SIAM J. APPL. ALGEBRA GEOMETRY
Vol. 1, pp. 374–387

Perturbation of Higher-Order Singular Values*

Wolfgang Hackbusch†, Daniel Kressner‡, and André Uschmajew§

Abstract. The higher-order singular values for a tensor of order d are defined as the singular values of the d different matricizations associated with the multilinear rank. When $d \geq 3$, the singular values are generally different for different matricizations but not completely independent. Characterizing the set of feasible singular values turns out to be difficult. In this work, we contribute to this question by investigating which first-order perturbations of the singular values for a given tensor are possible. We prove that, except for trivial restrictions, *any* perturbation of the singular values can be achieved for almost every tensor with identical mode sizes. This settles a conjecture from [W. Hackbusch and A. Uschmajew, *Numer. Math.,* 135 (2017), pp. 875–894] for the case of identical mode sizes. Our theoretical results are used to develop and analyze a variant of the Newton method for constructing a tensor with specified higher-order singular values or, more generally, with specified Gramians for the matricizations. We establish local quadratic convergence and demonstrate the robust convergence behavior with numerical experiments.

Key words. tensors, higher-order singular value decomposition, Newton method

AMS subject classifications. 15A18, 15A21, 15A69

DOI. 10.1137/16M1089873

1. Introduction.

Various types of matricizations (or flattenings) of a higher-order tensor $\mathbf{X}$ are connected with subspace-based decompositions for representing and compressing $\mathbf{X}$, such as the Tucker, the hierarchical Tucker, and the tensor train decompositions; see [9, 6, 7, 1] for surveys. In this work, we continue our study [8] of the singular values for matricizations associated with the Tucker decomposition. In particular, we address the question of whether these singular values can be moved in arbitrary directions by small perturbations of $\mathbf{X}$.

A tensor of order d has d different matricizations relevant for the Tucker decomposition. When $d = 2$, $\mathbf{X}$ becomes a matrix and these two matricizations correspond to the matrix itself and its transpose. In particular, these two matricizations have equal singular values. When $d \geq 3$, it is well known that the d vectors of singular values for the d different matricizations are generally not equal. This paper is concerned with the question of how different these singular value vectors may become. This is of practical importance because these singular values allow for quantifying the error committed when approximating $\mathbf{X}$ by tensors of lower multilinear rank. Specifically, the truncation of the *higher-order singular value decomposition* (HOSVD [10]), the workhorse for approximating tensors in a wide range of applications, chooses

*Received by the editors August 17, 2016; accepted for publication (in revised form) March 21, 2017; published electronically July 25, 2017.
http://www.siam.org/journals/siaga/1/M108987.html
†Max-Planck-Institut Mathematik in den Naturwissenschaften, 04103 Leipzig, Germany (wh@mis.mpg.de).
‡ANCHP, MATHICSE, EPFL, 1015 Lausanne, Switzerland (daniel.kressner@epfl.ch).
§Hausdorff Center for Mathematics & Institute for Numerical Simulation, University of Bonn, 53115 Bonn, Germany (uschmajew@ins.uni-bonn.de).

each of the multilinear ranks of the compressed tensor by neglecting small singular values. A deeper understanding of the relations between the d singular value vectors therefore allows for a deeper understanding of the compressibility of tensors and potentially has an impact on algorithmic design. Moreover, as we will show below, the perturbation technique presented in this paper allows us to construct tensors with prescribed singular values. This is a useful tool for testing and benchmarking algorithms.

1.1. Notation. Let us briefly recall the notation from [8]. Let $\mathbf{X} \in \mathbb{R}^{n_1 \times \cdots \times n_d} \cong \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ be a real tensor of order d with mode sizes $n_1, \ldots, n_d$. There are d *principal matricizations* (flattenings)

$$M_{\mathbf{X}}^{(j)} \in \mathbb{R}^{n_j \times n_j^c} \cong \mathbb{R}^{n_j} \otimes \left(\bigotimes_{i \neq j} \mathbb{R}^{n_i} \right),$$

where we have set $n_j^c = \prod_{i \neq j} n_i$. The kth row of $M_{\mathbf{X}}^{(j)}$ contains a vectorization (in some prespecified ordering) of the *slice* $\mathbf{X}(\cdots, k, \cdots)$ with k fixed at position j.

