GUARANTEES FOR EXISTENCE OF A BEST CANONICAL POLYADIC APPROXIMATION OF A NOISY LOW-RANK TENSOR*

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 [40, 4]. Tensor decompositions are powerful tools which can be used either to recover source signals responsible for measured data or to compress data [5].

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 low rank tensor is essentially unique [10, 25, 41]. 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 [23, 7]. In the case that the data has a natural tensor structure, CPD can allow for unique recovery of the source signals.

In applications one typically works with a measurement of a *low rank tensor*,

*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.

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. In particular, an arbitrary tensor may not have a best rank (less than or equal to) R approximation when $R > 1$ [27, 8]. Furthermore, even if a best rank $\mathbb{R}$ approximation exists, it can fail to be unique. In fact, when considering CPD over $\mathbb{R}$, nonexistence of a best low rank approximation occurs with positive probability [44, 49]. In contrast, it has recently been shown that a unique best rank R approximation exists with probability one when working over $\mathbb{C}$ [38]. See [14, 16] for further discussion of uniqueness of best approximations.[1]

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* [24, 45]. 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 [52].

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 [2, 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] Note that uniqueness of a best rank R approximation is a separate issue from uniqueness of the CPD of that approximation.

1.1. Brief statement of results. Given a tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, 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

$$\mathbf{T}(:,:,\ell)\mathbf{x} = \lambda_\ell \mathbf{y} \tag{1.1}$$

for all $\ell = 1,\dots,K$, and we call $\operatorname{span}(\boldsymbol{\lambda})$ the **joint generalized eigenvalue** of $\mathcal{T}$ corresponding to $\mathbf{x}$. In Section 2.4 we formally define appropriate notions of characteristic polynomials and algebraic multiplicity for JGE vectors. For now we roughly define the **algebraic multiplicity** of $\operatorname{span}(\boldsymbol{\lambda})$ to be the degree to which $\boldsymbol{\lambda}$ "divides" the characteristic polynomial of $\mathcal{T}$, and we define the **geometric multiplicity** of $\operatorname{span}(\boldsymbol{\lambda})$ to be the number of linearly independent JGE vectors corresponding to $\operatorname{span}(\boldsymbol{\lambda})$.

THEOREM 1.1. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ and assume there is some vector $\mathbf{v} \in \mathbb{K}^K$ such that $\mathcal{T} \cdot_3 \mathbf{v}$ is invertible. Then $\mathcal{T}$ has $\mathbb{K}$-rank R if and only if $\mathcal{T}$ has a basis of JGE eigenvectors.*

Furthermore, if $\mathcal{S} \in \mathbb{K}^{R\times R\times K}$ is a border $\mathbb{K}$-rank R tensor which has $\mathbb{K}$-rank strictly greater than R and there is a vector $\mathbf{u} \in \mathbb{K}^K$ such that $\mathcal{S} \cdot_3 \mathbf{u}$ is invertible, then $\mathcal{S}$ has a joint generalized eigenvalue with algebraic multiplicity strictly greater than geometric multiplicity.

Theorem 1.1 is a simplified statement of the upcoming Theorem 2.3. We remark that Theorem 1.1 can be viewed as a reformulation of algebraic-geometric tensor results in the language of the joint generalized eigenvalue problem, e.g. see [47, 29]. While this result is known in a different language, the reformulation in terms of joint generalized eigenvalues illustrates the connection between perturbation theory for joint generalized eigenvalues and existence of best rank R tensor approximations. Using this connection we give computable deterministic guarantees for the existence of best low rank tensor approximations.

THEOREM 1.2. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a tensor of $\mathbb{K}$-rank R and assume there is some vector $\mathbf{v} \in \mathbb{K}^K$ such that $\mathcal{T} \cdot_3 \mathbf{v}$ is invertible. For each $k = 1,\dots,\lfloor K/2\rfloor$, let $[\![\mathbf{A}_k, \mathbf{B}_k, \mathbf{C}_k]\!]$ be a CPD of the matrix subpencil $(\mathcal{T}(:,:,2k-1), \mathcal{T}(:,:,2k))$ of $\mathcal{T}$ where the columns of $\mathbf{C}_k$ have unit Euclidean norm. Define*

$$\epsilon_k := \frac{\sigma_{\min}(\mathbf{A}_k)\sigma_{\min}(\mathbf{B}_k)\min_{i\neq j}\chi(\mathbf{C}_k(:,i),\mathbf{C}_k(:,j))}{2} \qquad \text{for } k = 1,\dots,\lfloor K/2\rfloor$$

and set $\epsilon = \|(\epsilon_1,\dots,\epsilon_{\lfloor K/2\rfloor})\|_2$. If $\mathcal{W} \in \mathbb{K}^{R\times R\times K}$ satisfies $\|\mathcal{W}-\mathcal{T}\|_{\mathrm{F}} < \epsilon/2$ then $\mathcal{W}$ has a best $\mathbb{K}$-rank R approximation. Furthermore, any best $\mathbb{K}$-rank R approximation of $\mathcal{W}$ has a unique CPD, and, with probability one, the best $\mathbb{K}$-rank R approximation of $\mathcal{W}$ is unique. Here $\chi(\cdot,\cdot)$ denotes the chordal distance between unit vectors and $\sigma_{\min}(\cdot)$ denotes the smallest (possibly zero) singular value of a matrix.

The full statement of this result, Theorem 2.5, allows for more general collections of subpencils to be considered than those considered above. Additionally, a slight modification of this result may be applied to a measured tensor $\mathcal{M}'$ in an attempt to guarantee that $\mathcal{M}'$ has a best rank R approximation, see Theorem 2.7.

1.2. 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 $\mathcal{M}' = \mathcal{T}' + \mathcal{N}' \in \mathbb{K}^{I_1\times I_2\times I_3}$ where $\mathcal{T}'$ has rank R, then the Frobenius distance from $\mathcal{T}'$ to $\hat{\mathcal{T}}'$ cannot exceed $2\|\mathcal{N}'\|_{\mathrm{F}}$. 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.

We temporarily restrict our discussion to the case of $R \times R \times K$ tensors which are "slice mix invertible". In this setting, using the language of joint generalized eigenvalues, Theorem 2.3 shows that slice mix invertible tensors of size $R \times R \times K$ with border rank R and rank strictly greater than R have a joint generalized eigenvalue with algebraic multiplicity strictly greater than geometric multiplicity. We call such tensors defective. These defective tensors can be seen as a natural extension to the tensor setting of matrices which have a nontrivial Jordan block.

From this point our goal is to guarantee that a best border rank R approximation is nondefective, hence has rank equal to border rank. Mirroring the matrix case, a natural sufficient condition for a tensor to be nondefective is that the tensor has distinct joint generalized eigenvalues. This motivates the study and development of perturbation theoretic bounds for the joint generalized eigenvalue problem, see Theorems 2.8 and 4.3.

Theorem 2.8 in particular may be used to guarantee that a perturbation of a matrix pencil is nondefective. By applying this result to a collection of subpencils of the signal tensor, we obtain a full tensor based bound which may be used to guarantee the existence of a best rank R approximation of the measured tensor, see Theorems 1.2, 2.5, and 2.7.

The more general setting where $\mathcal{M}' = \mathcal{T}' + \mathcal{N}'$ has size $I_1 \times I_2 \times I_3$ and $\mathcal{T}'$ has rank $R \leq \min\{I_1, I_2\}$ is treated by a reduction to the $R \times R \times K$ setting. As above let $\hat{\mathcal{T}}'$ be a best border rank R approximation of $\mathcal{M}' = \mathcal{T}' + \mathcal{N}'$. Numerically, we wish to work with orthogonal compressions of $\mathcal{T}'$ and $\hat{\mathcal{T}}'$, as this puts us in the $R \times R \times K$ setting discussed above. However, there can be adverse effects if one independently computes the compressions of $\mathcal{T}'$ and $\hat{\mathcal{T}}'$.

Theorem 4.1 discusses how to orthogonally compress a *pair*[2] of tensors and shows that there exist orthogonal compressions $\mathcal{T}$ and $\hat{\mathcal{T}}$ of $\mathcal{T}'$ and $\hat{\mathcal{T}}'$, respectively, so that

$$\|\mathcal{T} - \hat{\mathcal{T}}\|_{\mathrm{F}} \leq \|\mathcal{T}' - \hat{\mathcal{T}}'\|_{\mathrm{F}}. \tag{1.2}$$

From this point we may use our results in the $R \times R \times K$ setting to guarantee that $\hat{\mathcal{T}}$ has rank R, hence $\hat{\mathcal{T}}'$ has rank R.

1.3. Connections to existing literature. We now briefly discuss the relationship between our results and existing literature.

1.3.1. Constrained existence results. While our results are novel in providing computable deterministic bounds for the existence of best unconstrained low rank tensor approximations, there are several known existence results in the setting of *constrained* tensor decompositions. For example, [33, Theorem 25] shows that if one places constraints on the coherence of the factor matrices for the approximation of $\mathcal{M}'$, then the constrained best low rank approximation problem always has a minimizer. As a special case, this result implies that best low rank approximations exist if one restricts to considering factor matrices which have orthogonal columns. Also see [19, 42] and references therein.

Another commonly considered constraint is nonnegativity. [32, Theorem 6.1], see also [37], shows that for nonnegative tensors, best low nonnegative rank approxima-

[2]Computing an orthogonal compression or a low multilinear rank approximation of a *single* tensor is a standard preprocessing in many applications, e.g., see [40].

tions always exist. That is, a best low rank approximation exists if one constrains their factors to be nonnegative.

As previously mentioned, it is in general known that if $\mathcal{M}'$ does not have a best rank R approximation then, near optimal approximations exhibit diverging components. Thus if one places constraints on the factors which prevent diverging components from occurring, then a best approximation with respect to those conditions will always exist. While in some applications it may be natural to consider such constraints, it is important to separately consider the *unconstrained* problem, as constraints will not always be applicable. Additionally, as discussed in [33], imposing constraints to "fix" ill-posedness may be misleading, as the problem at hand may in fact be ill-posed. For example, CPD may not be a proper data model for the problem, or the rank in consideration may be too low.

1.3.2. Existence in tubular neighborhoods. In the language of algebraic geometry, it is well known that given a smooth point on an algebraic variety, there is an open neighborhood around that point in which unique best approximations to the variety exist, e.g. see [15, 1, 43, 28]. Translating to this terminology, the "simple" (see Section 2.3) rank R tensors we consider are smooth points on the Rth secant variety of the Segre variety, thus there is necessarily some neighborhood around a simple rank R tensor in which best rank R approximations exist. This leads to a so called *tubular neighborhood* around the set of simple rank R tensors in which best rank R approximations are guaranteed to exist.

Though it was already known that there is *some neighborhood* around a given simple rank R tensor in which best rank R approximations are guaranteed to exist, until now there has been no known method for computing a lower bound for the *radius* of such a neighborhood. That is, until now there has been no known method for providing a *deterministic bound on noise levels* which may be used to guarantee that a measured tensor has a best rank R approximation.

1.3.3. The joint eigenvalue problem. A problem which is closely related to our joint generalized eigenvalue is the *joint eigenvalue problem* where one studies joint eigenvalues of a tuple of commuting matrices. In particular, by performing a modal multiplication in the first or second mode, a rank R slice mix invertible tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ can be transformed to a tensor whose frontal slices form a collection of commuting matrices.

There are indeed results in literature for the joint eigenvalue problem which may be used to show that if $(\mathbf{T}_1, \dots, \mathbf{T}_k) \in (\mathbb{K}^{R\times R})^K \cong \mathbb{K}^{R\times R\times K}$ is a simultaneously diagonalizable collection of matrices whose span contains a rank R matrix, then there exists *some neighborhood* around this tuple in which *other tuples of commuting matrices* are simultaneously diagonalizable, e.g., see [22, Theorem 3.5]; however, the results in this direction that we are aware of do not provide a method for computing a lower bound for the *radius* of such a neighborhood. Furthermore, though there is much literature, e.g. see [21], for the joint eigenvalue problem, there are two significant shortcomings in attempting to directly apply this literature to our setting.

First, we essentially wish to determine if some tensor $\hat{\mathcal{T}} \in \mathbb{K}^{R\times R\times K}$ is diagonalizable in the sense of the joint generalized eigenvalue problem based on its distance to a slice mix invertible tensor $\mathcal{T}$ which itself is diagonalizable in the sense of the joint generalized eigenvalue problem. Here $\hat{\mathcal{T}}$ is thought of as a best low rank approximation to some noisy observed tensor $\mathcal{T} + \mathcal{N}$ for which $\mathcal{T}$ is the signal portion. However, *a transformation which makes the frontal slices of $\mathcal{T}$ commute may not make the frontal slices of $\hat{\mathcal{T}}$ commute and there may exist no invertible transformation which simulta-*

neously converts both the frontal slices of $\mathcal{T}$ and the frontal slices of $\hat{\mathcal{T}}$ to tuples of commuting matrices. Thus, many tensors which can be treated with our results are missed by converting to the commuting case.

Second, treating a generalized eigenvalue problem[3] by transforming it to a classical eigenvalue problem can have adverse effects on stability. This is illustrated by the discussion in [17, Section 7.7.1]. These difficulties motivate studying perturbation theory for the joint generalized eigenvalue problem in its own right.

1.3.4. Perturbation theory for the generalized eigenvalue problem. One of our main results, Theorem 2.8, gives a lower bound on the radius of a neighborhood around a given diagonalizable matrix pencil in which all matrix pencils are diagonalizable. This result extends the "spectral variation" bound of [46, Chapter VI, Theorem 2.7] to the much stronger notion of a "matching distance" bound. Additionally, our extension is given in the tensor spectral norm. The proof of this extension may be unsurprising to experts in perturbation theory for classical generalized eigenvalues; however, the authors are unaware of an equivalent extension in the literature.

In general the literature on perturbation theory for generalized eigenvalues and generalized eigenvectors is rich, for example see [17, 35, 31]. We note that other bounds presented in the literature often use information from both the original matrix pencil and its perturbation. While such bounds can provide helpful insight into how generalized eigenvalues are affected by a specific perturbation, they are not well-suited to our needs.

1.4. Organization. Section 2 formally introduces our notation and definitions and gives formal statements of our main results. In Section 3 we discuss basic theory for the joint generalized eigenvalue problem. Section 4 discusses perturbation theoretic results for the joint generalized eigenvalue problem, while Section 5 gives deterministic bounds for existence of a best low rank tensor approximation using results from Sections 3 and 4. Numerical experiments illustrating our results are given in Section 6.

2. Notation, Definitions, and Main Results. We now give our basic definitions and notations as well as formal statements of the main results of the article. Let $\mathbb{K}$ denote $\mathbb{R}$ or $\mathbb{C}$. We denote scalars, vectors, matrices, and tensors with entries in $\mathbb{K}$ by lower case (m), bold lower case ($\mathbf{m}$), bold upper case ($\mathbf{M}$), and calligraphic script ($\mathcal{M}$), respectively. Additionally, script font ($\mathscr{M}$) is used to denote a subspace of $\mathbb{K}^I$. Given a matrix $\mathbf{M} \in \mathbb{K}^{I_1 \times I_2}$, let $\mathbf{M}^{\mathrm{T}}, \mathbf{M}^{\mathrm{H}}, \mathbf{M}^{\dagger}$ denote the transpose, Hermitian, and Moore-Penrose pseudo inverse of $\mathbf{M}$, respectively. If $\mathbf{M}$ is invertible, let $\mathbf{M}^{-1}$ and $\mathbf{M}^{-\mathrm{T}}$ denote the inverse and inverse transpose of $\mathbf{M}$. Additionally let $\|\mathbf{M}\|_2$ and $\|\mathbf{M}\|_{\mathrm{F}}$ denote the spectral and Frobenius norms of $\mathbf{M}$.

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 the linear form

$$\langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}} \rangle = \lambda_1 \gamma_1 + \lambda_2 \gamma_2 + \cdots + \lambda_K \gamma_K.$$

Here $\boldsymbol{\lambda}, \boldsymbol{\gamma} \in \mathbb{K}^K$ and $\bar{\boldsymbol{\gamma}}$ denotes the complex conjugate of $\boldsymbol{\gamma}$. Note that $\langle \cdot, \cdot \rangle$ denotes the usual inner product on $\mathbb{K}^K$ which is conjugate linear in the second term.[4]

[3] If $K = 2$, then our notions of joint generalized eigenvectors and joint generalized eigenvalues coincide exactly with the (classical) notions of generalized eigenvalues and generalized eigenvectors as discussed in [46].

[4] The linear form we are interested in is linear in $\boldsymbol{\gamma}$ while the inner product $\langle \cdot, \cdot \rangle$ is conjugate linear in the second argument. Hence, the complex conjugate of $\boldsymbol{\gamma}$ appears in the expression $\langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}} \rangle$. We use the notation $\langle \cdot, \cdot \rangle$ as we often consider some geometric relationship between $\boldsymbol{\lambda}$ and $\boldsymbol{\gamma}$.

