems. The data sets used in the experiments consist of

- (1) *MRI data*: This data set contains three 3-mode MRI images of size $240 \times 240 \times 155$ [1].
- (2) *Randomly generated data*: This data set contains 10 rank-10 4-mode tensors. Each test tensor is a $100 \times 100 \times 100 \times 100$ tensor that is created by adding 10 randomly generated rank-1 tensors. More specifically, each rank-10 tensor is generated according to

$$\mathcal{X}^{(m)} = \sum_{k=1}^r \bigodot_{j=1}^d \mathbf{y}_k^{(j)},$$

where $m \in [10]$, $r = 10$, $d = 4$ and $\mathbf{y}_k^{(j)} \in \mathbb{R}^{100}$. In the Gaussian case, each entry of $\mathbf{y}_k^{(j)}$ is drawn independently from the standard Gaussian distribution $\mathcal{N}(0, 1)$. In the case of coherent data, low-variance Gaussian noise is added to a constant, i.e., each entry $\mathbf{y}_{k,\ell}^{(j)}$ of $\mathbf{y}_k^{(j)}$ is set as $1 + \sigma g_{k,\ell}^{(j)}$ with $g_{k,\ell}^{(j)}$ being an i.i.d. standard Gaussian random variable defined above, and σ^2 denoting the desired variance. In the experiments of this section, $\sigma = \sqrt{0.1}$ is used. In both cases, the 2-norm of $\mathbf{y}_k^{(j)}$ is also normalized to 1.

The reason for running experiments on both Gaussian and coherent data is to show that although coherence requirements presented in section 3 are used to help get general theoretical results for a large class of modewise JL embeddings, they do not seem to be necessary in practice.

When JL embeddings are applied, experiments are performed using Gaussian JL matrices as well as Fast JL matrices. For Gaussian JL, $\mathbf{A}_j = \frac{1}{\sqrt{m}} \mathbf{G}$ is used for all $j \in [d]$, where m is the target dimension and each entry in \mathbf{G} is an i.i.d. standard Gaussian random variable $\mathbf{G}_{i,j} \sim \mathcal{N}(0, 1)$. For Fast JL, $\mathbf{A}_j = \frac{1}{\sqrt{m}} \mathbf{R} \mathbf{F} \mathbf{D}$ is used for all $j \in [d]$, where \mathbf{R} denotes the random restriction matrix, \mathbf{F} is the DFT matrix and \mathbf{D} is a diagonal matrix with Rademacher random variables forming its diagonal [36]. The compression on a test tensor \mathcal{X} is computed by

$$\mathcal{X}_p = \mathcal{X} \times_1 \mathbf{A}_1 \times \cdots \times_d \mathbf{A}_d,$$

where \mathcal{X}_p denotes the projected tensor.

5.1. Effect of JL Embeddings on Norm. In this section, numerical results have been presented, showing the effect of mode-wise JL embedding on the norm of 3 MRI 3-mode images treated as generic tensors, as well as randomly generated data.

The compression ratio for the j^{th} mode, denoted by $c_s^{(j)}$, is defined as the compression in the size of each of the mode- j fibers, i.e.,

$$c_s^{(j)} = \frac{m_j}{n_j}.$$

The target dimension m_j in JL matrices is chosen as $m_j = \lceil c_s n_j \rceil$ for all $j \in [d]$, to ensure that *at least* a fraction c_s of the ambient dimension in each mode is preserved. In the experiments, the compression ratio is set to be the same for all modes, i.e., $c_s^{(j)} = c_s$ for all $j \in [d]$.

Assuming \mathcal{X} and \mathcal{X}_p denote the original and projected tensors respectively, the relative norm of \mathcal{X} is defined by

$$c_{n,\mathcal{X}} = \frac{\|\mathcal{X}_p\|}{\|\mathcal{X}\|}.$$

The results of this section depict the interplay between $c_{n,\mathcal{X}}$ and c_s for both MRI and randomly generated data, where the numbers have been averaged over 1000 trials, as well as over all samples for each value of c_s . In the case of Figure 1, 1000 randomly generated JL matrices were applied to each mode of all 10 randomly generated tensors. In Figure 2, 1000 JL embedding choices have been averaged over each of the 3 MRI images as well as the 3 images themselves. As expected, it can be observed that increasing the compression ratio leads to better norm and distance preservation of the MRI data as the numbers on the vertical axes approach 1.