We denote by

- $\boldsymbol{\sigma}_{\mathbf{X}}^{(j)} \in \mathbb{R}^{n_j}$ the vector of singular values of $M_{\mathbf{X}}^{(j)}$ (arranged in nonincreasing order);
- $\boldsymbol{\Sigma}_{\mathbf{X}} = (\boldsymbol{\sigma}_{\mathbf{X}}^{(1)}, \ldots, \boldsymbol{\sigma}_{\mathbf{X}}^{(d)}) \in \mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_d}$ the tuple of *higher-order singular values* of $\mathbf{X}$;
- $G_{\mathbf{X}}^{(j)} = M_{\mathbf{X}}^{(j)} (M_{\mathbf{X}}^{(j)})^T$ the Gram matrix of the jth matricization;
- $\mathcal{S} = \{\mathbf{X} \in \mathbb{R}^{n_1 \times \cdots \times n_d} : \|\mathbf{X}\|_F = 1\}$ the unit sphere in $\mathbb{R}^{n_1 \times \cdots \times n_d}$ with respect to the Frobenius norm $\| \cdot \|_F$;
- $S_{\geq}^{(j)} = \{x \in \mathbb{R}^{n_j} : \|x\|_2 = 1,\ x_1 \geq x_2 \geq \cdots \geq x_{n_j} \geq 0\}$ the set of nonnegative and nonincreasing vectors of Euclidean norm one in $\mathbb{R}^{n_j}$;
- $\mathfrak{S}_{\geq} = S_{\geq}^{(1)} \times \cdots \times S_{\geq}^{(d)}$ the Cartesian product of above;
- $\mathbb{R}_{\text{sym}}^{n \times n}$ the space of real symmetric $n \times n$ matrices.

Additional notation will be introduced further in the text.

2. Problem statement. The singular values $\boldsymbol{\sigma}_{\mathbf{X}}^{(j)}$ of the matricizations $M_{\mathbf{X}}^{(j)}$ are not unrelated for different j. For example, if $\mathbf{X} \in \mathcal{S}$, then $\boldsymbol{\sigma}_{\mathbf{X}}^{(j)} \in S_{\geq}^{(j)}$ for all j. Beyond this simple fact, it is, however, not trivial to describe the relations between $\boldsymbol{\sigma}_{\mathbf{X}}^{(j)}$. For instance, it is not clear which combinations of singular values can actually occur. This amounts to the study of the following set.

Definition 2.1. *The set*

$$\mathfrak{F} = \mathfrak{F}(n_1, \ldots, n_d) \coloneqq \{(\boldsymbol{\sigma}_{\mathbf{X}}^{(1)}, \ldots, \boldsymbol{\sigma}_{\mathbf{X}}^{(d)}) : \mathbf{X} \in \mathcal{S}\}$$

is called the set of normalized feasible configurations *for higher-order singular values of* $\mathbb{R}^{n_1 \times \cdots \times n_d}$.

By the discussion above, $\mathfrak{F}$ is a subset of $\mathfrak{S}_{\geq}$. Since singular values depend continuously on matrix entries, $\mathfrak{F}$ is closed. An interesting, but apparently hard, problem is to decide for a given $\boldsymbol{\Sigma} \in \mathfrak{S}_{\geq}$ whether $\boldsymbol{\Sigma} \in \mathfrak{F}$, e.g., by constructing a tensor $\mathbf{X} \in \mathcal{S}$ with $\boldsymbol{\Sigma}_{\mathbf{X}} = \boldsymbol{\Sigma}$. Numerically, this can be tested using the alternating projection method from [8] or the Newton method introduced in section 5.

In this work, we focus on a different question regarding feasible configurations, namely whether $\mathfrak{F}$ contains interior points relative to $\mathfrak{S}_{\geq}$.

Problem 2.2. *Does $\mathfrak{F}$ contain interior points relative to $\mathfrak{S}_{\geq}$? For which $\mathbf{X} \in \mathcal{S}$ is $\mathbf{\Sigma_X}$ an interior point?*

