GUARANTEES FOR THE EXISTENCE OF BEST LOW RANK TENSOR APPROXIMATIONS*

ERIC EVERT† AND LIEVEN DE LATHAUWER†

Abstract. The canonical polyadic decomposition (CPD) of a low rank tensor plays a major role in data analysis and signal processing by allowing for unique recovery of underlying factors. However, it is well known that the low rank CPD approximation problem is ill-posed. That is, a tensor may fail to have a best rank R CPD approximation when $R > 1$.

This article gives deterministic bounds for the existence of best low rank tensor approximations over $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. More precisely, given a tensor $\mathcal{T} \in \mathbb{K}^{I\times I\times I}$ of rank $R \leq I$, we compute the radius of a Frobenius norm ball centered at $\mathcal{T}$ in which best $\mathbb{K}$-rank R approximations are guaranteed to exist. In addition we show that every $\mathbb{K}$-rank R tensor inside of this ball has a unique canonical polyadic decomposition. This neighborhood may be interpreted as a neighborhood of "mathematical truth" in with CPD approximation and computation is well-posed.

In pursuit of these bounds, we describe low rank tensor decomposition as a "joint generalized eigenvalue" problem. Using this framework, we show that, under mild assumptions, a low rank tensor which has rank strictly greater than border rank is defective in the sense of algebraic and geometric multiplicities for joint generalized eigenvalues. Bounds for existence of best low rank approximations are then obtained by establishing perturbation theoretic results for the joint generalized eigenvalue problem. In this way we establish a connection between existence of best low rank approximations and the tensor spectral norm. In addition we solve a "tensor Procrustes problem" which examines orthogonal compressions for pairs of tensors.

The main results of the article are illustrated by a variety of numerical experiments.

Key words. tensors, canonical polyadic decomposition, existence, low rank approximation, generalized eigenvalue, multilinear algebra, border rank, Bauer–Fike matching distance

AMS subject classifications. Primary: 15A69, 15A42

1. Introduction.

Tensors, i.e. multiindexed arrays, are natural generalizations of matrices and are commonplace in fields such as machine learning and signal processing where data sets often have inherent higher-order structure [25, 3]. Tensor decompositions are powerful tools which can be used either to recover source signals responsible for the data or to compress data, thus allowing for faster computations [4].

One of the most popular decompositions for tensors is the canonical polyadic decomposition (CPD) which expresses a given tensor as a sum of *rank one tensors*. A key property of this decomposition is that, under mild conditions, the CPD of a tensor is essentially unique [8, 16, 24]. This essential uniqueness makes CPD an indispensable tool in applications such as indpendent compenent analysis where one wants to recover source signals from measured data [15, 5]. In the case that the data has a natural tensor structure, CPD can allow for unique recovery of the source signals.

*Submitted to the editors November 17, 2020.

Funding: Research supported by: (1) Flemish Government: This work was supported by the Fonds de la Recherche Scientifique–FNRS and the Fonds Wetenschappelijk Onderzoek–Vlaanderen under EOS Project no 30468160 (SeLMA); (2) KU Leuven Internal Funds C16/15/059 and ID-N project no 3E190402. (3) This research received funding from the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" programme. (4) Work supported by Leuven Institute for Artificial Intelligence (Leuven.ai).

†KU Leuven, Dept. of Electrical Engineering ESAT/STADIUS, Kasteelpark Arenberg 10, bus 2446, B-3001 Leuven, Belgium and KU Leuven – Kulak, Group Science, Engineering and Technology, E. Sabbelaan 53, B-8500 Kortrijk, Belgium

(Eric.Evert, Lieven.DeLathauwer)@kuleuven.be.

In applications one typically works with a measurement of a *low rank tensor*, that is, a tensor which is representable by a small number of source signals. This measurement is always corrupted by measurement error, thus the measurement itself does not truly have low rank. As a consequence, in order to obtain the signals which represent the tensor, a practitioner must compute a low rank approximation to the measured tensor.

Although computation of a best low rank approximation is a natural problem, it is not necessarily a well-posed problem. That is, an arbitrary tensor may not have a best rank R approximation for some choice of R [18, 7] . In fact, when considering CPD over $\mathbb{R}$, nonexistence of a best low rank approximation occurs with positive probability [26, 31]. In contrast, it has recently been shown that best approximations exist with probability one when working over $\mathbb{C}$ [22].

The issue of potential nonexistence of best low rank approximations is not one which should be taken lightly. Should a tensor fail to have a best low rank approximation, one could of course choose to simply take a near optimal low rank approximation. However, when a best low rank approximation fails to exist, near optimal low rank approximations exhibit a phenomena known as *diverging components* [35, 27]. Roughly, the optimal way to represent a tensor which does not have a best low rank approximation is to express the tensor as destructive interference between signals which, in the limit, have infinite magnitude. Stated in an alternative fashion, computation of low rank approximations is unstable in a neighborhood around a tensor which does not have a best low rank approximation [33].

Although computation of a best rank R approximation can suffer from issues of nonexistence, tensor approximations typically work surprisingly well in practice. The reason for this is closely linked to the low rank structure of tensors which naturally arise in applications. In this article we explain this connection and quantify how much error can be present in the measurement of a low rank tensor before one encounters issues of nonexistence. Additionally, we further develop the perspective that low rank tensor approximation is in fact well-posed in a neighborhood of "mathematical truth" for tensors occurring in practical applications. In particular, we provide deterministic bounds for the existence of best low rank approximations for tensors of order three.

More precisely, given a low rank tensor of order three, we compute a deterministic lower bound for the radius of a Frobenius norm ball centered at the tensor in which best low rank approximations are guaranteed to exist and in which every rank R tensor has a unique CPD. While [1, Theorem 5.1] gives a probabilistic lower bound for the radius of such a neighborhood which may be applied to a restricted class of rank R symmetric tensors, the computation of a deterministic lower bound which may be applied to a large collection of rank R tensors is novel.

To obtain these bounds, we develop perturbation theoretic bounds for a "joint generalized eigenvalue decomposition" which, with mild assumptions, is equivalent to the CPD of a low rank tensor. In particular, we show that low border rank tensors which have rank strictly greater than border rank can be identified with matrix tuples which are defective in the sense of algebraic and geometric multiplicities for the joint generalized eigenvalue problem. Given a low rank tensor, we then compute lower bounds on the distance to a tensor which is defective in this sense.

1.1. Definitions and Notation. The remainder of the introduction gives our basic definitions and notations as well as statements of the main results of the article. Let $\mathbb{K}$ denote $\mathbb{R}$ or $\mathbb{C}$. Given a matrix $\mathbf{U} \in \mathbb{K}^{I_1 \times I_2}$, let $\mathbf{U}^T$ and $\mathbf{U}^H$ denote the transpose and Hermitian of $\mathbf{U}$, respectively. Let $\langle \cdot, \cdot \rangle$ denote the usual inner product on $\mathbb{K}^R$

which is conjugate linear in the second term. We make frequent use of polynomials of the form

$$p(\boldsymbol{\lambda}) = \lambda_1\gamma_1 + \lambda_2\gamma_2 + \cdots + \lambda_R\gamma_R.$$

We often express such a polynomial using $p(\boldsymbol{\lambda}) = \langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}} \rangle$.

A **tensor** is a multiindexed array $\mathcal{T}$ with entries in $\mathbb{K}$. In this article our primary interest lies with tensors of order 3. That is, we study tensors $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$. Define the **outer product** of vectors $\mathbf{a} \in \mathbb{K}^{I_1}, \mathbf{b} \in \mathbb{K}^{I_2}$ and $\mathbf{c} \in \mathbb{K}^{I_3}$, denoted $\mathbf{a} \otimes \mathbf{b} \otimes \mathbf{c}$, to be the $I_1 \times I_2 \times I_3$ tensor whose i, j, kth entry is given by

$$(\mathbf{a} \otimes \mathbf{b} \otimes \mathbf{c})_{ijk} = a_i b_j c_k.$$

A tensor which can be expressed as the outer product of nonzero vectors is called an **elementary tensor** or a **rank one tensor**.

Every tensor may be expressed as a sum of elementary tensors. That is, given a tensor $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ there exists a collection of vector tuples $\{(\mathbf{a}_r, \mathbf{b}_r, \mathbf{c}_r)\}_{r=1}^R \subset \mathbb{K}^{I_1} \times \mathbb{K}^{I_2} \times \mathbb{K}^{I_3}$ such that

$$\mathcal{T} = \sum_{r=1}^{R} \mathbf{a}_r \otimes \mathbf{b}_r \otimes \mathbf{c}_r. \tag{1.1}$$

In the case that R is as small as possible, the tensor $\mathcal{T}$ is said to have $\mathbb{K}$-**rank** R, and the decomposition in equation (1.1) is called a **canonical polyadic decomposition** (CPD) of $\mathcal{T}$. We let $\operatorname{rank}_{\mathbb{K}}(\mathcal{T})$ denote the rank of $\mathcal{T}$ over $\mathbb{K}$. A fundamental question in the study of tensors and in many applications is to determine the CPD of a given tensor over either $\mathbb{R}$ or $\mathbb{C}$. We emphasize that the field in which the entries of the tensor lie may differ from the field over which the decomposition is taken. In particular, if $\mathcal{T}$ is a tensor with real entries, then both real and complex rank are interesting to consider, and these ranks are not necessarily equal, e.g. see [30, 2, 17].

Let $\mathbf{A} \in \mathbb{K}^{I_1 \times R}$ be the matrix whose rth column is the vector $\mathbf{a}_r$, and similarly define $\mathbf{B} \in \mathbb{K}^{I_2 \times R}$ and $\mathbf{C} \in \mathbb{K}^{I_3 \times R}$. Write

$$\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$$

as a compact notation for the CPD of $\mathcal{T}$. The matrices $\mathbf{A}, \mathbf{B}$, and $\mathbf{C}$ are called the **factor matrices** of $\mathcal{T}$. As discussed in [25] one has

$$\mathcal{T}(:, :, k) = \mathbf{A} D_k(\mathbf{C}) \mathbf{B}^T \qquad \text{for all } k = 1, \ldots, I_3.$$

Here $D_k(\mathbf{C}) \in \mathbb{K}^{R \times R}$ is the diagonal matrix whose diagonal entries are given by the kth row of $\mathbf{C}$.

Say a CPD $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ is **unique** if for any other CPD $T = [\![\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]\!]$, there exist a permutation matrix Π and invertible diagonal matrices $\mathbf{D_A}, \mathbf{D_B}, \mathbf{D_C}$ such that

$$\mathbf{A}\mathbf{D_A}\Pi = \tilde{\mathbf{A}} \qquad \text{and} \qquad \mathbf{B}\mathbf{D_B}\Pi = \tilde{\mathbf{B}} \qquad \text{and} \qquad \mathbf{C}\mathbf{D_C}\Pi = \tilde{\mathbf{C}}.$$

That is, a CPD is unique if the only indeterminacies in the CPD are permutation and scaling/counter scaling of the columns of the factor matrices.

Central to the potential nonexistence of a best low rank tensor approximation is the fact that the set of tensors of rank less than or equal to R is, in general, not a closed set [18, 7]. For example, there are tensors having rank 3 which are a limit of rank 2 tensors. Of course, such a tensor cannot have a best rank 2 approximation. Say a $\mathcal{T}$ tensor has **border** $\mathbb{K}$**-rank** R if $\mathcal{T}$ is a limit of $\mathbb{K}$-rank R tensors and there is no integer $L < R$ such that $\mathcal{T}$ is a limit of $\mathbb{K}$-rank L tensors. That is, R should be the smallest integer such that $\mathcal{T}$ is a limit of $\mathbb{K}$-rank R tensors.

1.2. Multilinear rank and compression. In this article we are most interested in tensors which have low rank in a sense we now describe. Given $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$, define

$$\begin{aligned} R_1(\mathcal{T}) &:= \dim \operatorname{span}\{\mathcal{T}(:,k,\ell)\}_{\forall k,\ell}, \\ R_2(\mathcal{T}) &:= \dim \operatorname{span}\{\mathcal{T}(j,:,\ell)\}_{\forall j,\ell}, \\ R_3(\mathcal{T}) &:= \dim \operatorname{span}\{\mathcal{T}(j,k,:)\}_{\forall j,k}. \end{aligned} \tag{1.2}$$

The integer triple $(R_1(\mathcal{T}), R_2(\mathcal{T}), R_3(\mathcal{T}))$ is called the **multilinear rank** of $\mathcal{T}$. In the case where $R_i(\mathcal{T}) = I_i$ for each $i = 1, 2, 3$ we say $\mathcal{T}$ has **full multilinear rank**. The vectors in equation (1.2) used to define each $R_i(\mathcal{T})$ are called the **mode-i fibers** of $\mathcal{T}$, and we let $\mathcal{T}_{(i)}$ denote a matrix whose columns are equal to the mode-i fibers of $\mathcal{T}$. The matrix $\mathcal{T}_{(i)}$ is called a **mode-i unfolding** of $\mathcal{T}$.[1]

It is not difficult to show that one always has

$$\operatorname{rank}_{\mathbb{K}}(\mathcal{T}) \geq \max\{R_1(\mathcal{T}), R_2(\mathcal{T}), R_3(\mathcal{T})\}. \tag{1.3}$$

For generic tensors this inequality is strict; however, for many tensors occurring in application we have equality. Namely $\operatorname{rank}_{\mathbb{K}}(\mathcal{T}) \leq \min\{I_1, I_2, I_3\}$, and the factor matrices $\mathbf{A}, \mathbf{B}$, and $\mathbf{C}$ have full column rank. In this article we focus on the case where there exist indices $i_1 \neq i_2$ such that

$$\operatorname{rank}_{\mathbb{K}}(\mathcal{T}) = R_{i_1}(\mathcal{T}) = R_{i_2}(\mathcal{T}), \tag{1.4}$$

and it is in this sense that our tensors have low rank. For ease of notation and without loss of generality we take $i_1 = 1$ and $i_2 = 2$. All of our results may be applied to the other cases where equality (1.4) hold by a simple permutation of indices.

For a generic tensor $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$, its is straightforward to show that one has $R_i(\mathcal{T}) = I_i$ for each $i = 1, 2, 3$. However, it is again typical in applications to have strictly $R_i(\mathcal{T}) < I_i$ for each $i = 1, 2, 3$. In this case one may orthogonally compress $\mathcal{T}$ to a tensor $\mathcal{T}^c \in \mathbb{K}^{R_1(\mathcal{T}) \times R_2(\mathcal{T}) \times R_3(\mathcal{T})}$. The key feature of such mapping is that all geometric and algebraic information about $\mathcal{T}$ may be recovered from that of $\mathcal{T}^c$.

Define the **mode-1 product** between a tensor $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ and a matrix $\mathbf{U} \in \mathbb{K}^{J_1 \times I_1}$ to be the tensor $\mathcal{T} \cdot_1 \mathbf{U} \in \mathbb{K}^{J_1 \times I_2 \times I_3}$ with mode-1 fibers given by

$$(\mathcal{T} \cdot_1 \mathbf{U})(:,k,\ell) = \mathbf{U}(\mathcal{T}(:,k,\ell)) \qquad \text{for all } k, \ell.$$

The mode-2 and mode-3 products are defined analogously. The mode-i product is easily understood through its action on the factor matrices of $\mathcal{T}$. Suppose $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!] \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ and let $\mathbf{U}_i \in \mathbb{K}^{J_i \times I_i}$ for $i = 1, 2, 3$. Then one has

$$\mathcal{T} \cdot_1 \mathbf{U}_1 \cdot_2 \mathbf{U}_2 \cdot_3 \mathbf{U}_3 = [\![\mathbf{U}_1\mathbf{A}, \mathbf{U}_2\mathbf{B}, \mathbf{U}_3\mathbf{C}]\!] \in \mathbb{K}^{J_1 \times J_2 \times J_3}.$$

Having defined the mode-i product, one may orthogonally compress a tensor to a space with dimensions given by its multilinear rank as follows. For $i = 1, 2, 3$ let $\mathbf{V}_i \in \mathbb{K}^{I_i \times R_i(\mathcal{T})}$ be a matrix whose columns form an orthonormal basis for the span of the mode-i fibers of $\mathcal{T}$. Then the tensor $\mathcal{T} \cdot_1 \mathbf{V}_1^H \cdot_2 \mathbf{V}_2^H \cdot_3 \mathbf{V}_3^H \in \mathbb{K}^{R_1(\mathcal{T}) \times R_2(\mathcal{T}) \times R_3(\mathcal{T})}$ contains the same algebraic and geometric information as $\mathcal{T}$, and we call this tensor

[1] There is freedom in the ordering of the columns of $\mathcal{T}_{(i)}$, so in this sense we have not well-defined $\mathcal{T}_{(i)}$. However, the properties of mode-i unfoldings that we are most interested in are invariant under permutations of columns, so this ambiguity makes little difference for our purpose so long as one is consistent in the ordering of columns.

an **orthogonal compression** of $\mathcal{T}$ to its multilinear rank. One may recover $\mathcal{T}$ from $\mathcal{T}^c$ using

$$\mathcal{T} = (\mathcal{T} \cdot_1 \mathbf{V}_1^H \cdot_2 \mathbf{V}_2^H \cdot_3 \mathbf{V}_3^H) \cdot \mathbf{V}_1 \cdot \mathbf{V}_2 \cdot \mathbf{V}_3.$$

We briefly note that it is not necessary to use a column-wise orthonormal matrix to map a tensor to its multilinear rank; however, column-wise orthonormal matrices are useful in that they preserve certain geometric aspects of our problem. In addition, such an orthogonal compression preserves the smallest nonzero singular value of our factor matrices which plays an important role in our upcoming results.

To accomplish such an orthogonal compression it is common to take the core of a multilinear singular value decomposition of $\mathcal{T}$ [6]. Following the discussion above we typically restrict to the case where the multilinear rank of $\mathcal{T}$ is equal to the dimensions of $\mathcal{T}$.

1.3. The joint generalized eigenvalue problem. We now describe a type of eigenvalue decomposition for a matrix tuple which is equivalent to the CPD of a tensor in the case where $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ has $\mathbb{K}$-rank R and full multilinear rank. Note that we will interchangeably view $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ as either a tensor or a K-tuple of $R \times R$ matrices tuple via the map

$$\mathcal{T} \to (\mathcal{T}(:,:,1), \mathcal{T}(:,:,2), \cdots, \mathcal{T}(:,:,K)).$$

For notational convenience we denote $\mathbf{T}_k = \mathcal{T}(:,:,k)$.

Given a tensor (or a matrix tuple) $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, we call a nonzero vector $\mathbf{x} \in \mathbb{K}^R$ a **joint generalized eigenvector** (JGE vector) of $\mathcal{T}$ if there exists a $\boldsymbol{\lambda} \in \mathbb{K}^K$ and a nonzero vector $\mathbf{y} \in \mathbb{K}^R$ such that

$$\mathbf{T}_\ell \mathbf{x} = \lambda_\ell \mathbf{y} \tag{1.5}$$

for all $\ell = 1, \ldots, K$. Note that assuming $\mathbf{y} \neq 0$ implies that $\lambda_i = 0$ if and only if $\mathbf{T}_i\mathbf{x} = 0$.

There are natural scaling indeterminates[2] in equation (1.5). In particular, if the triple $(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y})$ is a solution to equation (1.5), then so is $(\alpha\mathbf{x}, \beta\boldsymbol{\lambda}, \gamma\mathbf{y})$ for any nonzero scalars $\alpha, \beta, \gamma \in \mathbb{K}$ satisfying $\alpha\beta\gamma = 1$. For this reason, if $(\boldsymbol{\lambda}, \mathbf{x})$ is a solution to (1.5), then we say the one-dimensional subspace $\boldsymbol{\Lambda} := \operatorname{span}(\boldsymbol{\lambda}) \subset \mathbb{K}^K$ is a **joint generalized eigenvalue** of $\mathcal{T}$ and that $(\boldsymbol{\Lambda}, \mathbf{x})$ is a **joint generalized eigenpair** of $\mathcal{T}$. We typically use $\boldsymbol{\Lambda} \subset \mathbb{K}^K$ to denote a JGE value, while $\boldsymbol{\lambda} \in \mathbb{K}^K$ typically denotes a vector of unit norm such that $\operatorname{span}(\boldsymbol{\lambda}) = \boldsymbol{\Lambda}$. Say a tensor has a **joint generalized eigenbasis** if its joint generalized eigenvectors form a basis for $\mathbb{K}^R$. We call the set of JGE values of a tensor the **spectrum** of the tensor. An equivalent formulation of joint generalized eigenvectors is used in [36] for blind source extraction.

We note that some authors have considered other notions of eigenvalues and eigenvectors for a tensor. These are not equivalent to and should not be confused with joint generalized eigenvalues. The eigenvalues and eigenvectors considered by other authors are intended to allow a spectral theory for general tensors. See [29, 20, 21].

1.3.1. JGE values vs. generalized eigenvalues. If $K = 2$, then our definition of a joint generalized eigenvector exactly coincides with the classical notion of a generalized eigenvector of a matrix pencil. However, our definitions lead to a slightly

[2]Scaling indeterminates in joint generalized eigenvalues are closely related to scaling indeterminates present in a CPD.

different notion of generalized eigenvalues. Namely, a generalized eigenvalue resulting from our definition is the orthogonal complement of a classical generalized eigenvalue.

This makes little difference for perturbation theoretic results, and one could extend the classical definition of generalized eigenvalues to general K; however, the resulting joint generalized eigenvalues would be subspaces of $\mathbb{K}^K$ of dimension $K-1$. As such, we feel the orthogonal setup leads to a more satisfying notion for joint generalized eigenvalues.

See [28, Chapter VI] for an overview of perturbation theory for the classical generalized eigenvalue problem. Also see [13, Chapter 7].

1.4. Multiplicities of JGE values. One of the main themes of this article is that low border rank tensors which have rank strictly greater than border rank can be identified with matrix tuples which are defective in the sense of algebraic and geometric multiplicities for joint generalized eigenvalues. These multiplicities we now define. To simplify the exposition, we only consider tensors of size $R \times R \times K$ when discussing multiplicities. While it is possible to define multiplicities of JGE values for general third order tensors, certain pathologies arise in the $I_1 \times I_2 \times I_3$ case which must be treated with care. By way of orthogonal compressions, we have no need to consider multiplicities in the general case.

Given a tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, define the polynomial $p_{\mathcal{T}}(\gamma)$ in the indeterminate $\gamma = (\gamma_1, \ldots, \gamma_K)$ by

$$p_{\mathcal{T}}(\gamma) := \det\left(\sum_{k=1}^{K} \gamma_i \mathbf{T}_k\right).$$

We call $p_{\mathcal{T}}(\gamma)$ the **characteristic polynomial** of $\mathcal{T}$. Given a nonzero $\boldsymbol{\lambda} \in \mathbb{K}^K$, the **algebraic multiplicity** of $\boldsymbol{\Lambda} = \operatorname{span}(\boldsymbol{\lambda})$ as a JGE value of the tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, denoted $\operatorname{am}(\boldsymbol{\Lambda})$, is the largest integer m such that $\langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}}\rangle^m$ divides the characteristic polynomial of $\mathcal{T}$.