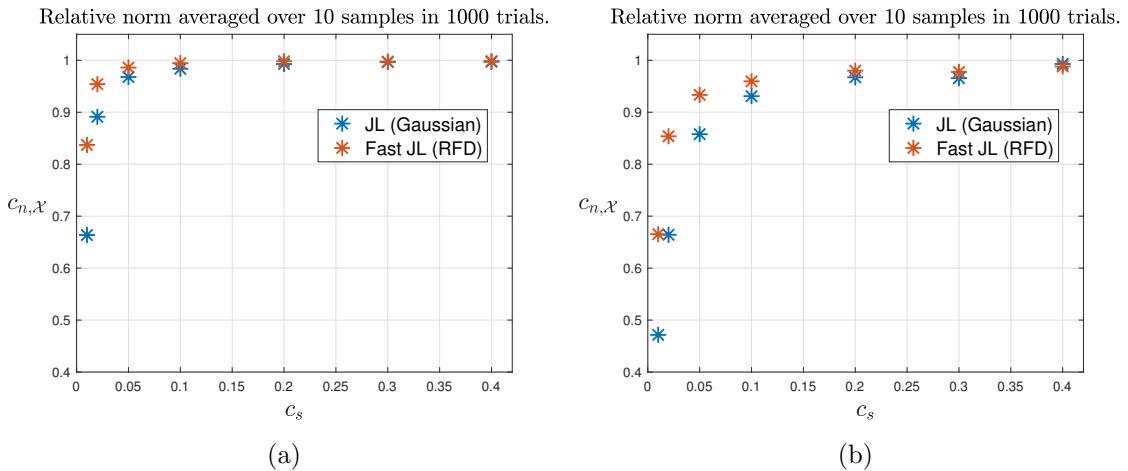


FIGURE 1. Relative norm of randomly generated data. (a) Gaussian data. (b) Coherent data.

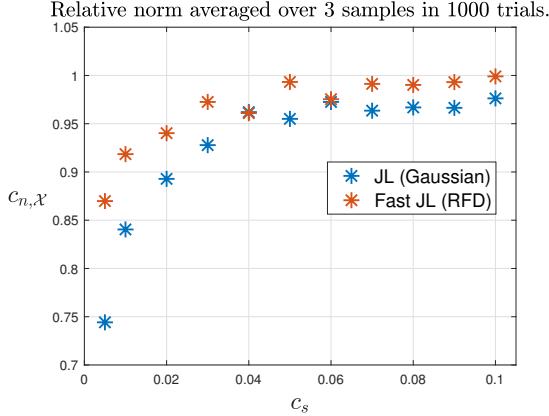


FIGURE 2. Average relative norm of 3 MRI data samples.

5.2. Effect of JL Embeddings on Least Squares Solutions. In this section, the first sample of the MRI data is used in the experiments. First, it is shown that this MRI sample has a relatively low-rank CP representation by calculating its CP reconstruction error for various values of rank. Next, the effect of modewise JL on least squares solutions is investigated by solving for the coefficients of the CP decomposition of the MRI sample in a least squares problem. This will be done by performing modewise JL on the data, which we call compressed least squares, and will be compared with the case where a regular least squares problem is solved.