An interesting point about this question is that if $\mathbf{\Sigma_X}$ is an interior point of $\mathfrak{F}$, then the higher-order singular values of $\mathbf{X}$ are locally independent in the sense that for small, but otherwise arbitrary, perturbations $\tilde{\mathbf{\Sigma}} = \mathbf{\Sigma} + O(\epsilon) \in \mathfrak{S}_{\geq}$, there exists a tensor $\tilde{\mathbf{X}}$ with $\mathbf{\Sigma}_{\tilde{\mathbf{X}}} = \tilde{\mathbf{\Sigma}}$. For instance, it is then possible to perturb only the singular values in one direction j while keeping the others fixed. Of course, this cannot hold for matrices (tensors of order $d = 2$), since a matrix and its transpose have identical singular values. In fact, $\mathfrak{F}(n_1, n_2)$ is the closure of a $\min(n_1 - 1, n_2 - 1)$-dimensional submanifold of $\mathfrak{S}_{\geq}(n_1, n_2)$ (which itself is of dimension $n_1 + n_2 - 2$) and hence contains no interior point (unless $n_1 = n_2 = 1$). Also, we need to exclude tensors with $n_j > n_j^c$ because the size $n_j \times n_j^c$ of the matricization $M_{\mathbf{X}}^{(j)}$ would then imply that some of the singular values are always zero. This leads us to the following conjecture from [8, Conjecture 3.5].

Conjecture 2.3. *For $d \geq 3$, let $n_1, \ldots, n_d$ satisfy the compatibility condition $n_j \leq n_j^c$ for $j = 1, \ldots, d$. Then for almost all $\mathbf{X} \in \mathcal{S}$ the higher-order singular value tuple $\mathbf{\Sigma_X}$ is an interior point of $\mathfrak{F}$ relative to $\mathfrak{S}_{\geq}$.*

Here "for almost all" refers to the standard Lebesgue measure on $\mathcal{S}$; that is, the claim is that the set of $\mathbf{X} \in \mathcal{S}$ for which $\mathbf{\Sigma_X}$ is not a relative interior point of $\mathfrak{F}$ has Lebesgue measure zero. The main theoretical result of this paper is a rigorous proof of this conjecture for $n \times \cdots \times n$ tensors.

Theorem 2.4. *Conjecture 2.3 is true for $n \times \cdots \times n$ tensors of order $d \geq 3$ and size $n \geq 2$.*

2.1. Equivalent formulation using Gram matrices. Our proof of Theorem 2.4 follows the strategy sketched in [8, Remark 3.6]. The main role is played by the map

$$\Phi : \mathbb{R}^{n_1 \times \cdots \times n_d} \to \mathbb{R}_{\text{sym}}^{n_1 \times n_1} \times \cdots \times \mathbb{R}_{\text{sym}}^{n_d \times n_d}, \quad \mathbf{X} \mapsto (G_{\mathbf{X}}^{(1)}, \ldots, G_{\mathbf{X}}^{(d)}),$$

which takes a tensor to the collection of Gram matrices of its principal matricizations. Compared to the mapping $\mathbf{X} \mapsto \mathbf{\Sigma_X}$, the function Φ is much simpler to study since it is quadratic. Note that Φ is also homogeneous of degree two, that is,

$$\Phi(t\mathbf{X}) = t^2 \Phi(\mathbf{X}) \tag{2.1}$$

for all $t \in \mathbb{R}$.

For $K \in R$, let us define

$$\mathcal{P}_K = \{(G^{(1)}, \ldots, G^{(d)}) \in \mathbb{R}_{\text{sym}}^{n_1 \times n_1} \times \cdots \times \mathbb{R}_{\text{sym}}^{n_d \times n_d} : \operatorname{tr}(G^{(j)}) = K \text{ for } j = 1, \ldots, d\}.$$

This set is a Cartesian product of affine hyperplanes in $\mathbb{R}_{\text{sym}}^{n_j \times n_j}$, and therefore

$$\dim(\mathcal{P}_K) = -d + \frac{1}{2} \sum_{j=1}^{d} n_j(n_j + 1).$$

Further, we consider the linear space

$$(2.2)\quad \mathcal{P} = \mathbb{R}\cdot\mathcal{P}_1 = \{(G^{(1)},\dots,G^{(d)}) \in \mathbb{R}^{n_1\times n_1}_{\mathrm{sym}} \times\cdots\times \mathbb{R}^{n_d\times n_d}_{\mathrm{sym}} : \operatorname{tr}(G^{(i)}) = \operatorname{tr}(G^{(j)}) \text{ for } i,j=1,\dots,d\},$$

which is of dimension

$$\dim(\mathcal{P}) = \dim(\mathcal{P}_K) + 1.$$

Given $\mathbf{X} \in \mathbb{R}^{n_1\times\cdots\times n_d}$, a simple but crucial observation is that $\operatorname{tr}(G^{(j)}_{\mathbf{X}}) = \|\mathbf{X}\|_F^2$ holds for every $j = 1,\dots,d$. Therefore, Φ maps into the linear space $\mathcal{P}$:

$$(2.3)\qquad \Phi(\mathbb{R}^{n_1\times\cdots\times n_d}) \subseteq \mathcal{P}.$$

Specifically, the unit sphere $\mathcal{S}$ is mapped into the affine plane $\mathcal{P}_1$:

$$(2.4)\qquad \Phi(\mathcal{S}) \subseteq \mathcal{P}_1.$$

It turns out that Problem 2.2 admits an equivalent formulation in terms of the map Φ.