Define the **geometric multiplicity** of $\boldsymbol{\Lambda}$ as a JGE value of $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, denoted $\operatorname{gm}(\boldsymbol{\Lambda})$, to be the dimension of the span of the set of JGE vectors of $\mathcal{T}$ which correspond to $\operatorname{span}(\boldsymbol{\lambda})$. That is,

$$\operatorname{gm}(\boldsymbol{\Lambda}) = \dim \operatorname{span}(\{\mathbf{x}|\ (\boldsymbol{\Lambda}, \mathbf{x}) \text{ is a JGE pair of } \mathcal{T}\}).$$

Similar to the cases of classical and generalized eigenvalues, the algebraic multiplicity of a JGE value is always greater than or equal to the geometric multiplicity of said JGE value, see Proposition 2.3. We say a tensor is **defective** if it has a JGE value $\boldsymbol{\Lambda}$ with $\operatorname{am}(\boldsymbol{\Lambda}) > \operatorname{gm}(\boldsymbol{\Lambda})$.

A degenerate case can occur in which $p_{\mathcal{T}}(\boldsymbol{\gamma})$ is identically equal to zero. In this case, $\langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}}\rangle^m$ divides $p_{\mathcal{T}}(\boldsymbol{\gamma})$ for all $\boldsymbol{\lambda} \in \mathbb{K}^K$ and for all integers m, hence we may reasonably say that $\operatorname{span}(\boldsymbol{\lambda})$ is a JGE value of $\mathcal{T}$ with infinite algebraic multiplicity for all $\boldsymbol{\lambda} \in \mathbb{K}^K$. In this sense, if $p_{\mathcal{T}}(\gamma) \equiv 0$ then the tensor $\mathcal{T}$ is necessarily defective. In the case $p_{\mathcal{T}}(\gamma)$ is not identically zero, we say $\mathcal{T}$ is **regular**. It is straightforward to show that $\mathcal{T}$ is regular if and only if the span of the frontal slices of $\mathcal{T}$ contains an invertible matrix. See Example 2.1 and Lemma 2.1 for further discussion. In the case that $\mathcal{T}$ has size $I_1 \times I_2 \times I_3$ and multilinear rank (R, R, K) we say $\mathcal{T}$ is regular if an orthogonal compression of $\mathcal{T}$ to its multilinear rank is regular.

We are now in position to formally state the connection between the joint generalized eigenvalue decomposition and the canonical polyadic decomposition.

THEOREM 1.1. *Let $\mathcal{T} \in \mathbb{R}^{I_1\times I_2\times I_3}$ and assume $\mathcal{T}$ has multilinear rank (R, R, K). Let $\mathcal{T}^c$ be an orthogonal compression of $\mathcal{T}$ to its multilinear rank. Then $\mathcal{T}$ has $\mathbb{K}$-rank*

R if and only if $\mathcal{T}^c$ has R JGE values in $\mathbb{K}^K$ counting geometric multiplicity. With these assumptions, if $\mathcal{T}^c = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ has $\mathbb{K}$-rank R then the JGE values of $\mathcal{T}^c$ are equal to the spans of columns of the factor matrix $\mathbf{C}$, and the JGE vectors of $\mathcal{T}^c$ are given by the columns of $\mathbf{B}^{-T}$. Furthermore, if $\mathcal{T}$ has border $\mathbb{K}$-rank R and is regular, then $\mathcal{T}^c$ has R JGE values in $\mathbb{K}^K$ counting algebraic multiplicity.

As an immediate consequence, if $\mathcal{T}$ is regular and has border $\mathbb{K}$-rank equal to R but $\mathbb{K}$-rank strictly greater than R, then $\mathcal{T}^c$ must have a JGE value which has algebraic multiplicity strictly greater than geometric multiplicity. That is, $\mathcal{T}$ is defective in the sense of joint generalized eigenvalues.

Proof. The proof of Theorem 1.1 is given in Section 2.4. ∎

Theorem 1.1 leads to a strategy for determining the radius of an open ball centered at a given rank R tensor in which best $\mathbb{K}$-rank R approximations are guaranteed to exist. In particular, it is sufficient to guarantee that nearby tensors are nondefective.

PROPOSITION 1.2. *Let $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ and assume $\mathcal{T}$ has rank R and multilinear rank (R, R, K), and let $\mathcal{T}^c$ be an orthogonal compression of $\mathcal{T}$ to its multilinear rank. If every border $\mathbb{K}$-rank R tensor in an open ball of Frobenius radius ϵ around $\mathcal{T}^c$ is regular and has R distinct JGE values counting algebraic multiplicity, then every tensor in an open ball of radius $\epsilon/2$ around $\mathcal{T}$ has a best $\mathbb{K}$-rank R approximation. In addition, every rank R tensor inside the open ball of Frobenius radius ϵ has a unique CPD.*

In particular, with these assumptions the ball of radius ϵ around $\mathcal{T}$ does not contain a tensor with border $\mathbb{K}$-rank R and $\mathbb{K}$-rank strictly greater than R.

Proof. See Section 3.2. ∎

1.5. Deterministic bounds for existence of a best rank R approximation. While Proposition 1.2 gives a strategy for determining the existence of a best rank R approximation, it provides little insight into how to actually compute ϵ. In the case that $K = 2$ we may use perturbation theory for the classical generalized eigenvalue problem to compute ϵ, see Theorem 3.5. This as well as general perturbation theory for the joint generalized eigenvalue problem is described in Section 3.

For general K, Proposition 4.1 shows that if a tensor is defective, than any subpencil of the tensor must be defective. Hence by examining any given subpencil, one may use Theorem 3.5 to compute the radius of a ball around a given tensor in which best $\mathbb{K}$-rank R approximations exist. A bound obtained from a single pencil may then be improved by combining information from a collection of non-overlapping subpencils of the tensor. In this way one arrives at the following multiple pencil based bound for the existence of a best $\mathbb{K}$-rank R approximation.

THEOREM 1.3. *Let $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ be a tensor of $\mathbb{K}$-rank R and multilinear rank (R, R, K), and let $\mathcal{T}^c$ be an orthogonal compression of $\mathcal{T}$ to its multilinear rank. Let $\mathbf{U} \in \mathbb{K}^{K \times K}$ be any unitary matrix and set $\mathcal{S} = \mathcal{T}^c \cdot_3 \mathbf{U}$. For each $i = 1, \ldots, \lfloor K/2 \rfloor$, let $\epsilon_i \geq 0$ be a nonnegative real number such that every matrix pencil in an open Frobenius norm ball of radius ϵ_i around the pencil*

$$(\mathbf{S}_{2i-1}, \mathbf{S}_{2i})$$

is regular and has R distinct generalized eigenvalues and set $\epsilon = ||(\epsilon_1, \ldots, \epsilon_{\lfloor K/2 \rfloor})||_2$. Then every tensor in an open ball of Frobenius radius $\epsilon/2$ centered at $\mathcal{T}$ has a best $\mathbb{K}$-rank R approximation and every rank R tensor inside the open ball of Frobenius radius ϵ centered at $\mathcal{T}$ has a unique CPD.

In particular, each ϵ_i may be computed as described in Theorem 3.5.

Proof. See Section 4.1. ∎

A slight variation on Theorem 1.3 may be used to check if a given multilinear rank (R, R, K) tensor has a best rank R approximation.

THEOREM 1.4. *Let $\tilde{\mathcal{T}} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ be a tensor of multilinear rank (R, R, K), and let $\tilde{\mathcal{T}}^c$ be an orthogonal compression of $\tilde{\mathcal{T}}$ to its multilinear rank. Let $\mathbf{U} \in \mathbb{K}^{K \times K}$ be any unitary matrix and set $\tilde{\mathcal{S}} = \tilde{\mathcal{T}}^c \cdot_3 \mathbf{U}$. For each $i = 1, \ldots, \lfloor K/2 \rfloor$, let $\epsilon_i \geq 0$ be a nonnegative real number such that every matrix pencil in an open Frobenius norm ball of radius ϵ_i around the pencil*

$$(\tilde{\mathbf{S}}_{2i-1}, \tilde{\mathbf{S}}_{2i})$$

is regular and has R distinct generalized eigenvalues and set $\epsilon = ||(\epsilon_1, \ldots, \epsilon_{\lfloor K/2 \rfloor})||_2$. If there exists some $\mathbb{K}$-rank R tensor $\mathcal{T}'$ such that $\|\tilde{\mathcal{T}} - \mathcal{T}'\|_F < \epsilon$, then $\tilde{\mathcal{T}}$ has a best $\mathbb{K}$-rank R approximation and any best $\mathbb{K}$-rank R approximation of $\tilde{\mathcal{T}}$ has a unique CPD.

In particular, each ϵ_i may be computed as described in Theorem 3.5.

Proof. The proof of Theorem 1.4 is nearly identical to the proof of that of Theorem 1.3. For completeness, details of the proof are given in the supplementary materials. ∎

In practice the tensor $\tilde{\mathcal{T}}$ appearing in the above theorem can be taken to be a multilinear rank (R, R, K) approximation of some measured tensor, and the tensor $\mathcal{T}'$ can simply be any $\mathbb{K}$-rank R approximation to $\tilde{\mathcal{T}}$. In this way Theorem 1.4 gives a method for determining if a multilinear rank (R, R, K) approximation to a measured tensor has a best $\mathbb{K}$-rank R approximation.[3]

To improve clarity, we emphasize the key differences between Theorems 1.3 and 1.4. Namely, Theorem 1.4 is applicable to any multilinear rank (R, R, K) tensor, while Theorem 1.3 is only applicable to tensors of $\mathbb{K}$-rank R. In addition, the multiple pencil based bound ϵ resulting from Theorem 1.4 need not be divided by two to obtain an existence bound. On the other hand, Theorem 1.3 guarantees that every tensor in some open neighborhood has a best rank R approximation, while Theorem 1.4 can only be applied to one tensor at a time.

1.6. An overview of our approach. The primary goal in this article is to give deterministic bounds for the existence of best low rank tensor approximations. A simple yet important idea for our approach is the following: If $\hat{\mathcal{T}}$ is a best border rank R approximation of a perturbation of some rank R tensor $\mathcal{T}$, then the distance from $\mathcal{T}$ to $\hat{\mathcal{T}}$ cannot be more than twice the magnitude of the perturbation. Therefore, to show that every tensor in a ball of radius $\epsilon/2$ around $\mathcal{T}$ has a best rank R approximation, it is sufficient to show that every border rank R tensor in a ϵ radius ball around $\mathcal{T}$ has rank R. Another important aspect of the argument is a WLOG reduction to considering tensors $\mathcal{T}$ and $\hat{\mathcal{T}}$ of size $R \times R \times K$. This reduction is made possible by Theorem 3.1 which considers computation of orthogonal compressions for a pair of tensors.

Using the joint generalized eigenvalue framework, Theorem 1.1 shows that regular tensors of size $R \times R \times K$ with border rank R and rank strictly greater than R are

[3] We caution the reader that the existence of a best rank R approximation for a multilinear rank (R, R, K) approximation of a tensor does not guarantee the existence of a best rank R approximation for the original tensor.

defective in the sense of multiplicities for joint generalized eigenvalues. Additionally Proposition 4.1 shows that if a tensor is defective in the sense of joint generalized eigenvalues, then every subpencil of the tensor is defective.

Theorem 3.5 gives a deterministic bound which may be used to determine that matrix pencils near a given regular matrix pencil are regular and nondefective. Therefore, Theorem 3.5 may be applied to a collection of subpencils of a given regular tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ to determine the radius ϵ of a ball centered at $\mathcal{T}$ in which every tensor is regular and has at least one subpencil which is nondefective. Using Proposition 4.1 together with Theorem 1.1, it follows that every border rank R tensor in this ϵ radius ball is regular and has rank R, hence every tensor in a ball of radius $\epsilon/2$ centered at $\mathcal{T}$ as a best $\mathbb{K}$-rank R approximation.

1.7. Organization. In Section 2 we discuss basic theory for the joint generalized eigenvalue problem. Section 3 discusses perturbation theoretic results for the joint generalized eigenvalue problem, while Section 4 gives deterministic bounds for existence of a best low rank tensor approximation using results from Sections 2 and 3. Numerical experiments illustrating our results are given in Section 5.

2. Basic results for the joint generalized eigenvalue problem. We now discuss several basic results for joint generalized eigenvalues which will be useful in the proofs of Theorem 1.1 and Proposition 1.2. In addition, these results will help illustrate similarities and differences between the classical (generalized) eigenvalue problem and the joint generalized eigenvalue problem.

2.1. A simplifying assumption. As a simplifying assumption throughout the section we restrict to the case where the tensor $\mathcal{T}$ of interest has size $R \times R \times K$ for some $K \leq R$. In addition, we typically assume that $\mathcal{T}$ is regular. The assumption that $\mathcal{T}$ is regular is closely related to the tensor having multilinear rank (R, R, K) and rank R, see Lemma 2.1.

As described in Section 1.2, we may always restrict a tensor to its multilinear rank while preserving geometry, so the assumption that $\mathcal{T}$ has multilinear rank $R \times R \times K$ is not restrictive. The assumption that $\mathcal{T}$ has a linear combination of frontal slices which is invertible holds for tensors with generic entries.

We remark that, without assuming $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ has rank R, the assumption that $\mathcal{T}$ has multilinear rank (R, R, K) is not sufficient for $\mathcal{T}$ to be regular. In fact, it is possible for a tensor to have multilinear rank (R, R, K) and border $\mathbb{K}$-rank R and not be regular. This is illustrated by the following example.

Example 2.1. Let $\mathcal{T} \in \mathbb{R}^{3\times 3\times 3}$ be the tensor with frontal slices

$$\mathbf{T}_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \mathbf{T}_2 = \begin{pmatrix} 0 & 2 & 0 \\ 0 & 0 & -2 \\ 0 & 0 & 0 \end{pmatrix} \quad \mathbf{T}_3 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Then $\mathcal{T}$ has multilinear rank $(3, 3, 3)$ and border $\mathbb{R}$-rank 3, but $\mathcal{T}$ is not regular. That is, the span of the frontal slices of $\mathcal{T}$ does not contain an invertible matrix. To see that $\mathcal{T}$ has border $\mathbb{R}$-rank 3, define the factor matrices

$$\mathbf{A}_n = \begin{pmatrix} 1 & 1 & 1 \\ 0 & \frac{1}{n} & \frac{1}{n^4} \\ 0 & 0 & \frac{1}{n^2} \end{pmatrix} \quad \mathbf{B}_n = \begin{pmatrix} 1 & 0 & 0 \\ -\frac{1}{n^2} & n & 0 \\ -n^2 & -n^2 & n^2 \end{pmatrix} \quad \mathbf{C}_n = \begin{pmatrix} 1 & \frac{1}{n^3} & 1 \\ \frac{4}{n} & \frac{2}{n} & \frac{6}{n} \\ \frac{1}{n^2} & 0 & 0 \end{pmatrix}$$

and set $\mathcal{T}^n = [\![\mathbf{A}_n, \mathbf{B}_n, \mathbf{C}_n]\!]$. Then $\lim \mathcal{T}^n = \mathcal{T}$, hence the border $\mathbb{R}$-rank of $\mathcal{T}$ is at most 3. That $\mathcal{T}$ has border $\mathbb{R}$-rank 3 then follows from the fact that $\mathcal{T}$ has multilinear

rank $(3,3,3)$. As a consequence of the forthcoming Lemma 2.1, $\mathcal{T}$ has rank strictly greater than 3.

LEMMA 2.1. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a tensor with $\mathbb{K}$-rank R and full multilinear rank, and let $\mathcal{T} = [\![\mathbf{A},\mathbf{B},\mathbf{C}]\!]$ be the CPD of $\mathcal{T}$. Then $\mathcal{T}$ is regular and the factor matrices $\mathbf{A},\mathbf{B} \in \mathbb{K}^{R\times R}$ are invertible.*

Proof. Assuming that $\mathcal{T}$ has multilinear rank (R,R,K) implies that each mode-i unfolding of $\mathcal{T}$ has full row rank. By choosing an appropriate ordering of columns, the mode-1 unfolding of $\mathcal{T}$ is given by

$$\mathbf{A}(\mathbf{C}\odot\mathbf{B})^T \in \mathbb{K}^{R\times RK},$$

where $\odot$ denotes the Khatri-Rao (column-wise Kronecker) product, e.g. see [3]. $\mathcal{T}$ having multilinear rank (R,R,K) implies that the mode-1 unfolding has rank R from which it follows that $\mathbf{A} \in \mathbb{K}^{R\times R}$ has rank R and is invertible. A similar argument shows that $\mathbf{B}$ is invertible.

We now show that $\mathcal{T}$ is regular. Using $\mathbf{T}_i = \mathbf{A}D_i(\mathbf{C})\mathbf{B}^T$ where $D_i(\mathbf{C})$ is the diagonal matrix whose entries are given by the ith row of $\mathbf{C}$, to show that $\mathcal{T}$ has a linear combination of frontal slices which is invertible, it is sufficient to show that there is a linear combination of rows of $\mathbf{C}$ which has no entry equal to zero. To this end, note that since $\mathcal{T}$ has $\mathbb{K}$-rank R, the factor matrix $\mathbf{C}$ cannot have a column with all entries equal to zero. Therefore, for each $r = 1,\dots,R$, the set of $\alpha \in \mathbb{K}^K$ such that

$$\sum_{k=1}^{K} \alpha_k \mathbf{C}(r,k) = 0$$

is a nontrivial Zariski closed subset of $\mathbb{K}^K$. It follows that for generic $\alpha \in \mathbb{K}^K$, one has

$$\sum_{k=1}^{K} \alpha_k \mathbf{C}(r,k) \neq 0 \text{ for all } r = 1,\dots,R$$

from which the result follows. ∎

2.2. Distinct JGE values correspond to linearly independent JGE vectors. A perhaps unsurprising yet important fact is that a set of joint generalized eigenvectors corresponding to a set of distinct joint generalized eigenvalues is linearly independent.

LEMMA 2.2. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ and assume $\mathcal{T}$ is regular. Let $\{(\mathbf{\Lambda}_j,\mathbf{x}_j)\}_{j=1}^{\ell}$ be a collection of JGE pairs of $\mathcal{T}$ such that the $\mathbf{\Lambda}_j$ are all distinct. Then $\{\mathbf{x}_j\}_{j=1}^{\ell}$ is a linearly independent set. As a consequence, if $\mathcal{T}$ has R distinct JGE values each with geometric multiplicity at least one, then $\mathcal{T}$ has a JGE basis.*

Proof. For $j = 1,\dots,\ell$ let $\boldsymbol{\lambda}_j \in \mathbb{K}^K$ be a vector such that $\operatorname{span}(\boldsymbol{\lambda}_j) = \mathbf{\Lambda}_j$, and let $\mathbf{L} \in \mathbb{K}^{K\times\ell}$ be the matrix whose jth column is equal to $\boldsymbol{\lambda}_j$. Since the $\mathbf{\Lambda}_j$ are distinct, the matrix $\mathbf{L}$ has Kruskal rank at least 2. That is, no pair of columns of $\mathbf{L}$ is linearly dependent. It follows that there exists a matrix $\mathbf{V} \in \mathbb{K}^{2\times K}$ such that $\mathbf{VL}$ has Kruskal rank 2.

Set $\mathbf{S}_i = \sum_{k=1}^{K} \mathbf{V}(i,k)\mathbf{T}_k$ for $i = 1,2$. It is straightforward to show that each $\mathbf{x}_j$ is a generalized eigenvector of the matrix pencil $(\mathbf{S}_1,\mathbf{S}_2)$ which corresponds to the generalized eigenvalue $\operatorname{span}(\mathbf{V}\boldsymbol{\lambda}_j)$. Furthermore, since the matrix $\mathbf{VL}$ has Kruskal rank 2, the generalized eigenvalues $\{\operatorname{span}(\mathbf{V}\boldsymbol{\lambda}_j)\}_{j=1}^{\ell}$ are all distinct. As a consequence

of [28, Chapter VI, Theorem 1.11], distinct generalized eigenvalues correspond to linearly independent eigenvectors, so it follows that $\{\mathbf{x}_j\}_{j=1}^{\ell}$ is a linearly independent set. ∎

It is possible for a joint generalized eigenvalue of a regular tensor to have nonzero algebraic multiplicity while having geometric multiplicity equal to zero, hence, the assumption that the joint generalized eigenvalues of $\mathcal{T}$ have nonzero geometric multiplicity in the second part of of Lemma 2.2 is necessary.

Example 2.2. Let $\mathcal{T} \in \mathbb{R}^{3\times 3\times 3}$ be the tensor with frontal slices

$$\mathbf{T}_1 = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 3 & 4 & 1 \end{pmatrix} \quad \mathbf{T}_2 = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 2 & 0 \\ 4 & 5 & -1 \end{pmatrix} \quad \mathbf{T}_3 = \begin{pmatrix} 1 & 0 & 0 \\ 5 & 3 & 0 \\ 7 & 8 & 5 \end{pmatrix}.$$

Then

$$p_{\mathcal{T}}(\gamma) = \det\left(\sum_{k=1}^{3} \gamma_k \mathbf{T}_k\right) = (\gamma_1+\gamma_2+\gamma_3)(\gamma_1+2\gamma_2+3\gamma_3)(\gamma_1-\gamma_2+5\gamma_3)$$

hence $\mathcal{T}$ is regular and has three JGE values counting algebraic multiplicity, namely span $(1,1,1)$, span $(1,2,3)$, and span $(1,-1,5)$. However, the JGE values span $(1,1,1)$ and span $(1,2,3)$ do not have a corresponding JGE vector.

In Example 2.2, one could reasonably consider left joint generalized eigenvalues of $\mathcal{T}$ in which case the JGE value span $(1,1,1)$ in Example 2.2 corresponds to the left JGE vector $(1,0,0) \in \mathbb{R}^3$. However, the JGE values span $(1,-1,5)$ and span $(1,2,3)$ do not have left JGE vectors. The JGE value span $(1,2,3)$ is of particular note since it neither has left nor right JGE vectors.

Assuming that a tensor $\mathcal{T}$ is regular, it is not difficult to show that $\mathcal{T}$ has a basis of right JGE vectors if and only if it has a basis of left JGE vectors. However, this fact and left JGE vectors in general have little impact on our results, so we do not discuss them further.

2.3. Algebraic and geometric multiplicities. We now more carefully examine algebraic and geometric multiplicities for joint generalized eigenvalues. In particular we show that algebraic multiplicity is always greater than or equal to geometric multiplicity of a joint generalized eigenvalue. As an immediate consequence, if the tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ has a JGE pair $(\mathbf{\Lambda}, \mathbf{x})$, then $\operatorname{am}(\mathbf{\Lambda}) \geq 1$.