5.2.1. CPD Reconstruction. Before the experimental results, a short description of the basic form of CPD calculation is presented as well as how the number of rank-1 tensors, r , is chosen. Given a tensor \mathcal{X} , assume r is known beforehand. The problem is now the calculation of $\mathbf{y}_k^{(j)}$ for $j \in [d]$ and $k \in [r]$ and $\boldsymbol{\alpha}$ in (17), i.e. the solution to

$$(33) \quad \min_{\hat{\mathcal{X}}} \|\mathcal{X} - \hat{\mathcal{X}}\| \text{ with } \hat{\mathcal{X}} = \sum_{k=1}^r \alpha_k \mathbf{y}_k^{(1)} \bigcirc \mathbf{y}_k^{(2)} \bigcirc \cdots \bigcirc \mathbf{y}_k^{(d)}.$$

As the Euclidean norm a d -mode tensor is equal to the Frobenius norm of its mode- j unfoldings for $j \in [d]$, by letting $\mathbf{y}_k^{(j)}$ be the k^{th} column of a matrix $\mathbf{Y}^{(j)} \in \mathbb{C}^{n_j \times r}$, the above minimization problem can be written as

$$\min_{\hat{\mathbf{Y}}^{(j)}} \left\| \mathbf{X}_{(j)} - \hat{\mathbf{Y}}^{(j)} \left(\mathbf{Y}^{(d)} \odot \cdots \odot \mathbf{Y}^{(j+1)} \odot \mathbf{Y}^{(j-1)} \odot \cdots \odot \mathbf{Y}^{(1)} \right)^{\top} \right\|_{\text{F}}$$

where $\hat{\mathbf{Y}}^{(j)} = \mathbf{Y}^{(j)} \text{diag}(\boldsymbol{\alpha})$, and \odot denotes the Khatri-Rao product defined as the columnwise matching Kronecker product. The operator $\text{diag}(\cdot)$ creates a diagonal matrix with $\boldsymbol{\alpha}$ as its diagonal. Once solved for, the columns of $\hat{\mathbf{Y}}^{(j)}$ can then be normalized and used to form the coefficients $\alpha_k = \prod_{j=1}^d \|\hat{\mathbf{y}}_k^{(j)}\|_2$ for $k \in [r]$, although this is optional, i.e., if the columns are not normalized, the coefficients α_k in the factorization will all be ones. This procedure is repeated iteratively until the fit ceases to improve (the objective function stops improving with respect to a tolerance) or the maximum number of iterations are exhausted. This procedure is known as CPD-ALS⁸ [35]. To choose the rank of the decomposition as well as obtaining the best estimates for $\mathbf{Y}^{(j)}$, a commonly used consistency diagnostic called CORCONDIA⁹ can be employed [11].

⁸Alternating Least Squares

⁹CORe CONSistency DIAgnostic

In the remainder of this section, the relative reconstruction error of CPD is calculated and plotted for various values of rank r . Assuming \mathcal{X} represents the data, this error is defined as

$$e_{cpd} = \frac{\|\mathcal{X} - \hat{\mathcal{X}}\|}{\|\mathcal{X}\|},$$

where $\hat{\mathcal{X}}$ denotes the reconstruction of \mathcal{X} . Figure 3 displays the results.

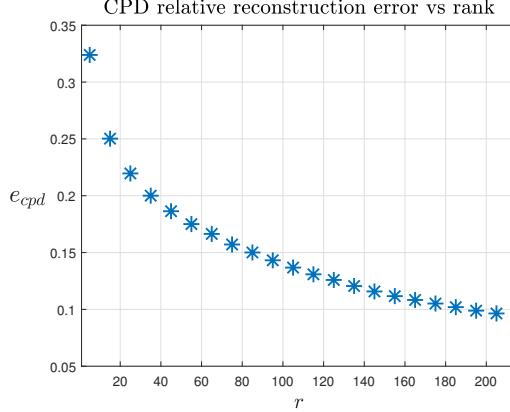


FIGURE 3. Relative reconstruction error of CPD calculated for different values of rank r for MRI data. As the rank increases, the error becomes smaller.