A **tensor** is a multiindexed array $\mathcal{M}'$ with entries in $\mathbb{K}$. In this article our primary interest lies with tensors of order 3. That is, we study tensors $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$. Often we will restrict our study to low rank (pencil-like) tensors having size $R \times R \times K$ and rank R. To the reader keep track of context, we label tensors of size $I_1 \times I_2 \times I_3$ with the symbol $'$, while tensors of size $R \times R \times K$ not labelled in this way. So for example, $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ while $\mathcal{M} \in \mathbb{K}^{R \times R \times K}$.

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 a **rank one tensor**.

Every tensor may be expressed as a sum of rank one tensors. That is, given a tensor $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ 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

$$\mathcal{M}' = \sum_{r=1}^{R} \mathbf{a}_r \otimes \mathbf{b}_r \otimes \mathbf{c}_r. \tag{2.1}$$

In the case that R is as small as possible, the tensor $\mathcal{M}'$ is said to have $\mathbb{K}$-**rank** R, and the decomposition in equation (2.1) is called a **canonical polyadic decomposition** (CPD) of $\mathcal{M}'$. We let $\operatorname{rank}_{\mathbb{K}}(\mathcal{M}')$ denote the rank of $\mathcal{M}'$ 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 $\mathcal{M}'$ 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 [50, 3, 26].

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

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

as a compact notation for the CPD of $\mathcal{M}'$. The matrices $\mathbf{A}, \mathbf{B}$, and $\mathbf{C}$ are called the **factor matrices** of $\mathcal{M}'$. As discussed in [40] one has

$$\mathcal{M}'(:, :, k) = \mathbf{A} D_k(\mathbf{C}) \mathbf{B}^{\mathrm{T}} \qquad \text{for all } k = 1, \ldots, I_3.$$

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 $\mathcal{M}' = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ is **unique** if for any other CPD $\mathcal{M}' = [\![\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]\!]$, 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 [27, 8]. 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 $\mathcal{M}'$ tensor has **border $\mathbb{K}$-rank** R if $\mathcal{M}'$ is a limit of $\mathbb{K}$-rank R tensors and there is no integer $L < R$ such that $\mathcal{M}'$ is a limit of $\mathbb{K}$-rank L tensors. That is, R should be the smallest integer such that $\mathcal{M}'$ is a limit of $\mathbb{K}$-rank R tensors. We let $\mathbb{K}$-$\underline{\text{rank}}(\mathcal{M}')$ denote the border $\mathbb{K}$-rank of $\mathcal{M}'$.

We will make frequent use of the fact that the set of border $\mathbb{K}$-rank R tensors is a closed set. This in particular implies that a tensor always has a (not necessarily unique) best border $\mathbb{K}$-rank R approximation.

2.1. Approximately low rank tensors. In this article we are mainly interested in tensors which have (approximately) low rank in the following sense. Stated briefly, for tensors $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ occuring in applications, one often has $\mathcal{M}' = \mathcal{T}' + \mathcal{N}'$ where $\mathcal{T}'$ is a rank R tensor with rank equal to multilinear rank and $\mathcal{N}'$ is noise. That is, $\mathcal{M}'$ is a noisy measurement of low rank signal tensor $\mathcal{T}'$. Our main result, Theorem 2.5, uses information from $\mathcal{T}'$ to compute an upper bound ϵ on the Frobenius norm of the noise $\|\mathcal{N}'\|_{\mathrm{F}}$ such that $\mathcal{M}'$ is guaranteed to have a best rank R approximation if $\|\mathcal{N}'\|_{\mathrm{F}} < \epsilon$. A detailed explanation of our perspective follows.

Given a tensor $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$, define

$$\begin{aligned} R_1(\mathcal{M}') &:= \dim \operatorname{span}\{\mathcal{M}'(:,k,\ell)\}_{\forall k,\ell}, \\ R_2(\mathcal{M}') &:= \dim \operatorname{span}\{\mathcal{M}'(j,:,\ell)\}_{\forall j,\ell}, \\ R_3(\mathcal{M}') &:= \dim \operatorname{span}\{\mathcal{M}'(j,k,:)\}_{\forall j,k}. \end{aligned} \tag{2.2}$$

The integer triple $(R_1(\mathcal{M}'), R_2(\mathcal{M}'), R_3(\mathcal{M}'))$ is called the **multilinear rank** of $\mathcal{M}'$. In the case where $R_i(\mathcal{M}') = I_i$ for each $i = 1, 2, 3$ we say $\mathcal{M}'$ has **full multilinear rank**. The vectors in equation (2.2) used to define each $R_i(\mathcal{M}')$ are called the **mode-i fibers** of $\mathcal{M}'$, and we let $\mathcal{M}'_{(i)}$ denote a matrix whose columns are equal to the mode-i fibers of $\mathcal{M}'$. The matrix $\mathcal{M}'_{(i)}$ is called a **mode-i unfolding** of $\mathcal{M}'$.[5]

It is not difficult to show that one always has

$$\operatorname{rank}_{\mathbb{K}}(\mathcal{M}') \geq \max\{R_1(\mathcal{M}'), R_2(\mathcal{M}'), R_3(\mathcal{M}')\}. \tag{2.3}$$

For generic tensors this inequality is strict; however, for the signal portion of many tensors occurring in applications we have equality. In particular, in applications the tensor $\mathcal{M}'$ is often a noisy measurement of some low rank signal tensor $\mathcal{T}'$ which satisfies $\operatorname{rank}_{\mathbb{K}}(\mathcal{T}') \leq \min\{I_1, I_2, I_3\}$ and, up to a permutation of indices,

$$R = \operatorname{rank}_{\mathbb{K}}(\mathcal{T}') = R_1(\mathcal{T}') = R_2(\mathcal{T}') \geq R_3(\mathcal{T}') = K. \tag{2.4}$$

That is, we think of $\mathcal{M}'$ as $\mathcal{M}' = \mathcal{T}' + \mathcal{N}'$ where $\mathcal{T}'$ has $\mathbb{K}$-rank R and multilinear rank (R, R, K) and $\mathcal{N}'$ is noise. Note that our results still hold if the multilinear rank of $\mathcal{T}'$ in one mode is less than R, and allowing this increases flexibility, hence $R_3(\mathcal{T}') = K$ is allowed to be less than or equal to R.

For later use we note that the multilinear ranks of a tensor not only lower bound the rank of the tensor, but also the border rank of the tensor. That is, one has

$$\mathbb{K}\text{-}\underline{\text{rank}}(\mathcal{M}') \geq \max\{R_1(\mathcal{M}'), R_2(\mathcal{M}'), R_3(\mathcal{M}')\}. \tag{2.5}$$

[5] There is freedom in the ordering of the columns of $\mathcal{M}'_{(i)}$, so in this sense we have not well-defined $\mathcal{M}'_{(i)}$. 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.

This fact follows from lower semi-continuity of matrix rank, as tensors in a sufficiently small neighborhood around $\mathcal{M}'$ must have multilinear rank greater than or equal to that of $\mathcal{M}'$. It follows that the multilinear rank of $\mathcal{M}'$ provides a lower bound for the rank of tensors in a neighborhood of $\mathcal{M}'$.

2.2. Orthogonal compressions and modal products. In the case where $R_i(\mathcal{T}') < I_i$ for some i the tensor $\mathcal{T}'$ maybe be orthogonally compressed to a tensor $\mathcal{T} \in \mathbb{K}^{R_1(\mathcal{T}')\times R_2(\mathcal{T}')\times R_3(\mathcal{T}')}$. In particular, 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

$$\mathcal{T}' = \mathcal{T}' \cdot_1 \mathbf{V}_1^{\mathrm{H}} \cdot_2 \mathbf{V}_2^{\mathrm{H}} \cdot_3 \mathbf{V}_3^{\mathrm{H}} \in \mathbb{K}^{R_1(\mathcal{T}')\times R_2(\mathcal{T}')\times R_3(\mathcal{T}')} \tag{2.6}$$

contains the same algebraic and geometric information as $\mathcal{T}$, and we call this tensor an **orthogonal compression** of $\mathcal{T}'$. One may recover $\mathcal{T}'$ from $\mathcal{T}$ using

$$\mathcal{T}' = (\mathcal{T}' \cdot_1 \mathbf{V}_1^{\mathrm{H}} \cdot_2 \mathbf{V}_2^{\mathrm{H}} \cdot_3 \mathbf{V}_3^{\mathrm{H}}) \cdot_1 \mathbf{V}_1 \cdot_2 \mathbf{V}_2 \cdot_3 \mathbf{V}_3. \tag{2.7}$$

In equations (2.6) and (2.7) above, the product $\mathcal{T}' \cdot_1 \mathbf{V}_1^{\mathrm{H}}$ denotes the **mode-1 product** between the tensor $\mathcal{T}'$ and the matrix $\mathbf{V}_1^{\mathrm{H}}$. In general, the mode-1 product between a tensor $\mathcal{M}' \in \mathbb{K}^{I_1\times I_2\times I_3}$ and a matrix $\mathbf{U} \in \mathbb{K}^{J_1\times I_1}$ is the tensor $\mathcal{M}' \cdot_1 \mathbf{U} \in \mathbb{K}^{J_1\times I_2\times I_3}$ with mode-1 fibers given by

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

The mode-2 and mode-3 products $\cdot_2$ and $\cdot_3$ are defined analogously.

A common first step in CPD computation in applications is to compute[6] the core of a multilinear $(R_1(\mathcal{T}'), R_2(\mathcal{T}'), R_3(\mathcal{T}'))$ rank approximation to $\mathcal{M}' = \mathcal{T}' + \mathcal{N}' \in \mathbb{K}^{I_1\times I_2\times I_3}$, and it will be convenient for us to have notation for such tensors. Recalling that we typically assume $\mathcal{T}'$ has multilinear rank $R\times R\times K$, we let $\mathcal{W}'$ denote a tensor of size $I_1 \times I_2 \times I_3$ which has multilinear rank (R, R, K) and we let $\mathcal{W} \in \mathbb{K}^{R\times R\times K}$ be an orthogonal compression of $\mathcal{W}'$. In particular, $\mathcal{W}$ is thought of as the core of a low multilinear rank approximation to $\mathcal{M}'$, while $\mathcal{W}'$ itself is thought of as a low multilinear rank approximation to $\mathcal{M}'$.

2.2.1. Implicit subpencils of tensors. We frequently consider subpencils of a given tensor. That is, we will consider pairs of matrices which are formed from the frontal slices of a given tensor. The subpencils considered in this article differ slightly from the standard terminology in that we allow subpencils which are "revealed" by an orthogonal change of basis in the third mode of the tensor. We call such a subpencil an **implicit subpencil**.

Defined precisely, an implicit subpencil of $\mathcal{S}$ is an $R \times R \times 2$ tensor of the form $\mathcal{S} = \mathcal{W} \cdot_3 \mathbf{V} \in \mathbb{K}^{R\times R\times 2}$ where $\mathbf{V} \in \mathbb{K}^{2\times K}$ has rank 2 and has orthonormal rows. Since we always work with implicit subpencils, we drop the distinction "implicit" in the remainder of the article.

2.3. The joint generalized eigenvalue problem. We now discuss the joint generalized eigenvalue problem in greater detail. Note that we will interchangeably

[6]To compute an orthogonal compression or an (approximately) best low multilinear rank approximation of a tensor it is common to take the truncated core of a multilinear singular value decomposition of the tensor [6].

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 $\mathbf{T}_k = \mathcal{T}(:,:,k)$.

Recall that 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 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

$$\mathbf{T}_\ell \mathbf{x} = \lambda_\ell \mathbf{y} \tag{2.8}$$

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

There are natural scaling indeterminates in equation (2.8). In particular, if the triple $(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y})$ is a solution to equation (2.8), then so is $(\alpha\mathbf{x}, \beta\boldsymbol{\lambda}, \gamma\mathbf{y})$ for any nonzero scalars $\alpha, \beta, \gamma \in \mathbb{K}$ satisfying $\alpha = \beta\gamma$. For this reason, if $(\boldsymbol{\lambda}, \mathbf{x})$ is a solution to (2.8), then we say the one-dimensional subspace $\mathscr{L} := \operatorname{span}(\boldsymbol{\lambda}) \subset \mathbb{K}^K$ is a **joint generalized eigenvalue** (JGE value) of $\mathcal{T}$ and that $(\mathscr{L}, \mathbf{x})$ is a **joint generalized eigenpair** of $\mathcal{T}$. We typically use $\mathscr{L} \subset \mathbb{K}^K$ to denote a JGE value, while $\boldsymbol{\lambda} \in \mathbb{K}^K$ typically denotes a vector of unit norm such that $\operatorname{span}(\boldsymbol{\lambda}) = \mathscr{L}$. 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 **(joint generalized) spectrum** of the tensor. An equivalent formulation of joint generalized eigenvectors is used in [54] 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 [48, 30, 36].

2.4. 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. For our purposes, it will not be necessary to consider multiplicities in the $I_1 \times I_2 \times I_3$ setting.

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

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

We call $p_{\mathcal{T}}(\boldsymbol{\gamma})$ the **characteristic polynomial** of $\mathcal{T}$. Given a nonzero $\boldsymbol{\lambda} \in \mathbb{K}^K$, the **algebraic multiplicity** of $\mathscr{L} = \operatorname{span}(\boldsymbol{\lambda})$ as a JGE value of the tensor $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, denoted $\operatorname{am}(\mathscr{L})$, is the largest integer m such that $\langle \boldsymbol{\lambda}, \bar{\boldsymbol{\gamma}} \rangle^m$ divides the characteristic polynomial of $\mathcal{T}$. If $p_{\mathcal{T}}$ factors as a product of R (not necessarily distinct) linear terms, then we say that $\mathcal{T}$ has R JGE values counting algebraic multiplicity.

Define the **geometric multiplicity** of $\mathscr{L}$ as a JGE value of $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$, denoted $\operatorname{gm}(\mathscr{L})$, 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,

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

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 3.3. However, we warn that it is possible for a JGE value to have nonzero algebraic multiplicity while having geometric multiplicity equal to zero, see Example 3.2.

DEFINITION 2.1. *Say a tensor $\mathcal{T} \in R \times R \times K$ is* **defective** *if it has a JGE value $\mathscr{L}$ with $\operatorname{am}(\mathscr{L}) > \operatorname{gm}(\mathscr{L})$.*

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 $p_{\mathcal{T}}(\boldsymbol{\gamma}) \equiv 0$ then the tensor $\mathcal{T}$ is necessarily defective.

It is straightforward to show that $p_{\mathcal{T}}(\gamma)$ is not identically zero if and only if the span of the frontal slices of $\mathcal{T}$ contains an invertible matrix. We call such a tensor a **slice mix invertible** tensor. See Example 3.1 and Lemma 3.1 for further discussion.

DEFINITION 2.2. *Say $\mathcal{T} \in R \times R \times K$ is* **simple** *if $\mathcal{T}$ is slice mix invertible and has R distinct JGE values each having algebraic multiplicity equal to 1. In other words, $\mathcal{T}$ is simple if $p_{\mathcal{T}}(\boldsymbol{\gamma})$ is not identically zero and factors as a product of R distinct linear terms.*

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

THEOREM 2.3. *Let $\mathcal{T}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ and assume $\mathcal{T}'$ has multilinear rank (R, R, K). Let $\mathcal{T} \in \mathbb{K}^{R \times R \times K}$ be an orthogonal compression of $\mathcal{T}'$. Then we have the following.*

1. *$\mathcal{T}'$ has $\mathbb{K}$-rank R if and only if the JGE vectors of $\mathcal{T}$ span $\mathbb{K}^R$. Furthermore, if $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ has $\mathbb{K}$-rank R, then the JGE values of $\mathcal{T}$ are equal to the spans of columns of the factor matrix $\mathbf{C}$, and the columns of $\mathbf{B}^{-\mathrm{T}}$ are JGE vectors of $\mathcal{T}$.*
2. *If $\mathcal{T}'$ has border $\mathbb{K}$-rank R and $\mathcal{T}$ is simple then $\mathcal{T}'$ has $\mathbb{K}$-rank R.*
3. *If $\mathcal{T}'$ has border $\mathbb{K}$-rank R but $\mathbb{K}$-rank strictly greater than R and if $\mathcal{T}$ is slice mix invertible, then $\mathcal{T}$ must have a JGE value $\mathscr{L} \subset \mathbb{K}^K$ which satisfies*

$$\operatorname{am}(\mathscr{L}) > \operatorname{gm}(\mathscr{L}) \geq 1.$$

In this case, $\mathcal{T}$ is defective in the sense of joint generalized eigenvalues.

Proof. The proof of Theorem 2.3 is given in Section 3.5. ∎

Theorem 2.3 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 border $\mathbb{K}$-rank R tensors are simple.