PROPOSITION 2.3. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ and assume $\mathcal{T}$ is regular. Then for any nonzero $\boldsymbol{\lambda} \in \mathbb{K}^K$ one has* $\operatorname{am}(\operatorname{span}(\boldsymbol{\lambda})) \geq \operatorname{gm}(\operatorname{span}(\boldsymbol{\lambda}))$. *In particular, if* $(\operatorname{span}(\boldsymbol{\lambda}), \mathbf{x})$ *is a JGE pair of $\mathcal{T}$, then there is an integer $m > 0$ such that the characteristic polynomial of $\mathcal{T}$ factors as*

$$p_{\mathcal{T}}(\gamma) = \langle \boldsymbol{\lambda}, \overline{\boldsymbol{\gamma}} \rangle^m g(\boldsymbol{\gamma}) \tag{2.1}$$

where g is a nonzero polynomial in the indeterminate γ.

Furthermore, if $\mathcal{T}$ has border $\mathbb{K}$-rank R, then $p_{\mathcal{T}}(\gamma)$ factors into a product of linear terms. That is, one has

$$p_{\mathcal{T}}(\gamma) = \Pi_{r=1}^{R} \langle \boldsymbol{\lambda}_r, \overline{\boldsymbol{\gamma}} \rangle.$$

Proof. Suppose $\operatorname{gm}(\boldsymbol{\lambda}) = m$ and let $\mathbf{x}_1, \ldots, \mathbf{x}_m$ be the corresponding JGE vectors. For $i = 1, \ldots, m$ let $\mathbf{y}_i$ be the nonzero vector corresponding to the eigenpair

$(\mathrm{span}(\boldsymbol{\lambda}), \mathbf{x}_i)$, as in equation (1.5). Note that the assumption that $\mathcal{T}$ is regular guarantees that $\{\mathbf{y}_1, \ldots, \mathbf{y}_m\}$ is a linearly independent set.

Now let $\mathbf{X}_1 \in \mathbb{R}^{R\times m}$ be a matrix whose ith column is $\mathbf{x}_i$, and let $\mathbf{X}_2 \in \mathbb{R}^{R\times(R-m)}$ be any matrix such that the block matrix $\mathbf{X} = (\mathbf{X}_1\ \mathbf{X}_2)$ is invertible. Similarly let $\mathbf{Y}_1 \in \mathbb{R}^{R\times m}$ be the matrix whose ith column is $\mathbf{y}_i$, and let $\mathbf{Y}_2 \in \mathbb{R}^{(R-m)\times R}$ be a matrix whose rows form a basis for the orthogonal complement of the range of $\mathbf{Y}_1$. Set

$$\mathbf{Y} = \begin{pmatrix} \mathbf{Y}_1^\dagger \\ \mathbf{Y}_2 \end{pmatrix}.$$

Then for each $k = 1, \ldots, K$ one has

$$\mathbf{Y}\mathbf{T}_k\mathbf{X} = \mathbf{Y}(\boldsymbol{\lambda}_k\mathbf{Y}_1\ \mathcal{T}_k\mathbf{X}_2) = \begin{pmatrix} \boldsymbol{\lambda}_k\mathbf{Y}_1^\dagger\mathbf{Y}_1 & \mathbf{Y}_1\mathcal{T}_k\mathbf{X}_2 \\ \boldsymbol{\lambda}_k\mathbf{Y}_2\mathbf{Y}_1 & \mathbf{Y}_2\mathcal{T}_k\mathbf{X}_2 \end{pmatrix} = \begin{pmatrix} \boldsymbol{\lambda}_k I_m & \mathbf{Y}_1\mathcal{T}_k\mathbf{X}_2 \\ \mathbf{0} & \mathbf{Y}_2\mathcal{T}_k\mathbf{X}_2 \end{pmatrix}.$$

The matrices $\mathbf{X}$ and $\mathbf{Y}$ are both invertible by construction so using

$$\det(\mathbf{Y})\det\left(\sum_{k=1}^K \lambda_k\mathbf{T}_k\right)\det(\mathbf{X}) = \det\left(\sum_{k=1}^K \lambda_k\mathbf{Y}\mathbf{T}_k\mathbf{X}\right)$$

shows that $p_{\mathcal{T}}(\gamma)$ factors as in equation (2.1). The fact that the polynomial $g(\gamma)$ is nonzero is a consequence of the assumption that $\mathcal{T}$ is regular, hence $p_{\mathcal{T}}$ is not identically zero.

To complete the proof, first assume that $\mathcal{T}$ has $\mathbb{K}$-rank R and let $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ be the CPD of $\mathcal{T}$. We have assumed $\mathcal{T}$ is regular, so Lemma 2.1 shows that $\mathbf{A}$ and $\mathbf{B}$ are invertible. Using $\mathbf{T}_k = \mathbf{A}D_k(\mathbf{C})\mathbf{B}^T$ shows that

$$p_{\mathcal{T}}(\boldsymbol{\gamma}) = \det\left(\mathbf{A}\sum_{k=1}^K \gamma_k D_k(\mathbf{C})\mathbf{B}^T\right) = \det(\mathbf{A})\left(\sum_{k=1}^K \gamma_k D_k(\mathbf{C})\right)\det(\mathbf{B}^T).$$

The final term in this equality easily factors as desired.

Now assume that $\mathcal{T}$ has border $\mathbb{K}$-rank R. Then $\mathcal{T} = \lim \mathcal{T}^n$ for some sequence $\{\mathcal{T}^n\}_{n=1}^\infty$ of $\mathbb{K}$-rank R tensors. Since $\mathcal{T}$ is regular, we may assume that each $\mathcal{T}^n$ is regular. Thus, as we have shown above, for each n we have

$$p_{\mathcal{T}^n}(\boldsymbol{\gamma}) = \alpha_n \Pi_{r=1}^R \langle \boldsymbol{\lambda}_{r,n}, \overline{\boldsymbol{\gamma}}\rangle.$$

Here the constant $\alpha_n \in \mathbb{K}$ has been added so that we may scale to have $||\boldsymbol{\lambda}_{r,n}||_2 = 1$ for each $r = 1, \ldots, R$ and for all $n \in \mathbb{N}$.

Since all the $\boldsymbol{\lambda}_{r,n}$ have norm 1, by passing to a common subsequence if necessary, the sequence $\{\boldsymbol{\lambda}_{r,n}\}_{n=1}^\infty$ converges to some unit norm $\boldsymbol{\lambda}_r \in \mathbb{K}^K$ for each $r = 1, \ldots, R$. We next show that the constants α_n are bounded in norm. To this end note that there must exist some unit vector $\boldsymbol{\gamma}_* \in \mathbb{K}^K$ such that

$$\Pi_{r=1}^R \langle \boldsymbol{\lambda}_{r,n}, \overline{\boldsymbol{\gamma}_*}\rangle \neq 0,$$

as otherwise would imply that the union of the orthogonal complements of the $\boldsymbol{\lambda}_r$ is dense in $\mathbb{K}^K$.

Since the $\mathcal{T}^n$ converge to $\mathcal{T}$ we must have $\lim_{n\to\infty} p_{\mathcal{T}^n}(\gamma) = p_{\mathcal{T}}(\gamma)$ from which it follows that

$$p_{\mathcal{T}}(\boldsymbol{\gamma}_*) = \lim_{n\to\infty} p_{\mathcal{T}^n}(\boldsymbol{\gamma}_*) = \lim_{n\to\infty} \alpha_n \Pi_{r=1}^R \langle \boldsymbol{\lambda}_{r,n}, \overline{\boldsymbol{\gamma}_*}\rangle.$$

Using this together with

$$\lim_{n\to\infty} \Pi_{r=1}^{R}\langle \boldsymbol{\lambda}_{r,n}, \overline{\boldsymbol{\gamma}_*}\rangle = \Pi_{r=1}^{R}\langle \boldsymbol{\lambda}_{r,n}, \overline{\boldsymbol{\gamma}_*}\rangle \neq 0$$

shows that the sequence $\{\alpha_n\}_{n=1}^{\infty}$ converges to some $\alpha \in \mathbb{K}$. It follows that for all $\boldsymbol{\gamma} \in \mathbb{K}^K$ we have

$$\lim_{n\to\infty} p_{\mathcal{T}^n}(\boldsymbol{\gamma}) = \lim_{n\to\infty} \alpha_n \Pi_{r=1}^{R}\langle \boldsymbol{\lambda}_{n,r}, \overline{\boldsymbol{\gamma}}\rangle = \alpha \Pi_{r=1}^{R}\langle \boldsymbol{\lambda}_{r}, \overline{\boldsymbol{\gamma}}\rangle = p_{\mathcal{T}}(\boldsymbol{\gamma}).$$

Furthermore, since $\mathcal{T}$ is regular, we must have $\alpha \neq 0$. We conclude that $p_{\mathcal{T}}(\boldsymbol{\gamma})$ factors as desired. ∎

2.4. Proof of Theorem 1.1. We are now in position to give the proof of our first main result, Theorem 1.1.

Proof. An orthogonal compression $\mathcal{T}^c$ of $\mathcal{T}$ has $\mathbb{K}$-rank R if and only if $\mathcal{T}$ has $\mathbb{K}$-rank R, so throughout the proof we WLOG assume that $\mathcal{T}$ has size $R \times R \times K$ and full multilinear rank, hence $\mathcal{T}^c = \mathcal{T}$.

First assume that $\mathcal{T}$ has R JGE values in $\mathbb{K}^K$ counting geometric multiplicity. That is, assume that $\mathcal{T}$ has a basis of JGE vectors. Then there exists an invertible matrix $\mathbf{X} \in \mathbb{K}^{R\times R}$, a collection of diagonal matrices $\{\boldsymbol{\Gamma}_k\}_{k=1}^{K}$, and matrix $\mathbf{Y} \in \mathbb{K}^{R\times R}$ such that

$$\mathbf{T}_k\mathbf{X} = \mathbf{Y}\boldsymbol{\Gamma}_k \qquad \text{for all } k = 1, \dots, K.$$

It follows that

$$\mathbf{T}_k = \mathbf{Y}\boldsymbol{\Gamma}_k\mathbf{X}^{-1} \qquad \text{for all } k = 1, \dots, K.$$

Set $\mathbf{A} = \mathbf{Y}$ and $\mathbf{B} = \mathbf{X}^{-T}$ and let $\mathbf{C} \in \mathbb{K}^{R\times K}$ be the matrix whose kth row is given by the diagonal entries of $\boldsymbol{\Gamma}_k$. Then $[\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!] = \mathcal{T}$ is an R-term CPD for $\mathcal{T}$, hence $\mathcal{T}$ has rank less than or equal to R. That $\mathcal{T}$ has rank R then follows from the assumption that $\mathcal{T}$ has multilinear rank (R, R, K) and the fact that the rank of a tensor dominates its multilinear rank.

Now, similar to the proof of Proposition 2.3, assume that $\mathcal{T}$ has rank R with CPD $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$. By assumption $\mathcal{T}$ has full multilinear rank, so Lemma 2.1 shows that the factor matrices $\mathbf{A}, \mathbf{B} \in \mathbb{K}^{R\times R}$ are invertible. Using $\mathbf{T}_k = \mathbf{A}D_k(\mathbf{C})\mathbf{B}^T$, we then have

$$\mathbf{T}_k\mathbf{B}^{-T} = \mathbf{A}D_k(\mathbf{C}) \qquad \text{for all } k = 1, \dots, K.$$

For $r = 1, \dots, R$, let $\mathbf{x}_r$ be the rth column of the matrix $\mathbf{B}^{-T}$ and let $\mathbf{y}_r$ be the rth column of $\mathbf{A}$. Then $\mathbf{x}_r, \mathbf{y}_r \in \mathbb{K}^K$ are nonzero vectors for all r and for each $r = 1, \dots, R$ and $k = 1, \dots, K$ we have

$$\mathbf{T}_k\mathbf{x}_r = c_{kr}\mathbf{y}_r$$

where c_{kr} denotes the (k, r) entry of $\mathbf{C}$. Setting $\boldsymbol{\Lambda}_r = \operatorname{span}\big(c_{1r}, \dots, c_{Kr}\big)$, it follows that $(\boldsymbol{\Lambda}_r, \mathbf{x}_r)$ is a joint generalized eigenpair of $\mathcal{T}$ for each $r = 1, \dots, R$. We conclude that $\mathcal{T}$ has a JGE basis, that the JGE values of $\mathcal{T}$ are given by the spans of the columns of $\mathbf{C}$, and that the JGE vectors of $\mathcal{T}$ are given by the columns of $\mathbf{B}^{-T}$, as claimed.

The fact that $\mathcal{T}$ being regular with border rank R and multilinear rank (R, R, K) implies that $\mathcal{T}^c$ has R joint generalized eigenvalues counting algebraic multiplicity is an immediate consequence of Proposition 2.3. ∎

2.5. Multiplicities for border rank R tensors. We end the section by examining geometric multiplicities of JGE values for regular border $\mathbb{K}$-rank R tensors. Namely we show that JGE values of a regular tensor with border $\mathbb{K}$-rank R have nonzero geometric multiplicity. This lemma plays an important role in the proof of Proposition 1.2, as it allows one to conclude that a regular border $\mathbb{K}$-rank R tensor having R distinct JGE values counting algebraic multiplicity also has R distinct JGE values counting geometric multiplicity. Additionally, the lemma shows that regular tensors which have border rank R and rank strictly greater than R can be seen as a generalization of matrices which have a nontrivial Jordan cell.

LEMMA 2.4. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a regular tensor and assume $\mathcal{T}$ has multilinear rank (R, R, K) and border $\mathbb{K}$-rank R. If $\mathbf{\Lambda} \subset \mathbb{R}^K$ is a JGE value of $\mathcal{T}$ with $\operatorname{am}(\mathbf{\Lambda}) = m \geq 1$, then $\operatorname{gm}(\mathbf{\Lambda}) \geq 1$.*

As a consequence, if $\mathcal{T}$ has $\mathbb{K}$-rank strictly greater than R, then $\mathcal{T}$ must have a JGE value with algebraic multiplicity strictly greater than one.

Proof. Let $\{\mathcal{T}^n\}_{n=1}^{\infty}$ be a sequence of $\mathbb{K}$-rank R tensors which converges to $\mathcal{T}$ and let $\boldsymbol{\lambda} \in \mathbb{K}^K$ be a unit vector such that $\operatorname{span}(\boldsymbol{\lambda}) = \mathbf{\Lambda}$. Proposition 2.3 shows that $p_{\mathcal{T}}(\boldsymbol{\gamma})$ and all the $p_{\mathcal{T}^n}(\boldsymbol{\gamma})$ factor as products of linear terms. Note that $\lim \mathcal{T}^n = \mathcal{T}$ implies that $\lim p_{\mathcal{T}^n} = p_{\mathcal{T}}$, so the normalized linear factors of the $p_{\mathcal{T}^n}$ converge to the linear factors of $p_{\mathcal{T}}$. It follows that there exists a sequence of unit vectors $\{\boldsymbol{\lambda}_n\}_{n=1}^{\infty} \subset \mathbb{K}^K$ such that $\operatorname{span}(\boldsymbol{\lambda}_n)$ is a JGE value of $\mathcal{T}^n$ for each n and such that

$$\lim_{n\to\infty} \boldsymbol{\lambda}_n = \boldsymbol{\lambda}.$$

As a consequence of Theorem 1.1, each $\operatorname{span}(\boldsymbol{\lambda}_n)$ has geometric multiplicity at least 1 as a JGE value of $\mathcal{T}^n$. For each n, let $\mathbf{x}_n \in \mathbb{K}^R$ be a unit norm JGE vector of $\mathcal{T}^n$ corresponding to $\boldsymbol{\lambda}_n$, and let $\mathbf{y}_n \in \mathbb{K}^R$ be the nonzero vector such that the tuple $(\mathcal{T}^n, \boldsymbol{\lambda}_n, \mathbf{x}_n, \mathbf{y}_n)$ satisfies equation (1.5) for each n.

By construction the sequences $\{\mathbf{x}_n\}$ and $\{\boldsymbol{\lambda}_n\}$ lie in a compact sets. Additionally, since the $\mathcal{T}^n$ converge to $\mathcal{T}$, the norms of the frontal slices of $\mathcal{T}^n$ are uniformly bounded, so since the $\boldsymbol{\lambda}_n$ all have norm 1, the sequence $\{\mathbf{y}_n\}$ must also lie in a compact set. Therefore, by passing to a common subsequence if necessary, the sequences $\{\mathbf{x}_n\}$ and $\{\mathbf{y}_n\}$ have limits $\mathbf{x}, \mathbf{y} \in \mathbb{K}^R$. For each $k = 1, \dots K$ we then have

$$\mathbf{T}_k\mathbf{x} = \lim_{n\to\infty} \mathbf{T}_k^n\mathbf{x}_n = \lim_{n\to\infty} \lambda_{n,k}\mathbf{y}^n = \lambda_k\mathbf{y} \tag{2.2}$$

where $\lambda_{n,k}$ denotes the kth entry of $\boldsymbol{\lambda}_n$.

By assumption, $\mathcal{T}$ is regular, so the frontal slices of $\mathcal{T}$ cannot have a common null space and the vector $\mathbf{y}$ cannot be equal to zero. We conclude that $(\mathbf{\Lambda}, \mathbf{x})$ is a joint generalized eigenpair of $\mathcal{T}$ and that $\operatorname{gm}(\mathbf{\Lambda}) \geq 1$.

It remains to show that if $\mathcal{T}$ has $\mathbb{K}$-rank strictly greater than R, then $\mathcal{T}$ must have a JGE value with algebraic multiplicity greater than one. Suppose toward a contradiction that every JGE value of $\mathcal{T}$ has algebraic multiplicity equal to one. Using Proposition 2.3, $p_{\mathcal{T}}$ factors as a product of R linear terms. Since each JGE of $\mathcal{T}$ has algebraic multiplicity one, it follows that $\mathcal{T}$ has R distinct JGE values counting algebraic multiplicity.

Using the first part of the result shows that each of the JGE values of $\mathcal{T}$ has nonzero geometric multiplicity. We conclude that $\mathcal{T}$ has R distinct JGE values counting geometric multiplicity. It follows from Theorem 1.1 that $\mathcal{T}$ has $\mathbb{K}$-rank R, a contradiction. ∎

We briefly clarify an important difference in the conclusions of Theorem 1.1 and Lemma 2.4. Both results imply that a regular border rank R tensor with rank strictly greater than border rank must be defective in the sense of JGE values. A priori, the conclusion of Theorem 1.1 leaves the possibility that such a tensor has a JGE value with algebraic multiplicity one and geometric multiplicity zero. Lemma 2.4 expands on this idea by showing that every JGE value of such a tensor has geometric multiplicity at least one. Hence, a regular border rank R tensor with rank strictly greater than border rank must have a JGE value with algebraic multiplicity strictly greater than one.

3. Perturbation bounds for joint generalized eigenvalues. We now examine how joint generalized eigenvalues are affected by a perturbation of the tensor of interest. Our main goal at this point is to provide a way to determine if border $\mathbb{K}$-rank R tensors near a given regular $\mathbb{K}$-rank R tensor have distinct JGE values. That is, we want to show that nearby border $\mathbb{K}$-rank R tensors are nondefective hence have rank equal to border rank. To achieve this goal, we provide a Bauer–Fike like bound for JGE eigenvalues which may be converted to a spectral matching distance bound in the case $K = 2$.

As in the previous section, it will be helpful to reduce to the case where $\mathcal{T}$ has size $R \times R \times K$ and full multilinear rank. To make this reduction we consider a "tensor Procrustes problem" which serves as an analogue of the classical matrix Procrustes problem [23]. In this tensor Procrustes problem, we take a perspective on orthogonal compressions which differs from the perspective most commonly taken in applications. This we now explain.

3.1. Orthogonal compressions of pairs of tensors: The tensor Procrustes problem. As previously discussed, the situation most common to application is that one has access to a measurement $\tilde{\mathcal{T}} = \mathcal{T} + \mathcal{N}$ of a rank R tensor $\mathcal{T}$ which is corrupted by measurement error $\mathcal{N}$. In order to approximate the factors underlying the tensor $\mathcal{T}$, the practitioner will compute a rank R approximation of $\tilde{\mathcal{T}}$. Typically the rank of $\mathcal{T}$ is significantly less than its dimensions, so to reduce the computational complexity of computing a rank R approximation of $\tilde{\mathcal{T}}$ one first computes a low multilinear rank approximation of $\tilde{\mathcal{T}}$. One then computes a R rank approximation of the low multilinear rank approximation.

While this approach is natural to use in applied settings, it can have adverse effects when searching for theoretical guarantees. To simplify to the $R \times R \times K$ setting, we must work with a core of $\mathcal{T}$ and a core of a best (border) rank R approximation of $\tilde{\mathcal{T}}$. While the above methodology yields a core of a rank R approximation of $\tilde{\mathcal{T}}$, if one independently computes this core and a core of $\mathcal{T}$, then it is possible that the distance between the resulting cores is greater than the distance between $\mathcal{T}$ and the computed approximation of $\tilde{\mathcal{T}}$. A key cause of the issue at hand is that the mode-j fibers of $\mathcal{T}$ cannot be expected to have the same span as the mode-j fibers of a best (border) rank R approximation of $\tilde{\mathcal{T}}$, hence an orthogonal compression of one of these tensors is not necessarily an orthogonal compression of the other.

To resolve this issue we first let $\hat{\mathcal{T}}$ be a best (border) rank R approximation of $\tilde{\mathcal{T}}$, then consider a tensor Procrustes problem for $\mathcal{T}$ and $\hat{\mathcal{T}}$ where the goal is compute cores for each tensor such that the distance between the cores is no more than the distance between the original tensors. Essentially our task is to compute orthogonal compressions such that the mode-j subspaces of the cores are well aligned.[4]

[4]In practice one rarely has access to both $\mathcal{T}$ and $\mathcal{T} + \mathcal{N}$, but this is immaterial to our present

THEOREM 3.1. *Let $\mathcal{T}, \hat{\mathcal{T}} \in \mathbb{K}^{I_1 \times \cdots \times I_\ell}$ be tensors having multilinear rank $(R_1, \cdots, R_\ell)$ and $(\hat{R}_1, \cdots, \hat{R}_\ell)$, respectively. Let $\mathbf{V} : \mathbb{K}^{R_j} \to \mathbb{K}^{I_j}$ be column-wise orthonormal with range equal to the subspace spanned by the mode-j fibers of $\mathcal{T}$. If $\tilde{R}_j \leq R_j$, then there exists a column-wise orthonormal matrix $\hat{\mathbf{V}} : \mathbb{R}^{R_j} \to \mathbb{R}^{I_j}$ with range containing the subspace spanned by the mode-j fibers of $\tilde{\mathcal{T}}$ such that*

$$\|\mathcal{T} \cdot_j \mathbf{V}^H - \hat{\mathcal{T}} \cdot_j \hat{\mathbf{V}}^H\|_F \leq \|\mathcal{T} - \hat{\mathcal{T}}\|_F.$$

As an immediate consequence, letting $R_j^* = \max\{R_j, \hat{R}_j\}$ for each j there exist orthogonal compressions $\mathcal{T}^c, \hat{\mathcal{T}}^c \in \mathbb{K}^{R_1^* \times \cdots \times R_\ell^*}$ of $\mathcal{T}$ and $\hat{\mathcal{T}}$, respectively, such that

$$\|\mathcal{T}^c - \hat{\mathcal{T}}^c\|_F \leq \|\mathcal{T} - \hat{\mathcal{T}}\|_F.$$

As will be come clear, we may use Theorem 3.1 to without loss of generality restrict to the case of $R \times R \times K$ tensors. Before proving the theorem we give a technical lemma.