Proposition 2.5. *Let $\mathbf{X} \in \mathcal{S}$. Then $\Sigma_{\mathbf{X}}$ is an interior point of $\mathfrak{F}$ relative to $\mathfrak{S}_\geq$ if and only if $\Phi(\mathbf{X})$ is an interior point of $\Phi(\mathcal{S})$ relative to $\mathcal{P}_1$.*

Therefore, if Conjecture 2.3 is true, it means that the Gram matrices of tensor matricizations are generically locally independent, that is, they can be moved in arbitrary directions within $\mathcal{P}$. The simplicity of the map Φ compared to $\mathbf{X} \mapsto \Sigma_{\mathbf{X}}$ also provides a convenient starting point for a Newton method, to be discussed in section 5.

Proof of Proposition 2.5. Suppose that $(G^{(1)}_{\mathbf{X}},\dots,G^{(d)}_{\mathbf{X}}) = \Phi(\mathbf{X})$ is an interior point of $\Phi(\mathcal{S})$ relative to $\mathcal{P}_1$. Then there exists an open (relative to $\mathcal{P}_1$) neighborhood $\mathcal{O} \subseteq \Phi(\mathcal{S})$ of $\Phi(\mathbf{X})$. Consider a spectral decomposition

$$G^{(j)}_{\mathbf{X}} = U^{(j)}_{\mathbf{X}}\Lambda^{(j)}_{\mathbf{X}}(U^{(j)}_{\mathbf{X}})^T$$

with $\Lambda^{(j)}_{\mathbf{X}} = \operatorname{diag}(\sigma^{(j)}_{\mathbf{X}})^2$. If $\boldsymbol{\sigma}^{(j)} \in S^{(j)}_{\geq}$ are vectors sufficiently close to $\boldsymbol{\sigma}^{(j)}_{\mathbf{X}}$, then the symmetric matrices $G^{(j)} = U^{(j)}_{\mathbf{X}}\operatorname{diag}(\boldsymbol{\sigma}^{(j)})^2(U^{(j)}_{\mathbf{X}})^T \in \mathcal{P}_1$ satisfy $(G^{(1)},\dots,G^{(d)}) \in \mathcal{O} \subseteq \Phi(\mathcal{S})$. Hence, there is a tensor $\mathbf{Y} \in \mathcal{S}$ such that $G^{(j)}_{\mathbf{Y}} = G^{(j)}$ for $j = 1,\dots,d$. Specifically, $\boldsymbol{\sigma}^{(j)}_{\mathbf{Y}} = \boldsymbol{\sigma}^{(j)}$. This shows that $\mathfrak{F}$ contains an open (relative to $\mathfrak{S}_\geq$) neighborhood of $\boldsymbol{\Sigma}_{\mathbf{X}}$; that is, $\boldsymbol{\Sigma}_{\mathbf{X}}$ is an interior point of $\mathfrak{F}$ relative to $\mathfrak{S}_\geq$.

To prove the reverse implication, let $\Sigma_{\mathbf{X}}$ be an interior point of $\mathfrak{F}$ relative to $\mathfrak{S}_\geq$. Then, by continuity of eigenvalues, there exists an open neighborhood $\mathcal{O} \subseteq \mathcal{P}_1$ of $\Phi(\mathbf{X})$ such that for any $(G^{(1)},\dots,G^{(d)}) \in \mathcal{O}$ the singular value vectors $\boldsymbol{\sigma}^{(j)}$ of $G^{(j)}$ are close enough to $\boldsymbol{\sigma}^{(j)}_{\mathbf{X}}$ to satisfy $(\boldsymbol{\sigma}^{(1)},\dots,\boldsymbol{\sigma}^{(d)}) \in \mathfrak{F}$. Hence, there exists a tensor $\mathbf{Y} \in \mathcal{S}$ with $\boldsymbol{\sigma}^{(j)}_{\mathbf{Y}} = \boldsymbol{\sigma}^{(j)}$ for $j = 1,\dots,d$. It remains to show that there is a tensor that has not only the correct singular values but also the correct Gramians. For this purpose, let $G^{(j)} = V^{(j)}\Lambda^{(j)}(V^{(j)})^T$ and $G^{(j)}_{\mathbf{Y}} = V^{(j)}_{\mathbf{Y}}\Lambda^{(j)}(V^{(j)}_{\mathbf{Y}})^T$, with $\Lambda^{(j)} = \operatorname{diag}(\sigma^{(j)})^2$, be spectral decompositions of $G^{(j)}$ and $G^{(j)}_{\mathbf{Y}}$, respectively. We then consider the orthogonally transformed tensor