5.2.2. *Compressed Least Squares Performance.* Let $\mathbf{y}_k^{(j)}$ be known in

$$\mathcal{X} \approx \sum_{k=1}^r \alpha_k \bigcirc_{j=1}^d \mathbf{y}_k^{(j)},$$

for $k \in [r]$ and $j \in [d]$. They can be obtained from a previous iteration in the CPD fitting procedure. Here, they come from the CPD of the data calculated in section 5.2.1. Also, assume these vectors have unit norms. In general, as stated in section 5.2.1, when $\mathbf{y}_k^{(j)}$ are obtained using a CPD algorithm, they do not necessarily have unit norms. Therefore, they are normalized and the norms are absorbed into the coefficients of CPD. In other words, $\alpha_k = \prod_{j=1}^d \|\mathbf{y}_k^{(j)}\|_2$ for $k \in [r]$. If the normalization of the vectors is not performed, $\alpha_k = 1$ for $k \in [r]$. The coefficients of the CPD fit are the solutions to the following least squares problem,

$$\boldsymbol{\alpha} = \arg \min_{\boldsymbol{\beta}} \left\| \mathcal{X} - \sum_{k=1}^r \beta_k \bigcirc_{j=1}^d \mathbf{y}_k^{(j)} \right\|.$$

As normalization of $\mathbf{y}_k^{(j)}$ was not performed when computing the CPD of the data in these experiments, the true solution will be $\boldsymbol{\alpha} = \mathbf{1}$. An approximate solution for the coefficients can be obtained by solving for

$$\boldsymbol{\alpha}_p = \arg \min_{\boldsymbol{\beta}} \left\| \mathcal{X} \bigtimes_{j=1}^d \mathbf{A}_j - \sum_{k=1}^r \beta_k \bigcirc_{j=1}^d \mathbf{A}_j \mathbf{y}_k^{(j)} \right\|.$$

where $\boldsymbol{\alpha}_p$ is the vector $\boldsymbol{\alpha}$ estimated for data randomly projected by JL matrices \mathbf{A}_j . This is in fact a way of demonstrating that solving (29) yields an approximate solution to (28) for a $(d-1)$ -mode tensor. Both of these problems can be solved using the vectorized versions of the tensors. Indeed,

for α_p , vectorization should be done after modewise random projection of \mathcal{X} and the rank-1 tensors, i.e.,

$$\alpha_p = \arg \min_{\beta} \|\mathbf{x}_p - \mathbf{B}\beta\|_2 = (\mathbf{B}^*\mathbf{B})^{-1} \mathbf{B}^* \mathbf{x}_p,$$

where $\mathbf{x}_p = \text{vect} \left(\mathcal{X} \bigtimes_{j=1}^d \mathbf{A}_j \right)$, and \mathbf{B} is a matrix whose k^{th} column is $\text{vect} \left(\bigodot_{j=1}^d \mathbf{A}_j \mathbf{y}_k^{(j)} \right)$ for $k \in [r]$.

The relative norm of coefficients, denoted by $c_{n,\alpha}$, is defined as

$$c_{n,\alpha} = \frac{\|\alpha_p\|_2}{\|\alpha\|_2},$$

and is plotted in Figure 4 for different values of c_s .