LEMMA 3.2. *Let $\mathbf{M}, \mathbf{N} \in \mathbb{K}^{I_1 \times I_2}$ be matrices of rank $R_{\mathbf{M}}$ and $R_{\mathbf{N}}$, respectively, where $R_{\mathbf{M}} \geq R_{\mathbf{N}}$. There exists a unitary matrix $\mathbf{W} \in \mathbb{K}^{I_1 \times I_1}$ such that* ran $\mathbf{M} \supseteq$ ran $\mathbf{WN}$ *and so that*

$$\|\mathbf{M} - \mathbf{WN}\|_F \leq \|\mathbf{M} - \mathbf{N}\|_F.$$

As a consequence, given any column-wise orthonormal matrix $\mathbf{Z}_{\mathbf{M}} \in \mathbb{K}^{I_1 \times R_{\mathbf{M}}}$ mapping onto ran $\mathbf{M}$*, there is a column-wise orthonormal matrix $\mathbf{Z}_{\mathbf{N}} \in \mathbb{K}^{I_1 \times R_{\mathbf{M}}}$ such that* ran $\mathbf{Z}_{\mathbf{N}} \supseteq$ ran $\mathbf{N}$ *and such that*

$$\|\mathbf{Z}_{\mathbf{M}}^H \mathbf{M} - \mathbf{Z}_{\mathbf{N}}^H \mathbf{N}\|_F \leq \|\mathbf{M} - \mathbf{N}\|_F.$$

Proof. The proof of Lemma 3.2 is a routine argument using the solution of the orthogonal Procrustes problem [23]. Details are given in the supplementary materials. ∎

We now give the proof of Theorem 3.1.

Proof. The main strategy of the proof is to apply Lemma 3.2 to each mode of the tensors $\mathcal{T}$ and $\hat{\mathcal{T}}$. To this end, fix j and let $\mathcal{T}_{(j)}$ and $\hat{\mathcal{T}}_{(j)}$ be mode-j unfoldings of $\mathcal{T}$ and $\hat{\mathcal{T}}$, respectively, with the same ordering of the fibers. Note that $\mathcal{T}_{(j)}$ has rank R_j and $\hat{\mathcal{T}}_{(j)}$ has rank $\hat{R}_j$, and WLOG assume $R_j \geq \hat{R}_j$.

Using Lemma 3.2, given any column-wise orthonormal matrix $\mathbf{V} \in \mathbb{R}^{I_j \times R_j}$ with range equal to ran $\mathcal{T}_{(j)}$, there exists an column-wise orthonormal matrix $\hat{\mathbf{V}} \in \mathbb{K}^{I_j \times R_j}$ such that the range of $\hat{\mathbf{V}}$ contains ran $\hat{\mathcal{T}}_{(j)}$ and such that

$$\|\mathbf{V}^H \mathcal{T}_{(j)} - \hat{\mathbf{V}}^H \hat{\mathcal{T}}_{(j)}\|_F \leq \|\mathcal{T}_{(j)} - \hat{\mathcal{T}}_{(j)}\|_F.$$

Since the ordering of the mode-j fibers is the same in $\mathcal{T}_{(j)}$ and $\hat{\mathcal{T}}_{(j)}$ we have

$$\|\mathcal{T}_{(j)} - \hat{\mathcal{T}}_{(j)}\|_F = \|\mathcal{T} - \hat{\mathcal{T}}\|_F \qquad \text{and} \qquad \|\mathbf{V}^H \mathcal{T}_{(j)} - \hat{\mathbf{V}}^H \hat{\mathcal{T}}_{(j)}\|_F = \|\mathcal{T} \cdot_j \mathbf{V}^H - \hat{\mathcal{T}} \cdot_j \hat{\mathbf{V}}^H\|_F$$

from which the first conclusion follows.

The second part of the result follows from a standard argument using induction on the mode of the fibers. ∎

discussion.

It is worth noting that the Procrustes cores $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ one arrives at using the methodology in the proof are not necessarily optimal. That is, there may exist cores $\mathcal{T}^*$ and $\hat{\mathcal{T}}^*$ such that

$$\|\mathcal{T}^* - \hat{\mathcal{T}}^*\|_F < \|\mathcal{T}^c - \hat{\mathcal{T}}^c\|_F \leq \|\mathcal{T} - \hat{\mathcal{T}}\|_F.$$

One may compute locally optimal cores by using an ALS approach for instance where one continues to iterate over the modes of the tensors and solves the matrix Procrustes problem for each corresponding unfolding. Methods for computing globally optimal cores are outside the scope of this article.

A similar issue occurs when computing a best low multilinear rank approximation of a given tensor $\tilde{\mathcal{T}}$. Here one wants to minimize $\|\tilde{\mathcal{T}} - \hat{\mathcal{T}}\|$ subject to a multilinear rank constraint on $\hat{\mathcal{T}}$. A common approach is to compute a MLSVD of $\tilde{\mathcal{T}}$ and truncate at the specified multilinear rank. While this approach is in a sense optimal "per mode", since the MLSVD treats each mode independently the result is rarely globally optimal when all modes are considered together.

3.2. Proof of Proposition 1.2. We are now in position to prove Proposition 1.2.

Proof. Let $\tilde{\mathcal{T}} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ be any tensor such that $\|\tilde{\mathcal{T}} - \mathcal{T}\| < \epsilon/2$, and let $\hat{\mathcal{T}}$ be a best border $\mathbb{K}$-rank R approximation of $\tilde{\mathcal{T}}$. Note that $\mathcal{T}$ has multilinear rank (R, R, K) by assumption, so $\mathcal{T}$ has border $\mathbb{K}$-rank R in addition to having $\mathbb{K}$-rank R. It follows that

$$\|\hat{\mathcal{T}} - \tilde{\mathcal{T}}\| \leq \|\mathcal{T} - \tilde{\mathcal{T}}\|, \qquad \text{hence} \qquad \|\mathcal{T} - \hat{\mathcal{T}}\| < \epsilon,$$

as otherwise $\mathcal{T}$ would be a strictly better border $\mathbb{K}$-rank R approximation of $\tilde{\mathcal{T}}$ than $\hat{\mathcal{T}}$. To show that $\tilde{\mathcal{T}}$ has a best $\mathbb{K}$-rank R approximation, it is sufficient to show that $\hat{\mathcal{T}}$ has $\mathbb{K}$-rank equal to R.

Since $\hat{\mathcal{T}}$ has border $\mathbb{K}$-rank R, the rank of each mode-i unfolding of $\hat{\mathcal{T}}$ is less than or equal to R. Therefore, using Theorem 3.1, we may simultaneously orthogonally compress $\mathcal{T}$ and $\hat{\mathcal{T}}$ to their multilinear rank without increasing the Frobenius norm between the tensors, so WLOG we may assume $\mathcal{T}$ and $\hat{\mathcal{T}}$ both have size $R \times R \times K^*$ for some $K^* \leq R$. After this restriction, our assumptions imply that $\hat{\mathcal{T}}$ is regular and has R distinct joint generalized eigenvalues counting algebraic multiplicity.

Since $\hat{\mathcal{T}}$ is regular and has border $\mathbb{K}$-rank R, Lemma 2.4 shows that each JGE value of $\hat{\mathcal{T}}$ with algebraic multiplicity at least one must also have geometric multiplicity at least one. We have shown $\hat{\mathcal{T}}$ has R distinct JGE values each having algebraic multiplicity one, so it follows that $\hat{\mathcal{T}}$ has R distinct JGE values each having geometric multiplicity one. Since $\hat{\mathcal{T}}$ is regular we may use Theorem 1.1, we conclude that $\hat{\mathcal{T}}$ has $\mathbb{K}$-rank equal to R.

It remains to prove that every rank R tensor in the open ball of Frobenius radius ϵ centered at $\mathcal{T}$ has a unique CPD. By assumption every rank R tensor in this ball is regular and has distinct joint generalized eigenvalues. Using Theorem 1.1 together with Lemma 2.2, it follows that every tensor inside this ball satisfies Kruskal's condition [16], hence every tensor in this ball has a unique CPD. ∎

3.3. Metrics for joint generalized eigenvalues. Before discussing our perturbation bounds for JGE values, we must first discuss the notion of distance we use to measure how much JGE values change under a perturbation of a tensor. Recalling the natural scaling invariance in the joint generalized eigenvalue problem, it is necessary to use a metric which is scale invariant when working with joint generalized

eigenvalues. Following work of Stewart and Sun [28] for classical generalized eigenvalues, we use the chordal metric to measure the distance between joint generalized eigenvalues.

Given two one-dimensional subspaces $\mathbf{\Lambda}_1, \mathbf{\Lambda}_2 \subseteq \mathbb{K}^K$, the **chordal metric** between $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$, denoted $\chi(\mathbf{\Lambda}_1, \mathbf{\Lambda}_2)$, is equal to the sine of the angle between $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$. Letting $\boldsymbol{\lambda}_i$ be a unit vector such that $\operatorname{span}(\boldsymbol{\lambda}_i) = \mathbf{\Lambda}_i$ for $i = 1, 2$, it is not difficult to show that the chordal metric may be computed as

$$\chi(\mathbf{\Lambda}_1, \mathbf{\Lambda}_2) = \|\mathbf{P}_{(\mathbf{\Lambda}_1)^\perp} \boldsymbol{\lambda}_2\|_2 = \max_{\substack{\gamma \in (\mathbf{\Lambda}_1)^\perp \\ \|\gamma\|_2 = 1}} |\langle \boldsymbol{\lambda}_2, \gamma \rangle| = \sqrt{1 - |\langle \boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2 \rangle|^2}. \tag{3.1}$$

Here $(\mathbf{\Lambda}_1)^\perp \subset \mathbb{K}^K$ is the orthogonal complement of $\mathbf{\Lambda}_1$ and $\mathbf{P}_{(\mathbf{\Lambda}_1)^\perp}$ denotes the projection onto $(\mathbf{\Lambda}_1)^\perp$. Note that the chordal metric is equal to the sine of the angle between the subspaces $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$, while $|\langle \boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2 \rangle|$ corresponds to the cosine of this angle.

We will also need a notion of separation between two sets of JGE values. We consider two different but closely related measures of separation. For $i = 1, 2$, let $\mathcal{T}_i \in \mathbb{K}^{R \times R \times K}$ be a tensor having R JGE values counting algebraic multiplicity, and let $\{\mathbf{\Lambda}_{i,k}\}_{k=1}^R$ be the spectrum of $\mathcal{T}_i$. Define the **spectral variation** from $\mathcal{T}_1$ to $\mathcal{T}_2$, denoted $\operatorname{sv}[\mathcal{T}_1, \mathcal{T}_2]$, by

$$\operatorname{sv}[\mathcal{T}_1, \mathcal{T}_2] = \max_i \min_j \chi(\mathbf{\Lambda}_{1,j}, \mathbf{\Lambda}_{2,i}).$$

Additionally define the **matching distance** between $\mathcal{T}_1$ and $\mathcal{T}_2$, denoted $\operatorname{md}[\mathcal{T}_1, \mathcal{T}_2]$, by

$$\operatorname{md}[\mathcal{T}_1, \mathcal{T}_2] = \min_{\pi \in \mathcal{S}_R} \max_{k=1,\ldots,K} \chi(\mathbf{\Lambda}_{1,k}, \mathbf{\Lambda}_{2,\pi(k)}).$$

Here $\mathcal{S}_R$ is group of permutations of $\{1, \ldots, R\}$.

The key difference between spectral variation and matching distance is that, in matching distance, each JGE value of $\mathcal{T}_2$ must be paired with a unique JGE value of $\mathcal{T}_1$, while in spectral variation the JGE values of $\mathcal{T}_2$ need not be paired with unique JGE values of $\mathcal{T}_1$. Due to the lack of unique pairings in spectral variation, the spectral variation from $\mathcal{T}_1$ to $\mathcal{T}_2$ need not equal the spectral variation from $\mathcal{T}_2$ to $\mathcal{T}_1$. Furthermore, the spectral variation from $\mathcal{T}_1$ to $\mathcal{T}_2$ may be equal to zero while the spectra of $\mathcal{T}_1$ and $\mathcal{T}_2$ are not equal. As a consequence, while matching distance defines a metric on sets of JGE values, spectral variation does not.

Example 3.1. Let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbb{K}^{R \times R \times 2}$. Suppose $\mathcal{T}_1$ has JGE values

$$\mathbf{\Lambda}_{1,1} = \operatorname{span}(\mathbf{e}_1) \qquad \text{and} \qquad \mathbf{\Lambda}_{1,2} = \operatorname{span}(\mathbf{e}_2)$$

and that $\mathcal{T}_2$ has JGE values

$$\mathbf{\Lambda}_{2,1} = \mathbf{\Lambda}_{2,2} = \operatorname{span}(\mathbf{e}_1).$$

Here $\mathbf{e}_1$ and $\mathbf{e}_2$ are the standard basis vectors in $\mathbb{K}^2$. Then one has

$$0 = \operatorname{sv}[\mathcal{T}_1, \mathcal{T}_2] \neq \operatorname{sv}[\mathcal{T}_2, \mathcal{T}_1] = 1.$$

Of course, matching distance bounds are more desirable to obtain than spectral variation bounds; however, they are also typically much more difficult to obtain. A standard strategy to obtain a matching distance bound is to first find a spectral variation bound which grows linearly with the magnitude of a perturbation.

3.4. Bauer–Fike Theorem for JGE values. We soon present our main perturbation theoretic bound for the spectral variation between two tensors of $\mathbb{K}$-rank R. Two bounds are presented. The first bound is given by the magnitude of the best rank one approximation for a transformation of the error tensor. Equivalently, this bound is equal to the spectral norm (i.e. the largest singular value) of a transformation of the error tensor. The spectral norm of a tensor $\mathcal{T} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$, denoted $\|\mathcal{T}\|_{\text{sp}}$ is given by

$$\|\mathcal{T}\|_{\text{sp}} := \max_{\|\mathbf{x}\|_2=\|\mathbf{y}\|_2=\|\mathbf{z}\|_2=1} |\mathcal{T} \cdot_1 \mathbf{x} \cdot_2 \mathbf{y} \cdot_3 \mathbf{z}|$$

where the vectors $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are of appropriate size. See [10, 11, 12] for more information about best rank 1 approximations of tensors, tensor singular values, and the spectral norm of tensors.

The second bound we present depends on the spectral norm of a mode-1 unfolding of the error tensor. That is, this bound depends on $\|\mathcal{E}_{(1)}\|_2$ where $\mathcal{E}$ is the error tensor. It is important to note that the spectral norm of a matrix is invariant under the ordering of columns, so this quantity is well-defined and does not depend on the ordering of columns in $\mathcal{E}_{(1)}$. The first bound is always tighter than the second bound; however the second bound has the advantage that one only need know the magnitude of the error tensor to compute it.

THEOREM 3.3. *Let $\mathcal{T}$ and $\tilde{\mathcal{T}}$ be $R \times R \times K$ tensors with entries in $\mathbb{K}$. Assume that $\mathcal{T}$ has $\mathbb{K}$-rank R with CPD $[\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ where the columns of $\mathbf{C}$ are normalized to have unit norm. Let $\{\mathbf{\Lambda}_k\}_{k=1}^K$ be the spectrum of $\mathcal{T}$ and set $\mathcal{E} = \mathcal{T} - \tilde{\mathcal{T}}$. If $\tilde{\mathbf{\Lambda}}$ is a JGE value of $\tilde{\mathcal{T}}$ with geometric multiplicity at least one, then*

$$\min_k \chi(\mathbf{\Lambda}_k, \tilde{\mathbf{\Lambda}}) \leq \sqrt{R} \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{\Lambda}}^\perp} \right\|_{\text{sp}} \leq \frac{\sqrt{R}\|\mathcal{E}_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}. \tag{3.2}$$

Here $\sigma_{\min}(\mathbf{A})$ and $\sigma_{\min}(\mathbf{B})$ denote the smallest[5] singular values of $\mathbf{A}$ and $\mathbf{B}$, respectively and $\mathbf{P}_{\tilde{\mathbf{\Lambda}}^\perp}$ denotes the projection onto the orthogonal complement of $\tilde{\mathbf{\Lambda}}$. As a consequence, if $\tilde{\mathcal{T}}$ is a regular border $\mathbb{K}$-rank R tensor, then

$$\text{sv}[\mathcal{T}, \tilde{\mathcal{T}}] \leq \sqrt{R} \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\text{sp}} \leq \frac{\sqrt{R}\|\mathcal{E}_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}. \tag{3.3}$$

In the following special cases the coefficient $\sqrt{R}$ may be dropped from equations (3.2) and (3.3):

1. *$K = 2$.*
2. *$\tilde{\mathcal{T}}$ is also regular and has $\mathbb{K}$-rank R with CPD $[\![\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]\!]$ where $\tilde{\mathbf{A}} = \mathbf{A}$ or $\tilde{\mathbf{B}} = \mathbf{B}$.*

That is, in these special cases one has

$$\text{sv}[\mathcal{T}, \tilde{\mathcal{T}}] \leq \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\text{sp}} \leq \frac{\|\mathcal{E}_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}. \tag{3.4}$$

Proof. The inequality

$$\min_k \chi(\mathbf{\Lambda}_k, \tilde{\mathbf{\Lambda}}) \leq \frac{\|\mathcal{E}_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}$$

[5]Lemma 2.1 shows that $\mathbf{A}$ and $\mathbf{B}$ are invertible, hence these singular values are nonzero.

in the $K = 2$ case is [28, Chapter VI, Theorem 2.7] and the $K > 2$ case follows a similar strategy. However, in the $K > 2$ case the orthogonal complement of a JGE value has dimension greater than one, hence a modification of the Stewart and Sun approach is required. This fact leads to the introduction of the constant $\sqrt{R}$ in equations (3.2) and (3.3) when $K > 2$.

We begin by proving the first inequality in equation (3.2). To this end, let $\tilde{\mathbf{\Lambda}} \subset \mathbb{K}^K$ be a JGE value of $\tilde{\mathcal{T}}$ and let $\tilde{\boldsymbol{\lambda}} \in \tilde{\mathbf{\Lambda}}$ and $\boldsymbol{\gamma} \in \tilde{\mathbf{\Lambda}}^{\perp}$ be vectors with unit two norm. Let $\mathbf{x} \in \mathbb{K}^R$ be a JGE vector of $\tilde{\mathcal{T}}$ corresponding to $\tilde{\mathbf{\Lambda}}$ and let $\mathbf{y} \in \mathbb{K}^R$ be a vector such that the tuple $(\tilde{\mathcal{T}}, \tilde{\boldsymbol{\lambda}}, \mathbf{x}, \mathbf{y})$ satisfies equation (1.5). Set $\mathbf{v} = \frac{\mathbf{B}^T\mathbf{x}}{\|\mathbf{B}^T\mathbf{x}\|_2}$ so that $\|\mathbf{v}\|_2 = 1$. Additionally, for $k = 1, \ldots, K$ let $\mathbf{E}_k$ denote the kth frontal slice of $\mathcal{E}$.