PROPOSITION 2.4. *Let $\mathcal{T}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ and assume $\mathcal{T}'$ has $\mathbb{K}$-rank R and multilinear rank (R, R, K), and let $\mathcal{T} \in \mathbb{K}^{R \times R \times K}$ be an orthogonal compression of $\mathcal{T}'$. If every border $\mathbb{K}$-rank R tensor in an open ball of Frobenius radius ϵ around $\mathcal{T}$ is simple, 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 $\mathbb{K}$-rank R tensor inside the open ball of Frobenius radius ϵ centered at $\mathcal{T}'$ has a unique CPD. Furthermore, the set of tensors inside this open ball for which the best rank R approximation is not unique has Lebesgue measure zero.*

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 4.2. ∎

As previously mentioned, the existence of some neighborhood in which best rank R approximations exist is known. The main motivation for Proposition 2.4 is that a lower bound for the radius of a ball centered at a simple rank R tensor in which border rank R-tensors are simple can be computed using perturbation theoretic results for joint generalized eigenvalues, e.g., see Theorems 2.5 and 2.8.

We also note that [38, 14] together show that the set of tensors which do not have a unique best border $\mathbb{K}$-rank R approximation has measure 0. The uniqueness of best $\mathbb{K}$-rank R approximations discussed in Proposition 2.4 follows quickly from these results after having shown existence of best $\mathbb{K}$-rank R approximations.[7]

2.5. Deterministic bounds for existence of a best rank R approximation. We now explain how to compute the ϵ appearing in Proposition 2.4. For general K, Proposition 5.1 shows that if a tensor is defective, then any subpencil of the tensor must be defective. Hence by examining any given subpencil, one may use Theorem 2.8 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 2.5. *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 $\mathcal{T} \in \mathbb{K}^{R \times R \times K}$ be an orthogonal compression of $\mathcal{T}'$. Let $\mathbf{U} \in \mathbb{K}^{K \times K}$ be a unitary matrix and set $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U}$. 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 simple 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. Furthermore, the set of tensors inside this open ball for which the best rank R approximation is not unique has Lebesgue measure zero.

In particular, each ϵ_i may be computed using Theorem 2.8.

Proof. See Section 5.1. ∎

Remark 2.6. As an alternative to Theorem 2.8, one could use [9, Corollary 3] to compute the individual ϵ_i appearing in Theorem 2.5. We observe in experiments that the ϵ_i computed by Theorem 2.8 are typically about the same or slightly larger than those computed using [9, Corollary 3]. Additionally, it is much more computationally intensive to obtain the ϵ_i using [9, Corollary 3]. However, for particularly ill-conditioned tensors, [9, Corollary 3] may outperform Theorem 2.8.

Though Theorem 2.5 explains how to compute a lower bound for the radius of a neighborhood around a given rank R tensor in which best rank R approximations are guaranteed to exist, this result may be difficult to directly use in applications as one will not have access to the low rank signal tensor one is interested in.

To remedy this difficulty, we present a variation on Theorem 2.8 which can be directly applied to (a low multilinear rank approximation of) a measured tensor. This

[7] In fact it is known that there is a neighborhood in which best rank R approximations exist *and are unique*. In regards to uniqueness, the conclusion of Proposition 2.4 is slightly weaker; however, the fact that a lower bound for ϵ is computable is a significant advantage.

variation, Theorem 2.7, only requires access to the measured tensor to use.

THEOREM 2.7. *Given $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ and an integer $R \leq \min\{I_1, I_2\}$, set $K = \min\{R_3(\mathcal{M}'), R\}$. Let $\mathcal{W}'$ be a multilinear rank (R, R, K) approximation of $\mathcal{M}'$ and let $\mathcal{W} \in \mathbb{K}^{R \times R \times K}$ be an orthogonal compression of $\mathcal{W}'$. Also let $\mathbf{U} \in \mathbb{K}^{K \times K}$ be a unitary matrix and set $\mathcal{S} = \mathcal{W} \cdot_3 \mathbf{U}$. 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 simple and set $\epsilon = ||(\epsilon_1, \ldots, \epsilon_{\lfloor K/2 \rfloor})||_2$. If there exists some $\mathbb{K}$-rank R tensor $\tilde{\mathcal{T}}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ such that

$$\|\mathcal{M}' - \tilde{\mathcal{T}}'\|_{\mathrm{F}} < \epsilon - \|\mathcal{M}' - \mathcal{W}'\|_{\mathrm{F}} \tag{2.9}$$

then $\mathcal{M}'$ has a best $\mathbb{K}$-rank R approximation and any best $\mathbb{K}$-rank R approximation of $\mathcal{M}'$ has a unique CPD. Furthermore, the set of tensors $\mathcal{M}'$ which satisfy these assumptions and do not have a unique best $\mathbb{K}$-rank R approximation has Lebesgue measure zero.

In particular, each ϵ_i may be computed using Theorem 2.8.

Proof. The proof of Theorem 2.7 is nearly identical to the proof of Theorem 2.5. For completeness, details of the proof are given in the supplementary materials. ∎

Theorem 2.7 is most naturally applicable to tensors which have low or approximately low multilinear rank. Indeed if this theorem is applied to a tensor which does not approximately have low multilinear rank, then the right hand side of equation (2.9) can be negative in which case no conclusion can be made.

However, as previously discussed, a common approach to computing a CPD in application is to first compute a core $\mathcal{W}$ of a low multilinear rank approximation of $\mathcal{M}'$, then to compute a low rank CPD approximation of $\mathcal{W}$. In this approach, a critical question to answer is whether or not the core tensor $\mathcal{W}$ has a best low rank approximation. To address this question, Theorem 2.7 can be directly applied to $\mathcal{W}$. Since $\mathcal{W}$ already has low multilinear rank, if one uses Theorem 2.7 in this way, then inequality (2.9) simplifies to checking if there is some $\mathbb{K}$-rank R tensor $\tilde{\mathcal{T}} \in \mathbb{K}^{R \times R \times K}$ such that

$$\|\mathcal{W} - \tilde{\mathcal{T}}\|_{\mathrm{F}} < \epsilon$$

where ϵ is computed as described in the theorem statement. The tensor $\tilde{\mathcal{T}}$ used in this theorem can for instance be obtained by applying any standard CPD algorithm to $\tilde{\mathcal{T}}$. The resulting approximation error being bounded by ϵ is sufficient to make sure the approximation problem for $\mathcal{W}$ is well-posed. The fact that $\tilde{\mathcal{T}}$ itself may not be optimal is not an issue.

Numerical experiments which illustrate Theorem 2.7 are given in Section 6.2.1.

2.6. Deterministic bounds when $K = 2$. The following theorem explains how to compute the ϵ_i appearing in Theorems 2.5 and 2.7. Before stating the result we give three brief definitions. Given two one-dimensional subspaces $\mathscr{L}_1, \mathscr{L}_2 \subseteq \mathbb{K}^K$, the **chordal metric**[8] between $\mathscr{L}_1$ and $\mathscr{L}_2$, denoted $\chi(\mathscr{L}_1, \mathscr{L}_2)$, is equal to the sine of the angle between $\mathscr{L}_1$ and $\mathscr{L}_2$.

[8]The chordal metric is a popular choice of metric for studying generalized eigenvalues as it accounts for the natural scaling invariance in the generalized eigenvalue problem. See e.g. [46, Chapter VI].

Next, for $i = 1, 2$ let $\mathcal{T}_i \in \mathbb{K}^{R\times R\times K}$ be a tensor which has R JGE values counting algebraic multiplicity and let $\{\mathscr{L}_{i,r}\}_{r=1}^{R}$ be the spectrum of $\mathcal{T}_i$. Define the **matching distance** between $\mathcal{T}_1$ and $\mathcal{T}_2$, denoted $\mathrm{md}[\mathcal{T}_1, \mathcal{T}_2]$, by

$$\mathrm{md}[\mathcal{T}_1, \mathcal{T}_2] = \min_{\pi\in S_R} \max_{k=1,\dots,K} \chi(\mathscr{L}_{1,k}, \mathscr{L}_{2,\pi(k)}). \tag{2.10}$$

Here S_R is the group of permutations of $\{1, \dots, R\}$. We note that the matching distance is a metric between the spectrum $\mathcal{T}_1$ and the spectrum of $\mathcal{T}_2$.

Finally, for a tensor $\mathcal{M}' \in K^{I_1\times I_2\times I_3}$, define the **spectral norm** of $\mathcal{M}'$, denoted $\|\mathcal{M}'\|_{\mathrm{sp}}$ by

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

where the vectors $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are of compatible dimension. Equivalently, the spectral norm of $\mathcal{M}'$ is the largest singular value of $\mathcal{M}'$ which is in turn given by the Frobenius norm of a best rank 1 approximation of $\mathcal{M}'$. See [12, 13, 14, 34] for more information about best rank 1 approximations of tensors, tensor singular values, and the spectral norm of tensors.

THEOREM 2.8. *Let $\mathcal{T}$ and $\mathcal{W}$ 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 Euclidean norm. Let $\{\mathscr{L}_r\}_{r=1}^{R}$ be the spectrum of $\mathcal{T}$. If*

$$\|\mathcal{T} - \mathcal{W}\|_{\mathrm{sp}} < \frac{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B}) \min_{i\neq j} \chi(\mathscr{L}_i, \mathscr{L}_j)}{2},$$

then $\mathcal{W}$ is a simple tensor of $\mathbb{K}$-rank R and satisfies

$$\mathrm{md}[\mathcal{T}, \mathcal{W}] \leq \frac{\|\mathcal{T} - \mathcal{W}\|_{\mathrm{sp}}}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}. \tag{2.11}$$

Since the Frobenius norm is an upper bound for the spectral norm of a tensor, Theorem 2.8 illustrates that if the pencil $(\mathbf{S}_{2i-1}, \mathbf{S}_{2i}) \cong \mathcal{S}_i$ appearing in Theorem 2.5 has CPD $\mathcal{S}_i = [\![\mathbf{A}_i, \mathbf{B}_i, \mathbf{C}_i]\!]$ where the columns of $\mathbf{C}_i$ have unit Euclidean norm, then one may take

$$\epsilon_i = \frac{\sigma_{\min}(\mathbf{A}_i)\sigma_{\min}(\mathbf{B}_i) \min_{k\neq j} \chi(\mathbf{C}_i(:,k), \mathbf{C}_i(:,j))}{2}.$$

Here $\mathbf{C}_i(:,k)$ is the kth column of $\mathbf{C}_i$. Note that one can use a generalized eigenvalue decomposition to determine if an $R \times R \times 2$ tensor is slice mix invertible and has rank R and, if so, compute a CPD of the tensor. This approach is often called Jennrich's algorithm [18].

In Theorem 4.3 we give a perturbation theoretic bound for the *joint* generalized eigenvalues of tensors $\mathcal{T}$ and $\mathcal{W}$ of size $R \times R \times K$. However, the measurement of variation between the joint generalized spectra of $\mathcal{T}$ and $\mathcal{W}$ used in Theorem 4.3 is notably weaker than the matching distance which is used in Theorem 2.8.

Remark 2.9. In the language of the classical generalized eigenvalue problem and matrix pencils, Theorem 2.8 gives a bound which guarantees that the matrix pencil $(\mathbf{W}_1, \mathbf{W}_2)$ is diagonalizable over $\mathbb{K}$.

3. Basic results for the joint generalized eigenvalue problem. We now develop several basic results for joint generalized eigenvalues. These results help reformulate CPD in terms of the joint generalized eigenvalue problem. This reformulation allows us to apply our upcoming perturbation theoretic results to the CPD setting.

3.1. 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 slice mix invertible. The assumption that $\mathcal{T}$ is slice mix invertible is closely related to the tensor having multilinear rank (R, R, K) and rank R, see Lemma 3.1.

As described in Section 2.1, orthogonal compressions preserve 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 slice mix invertible. In fact, it is possible for a tensor to have multilinear rank (R, R, K) and border $\mathbb{K}$-rank R and not be slice mix invertible. This is illustrated by the following example.

Example 3.1. 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 slice mix invertible. 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, for each n 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 let $\mathcal{T}^n$ be the $\mathbb{K}$-rank 3 tensor $\mathcal{T}^n = [\![\mathbf{A}_n, \mathbf{B}_n, \mathbf{C}_n]\!]$. Then $\lim \mathcal{T}^n = \mathcal{T}$, hence the border $\mathbb{K}$–rank of $\mathcal{T}$ is at most 3. That $\mathcal{T}$ has border $\mathbb{K}$–rank 3 then follows from the fact that $\mathcal{T}$ has multilinear rank $(3, 3, 3)$ together with equation (2.5). As a consequence of the forthcoming Lemma 3.1, $\mathcal{T}$ has $\mathbb{K}$-rank strictly greater than 3.

LEMMA 3.1. *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 a CPD of $\mathcal{T}$. Then $\mathcal{T}$ is slice mix invertible 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

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

where $\odot$ denotes the Khatri-Rao (column-wise Kronecker) product, e.g. see [4]. $\mathcal{T}$ 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 slice mix invertible. Using $\mathbf{T}_i = \mathbf{A}D_i(\mathbf{C})\mathbf{B}^{\mathrm{T}}$ 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 $\alpha \in \mathbb{K}^K$ outside of a set of measure zero, 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. ∎

3.2. 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 3.2. *Let $\mathcal{T} \in \mathbb{K}^{R \times R \times K}$ and assume $\mathcal{T}$ is slice mix invertible. Let $\{(\mathscr{L}_j, \mathbf{x}_j)\}_{j=1}^{\ell}$ be a collection of JGE pairs of $\mathcal{T}$ such that the $\mathscr{L}_j$ are all distinct. Then $\{\mathbf{x}_j\}_{j=1}^{\ell}$ is a linearly independent set. As a consequence, if $\mathcal{T}$ is simple and each of its JGE values has geometric multiplicity at least one, then $\mathcal{T}$ has a JGE basis.*

Proof. We treat the problem by reducing to the $K = 2$ setting. To this end, for $j = 1, \dots, \ell$ let $\boldsymbol{\lambda}_j \in \mathbb{K}^K$ be a vector such that $\operatorname{span}(\boldsymbol{\lambda}_j) = \mathscr{L}_j$, and let $\mathbf{L} \in \mathbb{K}^{K \times \ell}$ be the matrix whose jth column is equal to $\boldsymbol{\lambda}_j$. Since the $\mathscr{L}_j$ 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 $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{V} \in \mathbb{K}^{R \times R \times 2}$ and let $\mathbf{S}_1$ and $\mathbf{S}_2$ be the frontal slices of $\mathcal{S}$. 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 [46, 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 slice mix invertible 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 3.2 is necessary.

Example 3.2. 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 slice mix invertible and has three distinct JGE values counting algebraic multiplicity, namely $\operatorname{span}((1,1,1))$, $\operatorname{span}((1,2,3))$, and $\operatorname{span}((1,-1,5))$. However, the JGE values $\operatorname{span}((1,1,1))$ and $\operatorname{span}((1,2,3))$ do not have a corresponding JGE vector.

In Example 3.2, one could reasonably consider left joint generalized eigenvalues of $\mathcal{T}$ in which case the JGE value $\operatorname{span}((1,1,1))$ in Example 3.2 corresponds to the left JGE vector $(1,0,0) \in \mathbb{R}^3$. However, the JGE values $\operatorname{span}((1,-1,5))$ and $\operatorname{span}((1,2,3))$ do not have left JGE vectors. The JGE value $\operatorname{span}((1,2,3))$ is of particular note since it neither has left nor right JGE vectors.

Assuming that a tensor $\mathcal{T}$ is slice mix invertible, 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.

3.3. 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 $(\mathscr{L}, \mathbf{x})$, then $\operatorname{am}(\mathscr{L}) \geq 1$.

Before stating our result we recall that if $\mathcal{T}$ is not slice mix invertible, then $p_{\mathcal{T}}$ is identically zero in which case $\langle \boldsymbol{\lambda}, \boldsymbol{\gamma}\rangle^m$ divides $p_{\mathcal{T}}(\boldsymbol{\gamma})$ for all $\boldsymbol{\lambda} \in \mathbb{K}^K$ and all $m \in \mathbb{N}$. In this case one can say that $\operatorname{span}(\boldsymbol{\lambda})$ has infinite algebraic multiplicity as a JGE value of $\mathcal{T}$ for all $\boldsymbol{\lambda} \in \mathbb{K}^K$. For this reason, in the upcoming proposition it is sufficient to restrict to the case where $\mathcal{T}$ is slice mix invertible.

PROPOSITION 3.3. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ and assume $\mathcal{T}$ is slice mix invertible. 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*

$$p_{\mathcal{T}}(\boldsymbol{\gamma}) = \langle \boldsymbol{\lambda}, \overline{\boldsymbol{\gamma}}\rangle^m g(\boldsymbol{\gamma}) \tag{3.1}$$

where g is a nonzero polynomial in the indeterminate $\boldsymbol{\gamma}$.

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

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

hence $\mathcal{T}$ has R JGE values counting algebraic multiplicity.