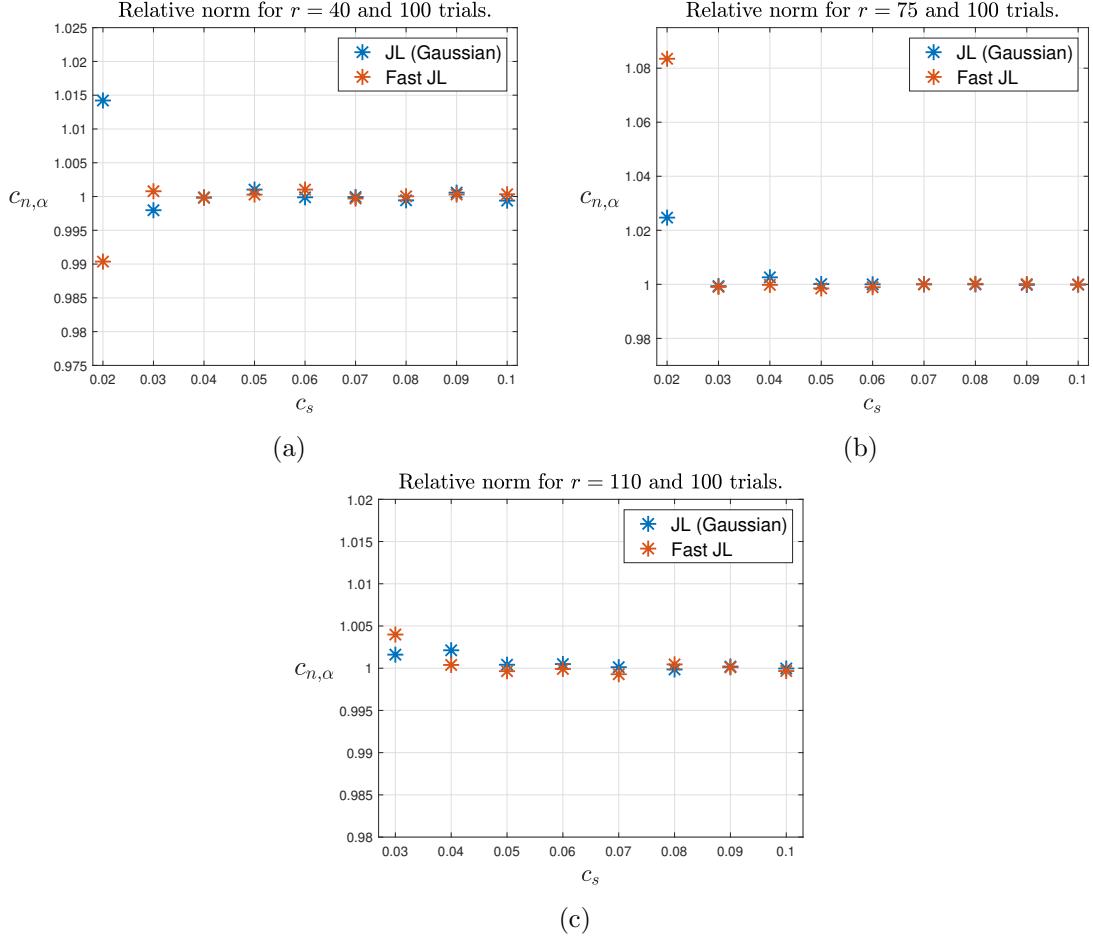


FIGURE 4. Effect of JL embeddings on the relative norm of least squares estimation of CPD coefficients. (a) $r = 40$. (b) $r = 75$. (c) $r = 110$. It can be observed that when the MRI sample is compressed to a very small tensor that is 0.03^3 of its original size, the coefficients are still very accurate in 2-norm.

6. CONCLUSION

We have proposed general modewise Johnson-Lindenstrauss (JL) subspace embeddings that are faster to generate and significantly smaller to store than traditional JL embeddings especially for

tensors in very large dimensions. We provided a subspace embedding result with improved space complexity bounds for embeddings of rank- r tensors in the setting of unknown basis tensors. This result also has applications in the vector setting, leading to general near-optimal oblivious subspace embedding constructions that require fewer random bits for subspaces spanned by basis vectors having special Kronecker structure. We also provided new fast JL embeddings for arbitrary r -dimensional subspaces using fewer random bits than standard methods. We showcase these results for applications including compressed least squares and fitting low-rank CP decompositions, while also confirming our results experimentally. There are several interesting future directions including the analysis of other randomly constructed embeddings, the construction of embeddings designed to maintain other types of structures (such as properties of the core tensor), and their effectiveness in reconstruction and inference tasks.

ACKNOWLEDGMENTS

M. Iwen was supported in part¹⁰ by NSF DMS 1912706 and NSF CCF 1615489, Deanna Needell and Elizaveta Rebrova by NSF CAREER DMS 1348721 and NSF BIGDATA 1740325, and Ali Zare by NSF CCF 1615489. Elizaveta Rebrova also acknowledges sponsorship by Capital Fund Management.