With this setup we have

$$\mathbf{0} = \langle \tilde{\boldsymbol{\lambda}}, \boldsymbol{\gamma} \rangle \mathbf{A}^{-1}\mathbf{y} = \sum_{k=1}^{K} \bar{\gamma}_k \mathbf{A}^{-1}\tilde{\mathbf{T}}_k\mathbf{B}^{-T}\mathbf{v} = \sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C})\mathbf{v} + \sum_{k=1}^{K} \bar{\gamma}_k \mathbf{A}^{-1}\mathbf{E}_k\mathbf{B}^{-T}\mathbf{v}.$$

It follows that

$$\sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C})\mathbf{v} = -\sum_{k=1}^{K} \bar{\gamma}_k \mathbf{A}^{-1}\mathbf{E}_k\mathbf{B}^{-T}\mathbf{v}.$$

Furthermore we have

(3.5)

$$\begin{aligned}
\max_{\substack{\boldsymbol{\gamma} \in \tilde{\mathbf{\Lambda}}^{\perp} \\ \|\boldsymbol{\gamma}\|_2 = 1}} \left\| \sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C})\mathbf{v} \right\|_2 &\leq \max_{\substack{\|\mathbf{w}\|_2 = 1 \\ \|\mathbf{z}\|_2 = 1}} \max_{\substack{\boldsymbol{\gamma} \in \tilde{\mathbf{\Lambda}}^{\perp} \\ \|\boldsymbol{\gamma}\|_2 = 1}} \left| \sum_{k=1}^{K} \bar{\gamma}_k \mathbf{w}^T \mathbf{A}^{-1}\mathbf{E}_k\mathbf{B}^{-T}\mathbf{z} \right| \\
&= \max_{\substack{\|\mathbf{w}\|_2 = 1 \\ \|\mathbf{z}\|_2 = 1 \\ \|\boldsymbol{\gamma}\|_2 = 1}} \left| \left(\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{\Lambda}}^{\perp}} \right) \cdot_1 \mathbf{w}^T \cdot_2 \mathbf{z}^T \cdot_3 \boldsymbol{\gamma}^T \right| \\
&= \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{\Lambda}}^{\perp}} \right\|_{\mathrm{sp}}.
\end{aligned}$$

In the $K = 2$ case, one may then follow the proof of [28, Chapter VI, Theorem 2.7] to obtain

$$\min_k \chi(\mathbf{\Lambda}_k, \tilde{\mathbf{\Lambda}}) \leq \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{\Lambda}}^{\perp}} \right\|_{\mathrm{sp}}.$$

To obtain the result for general K, we show that there exists some JGE value $\mathbf{\Lambda}$ of $\mathcal{T}$ such that

$$\frac{\chi(\mathbf{\Lambda}, \tilde{\mathbf{\Lambda}})}{\sqrt{R}} \leq \max_{\substack{\boldsymbol{\gamma} \in \tilde{\mathbf{\Lambda}}^{\perp} \\ \|\boldsymbol{\gamma}\|_2 = 1}} \left\| \sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C})\mathbf{v} \right\|_2, \tag{3.6}$$

from which the first inequality in equation (3.2) will follow. For $r = 1, \ldots, R$, let $\mathbf{c}_r = (c_{1r}, \ldots, c_{Kr})$ be the rth column of $\mathbf{C}$. Then c_{kr} is the rth diagonal entry of $D_k(\mathbf{C})$ for each k, and, as shown by Theorem 1.1, the set $\{\mathrm{span}(\mathbf{c}_r)\}_{r=1}^{R}$ is the spectrum of $\mathcal{T}$. By way of equation (3.1), to prove that equation (3.6) holds it is sufficient to show that there is an index r_* such that

$$\max_{\substack{\boldsymbol{\gamma} \in \tilde{\mathbf{\Lambda}}^{\perp} \\ \|\boldsymbol{\gamma}\|_2 = 1}} \frac{|\langle \mathbf{c}_{r_*}, \boldsymbol{\gamma} \rangle|}{\sqrt{R}} \leq \max_{\substack{\boldsymbol{\gamma} \in \tilde{\mathbf{\Lambda}}^{\perp} \\ \|\boldsymbol{\gamma}\|_2 = 1}} \left\| \sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C})\mathbf{v} \right\|_2. \tag{3.7}$$

To this end let $\mathbf{e}_r$ be the rth standard basis vector in $\mathbb{K}^R$ and observe

$$\left\| \sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C}) \mathbf{e}_r \right\|_2 = \left\| \sum_{k=1}^{K} \bar{\gamma}_k c_{kr} \mathbf{e}_r \right\|_2 = \left| \sum_{k=1}^{K} \bar{\gamma}_k c_{kr} \right| \|\mathbf{e}_r\|_2 = |\langle \mathbf{c}_r, \boldsymbol{\gamma} \rangle|. \tag{3.8}$$

Decompose $\mathbf{v} = \sum_{r=1}^{R} v_r \mathbf{e}_r$ and note that $\|\mathbf{v}\|_2 = 1$ implies $\sum_{r=1}^{R} |v_r|^2 = 1$. Then for each fixed $\boldsymbol{\gamma} \in \tilde{\boldsymbol{\Lambda}}^{\perp}$ with $\|\boldsymbol{\gamma}\|_2 = 1$ we have

$$\begin{aligned} \| \textstyle\sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C}) \mathbf{v} \|_2 &= \left\| \textstyle\sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C}) \left(\sum_{r=1}^{R} v_r \mathbf{e}_r \right) \right\|_2 = \left\| \textstyle\sum_{r=1}^{R} v_r \left(\sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C}) \mathbf{e}_r \right) \right\|_2 \\ &= \| \textstyle\sum_{r=1}^{R} v_r \langle \mathbf{c}_r, \boldsymbol{\gamma} \rangle \mathbf{e}_r \|_2 = \sqrt{\textstyle\sum_{r=1}^{R} |v_r|^2 |\langle \mathbf{c}_r, \boldsymbol{\gamma} \rangle|^2}. \end{aligned}$$

Now let r_* be an index such that $|v_{r_*}| \geq |v_r|$ for all $r = 1, \ldots, R$. Since $\|v\|_2 = 1$, we must have $|v_{r_*}| \geq \frac{1}{\sqrt{R}}$. Thus we have

$$\left\| \sum_{k=1}^{K} \bar{\gamma}_k D_k(\mathbf{C}) \mathbf{v} \right\|_2 = \sqrt{\sum_{r=1}^{R} |v_r|^2 |\langle \mathbf{c}_r, \boldsymbol{\gamma} \rangle|^2} \geq |v_{r_*}| |\langle \mathbf{c}_{r_*}, \boldsymbol{\gamma} \rangle| \geq \frac{|\langle \mathbf{c}_{r_*}, \boldsymbol{\gamma} \rangle|}{\sqrt{R}}.$$

Noting that the above inequality holds for all $\boldsymbol{\gamma} \in \tilde{\boldsymbol{\Lambda}}^{\perp}$ with $\|\boldsymbol{\gamma}\|_2 = 1$ gives equation (3.7) from which equation (3.6) follows. The desired inequality then follows by combining inequalities (3.5) and (3.6).

To prove the second inequality in equation (3.2) observe that

$$\left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\boldsymbol{\Lambda}}^{\perp}} \right\|_{\text{sp}} \leq \left\| \left(\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right)_{(1)} \right\|_2 = \left\| \mathbf{A}^{-1} \mathcal{E}_{(1)} (\mathbf{I}_K \otimes \mathbf{B}^{-T}) \right\|_2 \leq \|\mathbf{A}^{-1}\|_2 \|\mathbf{B}^{-1}\|_2 \|\mathcal{E}_{(1)}\|_2$$

from which the result follows.

It remains to show that the coefficient $\sqrt{R}$ is not necessary in the special case where $\tilde{\mathcal{T}}$ is regular with $\mathbb{K}$-rank R and CPD $[\![\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]\!]$ where $\tilde{\mathbf{A}} = \mathbf{A}$ or $\tilde{\mathbf{B}} = \mathbf{B}$. It is straightforward to check that the spectral variation $\text{sv}[\mathcal{T}, \tilde{\mathcal{T}}]$ is invariant under a simultaneous permutation of the first and second mode of $\mathcal{T}$ and $\tilde{\mathcal{T}}$, so it is sufficient to consider the case $\tilde{\mathbf{B}} = \mathbf{B}$.

In this case, Theorem 1.1 shows that the JGE vectors of both $\mathcal{T}$ and $\tilde{\mathcal{T}}$ are given by the columns of $\mathbf{B}^{-T}$, hence the vector $\mathbf{v}$ appearing the proof above is simply equal to $\mathbf{e}_r$ for some index $r \in 1, \ldots, R$. Using equation (3.5) together with equation (3.8) then shows that we have

$$\chi(\boldsymbol{\Lambda}_r, \tilde{\boldsymbol{\Lambda}}) = \max_{\substack{\boldsymbol{\gamma} \in \tilde{\boldsymbol{\Lambda}}^{\perp} \\ \|\boldsymbol{\gamma}\|_2 = 1}} |\langle \mathbf{c}_r, \boldsymbol{\gamma} \rangle| \leq \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\boldsymbol{\Lambda}}^{\perp}} \right\|_{\text{sp}} \leq \frac{\|\mathcal{E}_{(1)}\|_2}{\sigma_{\min}(\mathbf{A}) \sigma_{\min}(\mathbf{B})},$$

as claimed. ∎

3.4.1. Scaling of A and B. The quality of the bound obtained in Theorem 3.3 depends on the scaling of the factor matrices $\mathbf{A}$ and $\mathbf{B}$ in the CPD $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$. More precisely, for any diagonal matrix $\mathbf{D}$, one has that $\mathcal{T} = [\![\mathbf{AD}, \mathbf{BD}^{-1}, \mathbf{C}]\!]$ is a CPD of $\mathcal{T}$, and the quantity of $\sigma_{\min}(\mathbf{AD}) \sigma_{\min}(\mathbf{BD}^{-1})$ is dependent on the choice of $\mathbf{D}$. To obtain an optimal bound, one wants to minimize the right hand side of equation (3.3), hence one wants to maximize $\sigma_{\min}(\mathbf{AD}) \sigma_{\min}(\mathbf{BD}^{-1})$ over invertible diagonal matrices $\mathbf{D}$.

One strategy to compute the optimal $\mathbf{D}$ is to use an ALS approach in which one sets all but one of the diagonal entries of $\mathbf{D}$ to be equal to 1, then maximizes over the remaining diagonal entry of $\mathbf{D}$. By repeating this approach while varying which entry of $\mathbf{D}$ is not fixed, one is able to obtain a (local) optimum for equation (3.3).

As an alternative strategy, one may simply normalize $\mathbf{A}$ and $\mathbf{B}$ so that the rth column of $\mathbf{A}$ and $\mathbf{B}$ have the same norm for each $r = 1, \ldots, R$. Although this method cannot be expected to be optimal, it appears to perform well for randomly generated $\mathbf{A}$ and $\mathbf{B}$, and the associated computational cost is minimal.

3.4.2. On the coefficient $\sqrt{R}$. Knowing that there are various special cases in which the coefficient $\sqrt{R}$ may be dropped from equation (3.3), it is natural to wonder if there are examples where this coefficient is necessary, or if it is simply a relic of the proof. The following example shows that some constant is indeed necessary. That is, equation (3.3) does not hold for arbitrary pairs of regular $\mathbb{K}$-rank R tensors if the coefficient $\sqrt{R}$ is removed.

Example 3.2. Let $\mathcal{T}$ have CPD $[\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ where

$$\mathbf{A} = \mathbf{B} = \mathbf{I}_3 \qquad \text{and} \qquad \mathbf{C} = \begin{pmatrix} 1 & 0 & \frac{\sqrt{2}}{2} \\ 0 & 1 & \frac{\sqrt{2}}{2} \\ 0 & 0 & 0 \end{pmatrix}$$

and let $\tilde{\mathcal{T}}$ have CPD $[\![\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]\!]$ where

$$\tilde{\mathbf{A}} = \tilde{\mathbf{B}} = \begin{pmatrix} \frac{1}{10\sqrt{3}} & \frac{1}{\sqrt{3}} & 0 \\ \frac{1}{10\sqrt{3}} & -\frac{1}{2\sqrt{3}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{10\sqrt{3}} & -\frac{1}{2\sqrt{3}} & \frac{1}{\sqrt{2}} \end{pmatrix} \qquad \text{and} \qquad \tilde{\mathbf{C}} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$$

Then $\mathcal{T}$ and $\tilde{\mathcal{T}}$ are both regular $\mathbb{R}$-rank 3 tensors and the spectral variation $\mathrm{sv}[\mathcal{T}, \tilde{\mathcal{T}}]$ is equal to 1, as seen from the fact that the first column of $\tilde{\mathbf{C}}$ is orthogonal to all three columns of $\mathbf{C}$. However, setting $\mathcal{E} = \tilde{\mathcal{T}} - \mathcal{T}$ one has

$$\left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\mathrm{sp}} < 1$$

illustrating that equation (3.3) does not hold for this pair[6].

3.5. Matching distance bounds and the $K = 2$ case. While the spectral variation bound from Theorem 3.3 is interesting, as illustrated by Example 3.1, a spectral variation bound cannot be used to determine that a perturbation of a given tensor has distinct JGE values. To accomplish this we need to convert our spectral variation bound to a matching distance bound.

As briefly mentioned before, a spectral variation bound can often be converted to a matching distance bound using the following idea: Supposing temporarily that $\tilde{\mathcal{T}}$ is regular and has border $\mathbb{K}$-rank R, the bound given in Theorem 3.3 gives a union of regions around the JGE values of $\mathcal{T}$ in which the JGE values of $\tilde{\mathcal{T}}$ must lie. Additionally, it is not difficult to show that JGE values are continuous in a perturbation along the set of regular border rank R tensors. Since these regions grow

[6]Computation of the spectral norm is NP-hard [14]; however, the mode three unfolding $\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}$ has spectral norm ≈ 0.983 and this norm is as an upper bound for the tensor spectral norm. Using Tensorlab [34] suggests $\left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\mathrm{sp}} \approx 0.882$.

linearly in the distance between $\mathcal{T}$ and $\tilde{\mathcal{T}}$, and since JGE values are continuous, if $\tilde{\mathcal{T}}$ is regular and has border $\mathbb{K}$-rank R and if there is a path of regular border rank R tensors from $\mathcal{T}$ to $\tilde{\mathcal{T}}$ lying in the ball of radius $\|(\mathcal{T}-\tilde{\mathcal{T}})_{(1)}\|_2$ and if the regions around around the JGE values of $\mathcal{T}$ do not intersect, then each region must contain exactly one JGE value of $\tilde{\mathcal{T}}$, hence in this setting, the bound given by 3.3 is in fact a matching distance bound and may be used to guarantee the existence of distinct JGE values for $\tilde{\mathcal{T}}$.

The primary issue with such an argument is that, for general K, it is not known if there is a path of regular border rank R tensors connecting $\mathcal{T}$ and $\tilde{\mathcal{T}}$ which lies in the ball of radius $\|(\mathcal{T}-\tilde{\mathcal{T}})_{(1)}\|_2$ centered at $\mathcal{T}$. However, in the case that $K=2$, the condition needed to guarantee that the tensor regions around the JGE values of $\mathcal{T}$ do not intersect in fact guarantees the existence of such a path.

As discussed in [28, Chapter VI], if $\mathcal{T}$ is a regular tensor of size $R\times R\times 2$, then the (joint) generalized eigenvalue problem for $\mathcal{T}$ is equivalent to a classical eigenvalue problem. E.g. if $\mathbf{T}_1$ is invertible, then the JGE eigenvectors of $\mathcal{T}$ are exactly equal to the classical eigenvectors of $\mathbf{T}_1^{-1}\mathbf{T}_2$. Using this equivalence, we for free obtain basic facts about multiplicities of JGE values for regular $R\times R\times 2$ tensors. Namely, a regular $R\times R\times 2$ tensor always has R JGE values in $\mathbb{C}^2$ counting algebraic multiplicity, and each JGE value of such a tensor must have geometric multiplicity at least one. We will make frequent use of these facts in the upcoming results.

PROPOSITION 3.4. *Let $\mathcal{T}\in\mathbb{K}^{R\times R\times 2}$ be a tensor of $\mathbb{K}$-rank R with CPD $\mathcal{T}=[\![\mathbf{A},\mathbf{B},\mathbf{C}]\!]$, where $\mathbf{C}$ has columns of unit norm and let $\{\mathbf{\Lambda}_r\}_{r=1}^R$ be the set of JGE values of $\mathcal{T}$. Let $\mathcal{E}\in\mathbb{K}^{R\times R\times 2}$ be any tensor which satisfies*

$$\|\mathcal{E}_{(1)}\|_2 \leq \frac{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})\min_{i\neq j}\chi(\mathbf{\Lambda}_i,\mathbf{\Lambda}_j)}{2}.$$

Then $\mathcal{T}+\mathcal{E}$ is regular and has border $\mathbb{C}$-rank equal to R.

Proof. To simplify notation, set

$$\delta=\min_{i\neq j}\frac{\chi(\mathbf{\Lambda}_i,\mathbf{\Lambda}_j)}{2}.$$

By assumption the columns of $\mathbf{C}$ all have unit norm, therefore, by reindexing columns if necessary, Theorem 1.1 shows that for each $r=1,\dots,R$ the column $\mathbf{c}_r$ satisfies $\mathbf{\Lambda}_r=\operatorname{span}(\mathbf{c}_r)$. To show $\mathcal{T}+\mathcal{E}$ is regular, we first show that there is a unit vector $\omega\in\mathbb{K}^2$ such that $|\langle\mathbf{c}_r,\omega\rangle|\geq\delta$ for all r. Using this we will show that the matrix $\bar{\omega}_1(\mathbf{T}_1+\mathbf{E}_1)+\bar{\omega}_2(\mathbf{T}_2+\mathbf{E}_2)$ is invertible, hence $\mathcal{T}+\mathcal{E}$ is regular.

To this end let $\Psi\subset\mathbb{K}^2$ be a one-dimensional subspace of $\mathbb{K}^2$ such that $\chi(\mathbf{\Lambda}_1,\Psi)=\delta$. Then using the triangle inequality we must have

$$\chi(\mathbf{\Lambda}_r,\Psi)\geq\delta \qquad \text{for all } r=1,\dots,R, \tag{3.9}$$

as otherwise would imply that there is an index ℓ such that

$$\chi(\mathbf{\Lambda}_\ell,\Psi)<\delta \qquad \text{hence} \qquad \chi(\mathbf{\Lambda}_1,\mathbf{\Lambda}_\ell)\leq\chi(\mathbf{\Lambda}_1,\Psi)+\chi(\mathbf{\Lambda}_\ell,\Psi)<2\delta=\min_{i\neq j}\chi(\mathbf{\Lambda}_i,\mathbf{\Lambda}_j),$$

which is a contradiction.

Now set $\Omega=\Psi^\perp\subset\mathbb{K}^2$ and let $\omega\in\mathbb{K}^2$ be a unit vector such that $\Omega=\operatorname{span}(\omega)$. Using equation (3.9), for each r we have

$$\delta\leq\chi(\mathbf{\Lambda}_\ell,\Psi)=\sin(\angle(\mathbf{\Lambda}_r,\Psi))=\sin(\pi/2-\angle(\mathbf{\Lambda}_r,\Omega))=\cos(\angle(\mathbf{\Lambda}_r,\Omega))=|\langle\mathbf{c}_r,\omega\rangle|$$

where $\angle(\cdot,\cdot)$ denotes the angle between two one-dimensional subspaces of $\mathbb{K}^2$.

Letting $D_k(\mathbf{C}) \in \mathbb{K}^{R\times R}$ be the diagonal matrix whose diagonal entries are given by the kth row of $\mathbf{C}$ as usual, we have that $\bar{\omega}_1\mathcal{D}_1(\mathbf{C}) + \bar{\omega}_2 D_2(\mathbf{C})$ is a diagonal matrix with rth diagonal entry equal to $\langle \mathbf{c}_r, \omega\rangle$. If follows that

$$\sigma_{\min}(\bar{\omega}_1\mathcal{D}_1(\mathbf{C}) + \bar{\omega}_2 D_2(\mathbf{C})) = \min(\{|\langle \mathbf{c}_r, \omega\rangle|\}_{r=1}^R) \geq \delta. \tag{3.10}$$

Finally observe that

$$\begin{aligned}
\sigma_{\min}\big(\bar{\omega}_1(\mathbf{T}_1 + \mathbf{E}_1) + \bar{\omega}_2(\mathbf{T}_2 + \mathbf{E}_2)\big) &\geq \sigma_{\min}(\bar{\omega}_1\mathbf{T}_1 + \bar{\omega}_2\mathbf{T}_2) - \|\bar{\omega}_1\mathbf{E}_1 + \bar{\omega}_2\mathbf{E}_2\|_2 \\
&\geq \sigma_{\min}\Big(\mathbf{A}\Big(\bar{\omega}_1 D_1\big(\mathbf{C}\big) + \bar{\omega}_2 D_2(\mathbf{C})\Big)\mathbf{B}^T\Big) - \|(\mathbf{E}_1\ \mathbf{E}_2)\|_2 \left\| \begin{pmatrix} \bar{\omega}_1 \\ \bar{\omega}_2 \end{pmatrix} \right\|_2 \\
&\geq \sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})\sigma_{\min}\big(\bar{\omega}_1 D_1\big(\mathbf{C}\big) + \bar{\omega}_2 D_2(\mathbf{C})\big) - \|\mathcal{E}_{(1)}\|_2
\end{aligned}$$

Using inequality (3.10) together with the assumption that $\|\mathcal{E}_{(1)}\|_2 < \sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})\delta$ shows that $\mathcal{T} + \mathcal{E}$ is regular.

The fact that $\mathcal{T}+\mathcal{E}$ has border $\mathbb{C}$-rank equal to R is an immediate consequence of the regularity of $\mathcal{T}+\mathcal{E}$ together with the well know fact that tensors of size $R\times R\times 2$ have border rank less than or equal to R, e.g. see [19, Theorem 5.5.1.1]. ∎

We now present our matching distance bound for the spectra of two tensors in the case $K = 2$. Before proceeding we note that this result is heavily inspired by [28, Chapter VI, Theorem 2.7] and the proof may be unsurprising to experts in perturbation theory for classical generalized eigenvalues; however, the authors are unaware of an equivalent result in the literature.

THEOREM 3.5. *Let $\mathcal{T}$ and $\tilde{\mathcal{T}}$ be $R\times R\times 2$ tensors with entries in $\mathbb{K}$. Assume that $\mathcal{T}$ has $\mathbb{K}$-rank R with CPD $[\![\mathbf{A},\mathbf{B},\mathbf{C}]\!]$ where the columns of $\mathbf{C}$ are normalized to have unit norm. Let $\{\mathbf{\Lambda}_k\}_{k=1}^K$ be the spectrum of $\mathcal{T}$. If*

$$\|(\mathcal{T} - \tilde{\mathcal{T}})_{(1)}\|_2 < \frac{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})\min_{i\neq j}\chi(\mathbf{\Lambda}_i, \mathbf{\Lambda}_j)}{2},$$

then $\tilde{\mathcal{T}}$ is a regular tensor of $\mathbb{K}$-rank R and

$$\mathrm{md}[\mathcal{T}, \tilde{\mathcal{T}}] < \frac{\|(\mathcal{T} - \tilde{\mathcal{T}})_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}. \tag{3.11}$$

In addition, all JGE values of $\tilde{\mathcal{T}}$ have algebraic and geometric multiplicity equal to one.

Proof. Before proceeding we note that the proof of the $\mathbb{K} = \mathbb{R}$ case makes use of the result over $\mathbb{C}$, so we first consider the complex case. Now set

$$\epsilon = \frac{\|(\mathcal{T} - \tilde{\mathcal{T}})_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})} < \frac{\min_{i\neq j}\chi(\mathbf{\Lambda}_i, \mathbf{\Lambda}_j)}{2} \tag{3.12}$$

and for each $r = 1, \ldots, R$ let Δ^r be the collection of one-dimensional subspaces in $\mathbb{C}^2$ defined by

$$\Delta^r = \{\Psi \subset \mathbb{C}^2 |\ \dim(\Psi) = 1 \text{ and } \chi(\mathbf{\Lambda}_r, \Psi) \leq \epsilon\}.$$

That is, Δ^r is the ball of chordal radius ϵ centered at $\mathbf{\Lambda}_r$. Then our choice of ϵ together with the triangle inequality gives

$$\Delta^i \cap \Delta^j = \emptyset \qquad \text{for all } i \neq j.$$

To prove inequality (3.11), it is sufficient show that each Δ^r contains exactly one JGE value of $\tilde{\mathcal{T}}$.

To this end let $\mathcal{E} = \tilde{\mathcal{T}} - \mathcal{T}$ and for $\alpha \in [0,1]$ set $\mathcal{T}^\alpha = \mathcal{T} + \alpha\mathcal{E}$ so that $\mathcal{T} = \mathcal{T}^0$ and $\tilde{\mathcal{T}} = \mathcal{T}^1$. Note that Proposition 3.4 shows that $\mathcal{T}^\alpha$ is regular and has border $\mathbb{C}$-rank R for each $\alpha \in [0,1]$. It follows from Theorem 1.1 that each $\mathcal{T}^\alpha$ has R JGE values in $\mathbb{C}^2$ counting algebraic multiplicity.

Now define

$$\Delta^r(\alpha) = \{\Psi \subset \mathbb{C}^2 |\ \dim(\Psi) = 1 \text{ and } \chi(\Lambda_r, \Psi) \leq \alpha\epsilon\}.$$

Then $\Delta^r(\alpha) \subset \Delta^r$ for each α and r so $\Delta^i(\alpha) \cap \Delta^j(\alpha) = \emptyset$ whenever $i \neq j$. Furthermore, using Theorem 3.3 we have

$$\mathrm{sv}[\mathcal{T}, \mathcal{T}^\alpha] \leq \frac{\|\alpha\mathcal{E}_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})} = \alpha\epsilon.$$

It follows that the spectrum of $\mathcal{T}^\alpha$ is contained in $\cup_{r=1}^R \Delta^r(\alpha)$ for each $\alpha \in [0,1]$.