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 $(\operatorname{span}(\boldsymbol{\lambda}), \mathbf{x}_i)$, as in equation (2.8). Note that the assumption that $\mathcal{T}$ is slice mix invertible guarantees that $\{\mathbf{y}_1, \ldots, \mathbf{y}_m\}$ is a linearly independent set. To show that $p_{\mathcal{T}}$ factors as desired, we will use the $\mathbf{x}_i$ and $\mathbf{y}_i$ to construct matrices $\mathbf{X}$ and $\mathbf{Q}$ which simultaneously block upper triangularize the frontal slices $\mathbf{T}_k$ with $\lambda_k \mathbf{I}_m$ on the first diagonal block.

To simultaneously block upper diagonalize the $\mathbf{T}_k$, 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 $\mathbf{Q}_2 \in \mathbb{R}^{(R-m)\times R}$ be a matrix whose rows form a basis for the orthogonal complement of the range of $\mathbf{Y}_1$. Set

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

With this set up, for each $k = 1, \ldots, K$ one has

$$\mathbf{Q}\mathbf{T}_k\mathbf{X} = \mathbf{Q}(\lambda_k\mathbf{Y}_1\ \mathbf{T}_k\mathbf{X}_2) = \begin{pmatrix} \lambda_k\mathbf{Y}_1^\dagger\mathbf{Y}_1 & \mathbf{Y}_1^\dagger\mathbf{T}_k\mathbf{X}_2 \\ \lambda_k\mathbf{Q}_2\mathbf{Y}_1 & \mathbf{Q}_2\mathbf{T}_k\mathbf{X}_2 \end{pmatrix} = \begin{pmatrix} \lambda_k\mathbf{I}_m & \mathbf{Y}_1^\dagger\mathbf{T}_k\mathbf{X}_2 \\ \mathbf{0} & \mathbf{Q}_2\mathbf{T}_k\mathbf{X}_2 \end{pmatrix}.$$

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

$$\det(\mathbf{Q})\det\left(\sum_{k=1}^{K}\gamma_k\mathbf{T}_k\right)\det(\mathbf{X}) = \det\left(\sum_{k=1}^{K}\gamma_k\mathbf{Q}\mathbf{T}_k\mathbf{X}\right) = \det\left(\sum_{k=1}^{K}\gamma_k\begin{pmatrix}\lambda_k\mathbf{I}_m & \mathbf{Y}_1^\dagger\mathbf{T}_k\mathbf{X}_2\\ \mathbf{0} & \mathbf{Y}_2\mathbf{T}_k\mathbf{X}_2\end{pmatrix}\right)$$

shows that $p_{\mathcal{T}}(\boldsymbol{\gamma})$ factors as in equation (3.1). The fact that the polynomial $g(\boldsymbol{\gamma})$ is nonzero is a consequence of the assumption that $\mathcal{T}$ is slice mix invertible, 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 slice mix invertible, so Lemma 3.1 shows that $\mathbf{A}$ and $\mathbf{B}$ are invertible. Using $\mathbf{T}_k = \mathbf{A}D_k(\mathbf{C})\mathbf{B}^{\mathrm{T}}$ shows that

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

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 slice mix invertible, we may assume that each $\mathcal{T}^n$ is slice mix invertible. 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 $r = 1,\dots,R$ 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 $r = 1,\dots,R$. 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

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

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

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

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 slice mix invertible, we must have $\alpha \neq 0$. We conclude that $p_{\mathcal{T}}(\boldsymbol{\gamma})$ factors as desired. ∎

3.4. Multiplicities for border rank R tensors. We now examine geometric multiplicities of JGE values for slice mix invertible border $\mathbb{K}$-rank R tensors. Namely we show that JGE values of a slice mix invertible tensor with border $\mathbb{K}$-rank R have nonzero geometric multiplicity. This lemma plays an important role in the proof of Theorem 2.3 (2), as it allows one to conclude that a simple border $\mathbb{K}$-rank R tensor has a basis of JGE vectors.

LEMMA 3.4. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a slice mix invertible tensor and assume $\mathcal{T}$ has multilinear rank (R, R, K) and border $\mathbb{K}$-rank R. If $\mathscr{L} \subset \mathbb{K}^K$ is a JGE value of $\mathcal{T}$ with $\operatorname{am}(\mathscr{L}) = m \geq 1$, then $\operatorname{gm}(\mathscr{L}) \geq 1$.*

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 $\operatorname{span}(\boldsymbol{\lambda}) = \mathscr{L}$. Proposition 3.3 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 2.3, 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 (2.8) 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

$$\mathbf{T}_k\mathbf{x} = \lim_{n\to\infty} \mathbf{T}_k^n\mathbf{x}_n = \lim_{n\to\infty} \lambda_{n,k}\mathbf{y}^n = \lambda_k\mathbf{y} \tag{3.2}$$

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

By assumption, $\mathcal{T}$ is slice mix invertible, 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 $(\mathscr{L}, \mathbf{x})$ is a joint generalized eigenpair of $\mathcal{T}$ and that $\operatorname{gm}(\mathscr{L}) \geq 1$. ∎

3.5. Proof of Theorem 2.3. We are now in position to give the proof of our first main result, Theorem 2.3.

Proof. We first prove Item (1). Note that an orthogonal compression $\mathcal{T}$ 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 $\mathcal{T} = \mathcal{T}'$.

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[9] $\{\boldsymbol{\Lambda}_k\}_{k=1}^K$, and matrix $\mathbf{Y} \in \mathbb{K}^{R\times R}$ such that

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

[9]The matrices $\{\boldsymbol{\Lambda}_k\}_{k=1}^K$ are closely related to the JGE values of $\mathcal{T}$. In particular the subspace $\operatorname{span}((\boldsymbol{\Lambda}_1(r,r), \boldsymbol{\Lambda}_2(r,r), \dots, \boldsymbol{\Lambda}_K(r,r)) \subset \mathbb{K}^K$ is a JGE value of $\mathcal{T}$ for each $r = 1, \dots, R$. Here $\boldsymbol{\Lambda}_k(r,r)$ denotes the (r,r)th entry of $\boldsymbol{\Lambda}_k$.

It follows that

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

Set $\mathbf{A} = \mathbf{Y}$ and $\mathbf{B} = \mathbf{X}^{-\mathrm{T}}$ and let $\mathbf{C} \in \mathbb{K}^{R\times K}$ be the matrix whose kth row is given by the diagonal entries of $\mathbf{\Lambda}_k$. 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 3.3, 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 3.1 shows that the factor matrices $\mathbf{A}, \mathbf{B} \in \mathbb{K}^{R\times R}$ are invertible. Using $\mathbf{T}_k = \mathbf{A}D_k(\mathbf{C})\mathbf{B}^{\mathrm{T}}$, we then have

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

For $r = 1, \dots, R$, let $\mathbf{x}_r$ be the rth column of the matrix $\mathbf{B}^{-\mathrm{T}}$ 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 $\mathscr{L}_r = \operatorname{span}\big(c_{1r}, \dots, c_{Kr}\big)$, it follows that $(\mathscr{L}_r, \mathbf{x}_r)$ 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 $\mathbf{B}^{-\mathrm{T}}$, as claimed.

Item (2) follows immediately from Item (1) together with Lemma 3.4. To prove Item (3), suppose $\mathcal{T}'$ has border $\mathbb{K}$-rank R and assume towards a contradiction that $\operatorname{am}(\mathscr{L}) = \operatorname{gm}(\mathscr{L})$ for all JGE values $\mathscr{L}$ of $\mathcal{T}$. Proposition 3.3 shows that $p_{\mathcal{T}}$ factors as a product of linear terms, so in this case $\mathcal{T}$ must have a basis of JGE vectors which using Item (1) shows that $\mathcal{T}'$ has rank R, a contradiction. It follows that $\mathcal{T}$ has some JGE value with algebraic multiplicity strictly greater than geometric multiplicity. The proof is completed by Lemma 3.4 which shows that every JGE value of $\mathcal{T}$ has geometric multiplicity at least one. ∎

4. 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 slice mix invertible $\mathbb{K}$-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 [39].

In particular we will show that given two tensors $\mathcal{W}', \hat{\mathcal{W}}' \in \mathbb{K}^{I_1\times I_2\times I_3}$, both having multilinear rank bounded above by (R, R, K), there exist orthogonal compressions $\mathcal{W}, \hat{\mathcal{W}} \in \mathbb{K}^{R\times R\times K}$ of $\mathcal{W}'$ and $\hat{\mathcal{W}}'$ such that

$$\|\mathcal{W} - \hat{\mathcal{W}}\|_{\mathrm{F}} \leq \|\mathcal{W}' - \hat{\mathcal{W}}'\|_{\mathrm{F}}. \tag{4.1}$$

Though it is straightforward how to compute these orthogonal compressions in the special case that the mode-j fibers of $\mathcal{W}'$ and $\hat{\mathcal{W}}'$ span the same space for each $j = 1, 2, 3$, one should not expect an arbitrary pair of low multilinear rank tensors to have this special property.

In Section 4.1 we explain how to compute orthogonal compressions of $\mathcal{W}'$ and $\hat{\mathcal{W}}'$ which satisfy equation (4.1). Additionally, we explain our perspective on orthogonal compressions which differs slightly from the perspective most commonly taken in applications. While the tensor Procrustes problem plays an important role in our theoretical results, the solution may be unsurprising to readers familiar with the matrix Procrustes problem. Such readers can advance to Section 4.2.

4.1. 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 $\mathcal{M}' = \mathcal{T}' + \mathcal{N}'$ 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 $\mathcal{M}'$. 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 $\mathcal{M}'$ one first computes a low multilinear rank approximation $\mathcal{W}'$ of $\mathcal{M}'$. One then computes a R rank approximation of the low multilinear rank approximation $\mathcal{W}'$.

While this approach is natural to use in applied settings, it can have adverse effects when searching for theoretical guarantees. Letting $\hat{\mathcal{T}}'$ denote a best border rank R approximation of $\mathcal{M}'$, to simplify to the $R \times R \times K$ setting, we need orthogonal compressions of both $\mathcal{T}'$ and $\hat{\mathcal{T}}'$. If one independently computes orthogonal compressions $\mathcal{T}$ and $\hat{\mathcal{T}}$ of $\mathcal{T}'$ and $\hat{\mathcal{T}}'$, then it is possible that the distance between $\mathcal{T}$ and $\hat{\mathcal{T}}$ is greater than the distance between $\mathcal{T}'$ and $\hat{\mathcal{T}}'$. This difficulty must be treated with particular care if the span of the mode-j subspaces of $\mathcal{T}'$ is not the same as the span of the mode-j subspace of $\hat{\mathcal{T}}$.

Roughly speaking this difficulty is resolved by first applying an orthogonal transformation to $\hat{\mathcal{T}}$ which both does not increase the distance between $\mathcal{T}$ and $\hat{\mathcal{T}}$ and makes it so that the mode-j subspaces of the orthogonal transformation of $\hat{\mathcal{T}}$ are the same as those of $\mathcal{T}$. From this point it is straightforward how to compute the appropriate orthogonal compressions. As the (border) rank of $\hat{\mathcal{T}}'$ is invariant under orthogonal transformations, for our purposes there is no loss in studying an orthogonal transformation of $\hat{\mathcal{T}}'$ rather than $\hat{\mathcal{T}}'$ itself.[10]

THEOREM 4.1. *Let $\mathcal{W}', \hat{\mathcal{W}}' \in \mathbb{K}^{I_1 \times \cdots \times I_\ell}$ 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 $\mathcal{W}'$. If $\hat{R}_j \leq R_j$, 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 $\hat{\mathcal{W}}$ such that*

$$\|\mathcal{W}' \cdot_j \mathbf{V}^{\mathrm{H}} - \hat{\mathcal{W}}' \cdot_j \hat{\mathbf{V}}^{\mathrm{H}}\|_{\mathrm{F}} \leq \|\mathcal{W}' - \hat{\mathcal{W}}'\|_{\mathrm{F}}.$$

As an immediate consequence, letting $R_i^* = \max\{R_i, \hat{R}_i\}$ for each $i = 1, \ldots, \ell$ there exist orthogonal compressions $\mathcal{W}, \hat{\mathcal{W}} \in \mathbb{K}^{R_1^* \times \cdots \times R_\ell^*}$ of $\mathcal{W}'$ and $\hat{\mathcal{W}}'$, respectively, such that

$$\|\mathcal{W} - \hat{\mathcal{W}}\|_{\mathrm{F}} \leq \|\mathcal{W}' - \hat{\mathcal{W}}'\|_{\mathrm{F}}.$$

[10] In practice one rarely has access to both $\mathcal{T}'$ and $\hat{\mathcal{T}}'$, but this is immaterial to our present discussion.

As will become clear, we may use Theorem 4.1 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 4.2. *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* $\mathbf{U} \in \mathbb{K}^{I_1 \times I_1}$ *such that* $\operatorname{ran} \mathbf{M} \supseteq \operatorname{ran} \mathbf{UN}$ *and so that*

$$\|\mathbf{M} - \mathbf{UN}\|_{\mathrm{F}} \leq \|\mathbf{M} - \mathbf{N}\|_{\mathrm{F}}.$$

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

$$\|\mathbf{U}_{\mathbf{M}}^{\mathrm{H}}\mathbf{M} - \mathbf{U}_{\mathbf{N}}^{\mathrm{H}}\mathbf{N}\|_{\mathrm{F}} \leq \|\mathbf{M} - \mathbf{N}\|_{\mathrm{F}}.$$

Proof. The proof of Lemma 4.2 is a routine argument using the solution of the orthogonal Procrustes problem [39]. Details are given in the supplementary materials. ∎

We now give the proof of Theorem 4.1.

Proof. The main strategy of the proof is to apply Lemma 4.2 to each mode of the tensors $\mathcal{W}'$ and $\hat{\mathcal{W}}'$. To this end, fix j and let $\mathcal{W}'_{(j)}$ and $\hat{\mathcal{W}}'_{(j)}$ be mode-j unfoldings of $\mathcal{W}'$ and $\hat{\mathcal{W}}'$, respectively, with the same ordering of the fibers. Recall that $\mathcal{W}_{(j)}$ has rank R_j and $\hat{\mathcal{W}}'_{(j)}$ has rank $\hat{R}_j$, and WLOG assume $R_j \geq \hat{R}_j$.

Using Lemma 4.2, given any column-wise orthonormal matrix $\mathbf{V} \in \mathbb{R}^{I_j \times R_j}$ with range equal to ran $\mathcal{W}'_{(j)}$, there exists a 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 $\hat{\mathcal{W}}'_{(j)}$ and such that

$$\|\mathbf{V}^{\mathrm{H}}\mathcal{W}'_{(j)} - \hat{\mathbf{V}}^{\mathrm{H}}\hat{\mathcal{W}}'_{(j)}\|_{\mathrm{F}} \leq \|\mathcal{W}'_{(j)} - \hat{\mathcal{W}}'_{(j)}\|_{\mathrm{F}}.$$

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

$$\|\mathbf{V}^{\mathrm{H}}\mathcal{W}'_{(j)} - \hat{\mathbf{V}}^{\mathrm{H}}\hat{\mathcal{W}}'_{(j)}\|_{\mathrm{F}} = \|\mathcal{W}' \cdot_j \mathbf{V}^{\mathrm{H}} - \hat{\mathcal{W}}' \cdot_j \hat{\mathbf{V}}^{\mathrm{H}}\|_{\mathrm{F}} \qquad \text{and} \qquad \|\mathcal{W}'_{(j)} - \hat{\mathcal{W}}'_{(j)}\|_{\mathrm{F}} = \|\mathcal{W}' - \hat{\mathcal{W}}'\|_{\mathrm{F}}$$

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. ∎

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

$$\|\mathcal{W}^* - \hat{\mathcal{W}}^*\|_{\mathrm{F}} < \|\mathcal{W} - \hat{\mathcal{W}}\|_{\mathrm{F}} \leq \|\mathcal{W}' - \hat{\mathcal{W}}'\|_{\mathrm{F}}.$$

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 $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$. Here one wants to minimize $\|\mathcal{M}' - \mathcal{W}'\|$ subject to a multilinear rank constraint on $\mathcal{W}'$. A common approach is to compute a MLSVD of $\mathcal{M}'$ 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.

4.2. Proof of Proposition 2.4. We are now in position to prove Proposition 2.4.

Proof. The assumption that every border $\mathbb{K}$-rank R tensor in an open ϵ ball around $\mathcal{T}$ is simple in fact implies that this ϵ ball cannot contain a tensor with border $\mathbb{K}$-rank strictly less than R. We temporarily assume this implication holds and give a detailed proof at the end.