REFERENCES

- [1] Alzheimer’s disease neuroimaging initiative. <http://adni.loni.usc.edu/>.
- [2] D. Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *Journal of computer and System Sciences*, 66(4):671–687, 2003.
- [3] A. Ahmed and J. Romberg. Compressive multiplexing of correlated signals. *IEEE Transactions on Information Theory*, 61(1):479–498, 2014.
- [4] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- [5] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [6] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (2):218–233, 2003.
- [7] C. Battaglino, G. Ballard, and T. G. Kolda. A practical randomized cp tensor decomposition. *SIAM Journal on Matrix Analysis and Applications*, 39(2):876–901, 2018.
- [8] M. H. Beck, A. Jäckle, G. A. Worth, and H.-D. Meyer. The multiconfiguration time-dependent Hartree (MCTDH) method: a highly efficient algorithm for propagating wavepackets. *Physics reports*, 324(1):1–105, 2000.
- [9] J. A. Bengua, H. N. Phien, H. D. Tuan, and M. N. Do. Efficient tensor completion for color image and video recovery: Low-rank tensor train. *IEEE Transactions on Image Processing*, 26(5):2466–2479, 2017.
- [10] S. Bittens, R. Zhang, and M. A. Iwen. A deterministic sparse fft for functions with structured fourier sparsity. *Advances in Computational Mathematics*, 45(2):519–561, 2019.
- [11] R. Bro and H. A. Kiers. A new efficient method for determining the number of components in parafac models. *Journal of Chemometrics: A Journal of the Chemometrics Society*, 17(5):274–286, 2003.
- [12] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
- [13] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [14] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE T. Inform. Theory*, 51:4203–4215, 2005.
- [15] 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(3):283–319, 1970.
- [16] M. Charikar, K. Chen, and M. Farach-Colton. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 693–703. Springer, 2002.

¹⁰ Mark would also like to thank E.I. and D. M. for greatly accentuating his UCLA visit by squatting at his Airbnb Oct. 15 – 19, 2019, as well as a to commit a written act of dogeza to his near-optimal wife for agreeing to his being over 2000 miles away during E’s witching months. Mark also sends many thanks to E. S. for helping out with the baby in his place during his absence.

[17] G. Cormode and S. Muthukrishnan. What's hot and what's not: tracking most frequent items dynamically. *ACM Transactions on Database Systems (TODS)*, 30(1):249–278, 2005.

[18] A. Dasgupta, R. Kumar, and T. Sarlós. A sparse Johnson–Lindenstrauss transform. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 341–350. ACM, 2010.

[19] V. De Silva and L.-H. Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1084–1127, 2008.

[20] Y. C. Eldar and G. Kutyniok. *Compressed sensing: theory and applications*. Cambridge University Press, 2012.

[21] S. Foucart and H. Rauhut. A mathematical introduction to compressive sensing. *Applied and Numerical Harmonic Analysis*, 2013.

[22] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*. Springer, 2013.

[23] A. C. Gilbert, P. Indyk, M. Iwen, and L. Schmidt. Recent developments in the sparse fourier transform: A compressed fourier transform for big data. *IEEE Signal Processing Magazine*, 31(5):91–100, 2014.

[24] A. C. Gilbert, M. A. Iwen, and M. J. Strauss. Group testing and sparse signal recovery. In *2008 42nd Asilomar Conference on Signals, Systems and Computers*, pages 1059–1063. IEEE, 2008.

[25] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Physical review letters*, 105(15):150401, 2010.

[26] R. A. Harshman et al. Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multimodal factor analysis. 1970.

[27] M. Iwen and B. Ong. A distributed and incremental SVD algorithm for agglomerative data analysis on large networks. *SIAM Journal on Matrix Analysis and Applications*, 37(4):1699–1718, 2016.

[28] M. A. Iwen. Combinatorial sublinear-time fourier algorithms. *Foundations of Computational Mathematics*, 10(3):303–338, 2010.

[29] M. A. Iwen. Improved approximation guarantees for sublinear-time fourier algorithms. *Applied And Computational Harmonic Analysis*, 34(1):57–82, 2013.