In particular $\Delta^r(0)$ contains exactly one JGE value of $\mathcal{T}^0 = \mathcal{T}$ for each $r = 1, \ldots, R$. Since $\mathcal{T}^\alpha$ is regular for all $\alpha \in [0,1]$, [28, Chapter VI, Theorem 2.1] shows that the JGE values of $\mathcal{T}^\alpha$ are continuous in α for $\alpha \in [0,1]$. Therefore, because the $\Delta^r(\alpha)$ remain disjoint as α increases from 0 to 1, the continuity of the JGE values of $\mathcal{T}^\alpha$ in α implies that each $\Delta^r(\alpha)$ must contain exactly one JGE value of $\mathcal{T}^\alpha$ for each $\alpha \in [0,1]$ and for each $r = 1, \ldots, R$. In particular, $\Delta^r(1) = \Delta^r$ contains exactly one eigenvalue of $\mathcal{T}^1 = \tilde{\mathcal{T}}$ for each r. By definition, each Δ^r is a ball of chordal radius ϵ centered $\mathbf{\Lambda}_r$, so we conclude

$$\mathrm{md}[\mathcal{T}, \tilde{\mathcal{T}}] \leq \frac{\|(\mathcal{T} - \tilde{\mathcal{T}})_{(1)}\|_2}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})},$$

and that all the JGE value of $\tilde{\mathcal{T}}$ have algebraic multiplicity one, as claimed.

It remains to show that $\tilde{\mathcal{T}}$ has $\mathbb{C}$-rank R. As previously discussed, all JGE values of a regular $R \times R \times 2$ tensor must have geometric multiplicity at least 1. We have shown that $\tilde{\mathcal{T}}$ is regular and has distinct JGE values, so using Lemma 2.2 shows that $\tilde{\mathcal{T}}$ has a basis of JGE values in $\mathbb{C}$. Theorem 1.1 then shows that $\tilde{\mathcal{T}}$ has $\mathbb{C}$-rank R.

Now assume that $\mathbb{K} = \mathbb{R}$ and that $\mathcal{T}$ and $\tilde{\mathcal{T}}$ have real entries. Applying the result over the complex field shows that $\tilde{\mathcal{T}}$ must have distinct JGE values over $\mathbb{C}$ so it is sufficient to show that the JGE values of $\tilde{\mathcal{T}}$, hence the JGE vectors of $\tilde{\mathcal{T}}$, must all be real. Here, when we say that $\tilde{\mathcal{T}}$ has real JGE values, we mean that the JGE values of $\mathcal{T}$ can be expressed as the spans of real valued vectors. Equivalently, the classical generalized eigenvalues of the pencil $(\tilde{\mathbf{T}}_1, \tilde{\mathbf{T}}_2)$ are real valued.

Suppose towards a contradiction that $\tilde{\mathcal{T}}$ has a JGE value which is not real and set

$$\alpha_* = \inf_{\alpha \in [0,1]} \mathcal{T} + \alpha\mathcal{E} \text{ has a JGE value which is not real.}$$

Using the continuity for JGE values in α, it is straightforward to show that $\mathcal{T} + \alpha_0\mathcal{E}$ has a JGE value with algebraic multiplicity at least two. However, $\mathcal{T} + \alpha_0\mathcal{E}$ satisfies the assumptions of Theorem 3.5, so this is a contradiction as we have already shown that all JGE values of such a tensor must have algebraic multiplicity equal to one. We conclude that the JGE values and JGE vectors of $\tilde{\mathcal{T}}$ are all real. It follows from Theorem 1.1 that $\tilde{\mathcal{T}}$ has $\mathbb{R}$-rank R. ∎

Remark 3.6. In the language of the classical generalized eigenvalue problem and matrix pencils, Theorem 3.5 gives a bound which guarantees that the matrix pencil $(\tilde{\mathbf{T}}_1, \tilde{\mathbf{T}}_2)$ is diagonalizable over $\mathbb{K}$.

4. Deterministic bounds for existence of best rank R decompositions. We now turn to proving the deterministic bounds for existence of best $\mathbb{K}$-rank R approximations presented in the introduction of the article. While we do not have a matching distance bound in the case of general K, a key observation is that if the tensor $\mathcal{T}$ is defective in the sense of JGE values, then every subpencil of $\mathcal{T}$ must have a JGE value of algebraic multiplicity greater than one.

PROPOSITION 4.1. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a regular tensor. Also let $\mathbf{U} \in \mathbb{K}^{K\times K}$ be any invertible matrix and set $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U}$. If $\mathcal{T}$ has a JGE value of algebraic multiplicity m, then for any $i \neq j$ the subpencil*

$$(\mathbf{S}_i, \mathbf{S}_j)$$

must have a JGE value of algebraic multiplicity at least equal to m.

Before giving the proof we give a lemma which relates the joint generalized eigenvalues and eigenvectors of $\mathcal{T}$ and $\mathcal{T} \cdot_3 \mathbf{U}$.

LEMMA 4.2. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a regular tensor. Also let $\mathbf{U} \in \mathbb{K}^{K\times K}$ be any invertible matrix and set $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U}$. Then $\mathbf{S}$ is regular and $(\mathrm{span}(\boldsymbol{\lambda}), \mathbf{x})$ is a JGE pair of $\mathcal{T}$ if and only if $(\mathrm{span}(\mathbf{U}\boldsymbol{\lambda}), \mathbf{x})$ is a JGE pair of $\mathcal{S}$. Furthermore one has $\mathrm{am}(\boldsymbol{\lambda}) = \mathrm{am}(\mathbf{U}\boldsymbol{\lambda})$ and $\mathrm{gm}(\boldsymbol{\lambda}) = \mathrm{gm}(\mathbf{U}\boldsymbol{\lambda})$ where $\mathrm{span}(\boldsymbol{\lambda})$ and $\mathrm{span}(\mathbf{U}\boldsymbol{\lambda})$ are treated as JGE values of $\mathcal{T}$ and $\mathcal{S}$, respectively.*

Proof. Assume $\mathrm{span}(\boldsymbol{\lambda})$ is a JGE value of $\mathcal{T}$ with algebraic multiplicity m. View $\boldsymbol{\gamma} = (\boldsymbol{\gamma}_1, \ldots, \boldsymbol{\gamma}_K)$ as a vector in $\mathbb{K}^K$, and consider the change of variables $\boldsymbol{\gamma} \to \mathbf{U}^T\boldsymbol{\gamma}$. It is a straightforward to verify that

$$p_{\mathcal{T}\cdot_3\mathbf{U}}(\boldsymbol{\gamma}) = p_{\mathcal{T}}(\mathbf{U}^T\boldsymbol{\gamma}).$$

Since $\mathbf{U}$ is invertible, it follows that $p_{\mathcal{S}}(\boldsymbol{\gamma}) \not\equiv 0$, hence $\mathcal{S}$ is regular. Now set $\mathbf{U}^T\boldsymbol{\gamma} = \eta$. By assumption $p_{\mathcal{T}}(\eta)$ factors as

$$\langle \boldsymbol{\lambda}, \bar{\eta}\rangle^m g(\eta) = \langle \mathbf{U}\boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}}\rangle^m g(\mathbf{U}^T\boldsymbol{\gamma}),$$

so $\mathrm{span}(\mathbf{U}\boldsymbol{\lambda})$ is a JGE value of $\mathcal{T} \cdot_3 \mathbf{U}$ with algebraic multiplicity m.

The proof for geometric multiplicity is straightforward. ∎

We now give the proof of Proposition 4.1.

Proof. By assumption $\mathcal{T}$ is regular and has a JGE value of algebraic multiplicity m, so Lemma 4.2 shows that $\mathcal{S}$ is regular and has a JGE value with algebraic multiplicity m. Say this JGE value is equal to $\mathrm{span}(\boldsymbol{\lambda})$. It follows that the characteristic polynomial of $\mathcal{S}$ factors as

$$p_{\mathcal{S}}(\boldsymbol{\gamma}) = \langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}}\rangle^m g(\boldsymbol{\gamma})$$

where $g(\boldsymbol{\gamma})$ is a nonzero polynomial.

To complete the proof, assume $i < j$ and observe that the characteristic polynomial of the subpencil $(\mathbf{S}_i, \mathbf{S}_j)$ is given by

$$\begin{aligned} p_{(\mathbf{S}_i,\mathbf{S}_j)}((\gamma_i, \gamma_j)) &= p_{\mathcal{S}}(0, \ldots, 0, \gamma_i, 0 \ldots, 0, \gamma_j, 0 \ldots, 0) \\ &= (\bar{\gamma}_i\lambda_i + \bar{\gamma}_j\lambda_j)^m g(0, \ldots, 0, \gamma_i, 0 \ldots, 0, \gamma_j, 0 \ldots, 0) \end{aligned}$$

hence $\mathrm{span}((\lambda_i, \lambda_j))$ is a generalized eigenvalue of the subpencil $(\mathbf{S}_i, \mathbf{S}_j)$ with multiplicity at least m. ∎

4.1. Proof of Theorem 1.3. We are now in position to prove our multiple pencil based bound for the existence of a best $\mathbb{K}$-rank R approximation.

Proof. Let $\tilde{\mathcal{T}} \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ be any tensor such that

$$\|\mathcal{T} - \tilde{\mathcal{T}}\|_F < \epsilon/2$$

and let $\hat{\mathcal{T}}$ be a best border $\mathbb{K}$-rank R approximation of $\tilde{\mathcal{T}}$ in the Frobenius norm. Since $\mathcal{T}$ has rank R by assumption, we must have

$$\|\hat{\mathcal{T}} - \tilde{\mathcal{T}}\|_F \leq \|\mathcal{T} - \tilde{\mathcal{T}}\|_F \qquad \text{hence} \qquad \|\hat{\mathcal{T}} - \mathcal{T}\|_F < \epsilon.$$

Since $\hat{\mathcal{T}}$ has border $\mathbb{K}$-rank R, the multilinear rank of $\hat{\mathcal{T}}$ must be $(R, R, \hat{K})$ for some $\hat{K} \leq R$. It is easy to check that a tensor has $\mathbb{K}$-rank R if and only if an orthogonal compression to its multilinear rank has $\mathbb{K}$-rank R, so using Theorem 3.1, we may WLOG assume that $\mathcal{T}$ and $\hat{\mathcal{T}}$ both have size $R \times R \times \max\{K, \hat{K}\}$. There is no meaningful difference in the $K \geq \hat{K}$ and $K < \hat{K}$ cases, so for ease of notation we assume $K \geq \hat{K}$.

Now let $\mathbf{U} \in \mathbb{K}^{K \times K}$ be any unitary and set $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U}$ and $\hat{\mathcal{S}} = \hat{\mathcal{T}} \cdot_3 \mathbf{U}$. Since modal unitary multiplications preserve the Frobenius norm of a tensor we have $\|\mathcal{S} - \hat{\mathcal{S}}\|_F = \|\mathcal{T} - \hat{\mathcal{T}}\|_F < \epsilon$. Therefore, for some $i = 1, \ldots, \lfloor K/2 \rfloor$ we must have that

$$\|(\mathbf{S}_{2i-1} \;\; \mathbf{S}_{2i}) - (\hat{\mathbf{S}}_{2i-1} \;\; \hat{\mathbf{S}}_{2i})\|_2 \leq \|(\mathbf{S}_{2i-1}, \mathbf{S}_{2i}) - (\hat{\mathbf{S}}_{2i-1} \;\; \hat{\mathbf{S}}_{2i})\|_F < \epsilon_i,$$

as otherwise would imply

$$\|\mathcal{S} - \hat{\mathcal{S}}\|_F^2 \geq \sum_{k=1}^{\lfloor K/2 \rfloor} \|(\mathbf{S}_{2k-1} \;\; \mathbf{S}_{2k}) - (\hat{\mathbf{S}}_{2k-1} \;\; \hat{\mathbf{S}}_{2k})\|_F^2 \geq \sum_{k=1}^{\lfloor K/2 \rfloor} \epsilon_k^2 = \epsilon^2.$$

In particular, from our choice of ϵ_i the subpencil $(\hat{\mathbf{S}}_{2i-1}, \hat{\mathbf{S}}_{2i})$ must be regular and have R distinct generalized eigenvalues. Since $\hat{\mathcal{T}}$ has a regular subpencil, $\hat{\mathcal{T}}$ must be regular. Therefore, using Proposition 4.1 allows us to conclude that the tensor $\hat{\mathcal{T}}$ cannot have a JGE eigenvalue with algebraic multiplicity greater than 1. The result then follows from Proposition 1.2. ∎

One may notice that there is no guarantee that the tensor $\mathcal{T} \cdot_3 \mathbf{U}$ has a regular subpencil for a particular choice of $\mathbf{U}$, hence the bound given by Theorem 1.3 may be equal to 0 for a particular choice of $\mathbf{U}$. This does not cause serious difficulty: As a consequence of [9, Section 4.4 (ii)], if $\mathcal{T}$ is regular then for generic unitaries every subpencil of $\mathcal{T}$ is regular, hence generically the computed bound is nonzero.

This does point toward a greater issue, however. Different choices of unitaries lead to different values for the bound computed by Theorem 1.3. Of course, the result holds for each of these values, so one desires to use a unitary which gives the largest possible value of ϵ. Outside of a few special cases, is not known how to optimally pick the unitary $\mathbf{U}$. Determining how to optimally pick $\mathbf{U}$ is a direction for future research.

A subproblem of picking an optimal unitary is picking the optimal ordering of slices for a fixed unitary $\mathbf{U}$. E.g. if $\mathcal{T} \in \mathbb{K}^{4\times4\times4}$ and $\mathbf{U} \in K^{4\times4}$, then for $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U} \in \mathbb{K}^{4\times4\times4}$ one may arrive at different bounds by considering collections of subpencils of the form $\{(\mathbf{S}_1, \mathbf{S}_2), (\mathbf{S}_3, \mathbf{S}_4)\}$ or $\{(\mathbf{S}_1, \mathbf{S}_3), (\mathbf{S}_2, \mathbf{S}_4)\}$ or $\{(\mathbf{S}_1, \mathbf{S}_4), (\mathbf{S}_2, \mathbf{S}_3)\}$.

Let $\mathbf{E}_\mathbf{U}$ be the matrix with i, j entry equal to the bound given by Theorem 3.5 for the pencil $(\mathbf{S}_i, \mathbf{S}_j)$. Then choosing an optimal collection of subpencils (i.e. an optimal

ordering for the slices of $\mathcal{S}$) is equivalent to marking the entries of $\mathbf{E_U}$ so that there is at most one mark in each row and column of $\mathbf{E_U}$ and so that the norm of the vector of marked entries is maximal.

This problem is a variant of the longest path problem which is, in a sense, the "opposite" of the travelling salesman problem (where the total "distance" is minimized). The longest path problem is known to be NP-hard [32]; for our purposes and for small K, we may nevertheless consider trying all possible pencil collections/paths.

5. Numerical experiments. In this section, numerical experiments are used to examine the quality of the main perturbation bounds obtained in the paper, i.e. the spectral variation bound obtained in Theorem 3.3 and the existence bounds obtained in Theorem 1.3 and Theorem 1.4. In addition we illustrate the effect of a tensor Procrustes calculation as in Theorem 3.1 on the error between a rank R tensor and a rank R approximation of a perturbation of the tensor. Throughout the section, we define the signal to noise ratio (SNR) of a tensor $\mathcal{T}$ to a noise tensor $\mathcal{N}$ to be $20\log_{10}\left(\frac{\|\mathcal{T}\|_F}{\|\mathcal{N}\|_F}\right)$. The experiments discussed in this section all restrict to the real setting. That is, we restrict to considering real rank and real valued tensors.

In all of our experiments, rank R tensors of size $I \times I \times I$ are randomly generated by generating factor matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ of size $I \times R$ with entries independently sampled from a standard normal distribution. We then scale the factor matrices so that each rank R tensor has unit Frobenius norm. The entries of the noise tensors $\mathcal{N}$ are all independently sampled from a standard normal distribution and each $\mathcal{N}$ is scaled so that the resulting SNR is as specified.

5.1. Observed spectral variation vs. spectral variation bound. We first examine the spectral variation bound obtained from the spectral norm[7] of (a transformation of) the error tensor in Theorem 3.3 and compare it to the observed spectral variation occurring under a perturbation with a specified SNR. Three cases are treated: First, we consider the case of structured perturbations, i.e. perturbations only affecting the factor matrix $\mathbf{C}$. In this special case we may compare to the improved bound presented in equation (3.4). Second, we consider the challenging case of generic perturbations on a rank R tensor of size $R \times R \times R$. Finally we consider the situation most commonly occurring in applications of tensors with rank significantly smaller than dimension. That is, we examine spectral variation bounds for rank R tensors of size $I \times I \times I$ with $I \gg R$.

In each experiment we randomly generate rank R tensors $\mathcal{T}$ and perturbations $\mathcal{N}$ and present the mean observed spectral variation and either the mean computed spectral variation bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-2}\|$ or, when appropriate, the mean computed structured spectral variation bound $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-2}\|$.

5.1.1. Structured Perturbations. Figure 1 considers the structured perturbation case on tensors of rank R and size $R \times R \times R$ where $R = 4$ or $R = 10$. For various choices of SNR we randomly generate a rank R tensor $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ and an error matrix $\mathbf{E}$ with entries independently sampled from the standard normal distribution. The error matrix $\mathbf{E}$ is then scaled so that the SNR of $\mathbf{C}$ to $\mathbf{E}$ is as specified. A total of 20 trials are performed for each SNR and for each SNR we report the mean observed spectral variation between $\mathcal{T}$ and $\tilde{\mathcal{T}} := [\![\mathbf{A}, \mathbf{B}, \mathbf{C} + \mathbf{E}]\!]$ in addition to the

[7] As previously mentioned, computation of the spectral norm is NP-hard. Using Tensorlab [34], one may compute a (local) optimum for the best rank 1 approximation of the transformed error. The Frobenius norm of a best rank 1 approximation is equivalent to the spectral norm of the tensor [11].

mean computed value of $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$. Here $\mathcal{E} := \mathcal{T} - \hat{\mathcal{T}}$. We note that this experiment illustrates the first inequality in equation (3.4).

These figures illustrate that the error bound $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$ is highly predictive of the observed spectral variation when considering structured perturbations. We remind the reader that the error bound $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$ is guaranteed to be an upper bound in this special setting.

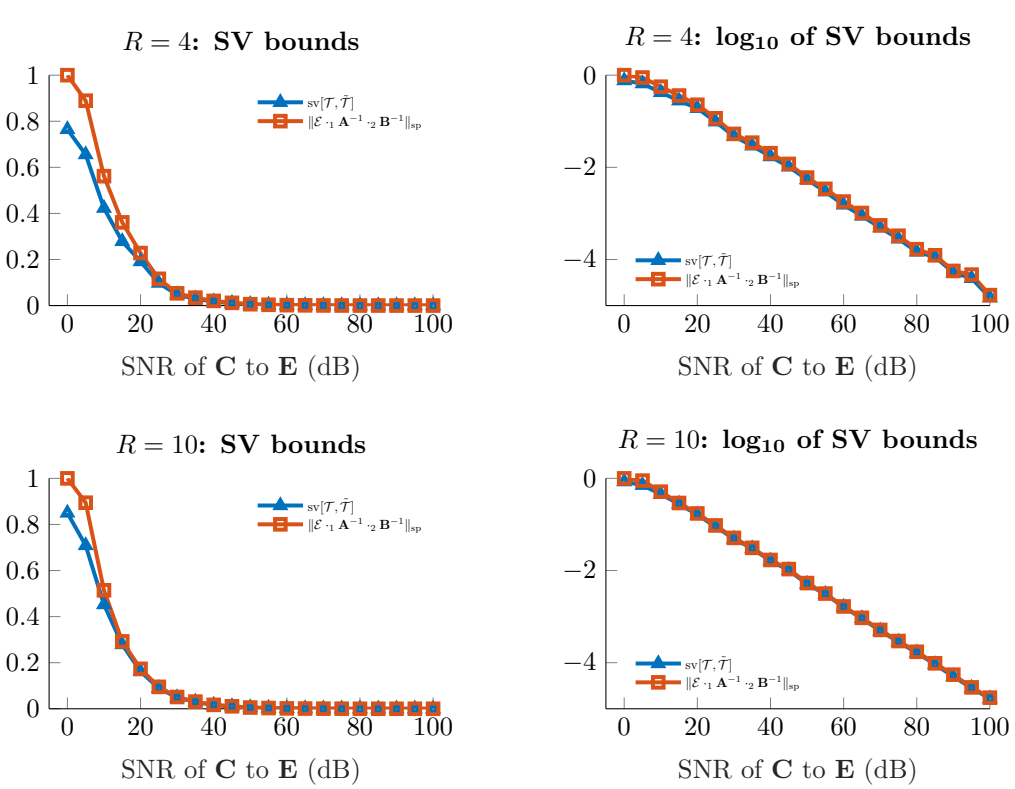


FIG. 1. *Observed spectral variation and computed bound between tensors of the form* $[\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ *and* $[\![\mathbf{A}, \mathbf{B}, \mathbf{C} + \mathbf{E}]\!]$ *with rank* $R = 4$ *or* 10. *The structured bound* $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$ *is observed to be highly predictive of the observed spectral variation.*

5.1.2. Generic Perturbations: Rank equal to dimensions. Figure 2 illustrates the case of generic perturbations on tensors of rank R and size $R \times R \times R$ where $R = 4$ or $R = 10$. In each trial we randomly generate a rank R tensor $\mathcal{T}$ and noise $\mathcal{N}$ at the specified SNR, then compute a rank R approximation $\hat{\mathcal{T}}$ of $\mathcal{T} + \mathcal{N}$ using Tensorlab [34]. A total of 50 trials are performed for each SNR, and for each SNR we report the mean observed spectral variation between $\mathcal{T}$ and $\hat{\mathcal{T}}$ as well as the mean computed value for $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$ where $\mathcal{E} := \hat{\mathcal{T}} - \mathcal{T}$. This experiment illustrates the first inequality in equation (3.3).

This case is particularly challenging since the perturbation is applied directly to the core of the tensor, which is not typically the case in applications. Furthermore, our bound serves as a worst case scenario bound which is unlikely to be achieved

under generic perturbations.