Let $\mathcal{M}' \in \mathbb{K}^{I_1 \times I_2 \times I_3}$ be any tensor such that $\|\mathcal{M}' - \mathcal{T}'\| < \epsilon/2$, and let $\hat{\mathcal{T}}'$ be a best border $\mathbb{K}$-rank R approximation of $\mathcal{M}'$. Note that $\mathcal{T}'$ has multilinear rank (R, R, K) by assumption, so inequality (2.5) shows that $\mathcal{T}'$ has border $\mathbb{K}$-rank R in addition to having $\mathbb{K}$-rank R. It follows that

$$\|\hat{\mathcal{T}}' - \mathcal{M}'\| \leq \|\mathcal{T}' - \mathcal{M}'\|, \qquad \text{hence} \qquad \|\mathcal{T}' - \hat{\mathcal{T}}'\| < \epsilon,$$

as otherwise $\mathcal{T}'$ would be a strictly better border $\mathbb{K}$-rank R approximation of $\mathcal{M}'$ than $\hat{\mathcal{T}}'$. To show that $\mathcal{M}'$ 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 4.1, we may simultaneously orthogonally compress $\mathcal{T}'$ and $\hat{\mathcal{T}}'$ to tensors of size $R \times R \times K^*$ for some $K^* \leq R$ without increasing the Frobenius distance between the tensors. Therefore we may WLOG we may assume $\mathcal{T}' = \mathcal{T}$ and $\hat{\mathcal{T}}' = \hat{\mathcal{T}}$ both have size $R \times R \times K^*$. After this restriction, our assumptions imply that $\hat{\mathcal{T}}$ is simple. An application of Theorem 2.3 (2) then shows that $\hat{\mathcal{T}}$ has rank R.

We next 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 border rank R tensor in this ball is simple. As previously mentioned, this assumption implies that the ϵ ball does not contain a tensor with border rank less than R, hence every rank R tensor in the ball is simple. Using Theorem 2.3 (1) together with Lemma 3.2, it follows that every tensor inside this ball satisfies Kruskal's condition [25], hence every tensor in this ball has a unique CPD.

We now show that the set of tensors contained in this open ball which do not have a unique best $\mathbb{K}$-rank R approximation has Lebesgue measure 0. For tensors with entries in $\mathbb{C}$ (equipped with Lebesgue measure over the complexes) this is an immediate consequence of [38]. For tensors with entries in $\mathbb{R}$ (equipped with Lebesgue measure over the reals) this follows from [14, Corollary 18] together with the first part of this result.

It remains to show that every tensor in an ϵ ball centered at $\mathcal{T}'$ has border $\mathbb{K}$-rank at least R. To this end, suppose towards a contradiction that there is some tensor $\mathcal{S}'$ such that $\|\mathcal{T}' - \mathcal{S}'\| < \epsilon$ and such that $\mathcal{S}'$ has border $\mathbb{K}$-rank $N < R$. Also fix $\delta_1, \delta_2 > 0$ such that $\|\mathcal{T}' - \mathcal{S}'\| + \delta_1 + \delta_2 < \epsilon$. Since $\mathcal{S}'$ has border $\mathbb{K}$-rank N, there must exist some tensor $\mathcal{W}' = [\![\mathbf{A}', \mathbf{B}', \mathbf{C}']\!]$ such that $\|\mathcal{W}' - \mathcal{S}'\| < \delta_1$ and such that $\mathbb{K}\text{-rank}(\mathcal{W}') = \mathbb{K}\text{-}\underline{\text{rank}}(\mathcal{W}') = N$. Furthermore, a routine argument shows that $\mathcal{W}$ can be chosen so that its factors $\mathbf{A}'$ and $\mathbf{B}'$ have full column rank.

One can then extend the factors $\mathbf{A}', \mathbf{B}', \mathbf{C}'$ to matrices

$$\tilde{\mathbf{A}}' = [\mathbf{A}' \;\; \mathbf{E}'_{\mathbf{A}}] \qquad \tilde{\mathbf{B}}' = [\mathbf{B}' \;\; \mathbf{E}'_{\mathbf{B}}] \qquad \tilde{\mathbf{C}}' = [\mathbf{C}' \;\; \mathbf{E}'_{\mathbf{C}}]$$

each having R columns where $\mathbf{E}'_{\mathbf{A}}, \mathbf{E}'_{\mathbf{B}}$ and $\mathbf{E}'_{\mathbf{C}}$ are chosen so that $\tilde{\mathbf{A}}'$ and $\tilde{\mathbf{B}}'$ have full column rank, $\tilde{\mathbf{C}}'$ has a repeated column, and $\|\mathcal{W}' - [\![\tilde{\mathbf{A}}', \tilde{\mathbf{B}}', \tilde{\mathbf{C}}']\!]\| < \delta_2$. Using Theorem 2.3 shows that the tensor $[\![\tilde{\mathbf{A}}', \tilde{\mathbf{B}}', \tilde{\mathbf{C}}']\!]$ has both $\mathbb{K}$-rank and $\mathbb{K}$-border rank

equal to R. Furthermore, by construction we have

$$\|\mathcal{T}' - [\![\tilde{\mathbf{A}}', \tilde{\mathbf{B}}', \tilde{\mathbf{C}}']\!]\| < \|\mathcal{T}' - \mathcal{S}'\| + \delta_1 + \delta_2 < \epsilon.$$

However, Theorem 2.3 also shows that this tensor is not simple, since $\tilde{\mathbf{C}}'$ has repeated columns. This contradicts our assumption that every border rank R tensor in an ϵ ball around $\mathcal{T}'$ is simple. ∎

4.3. Metrics for joint generalized eigenvalues. We now discuss in more detail the measures of distance we use for joint generalized eigenvalues. Recall the chordal metric between two one-dimensional subspaces $\mathscr{L}_1, \mathscr{L}_2 \subseteq \mathbb{K}^K$ is the sine of the angle between $\mathscr{L}_1$ and $\mathscr{L}_2$. Letting $\boldsymbol{\lambda}_i$ be a unit vector such that $\operatorname{span}(\boldsymbol{\lambda}_i) = \mathscr{L}_i$ for $i = 1, 2$, it is not difficult to show that the chordal metric may be computed as

$$\chi(\mathscr{L}_1, \mathscr{L}_2) = \|\mathbf{P}_{(\mathscr{L}_1)^\perp} \boldsymbol{\lambda}_2\|_2 = \max_{\substack{\boldsymbol{\gamma} \in (\mathscr{L}_1)^\perp \\ \|\boldsymbol{\gamma}\|_2 = 1}} |\langle \boldsymbol{\lambda}_2, \boldsymbol{\gamma} \rangle| = \sqrt{1 - |\langle \boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2 \rangle|^2}. \tag{4.2}$$

Here $(\mathscr{L}_1)^\perp \subset \mathbb{K}^K$ is the orthogonal complement of $\mathscr{L}_1 \subset \mathbb{K}^K$ and $\mathbf{P}_{(\mathscr{L}_1)^\perp}$ denotes the projection onto $(\mathscr{L}_1)^\perp$. Note that $|\langle \boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2 \rangle|$ is the cosine of the angle between the subspaces $\mathscr{L}_1$ and $\mathscr{L}_2$.

Also recall the matching distance defined in equation (2.10) which gives a metric between the spectra of two tensors. In addition to the matching distance we will need a weaker measure of separation between the spectra of two tensors. To this end, 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 $\{\mathscr{L}_{i,k}\}_{k=1}^R$ be the spectrum of $\mathcal{T}_i$. The **spectral variation** from $\mathcal{T}_1$ to $\mathcal{T}_2$, denoted $\operatorname{sv}[\mathcal{T}_1, \mathcal{T}_2]$, is defined by

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

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

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

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

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

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.

4.4. Bauer–Fike Theorem for JGE values. We now present our main perturbation theoretic bound for the spectral variation between two tensors of $\mathbb{K}$-rank R. The bound we present is in essence determined by the spectral norm of (a transformation of) the error tensor.

THEOREM 4.3. *Let $\mathcal{T}$ and $\mathcal{W}$ 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 $\{\mathscr{L}_k\}_{k=1}^K$ be the spectrum of $\mathcal{T}$ and set $\mathcal{E} = \mathcal{T} - \mathcal{W}$. If $\mathscr{W}$ is a JGE value of $\mathcal{W}$ with geometric multiplicity at least one, then*

$$\min_k \chi(\mathscr{L}_k, \mathscr{W}) \leq \sqrt{R} \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\mathscr{W}^\perp} \right\|_{\text{sp}} \leq \frac{\sqrt{R} \|\mathcal{E}\|_{\text{sp}}}{\sigma_{\min}(\mathbf{A}) \sigma_{\min}(\mathbf{B})}. \tag{4.3}$$

Here $\sigma_{\min}(\mathbf{A})$ and $\sigma_{\min}(\mathbf{B})$ denote the smallest[11] singular values of $\mathbf{A}$ and $\mathbf{B}$, respectively and $\mathbf{P}_{\mathscr{W}^\perp}$ denotes the projection onto the orthogonal complement of $\mathscr{W}$. As a consequence, if $\mathcal{W}$ is a slice mix invertible border $\mathbb{K}$-rank R tensor, then

$$\text{sv}[\mathcal{T}, \mathcal{W}] \leq \sqrt{R} \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\text{sp}} \leq \frac{\sqrt{R} \|\mathcal{E}\|_{\text{sp}}}{\sigma_{\min}(\mathbf{A}) \sigma_{\min}(\mathbf{B})}. \tag{4.4}$$

In the following special cases the coefficient $\sqrt{R}$ may be dropped from equations (4.3) and (4.4):

1. *$K = 2$.*
2. *$\mathcal{W}$ is also slice mix invertible and has $\mathbb{K}$-rank R with CPD $[\![\mathbf{A}_\mathcal{W}, \mathbf{B}_\mathcal{W}, \mathbf{C}_\mathcal{W}]\!]$ where $\mathbf{A}_\mathcal{W} = \mathbf{A}$ or $\mathbf{B}_\mathcal{W} = \mathbf{B}$.*

That is, in these special cases one has

$$\text{sv}[\mathcal{T}, \mathcal{W}] \leq \left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\text{sp}} \leq \frac{\|\mathcal{E}\|_{\text{sp}}}{\sigma_{\min}(\mathbf{A}) \sigma_{\min}(\mathbf{B})}. \tag{4.5}$$

In the general case, the proof of Theorem 4.3 follows a similar strategy to the proof of [46, Chapter VI, Theorem 2.7] with modifications to address issues arising when $K > 2$. In particular, [46, Chapter VI, Theorem 2.7] shows

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

in the $K = 2$ case. As the proof in the general case is similar to Stewart and Sun's argument, we give it in the supplementary materials. However, if we additionally assume that $\mathcal{W}$ has $\mathbb{K}$-rank R, then a new, simplified proof which is more intuitive can be given. This we now describe.

Proof. Assume that $\mathcal{W} \in \mathbb{K}^{R \times R \times R}$ has rank R and is slice mix invertible and let $[\![\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]\!]$ be a CPD of $\mathcal{W} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}$ where the columns of $\tilde{\mathbf{C}}$ have unit norm.[12] Additionally set $\mathcal{E} := \mathcal{W} - \mathcal{T}$. As a consequence of Theorem 2.3 (1), the spans of the columns of $\mathbf{C}$ and $\tilde{\mathbf{C}}$ give the JGE values of $\mathcal{T}$ and $\mathcal{W}$, respectively. For each column $\tilde{\mathbf{c}}_j$ of $\tilde{\mathbf{C}}$ we will show that the spectral norm of $\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}$ gives an upper bound for

$$\min_r \chi(\tilde{\mathbf{c}}_j, \mathbf{c}_r) = \min_r \|\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \mathbf{c}_r\|_2$$

[11] Lemma 3.1 shows that $\mathbf{A}$ and $\mathbf{B}$ are invertible, hence these singular values are nonzero.

[12] Up to permutation and scaling one has $\tilde{\mathbf{A}} = \mathbf{A}^{-1}\mathbf{A}_\mathcal{W}$ and $\tilde{\mathbf{B}} = \mathbf{B}^{-1}\mathbf{B}_\mathcal{W}$ and $\tilde{\mathbf{C}} = \mathbf{C}_\mathcal{W}$.

from which the result follows. Here $\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}$ is the orthogonal projection on the orthogonal complement of $\operatorname{span}(\tilde{\mathbf{c}}_j)$.

Expanding out the spectral norm of the transformed error tensor, we have

(4.6)
$$\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}\|_{\mathrm{sp}} = \max_{\|\mathbf{u}\|_2=\|\mathbf{v}\|_2=\|\mathbf{z}\|_2=1} \left| \left(\sum_{r\neq j} \tilde{\mathbf{a}}_r \otimes \tilde{\mathbf{b}}_r \otimes \left(\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \tilde{\mathbf{c}}_r\right) - \sum_{r=1}^{R} \mathbf{e}_r \otimes \mathbf{e}_r \otimes \left(\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \mathbf{c}_r\right) \right) \cdot_1 \mathbf{u} \cdot_2 \mathbf{v} \cdot_3 \mathbf{z} \right|.$$

Note that since $\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \tilde{\mathbf{c}}_j = \mathbf{0}$, the tensor $\tilde{\mathcal{T}} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \cdot_3 \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}$ which appears on the left hand side of the difference in equation (4.6) has rank at most $R-1$. It follows that there is some vector $\mathbf{y} \in \mathbb{K}^R$ which is orthogonal to the subspace spanned by the mode-2 fibers of this tensor. In particular, let $\mathbf{y} \in \mathbb{K}^R$ be a unit vector such that $\langle \mathbf{y}, \tilde{\mathbf{b}}_r \rangle = 0$ for all $r \neq j$. With this choice of $\mathbf{y}$ we have

$$\sum_{r\neq j} \langle \mathbf{y}, \tilde{\mathbf{b}}_r \rangle \tilde{\mathbf{a}}_r \otimes \left(\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \tilde{\mathbf{c}}_r\right) = 0$$

from which it follows that

(4.7)
$$\begin{aligned} \|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\mathbf{y}^{\mathrm{H}} \cdot_3 \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}\|_{\mathrm{sp}} &= \max_{\|\mathbf{u}\|_2=\|\mathbf{z}\|_2=1} \left| \left(\textstyle\sum_{r=1}^{R} \langle \mathbf{y}, \mathbf{e}_r \rangle \mathbf{e}_r \otimes \left(\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \mathbf{c}_r\right) \right) \cdot_1 \mathbf{u} \cdot_3 \mathbf{z} \right| \\ &= \left\| \textstyle\sum_{r=1}^{R} \langle \mathbf{y}, \mathbf{e}_r \rangle \mathbf{e}_r \otimes \left(\mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \mathbf{c}_r\right) \right\|_{\mathrm{sp}} \\ &\geq \max_r \left\| \langle \mathbf{y}, \mathbf{e}_r \rangle \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \mathbf{c}_r \right\|_2. \end{aligned}$$

Recalling that $\mathbf{y}$ is a unit vector, there must exist at least one index r_* such that $|\langle \mathbf{y}, \mathbf{e}_r \rangle| \geq 1/\sqrt{R}$, from which we obtain

$$\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\mathbf{y}^{\mathrm{H}} \cdot_3 \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}\|_{\mathrm{sp}} \geq \|\langle \mathbf{y}, \mathbf{e}_{r_*} \rangle \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp} \mathbf{c}_{r_*}\|_2 \geq \frac{\chi(\tilde{\mathbf{c}}_j, \mathbf{c}_{r_*})}{\sqrt{R}} \geq \frac{\min_r \chi(\tilde{\mathbf{c}}_j, \mathbf{c}_r)}{\sqrt{R}}.$$

We conclude

$$\sqrt{R}\|\mathbf{A}^{-1}\|_2 \|\mathbf{B}^{-1}\|_2 \|\mathcal{E}\|_{\mathrm{sp}} \geq \sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\mathbf{y}^{\mathrm{H}} \cdot_3 \mathbf{P}_{\tilde{\mathbf{c}}_j^\perp}\|_{\mathrm{sp}} \geq \min_r \chi(\tilde{\mathbf{c}}_j, \mathbf{c}_r).$$

The result in this special case follows from combining the above inequality with Theorem 2.3 (1). If additionally one has $\mathbf{B}_{\mathcal{W}} = \mathbf{B}$, then one can take $\tilde{\mathbf{B}} = \mathbf{I}$. The fact that the coefficient $\sqrt{R}$ may be dropped in this case easily follows. The case $\mathbf{A}_{\mathcal{W}} = \mathbf{A}$ is similar. ∎

4.4.1. Scaling of A and B. The quality of the bound obtained in Theorem 4.3 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 (4.4), 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 (4.4).

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.

4.4.2. On the coefficient $\sqrt{R}$. Knowing that there are various special cases in which the coefficient $\sqrt{R}$ may be dropped from equation (4.4), 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 (4.4) does not hold for arbitrary pairs of slice mix invertible $\mathbb{K}$-rank R tensors if the coefficient $\sqrt{R}$ is removed.

Example 4.2. 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 $\mathcal{W}$ 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 $\mathcal{W}$ are both slice mix invertible $\mathbb{R}$-rank 3 tensors and the spectral variation $\mathrm{sv}[\mathcal{T}, \mathcal{W}]$ 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 $\mathcal{E} = \mathcal{W} - \mathcal{T}$ one has

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

illustrating that equation (4.4) does not hold for this pair[13].