[30] M. A. Iwen. Compressed sensing with sparse binary matrices: Instance optimal error guarantees in near-optimal time. *Journal of Complexity*, 30(1):1–15, 2014.

[31] R. Jin, T. G. Kolda, and R. Ward. Faster Johnson–Lindenstrauss transforms via kronecker products. *arXiv preprint arXiv:1909.04801*, 2019.

[32] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.

[33] D. M. Kane and J. Nelson. Sparser Johnson–Lindenstrauss transforms. *Journal of the ACM*, 61(1):4, 2014.

[34] T. G. Kolda. Orthogonal tensor decompositions. *SIAM Journal on Matrix Analysis and Applications*, 23(1):243–255, 2001.

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

[36] F. Krahmer and R. Ward. New and improved Johnson–Lindenstrauss embeddings via the restricted isometry property. *SIAM Journal on Mathematical Analysis*, 43(3):1269–1281, 2011.

[37] K. G. Larsen and J. Nelson. Optimality of the Johnson–Lindenstrauss lemma. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 633–638. IEEE, 2017.

[38] X. Li, J. Haupt, and D. Woodruff. Near optimal sketching of low-rank tensor regression. In *Advances in Neural Information Processing Systems*, pages 3466–3476, 2017.

[39] J. Liu, P. Musalski, P. Wonka, and J. Ye. Tensor completion for estimating missing values in visual data. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):208–220, 2012.

[40] C. Lubich. *From quantum to classical molecular dynamics: reduced models and numerical analysis*. European Mathematical Society, 2008.

[41] O. A. Malik and S. Becker. Guarantees for the Kronecker fast Johnson–Lindenstrauss transform using a coherence and sampling argument, 2019.

[42] S. Merhi, R. Zhang, M. A. Iwen, and A. Christlieb. A new class of fully discrete sparse fourier transforms: Faster stable implementations with guarantees. *Journal of Fourier Analysis and Applications*, 25(3):751–784, 2019.

[43] H. Rauhut, R. Schneider, and Ž. Stojanac. Low rank tensor recovery via iterative hard thresholding. *Linear Algebra and its Applications*, 523:220–262, 2017.

[44] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.

[45] B. Romera-Paredes, H. Aung, N. Bianchi-Berthouze, and M. Pontil. Multilinear multitask learning. In *International Conference on Machine Learning*, pages 1444–1452, 2013.

[46] B. Segal and M. Iwen. Improved sparse fourier approximation results: faster implementations and stronger guarantees. *Numerical Algorithms*, 63(2):239–263, 2013.

[47] Y. Sun, Y. Guo, J. A. Tropp, and M. Udell. Tensor random projection for low memory dimension reduction. In *NeurIPS Workshop on Relational Representation Learning*, 2018.

- [48] G. Tsitsikas and E. E. Papalexakis. The core consistency of a compressed tensor. *arXiv preprint arXiv:1811.07428*, 2018.
- [49] L. R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966.
- [50] N. Vannieuwenhoven, R. Vandebril, and K. Meerbergen. A new truncation strategy for the higher-order singular value decomposition. *SIAM Journal on Scientific Computing*, 34(2):1027–1052, 2012.
- [51] M. A. O. Vasilescu and D. Terzopoulos. Multilinear independent components analysis. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 547–553. IEEE, 2005.
- [52] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.
- [53] R. Vershynin. Concentration inequalities for random tensors. *arXiv preprint arXiv:1905.00802*, 2019.
- [54] A. Zare, A. Ozdemir, M. A. Iwen, and S. Aviyente. Extension of PCA to higher order data structures: An introduction to tensors, tensor decompositions, and tensor PCA. *Proceedings of the IEEE*, 106(8):1341–1358, 2018.
- [55] H. Zhang, W. He, L. Zhang, H. Shen, and Q. Yuan. Hyperspectral image restoration using low-rank matrix recovery. *IEEE Transactions on Geoscience and Remote Sensing*, 52(8):4729–4743, 2013.