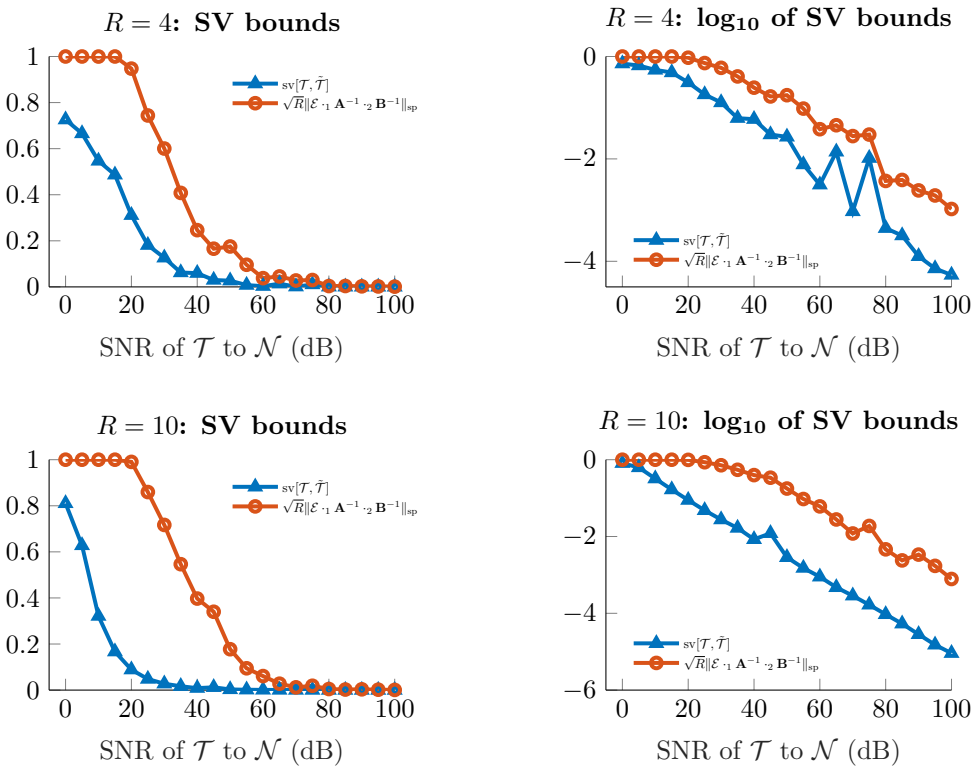


FIG. 2. *Observed spectral variation and computed bound for a rank R tensor $\mathcal{T}$ and a rank R approximation to $\mathcal{T}+\mathcal{N}$ where $R = 4$ or 10. The bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\mathrm{sp}}$ typically exceeds the observed spectral variation by about one order of magnitude when $R = 4$ and by about two orders of magnitude when $R = 10$. For perspective we note that $\log_{10}(\sqrt{4}) \approx 0.3$ while $\log_{10}(\sqrt{10}) = 0.5$.*

5.1.3. Generic Perturbations: Rank less than dimensions. Figures 3 and 4 consider the case which most commonly occurs in applications where the rank of the tensor of interest is significantly less than dimensions of the tensor. Here we work with tensors of rank R and dimensions $I \times I \times I$ where either $R = 4$ and $I = 10, 20,$ or 100 or $R = 10$ and $I = 20$ or 100.

To directly apply the spectral variation bound obtained from the spectral norm in Theorem 3.3 to a pair of rank R tensors, the factor matrices of at least one tensor should be invertible. In particular the tensors should have dimensions equal to their rank. Thus it is most appropriate to apply this bound to core tensors $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ of size $R \times R \times R$ satisfying $\|\mathcal{T}^c - \hat{\mathcal{T}}^c\|_F \leq \|\mathcal{T} - \hat{\mathcal{T}}\|_F$ as described in the tensor Procrustes problem of Theorem 3.1. Write $\mathcal{T}^c = [\![\mathbf{A}^c, \mathbf{B}^c, \mathbf{C}^c]\!]$. For each SNR we report the mean observed spectral variation between the core tensors $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ as well as the mean computed value for $\sqrt{R}\|\mathcal{E}^c \cdot_1 (\mathbf{A}^c)^{-1} \cdot_2 (\mathbf{B}^c)^{-1}\|_{\mathrm{sp}}$ where $\mathcal{E}^c := \hat{\mathcal{T}}^c - \mathcal{T}^c$. For each choice of I and R the mean is computed over 20 trials.

In this setting we find that the error bound $\sqrt{R}\|\mathcal{E}^c \cdot_1 (\mathbf{A}^c)^{-1} \cdot_2 (\mathbf{B}^c)^{-1}\|_{\mathrm{sp}}$ from

the first inequality in equation (3.3) is highly predictive of the observed spectral variation, with the error bound typically within a half order of magnitude of the observed spectral variation. Furthermore, we find that the observed spectral variation and spectral variation bound become remarkably small as I grows, even for low SNR. It is well known that computation of the core of an observed tensor can significantly increase the effective SNR. Our results further suggest that a significant portion of the error remaining after computing a rank R approximation of an observed tensor $\mathcal{T}+\mathcal{N}$ is in the form of rotational error (i.e. error accounted for by unitary rotations of the mode-j subspaces of the tensor) which disappears after a Tensor Procrustes computation.

While this rotational error is still meaningful from the perspective of quality of the estimated factors, rotational error cannot lead to nonexistence of a best rank R approximation and cannot negatively impact the condition number of computing a CPD (see [33] for details on this condition number), as these issues are invariant under unitary changes of basis. The observation that, in the setting of independent and identically distributed (i.i.d) noise, a significant portion error observed between the original tensor $\mathcal{T}$ and a rank R approximation to the observed tensor $\mathcal{T}+\mathcal{N}$ is merely rotational error further explains the success of CPD in many large scale settings. In particular, this observation illustrates that CPD computation in practice is likely to be more stable than one might naively expect. See Section 5.3 for experiments in this direction.

5.2. Bounds for existence of best rank R approximations. We now present the mean of the bound ϵ obtained using Theorem 1.3 for the radius of a neighborhood around a randomly generated rank R tensor of size $R\times R\times R$ of unit Frobenius norm in which best rank R approximations are guaranteed to exist. For each choice of R, the bound ϵ is averaged over 20 trials to compute the mean bound $\bar{\epsilon}$. The mean bound $\bar{\epsilon}$ is plotted in dB as $-20\log_{10}(\bar{\epsilon})$. Thus, on average for the tensors in our experiment, a best rank R approximation is guaranteed to exist if one has SNR greater than the reported value. The choices of R considered are $R=2$ through 10.

As previously discussed, the bound obtained from Theorem 1.3 depends on the choice of unitary $\mathbf{U}$ chosen when setting $\mathcal{S}=\mathcal{T}\cdot_3\mathbf{U}$. While finding an optimal unitary is computationally challenging, as evidenced by its relationship to the longest path problem, one may simply try many randomly generated unitaries. For each tensor $\mathcal{T}$ we try a total of $1,000$ random unitary matrices and record the largest bound found after $1, 10, 100$, and 1000 unitaries tried. As a further step, one may find the optimal ordering of tensor slices for each unitary. Bounds computed without reordering slices are presented in the left image of Figure 5 while bounds computed with finding the optimal slice ordering are presented in the right image of the same figure. The collection of tensors and unitaries tried in each case are the same.

5.2.1. Bounds for existence in the case of rank less than dimensions. We now examine the setting more common in applications where a tensor has rank significantly less than its dimensions. We first note that since the bound obtained from Theorem 1.3 is a deterministic bound, the bound must take into account a worst case scenario and does not improve as tensor dimensions increase. A worst case scenario achievable for the core of a tensor is always achievable on the full tensor by applying a perturbation directly to the core of the tensor.

Although the deterministic bound given by Theorem 1.3 does not improve as tensor dimensions grow, *the probability of a worst case scenario occurring dramatically decreases.* This is illustrated in Figure 6 with the following experiment.

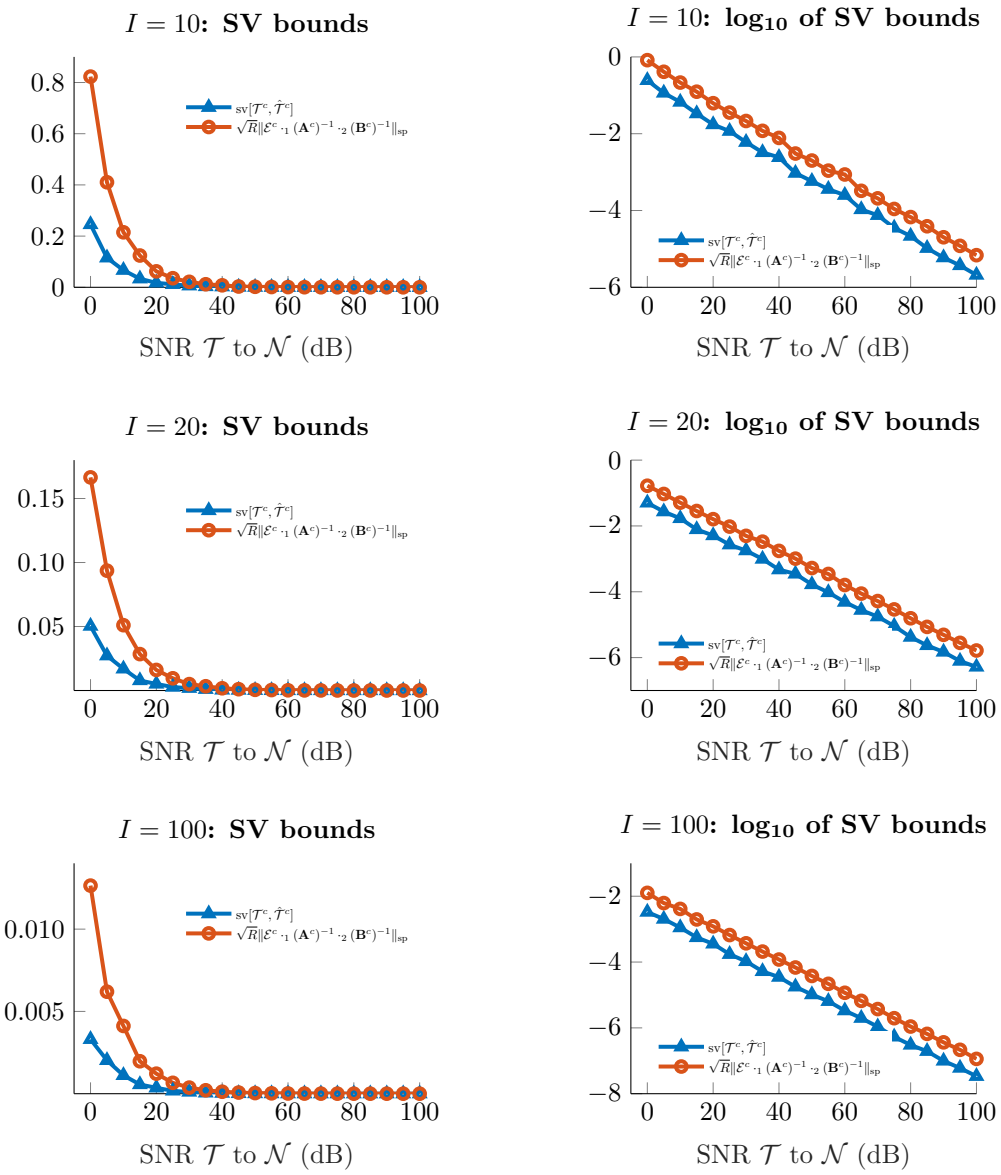


FIG. 3. *Observed spectral variation and computed bound for cores* $\mathcal{T}^c$ *and* $\hat{\mathcal{T}}^c$ *where* $\mathcal{T}$ *is a rank* 4 *tensor of size* $I \times I \times I$ *with* $I = 10, 20$ *or* 100 *and* $\hat{\mathcal{T}}$ *is a rank* 4 *approximation to* $\mathcal{T} + \mathcal{N}$*. The bound* $\sqrt{R}\|\mathcal{E}^c \cdot_1 (\mathbf{A}^c)^{-1} \cdot_2 (\mathbf{B}^c)^{-1}\|_{\mathrm{sp}}$ *provides a strong estimate for the observed spectral variation. The spectral variation bound and observed spectral variation both decrease as* I *grows.*

For various choices of R and I and SNR, we randomly generate a rank R tensor $\mathcal{T}$ and a perturbation $\mathcal{N}$ of size $I \times I \times I$ such that the SNR of $\mathcal{T}$ to $\mathcal{N}$ is as specified. Then, as is commonly done in practice, we compute an MLSVD of $\mathcal{T}+\mathcal{N}$ and truncate the MLSVD to obtain a multilinear rank $R \times R \times R$ approximation of $\mathcal{T} + \mathcal{N}$. We let

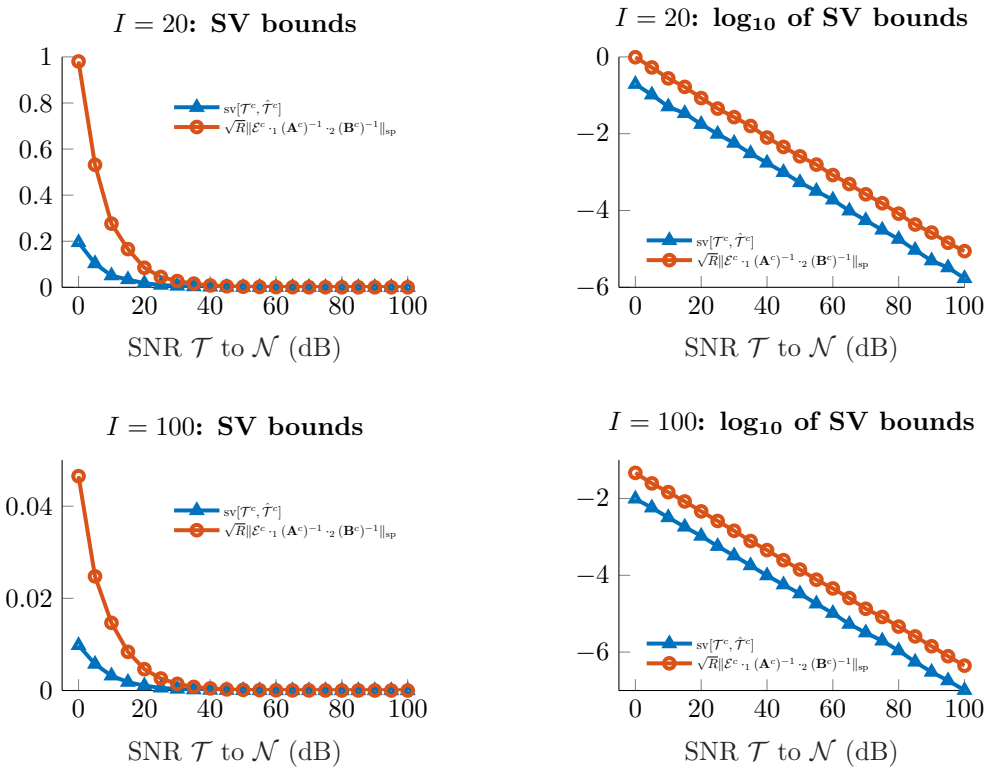


FIG. 4. *Observed spectral variation and computed bound for cores $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ where $\mathcal{T}$ is a rank 10 tensor of size $I \times I \times I$ with $I = 20$ or 100 and $\hat{\mathcal{T}}$ is a rank 10 approximation to $\mathcal{T} + \mathcal{N}$. The bound $\sqrt{R}\|\mathcal{E}^c \cdot_1 (\mathbf{A}^c)^{-1} \cdot_2 (\mathbf{B}^c)^{-1}\|_{\mathrm{sp}}$ provides a strong estimate for the observed spectral variation. The spectral variation bound and observed spectral variation both decrease as I grows.*

$(\mathcal{T}+\mathcal{N})^c$ denote the truncated MLSVD of $\mathcal{T}+\mathcal{N}$.

Theorem 1.4 is used to compute a bound[8] ϵ such that if there exists a rank R tensor $\mathcal{T}'$ satisfying $\|(\mathcal{T}+\mathcal{N})^c - \mathcal{T}'\|_F < \epsilon$, then $(\mathcal{T}+\mathcal{N})^c$ is guaranteed to have a best rank R approximation. We then compute a rank R approximation $\hat{\mathcal{T}}$ of $(\mathcal{T}+\mathcal{N})^c$ and check if $\|(\mathcal{T}+\mathcal{N})^c - \hat{\mathcal{T}}\|_F < \epsilon$, in which case we are able to guarantee that $(\mathcal{T}+\mathcal{N})^c$ has a best rank R approximation. For each choice of I, R, and SNR, a total of 10 trials are performed and the proportion of truncated MLSVDs guaranteed to have a best rank R approximation is reported.

As expected, we observe that the SNR needed for $(\mathcal{T}+\mathcal{N})^c$ to be guaranteed to have a best rank R approximation greatly decreases as I increases. Additionally, for each choice of I and R, we observe a sharp transition from being unable to guarantee the existence of a best rank R approximation of $(\mathcal{T}+\mathcal{N})^c$ to expecting to be able to guarantee the existence of a best rank R approximation of $(\mathcal{T}+\mathcal{N})^c$.

[8]For each $(\mathcal{T}+\mathcal{N})^c$, a total of 1000 randomly generated unitaries are tried when computing this radius. To limit computational complexity, we do not compute the optimal slice orderings. It is expected that either increasing the number of unitaries tried or computing the optimal slice ordering for each unitary would cause the curves in Figure 6 to shift slightly to the left.

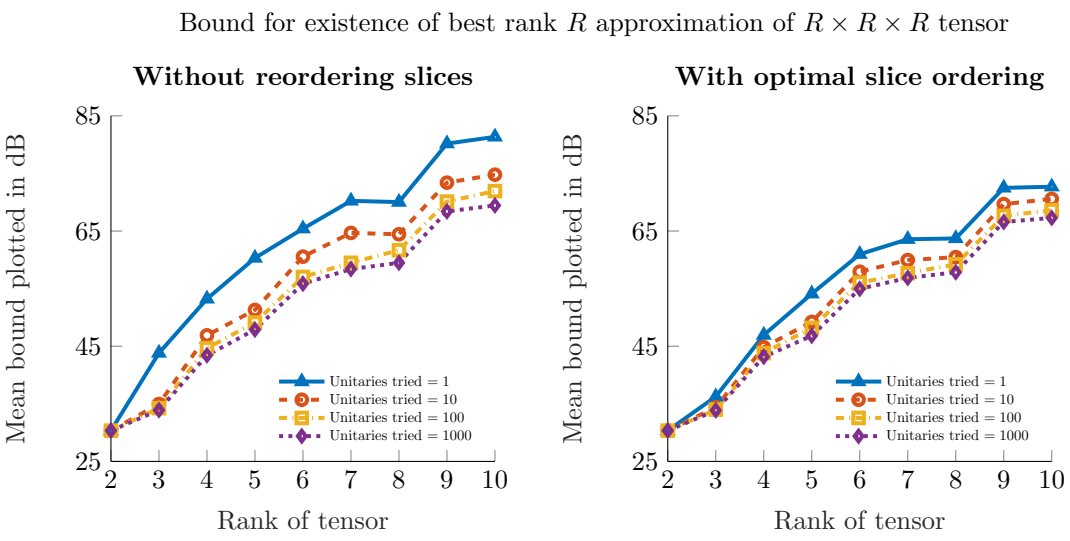


FIG. 5. *Mean bound for radius of a ball centered at a regular rank R tensor of unit Frobenius norm in which best rank R approximations are guaranteed to exist. The minimum SNR needed for existence of a best low rank approximation to be guaranteed increases as R increases from 2 to 10. Increasing the number of unitaries tried can notably improve the bound; however, there are diminishing returns as the number of unitaries tried increases.*

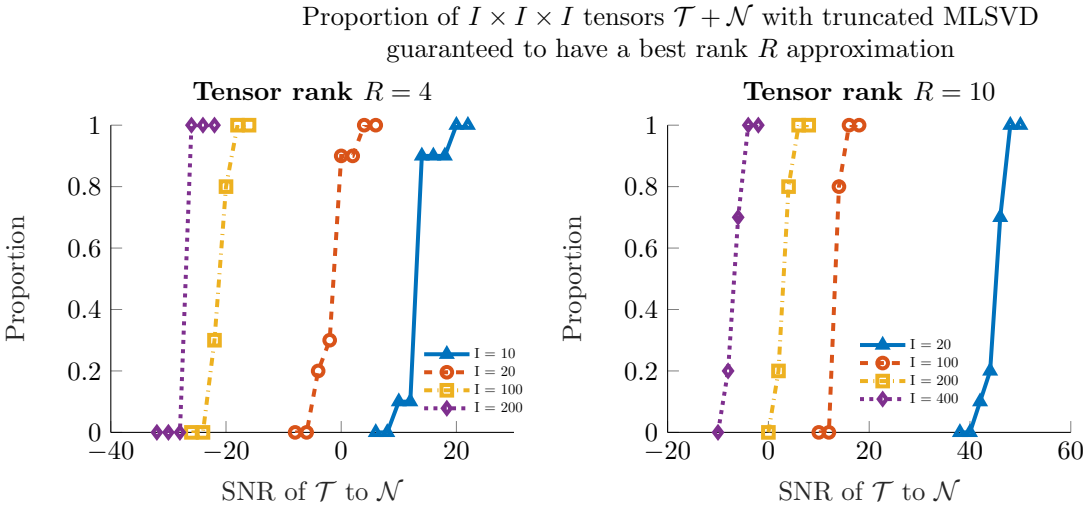


FIG. 6. *Proportion of $I \times I \times I$ tensors $\mathcal{T} + \mathcal{N}$ for which the truncated MLSVD of size $R \times R \times R$ is guaranteed to have a best rank R approximation. The necessary SNR decreases greatly as tensor dimensions grow. For each choice of I and R, a sharp transition point is observed. When the SNR exceeds the transition point, the truncated MLSVD is expected to have a best rank R approximation.*

5.3. 0 dB SNR i.i.d noise is not 0 dB SNR noise on the core. In this section we perform an experiment which is intended to help answer the following question: "Given a tensor $\mathcal{T}$ and a perturbation $\mathcal{N}$, what is the expected distance between

a pair of Procrustes cores computed for $\mathcal{T}$ and a best border rank R approximation to $\mathcal{T}+\mathcal{N}$?" As we now argue, this is a critical measure of distance when considering the existence of a best rank R approximation of $\mathcal{T}+\mathcal{N}$.

Recall the main idea in this article for determining if a tensor $\mathcal{T}+\mathcal{N}$ has a best rank R approximation: The tensor $\mathcal{T}+\mathcal{N}$ is guaranteed to have a best border rank R approximation $\hat{\mathcal{T}}$, so we need only verify that $\hat{\mathcal{T}}$ has rank R. To determine if $\hat{\mathcal{T}}$ has rank R, we check if $\hat{\mathcal{T}}$ lies in a neighborhood around $\mathcal{T}$ whose radius is determined by Theorem 1.3.