4.5. Matching distance bounds and the $K = 2$ case. While the spectral variation bound from Theorem 4.3 provides insight as to how a CPD is affected by a perturbation, as illustrated by Example 4.1, 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 use a matching distance bound such as Theorem 2.8.

As briefly mentioned before, a spectral variation bound can often be converted to a matching distance bound using the following idea: Supposing temporarily that $\mathcal{W}$ is slice mix invertible and has border $\mathbb{K}$-rank R and letting $\{\mathscr{L}_r\}_{r=1}^R$ be the spectrum of $\mathcal{T}$, the bound given in Theorem 4.3 gives a collection of neighborhoods $\{\mathfrak{N}_\epsilon(\mathscr{L}_r)\}_{r=1}^R$ in which the JGE values of $\mathcal{W}$ must lie. Additionally, it is not difficult to show that JGE values are continuous in a perturbation along the set of slice mix invertible border rank R tensors.[14]

More precisely, if $\mathfrak{T} : [0,1] \to \mathbb{K}^{R \times R \times K}$ is a continuous function such that $\mathfrak{T}(0) = \mathcal{T}$ and $\mathfrak{T}(t)$ has border $\mathbb{K}$-rank R and is slice mix invertible for all $t \in [0,1]$, then there exists a continuous function $\mathfrak{L} : [0,1] \to [\mathrm{Gr}(1, \mathbb{K}^K)]^R$ such that $\mathfrak{L}(t)$ is equal to the spectrum of $\mathfrak{T}(t)$ for all $t \in [0,1]$. Here $\mathrm{Gr}(1, \mathbb{K}^K)$ denotes the set of one-dimensional subspaces of $\mathbb{K}^K$ and $[\mathrm{Gr}(1, \mathbb{K}^K)]^R$ denotes a collection of R elements of $\mathrm{Gr}(1, \mathbb{K}^K)$ with repetition allowed.[15]

[13] Computation of the spectral norm is NP-hard [20]; 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 [53] suggests $\left\| \mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1} \right\|_{\mathrm{sp}} \approx 0.882$.

[14] For example, one may use the fact that the characteristic polynomial of a slice mix invertible border $\mathbb{K}$-rank R tensor factors as a product of R linear terms to show this.

[15] The collection of one-dimensional subspaces of $\mathbb{K}^K$ is a metric space under the chordal metric. In general, the collection of d-dimensional subpsaces of $\mathbb{K}^K$ is called a Grassmanian.

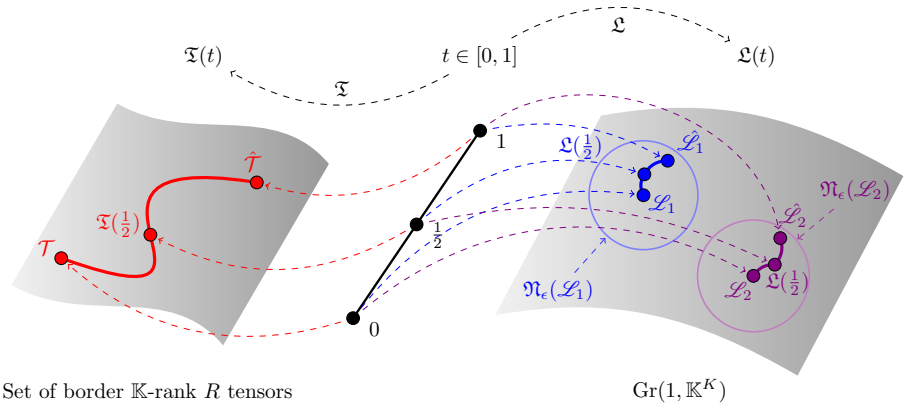


FIG. 1. *Illustration of continuous functions $\mathfrak{T}$ and $\mathfrak{L}$ defined on $[0,1]$. The function $\mathfrak{T}$ maps to a path of border $\mathbb{K}$-rank R tensors from $\mathcal{T}$ to $\hat{\mathcal{T}}$. The function $\mathfrak{L}$ is defined so that $\mathfrak{L}(t)$ is equal to the joint generalized spectrum of $\mathfrak{T}(t)$ for all $t \in [0,1]$. All JGE values of $\hat{\mathcal{T}}$ are contained in the union of the regions $\mathfrak{N}_\epsilon(\mathscr{L}_1)$ and $\mathfrak{N}_\epsilon(\mathscr{L}_2)$. Furthermore the neighborhood $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ contains a JGE value of $\hat{\mathcal{T}}$ for each $r = 1, 2$.*

The above discussion implies that if $\mathcal{W}$ is slice mix invertible and has border $\mathbb{K}$-rank R and if there is a path of slice mix invertible border $\mathbb{K}$-rank R tensors from $\mathcal{T}$ to $\mathcal{W}$ lying in the ball of radius $\|\mathcal{T} - \mathcal{W}\|_{\text{sp}}$ (i.e. if $\mathfrak{T}$ can be chosen so $\mathfrak{T}(0) = \mathcal{T}$ and $\mathfrak{T}(1) = \mathcal{W}$ and $\|\mathcal{T} - \mathfrak{T}(t)\|_{\text{sp}} \leq \|\mathcal{T} - \mathcal{W}\|_{\text{sp}}$ for all $t \in [0,1]$) and if the $\mathfrak{N}_\epsilon(\mathscr{L}_i) \cap \mathfrak{N}_\epsilon(\mathscr{L}_j) = \emptyset$ for each $i \neq j$, then each $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ must contain exactly one JGE value of $\mathcal{W}$. As a consequence, if these assumptions hold, then the bound given by 4.3 is in fact a matching distance bound and may be used to guarantee the existence of distinct JGE values for $\mathcal{W}$. See Figure 1 for an illustration of the functions $\mathfrak{T}$ and $\mathfrak{L}$ and of the neighborhoods $\mathfrak{N}_\epsilon(\mathscr{L}_r)$.

The primary issue with using such an argument for general K is that, it is not known if there is a continuous function $\mathfrak{T}$ which satisfies the necessary assumptions. That is, it is not known if there is a path of slice mix invertible border rank R tensors from $\mathcal{T}$ to $\mathcal{W}$ which lies in the ball of radius $\|\mathcal{T} - \mathcal{W}\|_{\text{sp}}$ centered at $\mathcal{T}$. However, in the case that $K = 2$, the condition needed to guarantee that the $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ do not intersect in fact guarantees the existence of such a path.

As discussed in [46, Chapter VI], if $\mathcal{T}$ is a slice mix invertible 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 slice mix invertible $R \times R \times 2$ tensors. Namely, a slice mix invertible $R \times R \times 2$ 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 4.4. *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 $\{\mathscr{L}_r\}_{r=1}^R$ 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

$$\|\mathcal{E}\|_{\text{sp}} \leq \frac{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B}) \min_{i\neq j} \chi(\mathscr{L}_i, \mathscr{L}_j)}{2}.$$

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

Proof. To simplify notation, set

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

By assumption the columns of $\mathbf{C}$ all have unit norm, therefore, by reindexing columns if necessary, Theorem 2.3 shows that for each $r = 1, \ldots, R$ the column $\mathbf{c}_r$ satisfies $\mathscr{L}_r = \text{span}(\mathbf{c}_r)$. To show $\mathcal{T} + \mathcal{E}$ is slice mix invertible, we first show that there is a unit vector $\mathbf{q} \in \mathbb{K}^2$ such that $|\langle \mathbf{c}_r, \mathbf{q}\rangle| \geq \delta$ for all r. Using this we will show that the matrix $\bar{q}_1(\mathbf{T}_1 + \mathbf{E}_1) + \bar{q}_2(\mathbf{T}_2 + \mathbf{E}_2)$ is invertible, hence $\mathcal{T} + \mathcal{E}$ is slice mix invertible.

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

$$\chi(\mathscr{L}_r, \mathscr{Z}) \geq \delta \qquad \text{for all } r = 1, \ldots, R, \tag{4.8}$$

as otherwise would imply that there is an index $\ell \neq 1$ such that

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

which is a contradiction.

Now set $\mathscr{Q} = \mathscr{Z}^\perp \subset \mathbb{K}^2$ and let $\mathbf{q} \in \mathbb{K}^2$ be a unit vector such that $\mathscr{Q} = \text{span}(\mathbf{q})$. Using equation (4.8), for each r we have

$$\delta \leq \chi(\mathscr{L}_\ell, \mathscr{Z}) = \sin(\angle(\mathscr{L}_r, \mathscr{Z})) = \sin(\pi/2 - \angle(\mathscr{L}_r, \mathscr{Q})) = \cos(\angle(\mathscr{L}_r, \mathscr{Q})) = |\langle \mathbf{c}_r, \mathbf{q}\rangle|$$

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 $\bar{q}_1 D_1(\mathbf{C}) + \bar{q}_2 D_2(\mathbf{C})$ is a diagonal matrix with rth diagonal entry equal to $\langle \mathbf{c}_r, \mathbf{q}\rangle$. If follows that

$$\sigma_{\min}(\bar{q}_1 D_1(\mathbf{C}) + \bar{q}_2 D_2(\mathbf{C})) = \min(\{|\langle \mathbf{c}_r, \mathbf{q}\rangle|\}_{r=1}^R) \geq \delta. \tag{4.9}$$

Finally observe that

$$\begin{aligned}
\sigma_{\min}\big(\bar{q}_1(\mathbf{T}_1 + \mathbf{E}_1) + \bar{q}_2(\mathbf{T}_2 + \mathbf{E}_2)\big) &\geq \sigma_{\min}(\bar{q}_1\mathbf{T}_1 + \bar{q}_2\mathbf{T}_2) - \|\bar{q}_1\mathbf{E}_1 + \bar{q}_2\mathbf{E}_2\|_2 \\
&\geq \sigma_{\min}\Big(\mathbf{A}\Big(\bar{q}_1 D_1\big(\mathbf{C}\big) + \bar{q}_2 D_2(\mathbf{C})\Big)\mathbf{B}^{\mathrm{T}}\Big) - \|\mathcal{E}\|_{\text{sp}} \\
&\geq \sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})\sigma_{\min}\big(\bar{q}_1 D_1\big(\mathbf{C}\big) + \bar{q}_2 D_2(\mathbf{C})\big) - \|\mathcal{E}\|_{\text{sp}}
\end{aligned}$$

Using inequality (4.9) together with the assumption that $\|\mathcal{E}\|_{\text{sp}} < \sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})\delta$ shows that $\mathcal{T} + \mathcal{E}$ is slice mix invertible.

Having shown $\mathcal{T} + \mathcal{E}$ is slice mix invertible, it follows that $\mathcal{T} + \mathcal{E}$ has multilinear rank $(R, R, \cdot)$. From equation (2.5) we conclude that the border rank of $\mathcal{T} + \mathcal{E}$ is at least R. Combining this with the well-known fact that tensors of size $R \times R \times 2$ have border rank less than or equal to R, e.g. see [28, Theorem 5.5.1.1], shows that $\mathcal{T} + \mathcal{E}$ has border rank R. ∎

4.6. Proof of Theorem 2.8. We now give the proof of Theorem 2.8.

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

$$\epsilon = \frac{\|\mathcal{T} - \mathcal{W}\|_{\text{sp}}}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})} < \frac{\min_{i \neq j} \chi(\mathscr{L}_i, \mathscr{L}_j)}{2} \tag{4.10}$$

and for each $r = 1, \ldots, R$ let $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ be the collection of one-dimensional subspaces in $\mathbb{C}^2$ defined by

$$\mathfrak{N}_\epsilon(\mathscr{L}_r) = \{\mathscr{Z} \subset \mathbb{C}^2 |\ \dim(\mathscr{Z}) = 1 \text{ and } \chi(\mathscr{L}_r, \mathscr{Z}) \leq \epsilon\}.$$

That is, $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ is the ball of chordal radius ϵ centered at $\mathscr{L}_r$. Then our choice of ϵ together with the triangle inequality gives

$$\mathfrak{N}_\epsilon(\mathscr{L}_i) \cap \mathfrak{N}_\epsilon(\mathscr{L}_j) = \emptyset \qquad \text{for all } i \neq j.$$

To prove inequality (2.11), it is sufficient show that each $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ contains exactly one JGE value of $\mathcal{W}$.

To this end let $\mathcal{E} = \mathcal{W} - \mathcal{T}$ and for $\alpha \in [0, 1]$ set $\mathcal{T}^\alpha = \mathcal{T} + \alpha\mathcal{E}$ so that $\mathcal{T} = \mathcal{T}^0$ and $\mathcal{W} = \mathcal{T}^1$. Note that Proposition 4.4 shows that $\mathcal{T}^\alpha$ is slice mix invertible and has border $\mathbb{C}$-rank R for each $\alpha \in [0, 1]$. It follows from Theorem 2.3 that each $\mathcal{T}^\alpha$ has R JGE values in $\mathbb{C}^2$ counting algebraic multiplicity.

Now define

$$\mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_r) = \{\mathscr{Z} \subset \mathbb{C}^2 |\ \dim(\mathscr{Z}) = 1 \text{ and } \chi(\mathscr{L}_r, \mathscr{Z}) \leq \alpha\epsilon\}.$$

Then $\mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_r) \subset \mathfrak{N}_\epsilon(\mathscr{L}_r)$ for each $\alpha \in [0, 1]$ and r so $\mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_i) \cap \mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_j) = \emptyset$ whenever $i \neq j$. Furthermore, using Theorem 4.3 we have

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

It follows that the spectrum of $\mathcal{T}^\alpha$ is contained in $\cup_{r=1}^R \mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_r)$ for each $\alpha \in [0, 1]$.

In particular $\mathfrak{N}_0(\mathscr{L}_r) = \{\mathscr{L}_r\}$ contains exactly one JGE value of $\mathcal{T}^0 = \mathcal{T}$ for each $r = 1, \ldots, R$. Since $\mathcal{T}^\alpha$ is slice mix invertible for all $\alpha \in [0, 1]$, [46, 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 $\mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_r)$ remain disjoint as α increases from 0 to 1, the continuity of the JGE values of $\mathcal{T}^\alpha$ in α implies that each $\mathfrak{N}_{\alpha\epsilon}(\mathscr{L}_r)$ 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, $\mathfrak{N}_{1\epsilon}(\mathscr{L}_r) = \mathfrak{N}_\epsilon(\mathscr{L}_r)$ contains exactly one eigenvalue of $\mathcal{T}^1 = \mathcal{W}$ for each r. By definition, each $\mathfrak{N}_\epsilon(\mathscr{L}_r)$ is a ball of chordal radius ϵ centered $\mathscr{L}_r$, so we conclude

$$\text{md}[\mathcal{T}, \mathcal{W}] \leq \epsilon = \frac{\|\mathcal{T} - \mathcal{W}\|_{\text{sp}}}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})},$$

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

It remains to show that $\mathcal{W}$ has $\mathbb{C}$-rank R. As previously discussed, all JGE values of a slice mix invertible $R \times R \times 2$ tensor must have geometric multiplicity at least 1. It follows that $\mathcal{W}$ is simple, so Theorem 2.3 (2) shows that $\mathcal{W}$ has $\mathbb{C}$-rank R.

Now assume that $\mathbb{K} = \mathbb{R}$ and that $\mathcal{T}$ and $\mathcal{W}$ have real entries. Applying the result over the complex field shows that $\mathcal{W}$ must have distinct JGE values over $\mathbb{C}$ so it is

sufficient to show that the JGE values of $\mathcal{W}$, hence the JGE vectors of $\mathcal{W}$, must all be real. Here, when we say that $\mathcal{W}$ 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 $(\mathbf{W}_1, \mathbf{W}_2)$ are real valued.

Suppose towards a contradiction that $\mathcal{W}$ 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 2.8, 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 $\mathcal{W}$ are all real. It follows from Theorem 2.3 (1) that $\mathcal{W}$ has $\mathbb{R}$-rank R. ∎

5. 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 5.1. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a slice mix invertible 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 5.2. *Let $\mathcal{T} \in \mathbb{K}^{R\times R\times K}$ be a slice mix invertible tensor and let* $\operatorname{span}(\boldsymbol{\lambda})$ *be a JGE value of $\mathcal{T}$. Also let $\mathbf{U} \in \mathbb{K}^{K\times K}$ be an invertible matrix and set $\mathcal{S} = \mathcal{T} \cdot_3 \mathbf{U}$. Then $\mathcal{S}$ is slice mix invertible and* $\operatorname{span}(\mathbf{U}\boldsymbol{\lambda})$ *is a JGE value of $\mathcal{S}$ which satisfies* $\operatorname{am}(\boldsymbol{\lambda}) = \operatorname{am}(\mathbf{U}\boldsymbol{\lambda})$ *where* $\operatorname{span}(\boldsymbol{\lambda})$ *and* $\operatorname{span}(\mathbf{U}\boldsymbol{\lambda})$ *are treated as JGE values of $\mathcal{T}$ and $\mathcal{S}$, respectively.*

Furthermore, if $\operatorname{span}(\boldsymbol{\lambda})$ *has a corresponding JGE vector* $\mathbf{x}$*, then* $(\operatorname{span}(\mathbf{U}\boldsymbol{\lambda}), \mathbf{x})$ *is a JGE pair of $\mathcal{S}$. As a consequence one has* $\operatorname{gm}(\boldsymbol{\lambda}) = \operatorname{gm}(\mathbf{U}\boldsymbol{\lambda})$.