Moreover, since tensor (border) rank is invariant under modal unitary multiplication, it is sufficient to check that if one computes Procrustes cores $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ for the pair $(\mathcal{T},\hat{\mathcal{T}})$, then $\hat{\mathcal{T}}^c$ lies in the neighborhood around $\mathcal{T}^c$ given by Theorem 1.3. The key point is that *the existence of a best rank R approximation of $\mathcal{T}+\mathcal{N}$ can be determined by checking the distance between $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$.*

The experiment shown in Figure 7 is performed as follows. For various choices of I and R we randomly generate rank R tensors of size $I\times I\times I$ with unit Frobenius norm. For each tensor we generate a noise tensor $\mathcal{N}$ and scale $\mathcal{N}$ so that the SNR of $\mathcal{T}$ to $\mathcal{N}$ is as specified. After generating $\mathcal{T}$ and $\mathcal{N}$, we compute a rank R approximation[9] $\hat{\mathcal{T}}$ of $\mathcal{T}+\mathcal{N}$ and the distance $\|\mathcal{T}-\hat{\mathcal{T}}\|_F$ between the tensors $\hat{\mathcal{T}}$ and $\mathcal{T}$. We then compute Procrustes cores $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ of $\mathcal{T}$ and $\hat{\mathcal{T}}$, respectively, for which it is guaranteed that $\|\mathcal{T}^c-\hat{\mathcal{T}}^c\|_F \le \|\mathcal{T}-\hat{\mathcal{T}}\|_F$ and record the distance $\|\mathcal{T}^c-\hat{\mathcal{T}}^c\|_F$ between the cores. For each choice of I and R, a total of 10 trials are performed and the mean relative distance between the relevant pairs at each step is reported in dB. That is, for a mean relative distance $\bar{d}$ we report $-20\log_{10}(\bar{d})$.

Figure 7 illustrates that the distance between the core tensors $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ is significantly smaller than the distance between $\mathcal{T}$ and the perturbed tensor $\mathcal{T}+\mathcal{N}$. We observe that the distance between $\mathcal{T}$ and $\hat{\mathcal{T}}$ is significantly less than that between $\mathcal{T}$ and $\mathcal{T}+\mathcal{N}$, and that the distance between $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ is even smaller still. Furthermore, consistent with Figure 6, this experiment shows that when $R \ll I$, issues of nonexistence of a best rank R approximation are unlikely to occur (given our experimental setup of i.i.d factors and i.i.d noise) unless the SNR between $\mathcal{T}$ and $\mathcal{N}$ is significantly lower than suggested by Figure 5.

6. Conclusion. We have given a deterministic bound on the radius of an open Frobenius norm ball centered at a regular rank R tensor in which best rank R approximations are guaranteed to exist and in which every rank R tensor has a unique CPD. This illustrates the existence of a neighborhood of "mathematical truth" around tensors occurring in application and gives a quantitative explanation for the numerical success of the CPD in practical settings. In addition we solved a tensor Procrustes problem which provides a method for computing orthogonal compressions for a pair of tensors.

The bound for existence of best rank R approximations was computed by connecting the CPD to a JGE value decomposition for a tuple of matrices. Several basic results were established for JGE values, including an examination of algebraic and geometric multiplicities. We showed that a regular tensor which has rank greater than border must be defective in the sense of algebraic and geometric multiplicities for JGE values. Furthermore, we developed perturbation theoretic bounds which may be used to guarantee that tensors in a neighborhood of a given regular rank R tensor are nondefective, thus giving a neighborhood in which every tensor has a best rank

[9] For the purpose of this experiment, it is not necessary to compute a best rank R approximation.

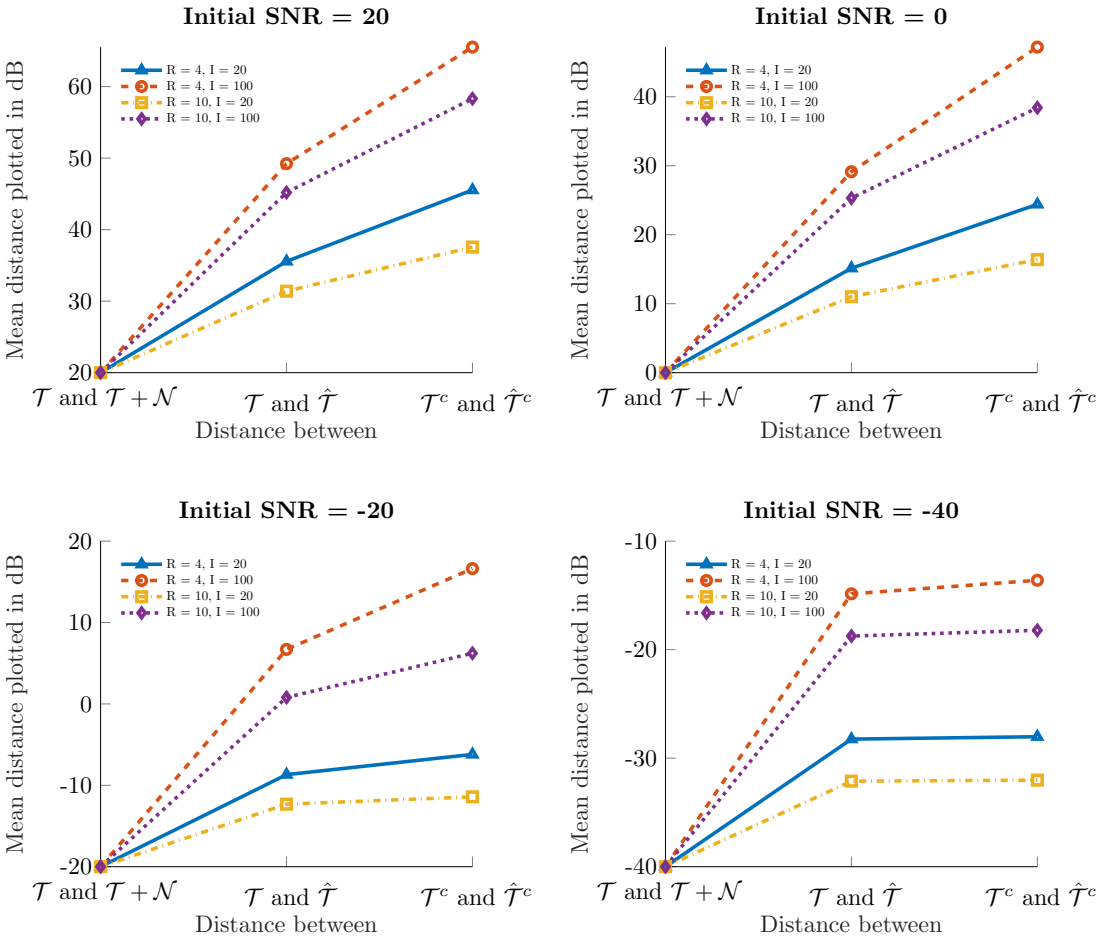


FIG. 7. *Distance between $\mathcal{T}$ and $\mathcal{T}+\mathcal{N}$, between $\mathcal{T}$ and $\hat{\mathcal{T}}$, and between $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ plotted in dB. The mean distance between $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ is observed to be significantly smaller than that between $\mathcal{T}$ and $\mathcal{T}+\mathcal{N}$, and the difference between these distances increases as the ratio of I to R increases. The mean distance between $\mathcal{T}^c$ and $\hat{\mathcal{T}}^c$ is smaller than that between $\mathcal{T}$ and $\hat{\mathcal{T}}$; however, the order of magnitude of the gap depends on the initial SNR of $\mathcal{T}$ to $\mathcal{N}$ with smaller gaps when the initial SNR is lower.*

R approximation.

REFERENCES

[1] A. ANANDKUMAR, R. GE, D. HSU, S. M. KAKADE, AND M. TELGARSKY, *Tensor decompositions for learning latent variable models*, " J. Mach. Learn. Res., 15 (2014), p. 2773–2832.
[2] G. BERGQVIST, *Exact probabilities for typical ranks of* $2 \times 2 \times 2$ *and* $3 \times 3 \times 2$ *tensors*, Linear Algebra Appl., 438 (2013), pp. 663–667.
[3] A. CICHOCKI, D. MANDIC, L. DE LATHAUWER, G. ZHOU, Q. ZHAO, C. CAIAFA, AND H. A. PHAN, *Tensor decompositions for signal processing applications: From two-way to mul-*

tiway component analysis, IEEE Signal Process. Mag., 32 (2015), pp. 145–163.
[4] P. Comon and C. Jutten, *Handbook of Blind Source Separation: Independent Component Analysis and Applications*, Academic Press, 2010.
[5] L. De Lathauwer, B. De Moor, and J. Vandewalle, *An introduction to independent component analysis*, J. Chemometrics, 14 (2000), pp. 123–149.
[6] L. De Lathauwer, B. De Moor, and J. Vandewalle, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1253–1278.
[7] V. de Silva and L. Lim, *Tensor rank and the ill-posedness of the best low-rank approximation problem*, SIAM J. Matrix Anal. & Appl., 30 (2008), pp. 1084–1127.
[8] I. Domanov and L. De Lathauwer, *On the uniqueness of the canonical polyadic decomposition of third-order tensors—part ii: Uniqueness of the overall decomposition*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 876–903.
[9] I. Domanov and L. De Lathauwer, *Canonical polyadic decomposition of third-order tensors: Reduction to generalized eigenvalue decomposition*, SIAM J Matrix Anal Appl., 34 (2014), pp. 636–660.
[10] S. Friedland and L.-H. Lim, *Nuclear norm of higher-order tensors*, Math. Comp., 87 (2018), pp. 1255–1281.
[11] S. Friedland, V. Mehrmann, R. Pajarola, and S. Suter, *On best rank one approximation of tensors*, Numerical Linear Algebra with Applications, 20 (2013), pp. 942–955.
[12] S. Friedland and G. Ottaviani, *The number of singular vector tuples and uniqueness of best rank-one approximation of tensors*, Foundations of Computational Mathematics, 14 (2014), pp. 1209–1242.
[13] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, 4th ed., 2013.
[14] C. J. Hillar and L.-H. Lim, *Most tensor problems are np-hard*, J. ACM, 60 (2013), https://doi.org/10.1145/2512329, https://doi.org/10.1145/2512329.
[15] T. Kolda and B. Bader, *Tensor decompositions and applications*, SIAM Review, 51 (2009), pp. 455–500.
[16] J. Kruskal, *Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics*, Linear Algebra Appl., 18 (1977), pp. 95–138.
[17] J. B. Kruskal, *Rank, Decomposition, and Uniqueness for 3-Way and n-Way Arrays*, North-Holland Publishing Co., NLD, 1989, p. 7–18.
[18] J. B. Kruskal, R. A. Harshman, and M. E. Lundy, *How 3-MFA Data Can Cause Degenerate Parafac Solutions, among Other Relationships*, North-Holland Publishing Co., NLD, 1989, p. 115–122.
[19] J. Landsberg, *Tensors: Geometry and Applications*, American Mathematical Society, Providence, RI, 2012.
[20] Lek-Heng Lim, *Singular values and eigenvalues of tensors: a variational approach*, in 1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, 2005., 2005, pp. 129–132.
[21] L. Qi, H. Chen, and Y. Chen, *Tensor Eigenvalues and Their Applications*, Springer, 2018.
[22] Y. Qi, M. Michałek, and L.-H. Lim, *Complex best r-term approximations almost always exist in finite dimensions*, Applied and Computational Harmonic Analysis, (2019), https://doi.org/10.1016/j.acha.2018.12.003.
[23] P. Schönemann, *A generalized solution of the orthogonal procrustes problem*, Psychometrika, 31 (1966), pp. 1–10, https://doi.org/10.1007/BF02289451.
[24] N. Sidiropoulos and R. Bro, *On the uniqueness of multilinear decomposition of n-way arrays*, Journal of Chemometrics, 14 (2000), pp. 229–239.
[25] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, *Tensor decomposition for signal processing and machine learning*, IEEE Trans. Signal Process., 65 (2017), pp. 3551–3582.
[26] A. Stegeman, *Degeneracy in candecomp/parafac explained for $p \times p \times 2$ arrays of rank $p + 1$ or higher*, Psychometrika, 71 (2006), pp. 483–501.
[27] A. Stegeman, *Low-rank approximation of generic $p \times q \times 2$ arrays and diverging components in the candecomp/parafac model*, SIAM J. Matrix Anal. & Appl., 30 (2008), pp. 988–1007.
[28] G. W. Stewart and J. Sun, *Matrix Perturbation Theory*, Academic Press, 1990.
[29] B. Sturmfels, *Tensors and their eigenvectors*, Not. Am. Math. Soc., 63 (2016), pp. 604–606.
[30] J. ten Berge and H. Kiers, *Simplicity of core arrays in three-way principal component analysis and the typical rank of $p \times q \times 2$ arrays*, Linear Algebra Appl., 294 (1999), pp. 169—179.
[31] J. Ten Berge, H. Kiers, and J. De Leeuw, *Explicit candecomp/parafac solutions for a contrived $2 \times 2 \times 2$ array of rank three*, Psychometrika, 53 (1988), pp. 579–584.
[32] R. Uehara and Y. Uno, *Efficient algorithms for the longest path problem*, in Algorithms and

Computation, R. Fleischer and G. Trippen, eds., Berlin, Heidelberg, 2005, Springer Berlin Heidelberg, pp. 871–883.
[33] N. Vannieuwenhoven, *Condition numbers for the tensor rank decomposition*, Linear Algebra Appl., 535 (2017), pp. 35–86.
[34] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer, *Tensorlab 3.0*, Mar. 2016, https://www.tensorlab.net/ (accessed 2015/05/14). Version 5.2.2.
[35] A. S. W.P. Krijnen, T.K. Dijkstra, *On the non-existence of optimal solutions and the occurrence of "degeneracy" in the candecomp/parafac model*, Psychometrika, 73 (2008), pp. 431–439.
[36] M. Zhang, X. Li, and J. Peng, *Using joint generalized eigenvectors of a set of covariance matrix pencils for deflationary blind source extraction*, IEEE Trans. Signal Process., 66 (2018), pp. 2892–2904.

7. Supplementary Proofs. The supplementary materials give the proofs of Theorem 1.4 and Lemma 3.2.

7.1. Proof of Theorem 1.4. As previously mentioned, the proof of Theorem 1.4 is nearly identical to that of Theorem 1.3, with all notable differences occurring at the start of the proof. For the sake of completeness, a detailed proof is given.

Proof. Let $\tilde{\mathcal{T}}$ be as in the statement of the theorem. Assume there exists some $\mathbb{K}$-rank R tensor $\mathcal{T}'$ such that

$$\|\tilde{\mathcal{T}} - \mathcal{T}'\|_F < \epsilon$$

and $\hat{\mathcal{T}}$ be a best border $\mathbb{K}$-rank R approximation of $\tilde{\mathcal{T}}$ in the Frobenius norm. Since $\mathcal{T}'$ has rank R by assumption, we must have

$$\|\tilde{\mathcal{T}} - \hat{\mathcal{T}}\|_F \leq \|\tilde{\mathcal{T}} - \mathcal{T}'\|_F < \epsilon.$$

Since $\hat{\mathcal{T}}$ has border $\mathbb{K}$-rank R, the multilinear rank of $\hat{\mathcal{T}}$ must be $(R, R, \hat{K})$ for some $\hat{K} \leq R$. It is easy to check that a tensor has $\mathbb{K}$-rank R if and only if an orthogonal compression to its multilinear rank has $\mathbb{K}$-rank R, so using Theorem 3.1 we may WLOG assume that $\hat{\mathcal{T}}$ and $\tilde{\mathcal{T}}$ both have size $R \times R \times \max\{K, \hat{K}\}$. There is no meaningful difference in the $K \geq \hat{K}$ and $K < \hat{K}$ cases, so for ease of notation we assume $K \geq \hat{K}$.

Now let $\mathbf{U} \in \mathbb{K}^{K \times K}$ be any unitary and set $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U}$ and $\hat{\mathcal{S}} = \hat{\mathcal{T}} \cdot_3 \mathbf{U}$. Since modal unitary multiplications preserve the Frobenius norm of a tensor we have $\|\tilde{\mathcal{S}} - \hat{\mathcal{S}}\|_F = \|\tilde{\mathcal{T}} - \hat{\mathcal{T}}\|_F < \epsilon$. Therefore, for some $i = 1, \ldots, \lfloor K/2 \rfloor$ we must have that

$$\|(\tilde{\mathbf{S}}_{2i-1} \;\; \tilde{\mathbf{S}}_{2i}) - (\hat{\mathbf{S}}_{2i-1} \;\; \hat{\mathbf{S}}_{2i})\|_2 \leq \|(\tilde{\mathbf{S}}_{2i-1}, \tilde{\mathbf{S}}_{2i}) - (\hat{\mathbf{S}}_{2i-1} \;\; \hat{\mathbf{S}}_{2i})\|_F < \epsilon_i,$$

as otherwise would imply

$$\|\tilde{\mathcal{S}} - \hat{\mathcal{S}}\|_F^2 \geq \sum_{k=1}^{\lfloor K/2 \rfloor} \|(\tilde{\mathbf{S}}_{2k-1} \;\; \tilde{\mathbf{S}}_{2k}) - (\hat{\mathbf{S}}_{2k-1} \;\; \hat{\mathbf{S}}_{2k})\|_F^2 \geq \sum_{k=1}^{\lfloor K/2 \rfloor} \epsilon_k^2 = \epsilon^2.$$

In particular, from our choice of ϵ_i the subpencil $(\hat{\mathbf{S}}_{2i-1}, \hat{\mathbf{S}}_{2i})$ must be regular and have R distinct generalized eigenvalues. Since $\hat{\mathcal{T}}$ has a regular subpencil, $\hat{\mathcal{T}}$ must be regular. Therefore, using Proposition 4.1 allows us to conclude that the tensor $\hat{\mathcal{T}}$ cannot have a JGE eigenvalue with algebraic multiplicity greater than 1 and Lemma 2.4 shows that every JGE value of $\hat{\mathcal{T}}$ has geometric multiplicity at least 1. We conclude that $\hat{\mathcal{T}}$ is regular and has a basis of JGE vectors, hence the result then follows from Theorem 1.1.

The fact that $\hat{\mathcal{T}}$ has a unique CPD is a straight forward consequence of Kruskal's condition [16]. ∎

7.2. Proof of Lemma 3.2. We now give the proof of Lemma 3.2.

Proof. Consider the orthogonal Procrustes problem:

$$\min_{\Psi \in O(I_1)} \|\mathbf{M} - \Psi \mathbf{N}\|_F. \tag{7.1}$$

where $O(I_1)$ denotes the group of $I_1 \times I_1$ unitary matrices. [23] shows that $\mathbf{W} \in O(I_1)$ is a minimizer of (7.1) if and only if $\mathbf{W}$ can be written $\mathbf{W} = \mathbf{U}\mathbf{V}^H$ where $\mathbf{U}, \mathbf{V} \in O(I_1)$ are unitary matrices such that

$$\mathbf{M}\mathbf{N}^H = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H \tag{7.2}$$

gives a singular value decomposition for $\mathbf{M}\mathbf{N}^H$ for an appropriate choice of $\mathbf{\Sigma}$.

It is straightforward to show that the unitary matrices $\mathbf{U}$ and $\mathbf{V}$ in equation (7.2) may be chosen so that

$$\mathbf{U} = (\mathbf{U}_1\ \mathbf{U}_2) \qquad \text{and} \qquad \mathbf{V} = (\mathbf{V}_1\ \mathbf{V}_2)$$

where $\mathbf{U}_1, \mathbf{V}_1 \in \mathbb{K}^{I_1 \times R_\mathbf{M}}$ and $\mathbf{U}_2, \mathbf{V}_2 \in \mathbb{K}^{I_1 \times (I_1 - R_\mathbf{M})}$ are column-wise orthonormal matrices satisfying

$$\operatorname{ran}\ \mathbf{U}_1 = \operatorname{ran}\ \mathbf{M} \qquad \text{and} \qquad \operatorname{ran}\ \mathbf{V}_1 \supseteq \operatorname{ran}\ \mathbf{N}$$

$$\operatorname{ran}\ \mathbf{U}_2 = (\operatorname{ran}\ \mathbf{M})^\perp \qquad \text{and} \qquad \operatorname{ran}\ \mathbf{V}_2 \subseteq (\operatorname{ran}\ \mathbf{N})^\perp.$$

Using this choice of $\mathbf{U}$ and $\mathbf{V}$ we have $\mathbf{V}_2^H \mathbf{N} = \mathbf{0}$, hence

$$\mathbf{U}\mathbf{V}^H\mathbf{N} = (\mathbf{U}_1\ \mathbf{U}_2)(\mathbf{V}_1\ \mathbf{V}_2)^H\mathbf{N} = \mathbf{U}_1\mathbf{V}_1^H\mathbf{N}.$$

It follows that $\operatorname{ran}\ \mathbf{M} \supseteq \operatorname{ran}\ \mathbf{U}\mathbf{V}^H\mathbf{N}$ which completes the first part of the proof.

To prove the second claim let $\mathbf{Z}_\mathbf{M} \in \mathbb{R}^{I_1 \times R_\mathbf{M}}$ be any column-wise orthonormal matrix mapping onto $\operatorname{ran}\ \mathbf{M}$, and set $\mathbf{Z}_\mathbf{N} = \mathbf{V}\mathbf{U}^H\mathbf{Z}_\mathbf{M}$ with $\mathbf{U}$ and $\mathbf{V}$ as above. Then we have

$$\|\mathbf{Z}_\mathbf{M}^H\mathbf{M} - \mathbf{Z}_\mathbf{N}^H\mathbf{N}\|_F \leq \|\mathbf{Z}_\mathbf{M}^H\|_2 \|\mathbf{M} - \mathbf{U}\mathbf{V}^H\mathbf{N}\|_F \leq \|\mathbf{M} - \mathbf{N}\|_F,$$

To see that $\operatorname{ran}\ \mathbf{Z}_\mathbf{N} \supseteq \operatorname{ran}\ \mathbf{N}$, we use $\mathbf{U}_2^H\mathbf{M} = \mathbf{0}$ to find

$$\mathbf{V}\mathbf{U}^H\mathbf{Z}_\mathbf{M} = (\mathbf{V}_1\ \mathbf{V}_2)(\mathbf{U}_1\ \mathbf{U}_2)^H\mathbf{Z}_\mathbf{M} = \mathbf{V}_1\mathbf{U}_1^H\mathbf{Z}_\mathbf{M}.$$

Since $\mathbf{Z}_\mathbf{M} \in \mathbb{K}^{I_1 \times R_\mathbf{M}}$ and $\mathbf{U}_1 \in \mathbb{R}^{I_1 \times R_\mathbf{M}}$ are column-wise orthonormal matrices with the same range, it follows that $\mathbf{U}_1^H\mathbf{Z}_\mathbf{M}$ is an invertible map on $\mathbb{K}^{R_\mathbf{M}}$. Therefore

$$\operatorname{ran}\ \mathbf{Z}_\mathbf{N} = \operatorname{ran}\ \mathbf{V}_1\mathbf{U}_1^H\mathbf{Z}_\mathbf{M} = \operatorname{ran}\ \mathbf{V}_1 \supseteq \operatorname{ran}\ \mathbf{N}$$

as claimed. ∎