Proof. Assume $\operatorname{span}(\boldsymbol{\lambda})$ is a JGE value of $\mathcal{T}$ with algebraic multiplicity m. View $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_K)$ 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 slice mix invertible. Now set $\mathbf{U}^T\boldsymbol{\gamma} = \eta$. By assumption $p_{\mathcal{T}}(\eta) = p_{\mathcal{T}\cdot_3\mathbf{U}}(\boldsymbol{\gamma})$ factors as

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

so $\operatorname{span}(\mathbf{U}\boldsymbol{\lambda})$ 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 5.1.

Proof. By assumption $\mathcal{T}$ is slice mix invertible and has a JGE value of algebraic multiplicity m, so Lemma 5.2 shows that $\mathcal{S}$ is slice mix invertible and has a JGE value with algebraic multiplicity m. Say this JGE value is equal to $\operatorname{span}(\boldsymbol{\lambda})$. 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

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

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

5.1. Proof of Theorem 2.5. 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 $\mathcal{M}' \in \mathbb{K}^{I_1\times I_2\times I_3}$ be any tensor such that

$$\|\mathcal{T}' - \mathcal{M}'\|_{\mathrm{F}} < \epsilon/2$$

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

$$\|\hat{\mathcal{T}}' - \mathcal{M}'\|_{\mathrm{F}} \le \|\mathcal{T}' - \mathcal{M}'\|_{\mathrm{F}} \qquad \text{hence} \qquad \|\hat{\mathcal{T}}' - \mathcal{T}'\|_{\mathrm{F}} < \epsilon.$$

Since $\hat{\mathcal{T}}'$ has border $\mathbb{K}$-rank R, equation (2.5) shows that the multilinear rank of $\hat{\mathcal{T}}'$ must be $(R,R,\hat{K})$ for some $\hat{K} \le R$. It is easy to check that a tensor has $\mathbb{K}$-rank R if and only if an orthogonal compression of the tensor has $\mathbb{K}$-rank R, so using Theorem 4.1, we may WLOG assume that $\mathcal{T}' = \mathcal{T}$ and $\hat{\mathcal{T}}' = \hat{\mathcal{T}}$ both have size $R \times R \times \max\{K,\hat{K}\}$. There is little difference in the $K \ge \hat{K}$ and $K < \hat{K}$ cases, so for ease of notation we assume $K \ge \hat{K}$.

Now let $\mathbf{U} \in \mathbb{K}^{K\times K}$ be a 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 $\|\mathcal{S} - \hat{\mathcal{S}}\|_{\mathrm{F}} = \|\mathcal{T} - \hat{\mathcal{T}}\|_{\mathrm{F}} < \epsilon$. Therefore, for some $i = 1,\dots,\lfloor K/2 \rfloor$ we must have that

$$\|(\mathbf{S}_{2i-1},\mathbf{S}_{2i}) - (\hat{\mathbf{S}}_{2i-1},\hat{\mathbf{S}}_{2i})\|_{\mathrm{F}} < \epsilon_i,$$

as otherwise would imply

$$\|\mathcal{S} - \hat{\mathcal{S}}\|_{\mathrm{F}}^2 \ge \sum_{k=1}^{\lfloor K/2 \rfloor} \|(\mathbf{S}_{2k-1},\mathbf{S}_{2k}) - (\hat{\mathbf{S}}_{2k-1},\hat{\mathbf{S}}_{2k})\|_{\mathrm{F}}^2 \ge \sum_{k=1}^{\lfloor K/2 \rfloor} \epsilon_k^2 = \epsilon^2.$$

In particular, from our choice of ϵ_i the subpencil $(\hat{\mathbf{S}}_{2i-1},\hat{\mathbf{S}}_{2i})$ must be simple. Using Proposition 5.1 allows us to conclude that the tensor $\hat{\mathcal{T}}$ is simple. The result then follows from Proposition 2.4. ∎

One may notice that there is no guarantee that the tensor $\mathcal{T} \cdot_3 \mathbf{U}$ has a slice mix invertible subpencil for a particular choice of $\mathbf{U}$, hence the bound given by Theorem 2.5 may be equal to 0 for a particular choice of $\mathbf{U}$. This does not cause serious difficulty: As a consequence of [11, Section 4.4 (ii)], if $\mathcal{T}$ is slice mix invertible then for generic unitaries every subpencil of $\mathcal{T}$ is slice mix invertible, 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 2.5. 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 2.8 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 $\mathbf{E}_\mathbf{U}$ so that there is at most one mark in each row and column of $\mathbf{E}_\mathbf{U}$ 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 [51]; for our purposes and for small K, we may nevertheless consider trying all possible pencil collections/paths.

6. 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 4.3 and the existence bounds obtained in Theorem 2.5 and Theorem 2.7. In addition we illustrate the effect of a tensor Procrustes calculation as in Theorem 4.1 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 signal tensor $\mathcal{T}'$ to a noise tensor $\mathcal{N}'$ to be $20 \log_{10}\left(\frac{\|\mathcal{T}'\|_\mathrm{F}}{\|\mathcal{N}'\|_\mathrm{F}}\right)$. The experiments discussed in this section all restrict to the real setting. That is, we restrict to considering real valued tensors and real rank.

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.

6.1. Observed spectral variation vs. spectral variation bound. We first examine the spectral variation bound obtained from the spectral norm[16] of (a transformation of) the error tensor in Theorem 4.3 and compare it to the observed spectral variation occurring under a perturbation with a specified SNR. Three cases are

[16] As previously mentioned, computation of the spectral norm is NP-hard. Using Tensorlab [53], 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 to a tensor is equal to the spectral norm of the tensor [13].

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 (4.5). 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 present the mean observed spectral variation and either the mean computed spectral variation bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-2}\|_{\text{sp}}$ or, when appropriate, the mean computed structured spectral variation bound $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-2}\|_{\text{sp}}$.

6.1.1. Structured Perturbations. Figure 2 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 $\hat{\mathcal{T}} := [\![\mathbf{A}, \mathbf{B}, \mathbf{C} + \mathbf{E}]\!]$ in addition to the mean computed value of $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$. Here $\mathcal{E} := \hat{\mathcal{T}} - \mathcal{T}$. This experiment illustrates the first inequality in equation (4.5).

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.

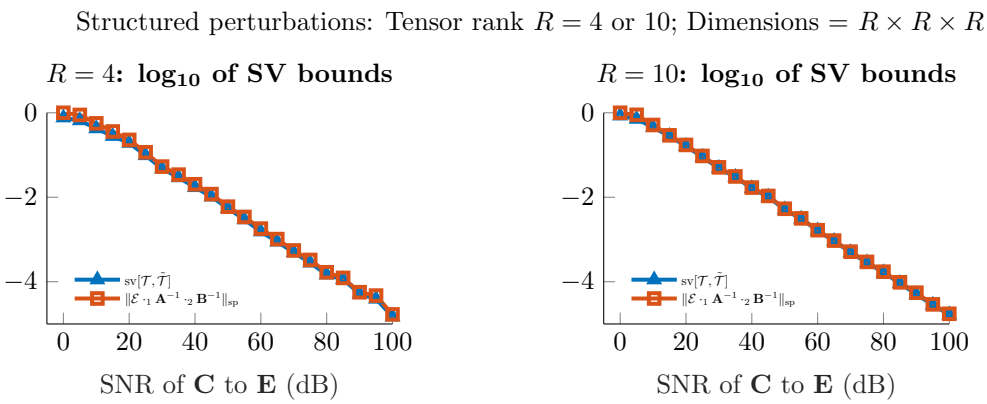


FIG. 2. *Observed spectral variation and computed bound between tensors of the form* $[\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$ *and* $[\![\mathbf{A}, \mathbf{B}, \mathbf{C} + \mathbf{E}]\!]$ *with rank* $R = 4$ *or* 10. *The structured bound* $\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\text{sp}}$ *is observed to be highly predictive of the observed spectral variation.*

6.1.2. Generic Perturbations: Rank equal to dimensions. Figure 3 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 [53]. 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 (4.4).

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.

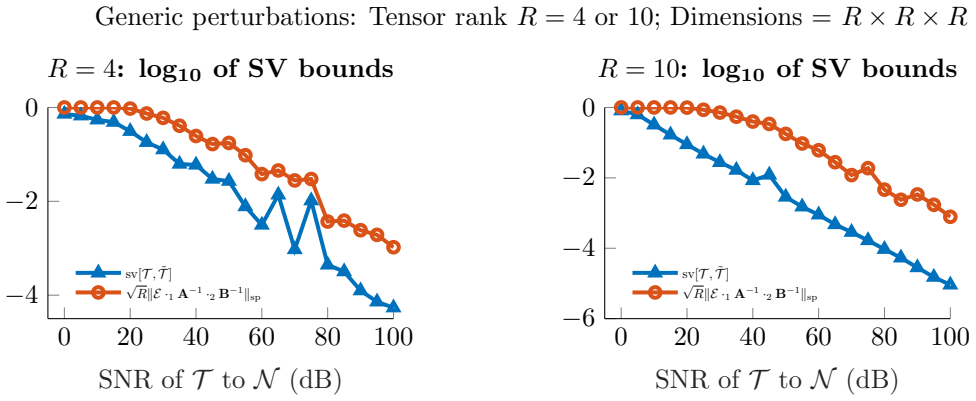


FIG. 3. *Observed spectral variation and computed bound for a rank R tensor $\mathcal{T}$ and a rank R approximation to $\mathcal{T}+\mathcal{N}$ where $R = 4$ or 10. The bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\mathrm{sp}}$ typically exceeds the observed spectral 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 $\log_{10}(\sqrt{10}) = 0.5$.*

6.1.3. Generic Perturbations: Rank less than dimensions. Figures 4 and 5 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 $\mathcal{T}'$ 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.

Let $\hat{\mathcal{T}}'$ denote a rank R approximation of $\mathcal{T}' + \mathcal{N}'$. To directly apply the spectral variation bound obtained from the spectral norm in Theorem 4.3 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 $\mathcal{T}$ and $\hat{\mathcal{T}}$ of size $R \times R \times R$ satisfying $\|\mathcal{T} - \hat{\mathcal{T}}\|_{\mathrm{F}} \leq \|\mathcal{T}' - \hat{\mathcal{T}}'\|_{\mathrm{F}}$ as described in the tensor Procrustes problem of Theorem 4.1. Write $\mathcal{T} = [\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$. For each SNR we report the mean observed spectral variation between the core tensors $\mathcal{T}$ and $\hat{\mathcal{T}}$ as well as the mean computed value for $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\mathrm{sp}}$ where $\mathcal{E} := \hat{\mathcal{T}} - \mathcal{T}$. For each choice of I and R the mean is computed over 20 trials.

In this setting we find that the error bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\mathrm{sp}}$ from the first inequality in equation (4.4) 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 [52] 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 6.3 for experiments in this direction.

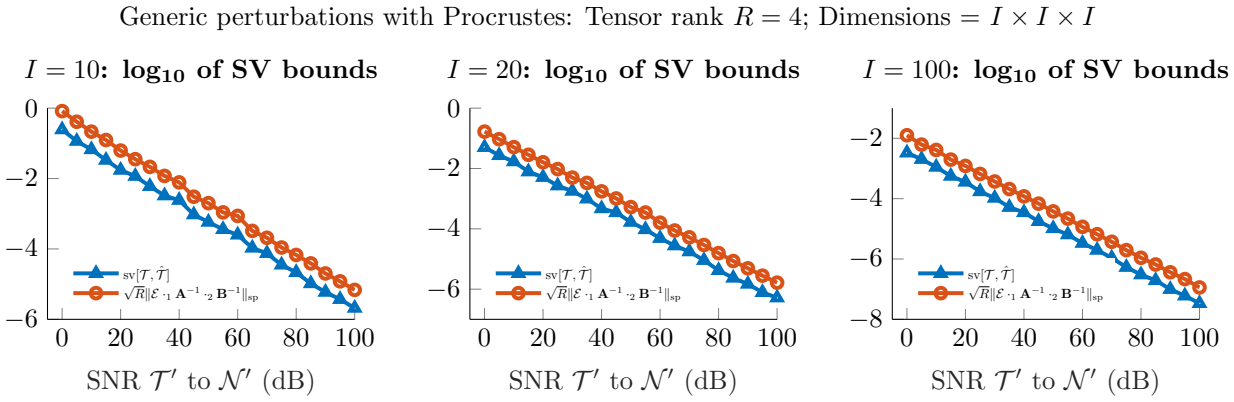


FIG. 4. *Observed spectral variation and computed bound for cores $\mathcal{T}$ and $\hat{\mathcal{T}}$ where $\mathcal{T}'$ is a rank 4 tensor of size $I \times I \times I$ with $I = 10, 20$ or 100 and $\hat{\mathcal{T}}'$ is a rank 4 approximation to $\mathcal{T}' + \mathcal{N}'$. The bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\mathrm{sp}}$ provides a strong estimate for the observed spectral variation. The spectral variation bound and observed spectral variation both decrease as I grows.*

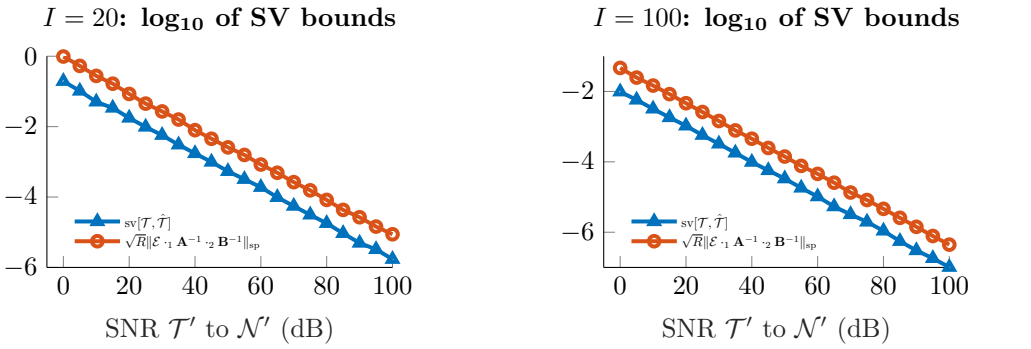


FIG. 5. *Observed spectral variation and computed bound for cores $\mathcal{T}$ and $\hat{\mathcal{T}}$ where $\mathcal{T}'$ is a rank 10 tensor of size $I \times I \times I$ with $I = 20$ or 100 and $\hat{\mathcal{T}}'$ is a rank 10 approximation to $\mathcal{T}' + \mathcal{N}'$. The bound $\sqrt{R}\|\mathcal{E} \cdot_1 \mathbf{A}^{-1} \cdot_2 \mathbf{B}^{-1}\|_{\mathrm{sp}}$ provides a strong estimate for the observed spectral variation. The spectral variation bound and observed spectral variation both decrease as I grows.*

6.2. Bounds for existence of best rank R approximations. We now present the mean of the bound ϵ obtained using Theorem 2.5 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 ϵ_{avg}. The mean bound ϵ_{avg} is plotted in dB as $-20\log_{10}(\epsilon_{\text{avg}})$. 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 2.5 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 6 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.

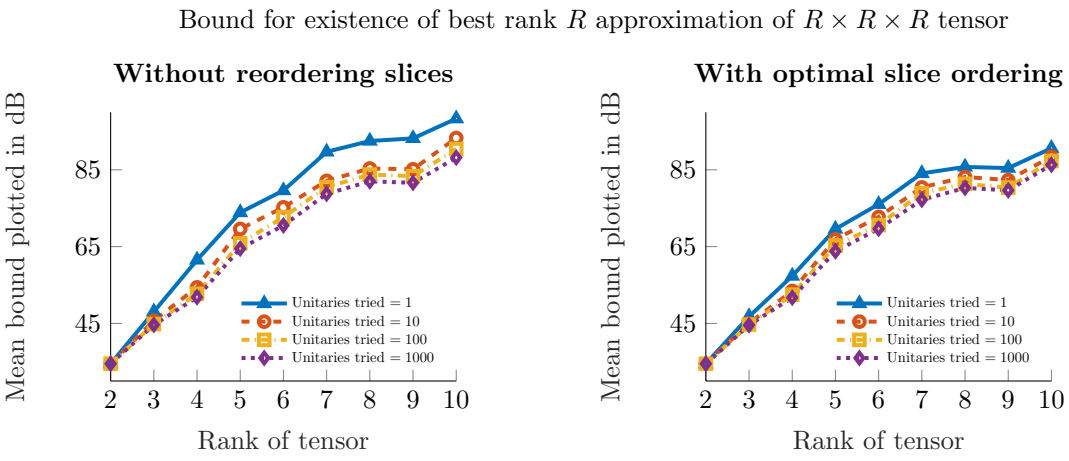


FIG. 6. *Mean bound for radius of a ball centered at a slice mix invertible rank R tensor of unit Frobenius norm in which best rank R approximations are guaranteed to exist. The minimum SNR needed for existence of a best low rank approximation to be guaranteed increases as R increases from* 2 *to* 10. *Increasing the number of unitaries tried can notably improve the bound; however, there are diminishing returns as the number of unitaries tried increases.*

6.2.1. 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 2.5 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 2.5 does not improve as tensor dimensions grow, *the probability of a worst case scenario occurring dramatically*

decreases. This is illustrated in Figure 7 with the following experiment.

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 $\mathcal{W}$ denote the truncated MLSVD of $\mathcal{T}' + \mathcal{N}'$.

Theorem 2.7 is used to compute a bound[17] ϵ such that if there exists a rank R tensor within ϵ Frobenius distance of $\mathcal{W}$, then $\mathcal{W}$ is guaranteed to have a best rank R approximation. We then compute a rank R approximation $\hat{\mathcal{T}}$ of $\mathcal{W}$ and check if $\|\mathcal{W} - \hat{\mathcal{T}}\|_{\mathrm{F}} < \epsilon$, in which case we are able to guarantee that $\mathcal{W}$ 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 $\mathcal{W}$ 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 $\mathcal{W}$ to expecting to be able to guarantee the existence of a best rank R approximation of $\mathcal{W}$.

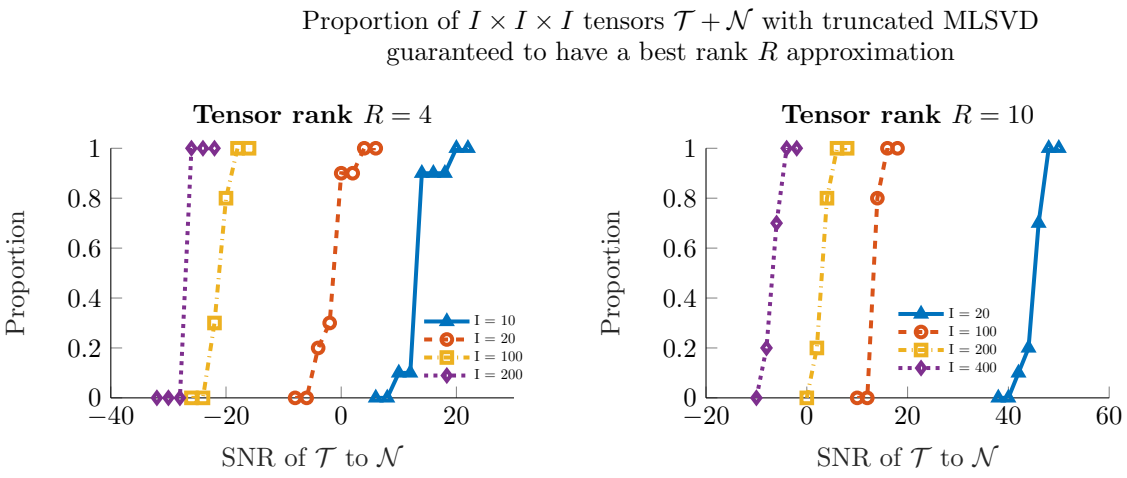


FIG. 7. *Proportion of $I \times I \times I$ tensors $\mathcal{T}' + \mathcal{N}'$ for which the truncated MLSVD $\mathcal{W}$ of size $R \times R \times R$ is guaranteed to have a best rank R approximation. The necessary SNR decreases greatly as tensor dimensions grow. For each choice of I and R, a sharp transition point is observed. When the SNR exceeds the transition point, the truncated MLSVD is expected to have a best rank R approximation.*

6.3. 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 approx-

[17] For each $\mathcal{W}$, 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 7 to shift slightly to the left.

imation 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 2.5.

Moreover, since tensor (border) rank is invariant under modal unitary multiplication, it is sufficient to check that if one computes Procrustes cores $\mathcal{T}$ and $\hat{\mathcal{T}}$ for the pair $\mathcal{T}', \hat{\mathcal{T}}'$, then $\hat{\mathcal{T}}$ lies in the neighborhood around $\mathcal{T}$ given by Theorem 2.5. 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 $\mathcal{T}$ and $\hat{\mathcal{T}}$.*

The experiment shown in Figure 8 is performed as follows. For various choices of I and R we randomly generate rank R tensors $\mathcal{T}'$ 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[18] $\hat{\mathcal{T}}'$ of $\mathcal{T}' + \mathcal{N}'$ and the distance $\|\mathcal{T}' - \hat{\mathcal{T}}'\|_{\mathrm{F}}$ between the tensors $\hat{\mathcal{T}}'$ and $\mathcal{T}'$. We then compute Procrustes cores $\mathcal{T}$ and $\hat{\mathcal{T}}$ of $\mathcal{T}'$ and $\hat{\mathcal{T}}'$, respectively, for which it is guaranteed that $\|\mathcal{T} - \hat{\mathcal{T}}\|_{\mathrm{F}} \leq \|\mathcal{T}' - \hat{\mathcal{T}}'\|_{\mathrm{F}}$ and record the distance $\|\mathcal{T} - \hat{\mathcal{T}}\|_{\mathrm{F}}$ between the cores. For each choice of I and R, a total of 10 trials are performed and the mean distance between the relevant pairs at each step is reported.

Figure 8 illustrates that the distance between the core tensors $\mathcal{T}$ and $\hat{\mathcal{T}}$ 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 $\mathcal{T}$ and $\hat{\mathcal{T}}$ is even smaller still. Furthermore, consistent with Figure 7, 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 6.

7. Conclusion. We have given a deterministic bound on the radius of an open Frobenius norm ball centered at a simple 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 the signal portion of 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 slice mix invertible tensor of size $R \times R \times K$ which has border $\mathbb{K}$-rank R and $\mathbb{K}$-rank strictly greater than R 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 simple rank R tensor are simple, thus giving a neighborhood in which every tensor has a best rank R approximation.

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

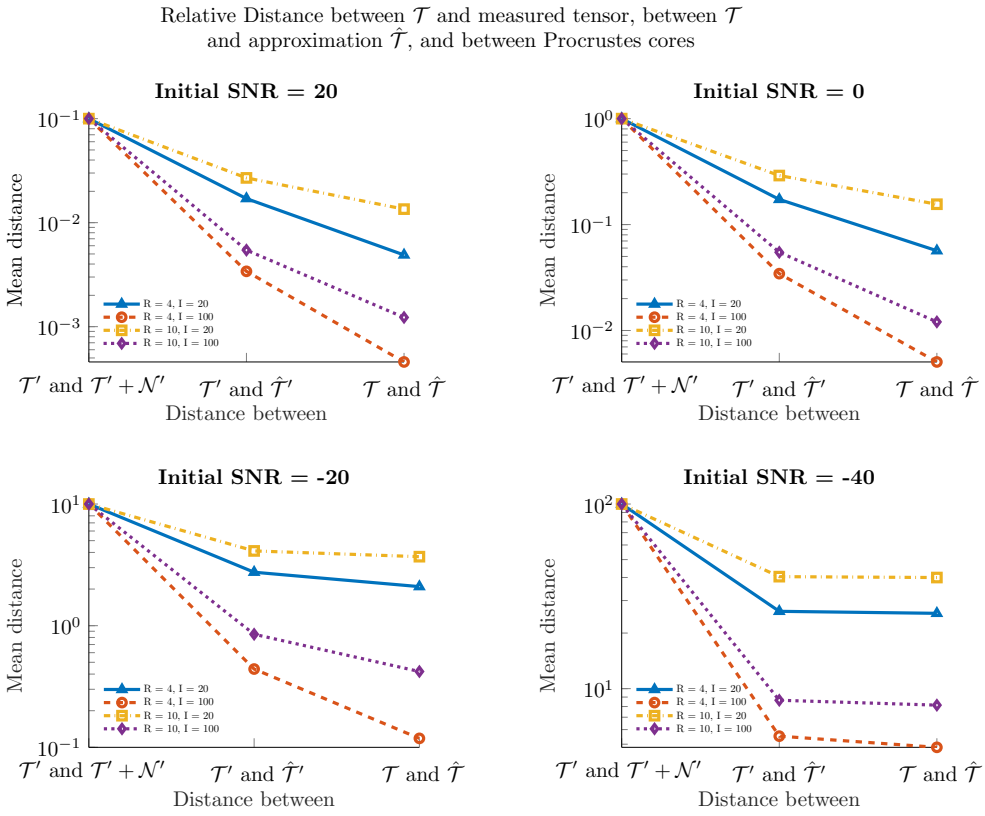


FIG. 8. *Distance between $\mathcal{T}'$ and $\mathcal{T}'+\mathcal{N}'$, between $\mathcal{T}'$ and $\hat{\mathcal{T}}'$, and between $\mathcal{T}$ and $\hat{\mathcal{T}}$. The mean distance between $\mathcal{T}$ and $\hat{\mathcal{T}}$ is observed to be significantly smaller than that between $\mathcal{T}'$ and $\mathcal{T}'+\mathcal{N}'$, and the difference between these distances increases as the ratio of I to R increases. The mean distance between $\mathcal{T}$ and $\hat{\mathcal{T}}$ 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.*

8. Acknowledgements. The authors thank the anonymous referees for their comments which helped improve clarity and exposition in the article.

REFERENCES

[1] T. ABATZOGLOU, *Unique best approximation from a C^2-manifold in Hilbert space*, Pacific Jounal of Mathematics, 87 (1980), pp. 233–244.
[2] 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), pp. 2773–2832.
[3] G. BERGQVIST, *Exact probabilities for typical ranks of $2 \times 2 \times 2$ and $3 \times 3 \times 2$ tensors*, Linear Algebra Appl., (2013), pp. 663–667.
[4] A. CICHOCKI, D. P. MANDIC, L. DE LATHAUWER, G. ZHOU, Q. ZHAO, C. F. CAIAFA, AND A.-H. PHAN, *Tensor decompositions for signal processing applications: From two-way to multiway component analysis*, IEEE Signal Process. Mag., 32 (2015), pp. 145–163.
[5] P. COMON AND C. JUTTEN, *Handbook of blind source separation: Independent component analysis and applications*, Academic press, 2009.
[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] L. DE LATHAUWER, B. D. MOOR, AND J. VANDEWALLE, *An introduction to independent com-*

ponent analysis, J. Chemometrics, 14 (2000), pp. 123–149.
[8] 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.
[9] J. W. Demmel and B. Kågström, *Computing stable eigendecompositions of matrix pencils*, Linear Algebra and its Applications, 88-89 (1987), pp. 139 – 186.
[10] I. Domanov and L. De Lathauwer, *On the uniqueness of the canonical polyadic decomposition of third-order tensors — Part I: Basic results and uniqueness of one factor matrix*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 855–875.
[11] I. Domanov and L. De Lathauwer, *Canonical polyadic decomposition of third-order tensors: Reduction to generalized eigenvalue decomposition*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 636–660.
[12] S. Friedland and L.-H. Lim, *Nuclear norm of higher-order tensors*, Math. Comp., 87 (2018), pp. 1255–1281.
[13] 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.
[14] 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.
[15] S. Friedland and M. Stawiska, *Best approximation on semi-algebraic sets and k-border rank approximation of symmetric tensors*, ArXiv, (2013), https://arxiv.org/abs/1311.1561.
[16] S. Friedland and M. Stawiska, *Some approximation problems in semi-algebraic geometry*, Banach Center Publ., 107 (2015), pp. 133–147.
[17] G. Golub and C. V. Loan, *Matrix Computations*, The Johns Hopkins University Press, 1996.
[18] R. A. Harshman, *Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multi-modal factor analysis*, UCLA Working Papers in Phonetics, 16 (1970), pp. 1–84.
[19] R. A. Harshman and M. E. Lundy, *Data preprocessing and the extended PARAFAC model*, in Research methods for multimode data analysis., H. G. Law, C. W. Snyder Jr, J. A. Hattie, and R. P. McDonald, eds., Praeger, New York, 1984, pp. 216–284.
[20] C. J. Hillar and L.-H. Lim, *Most tensor problems are NP-hard*, J. ACM, 60 (2013), pp. 45:1–45:39.
[21] R. D. Katz, H. Schneider, and S. Sergeev, *On commuting matrices in max algebra and in classical nonnegative algebra*, Linear Algebra and its Applications, 436 (2012), pp. 276–292. Special Issue devoted to the Applied Linear Algebra Conference (Novi Sad 2010).
[22] T. Košir and B. Plestenjak, *On stability of invariant subspaces of commuting matrices*, Linear Algebra and its Applications, 342 (2002), pp. 133–147.
[23] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM Rev., 51 (2009), pp. 455–500.
[24] W. P. Krijnen, T. K. Dijkstra, and A. Stegeman, *On the non-existence of optimal solutions and the occurrence of "degeneracy" in the CANDECOMP/PARAFAC model*, Psychometrika, 73 (2008), pp. 431–439.
[25] J. B. 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.
[26] J. B. Kruskal, *Rank, Decomposition, and Uniqueness for 3-Way and n-Way Arrays*, North-Holland Publishing Co., NLD, 1989, p. 7–18.
[27] J. B. Kruskal, R. A. Harshman, and M. E. Lundy, *How 3-MFA data can cause degenerate PARAFAC solutions, among other relationships*, in Multiway Data Analysis, R. Coppi and S. Bolasco, eds., North-Holland Publishing Co., Amsterdam, The Netherlands, 1989, pp. 115–121.
[28] J. Landsberg, *Tensors: Geometry and applications*, American Mathematical Society, 2012.
[29] J. Landsberg and M. Michałek, *Abelian tensors*, Journal de Mathématiques Pures et Appliquées, 108 (2017), pp. 333–371.
[30] 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.
[31] R.-C. Li, *On perturbations of matrix pencils with real spectra, a revisit*, Math. Comput., 72 (2003), p. 715–728.
[32] L.-H. Lim and P. Comon, *Nonnegative approximations of nonnegative tensors*, J. Chemometrics, 23 (2009), pp. 432–441.
[33] L.-H. Lim and P. Comon, *Blind multilinear identification*, IEEE Trans. Inf. Theory, 60 (2014), pp. 1260–1280.

[34] Y. X. Liqun Qi, Shenglong Hu, *Spectral norm and nuclear norm of a third order tensor*, Journal of Industrial & Management Optimization, (2021).
[35] Y. NAKATSUKASA, *Gerschgorin's theorem for generalized eigenvalue problems in the euclidean metric*, Mathematics of Computation, 80 (2011), pp. 2127–2142.
[36] L. Qi, H. Chen, and Y. Chen, *Tensor Eigenvalues and Their Applications*, Springer, 2018.
[37] Y. Qi, P. Comon, and L.-H. Lim, *Semialgebraic geometry of nonnegative tensor rank*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 1556–1580.
[38] Y. Qi, M. Michałek, and L. Lim, *Complex best r-term approximations almost always exist in finite dimensions*, Applied and Computational Harmonic Analysis, (2019).
[39] P. Schönemann, *A generalized solution of the orthogonal procrustes problem*, Psychometrika, 31 (1966), pp. 1–10.
[40] N. 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.
[41] N. D. Sidiropoulos and R. Bro, *On the uniqueness of multilinear decomposition of N-way arrays*, J. Chemometrics, 14 (2000), pp. 229–239.
[42] M. Sørensen, L. De Lathauwer, P. Comon, S. Icart, and L. Deneire, *Canonical polyadic decomposition with a columnwise orthonormal factor matrix*, SIAM J. Matrix Anal. Appl., 33 (2012), pp. 1190–1213.
[43] F. Sottile, *Real algebraic geometry for geometric constraints*, in Handbook of geometric constraint systems principles (Discrete Mathematics and Its Applications), M. Sitharam, A. S. John, and J. Sidman, eds., Chapman and Hall/CRC, Boca Raton, 2016, ch. 12, pp. 273–285.
[44] 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.
[45] 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.
[46] G. W. Stewart and J. Sun, *Matrix Perturbation Theory*, Academic Press, 1990.
[47] V. Strassen, *Rank and optimal computation of generic tensors*, Linear Algebra Appl., 52 (1983), pp. 645–685.
[48] B. Sturmfels, *Tensors and their eigenvectors*, Not. Am. Math. Soc., 63 (2016), pp. 604–606.
[49] 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.
[50] J. M. F. Ten Berge and H. A. L. 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.
[51] 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.
[52] N. Vannieuwenhoven, *Condition numbers for the tensor rank decomposition*, Linear Algebra Appl., 535 (2017), pp. 35–86.
[53] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer, *Tensorlab 3.0*, Mar. 2016.
[54] 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.