$$\tilde{\mathbf{Y}} = \left(V^{(1)}(V^{(1)}_{\mathbf{Y}})^T \otimes\cdots\otimes V^{(d)}(V^{(d)}_{\mathbf{Y}})^T\right)\cdot\mathbf{Y}.$$

Here, the application of the tensor product operator is to be understood in the usual sense. For instance, using the j-mode matrix product $\times_j$ (see [9]), the formula becomes $\tilde{\mathbf{Y}} = \mathbf{Y} \times_1 (V^{(1)}_{\mathbf{Y}})^T V^{(1)} \times_2 \cdots \times_d (V^{(d)}_{\mathbf{Y}})^T V^{(d)}$. In particular, the jth Gram matrix of this tensor is given by

$$G^{(j)}_{\tilde{\mathbf{Y}}} = V^{(j)}(V^{(j)}_{\mathbf{Y}})^T M^{(j)}_{\mathbf{Y}} (M^{(j)}_{\mathbf{Y}})^T V^{(j)}_{\mathbf{Y}} (V^{(j)})^T = V^{(j)} \Lambda^{(j)} (V^{(j)})^T = G^{(j)}.$$

Since $(G^{(1)}, \dots, G^{(d)}) \in \mathcal{O}$ is arbitrary, this shows $\mathcal{O} \subseteq \Phi(\mathcal{S})$, thereby completing the proof. ∎

2.2. Sufficient conditions. Due to (2.4), $\Phi(\mathbf{X})$ will be an interior point of $\mathcal{P}_1$ if its derivative

$$\Phi'(\mathbf{X}) : \mathbb{R}^{n_1 \times \cdots \times n_d} \to \mathbb{R}^{n_1 \times n_1}_{\text{sym}} \times \cdots \times \mathbb{R}^{n_d \times n_d}_{\text{sym}}$$

has rank $\dim(\mathcal{P}_1)$ on the tangent space $T_{\mathcal{S}}(\mathbf{X})$ of $\mathcal{S}$ at $\mathbf{X}$. In fact, this would mean that the restriction $\Phi|_{\mathcal{S}}$ as a map to $\mathcal{P}_1$ is a submersion at $\mathbf{X}$ and, as such, maps open (in $\mathcal{S}$) neighborhoods of $\mathbf{X}$ to open (in $\mathcal{P}_1$) neighborhoods of $\Phi(\mathbf{X})$; see, e.g., [4, section 16.7.5]. Note that $\Phi'(\mathbf{X})$ maps $T_{\mathcal{S}}(\mathbf{X})$ into the *linear* space $\mathcal{P}_0$. We obtain the following sufficient condition.

Proposition 2.6. *If $\mathbf{X} \in \mathcal{S}$ satisfies*

$$\operatorname{rank}(\Phi'(\mathbf{X})) = \dim(\mathcal{P}) \tag{2.5}$$

or, equivalently,

$$\Phi'(\mathbf{X})[T_{\mathcal{S}}(\mathbf{X})] = \mathcal{P}_0, \tag{2.6}$$

then $\Phi(\mathbf{X})$ is an interior point of $\mathcal{P}_1$.

Proof. As explained above, (2.6) implies that $\Phi(\mathbf{X})$ is an interior point of $\mathcal{P}_1$. Since $\dim(\mathcal{P}) = \dim(\mathcal{P}_0) + 1$ and $\Phi'(\mathbf{X})[T_{\mathcal{S}}(\mathbf{X})] \subseteq \mathcal{P}_0$, it is also clear that (2.5) implies (2.6). To prove that (2.6) implies (2.5), it is enough to note that, by homogeneity (2.1), $\Phi'(\mathbf{X})$ maps the one-dimensional subspace spanned by $\mathbf{X}$, which is orthogonal to $T_{\mathcal{S}}(\mathbf{X})$, to the one-dimensional subspace spanned by $\Phi(\mathbf{X})$, which is linearly independent from $\mathcal{P}_0$, since $\Phi(\mathbf{X}) \in \mathcal{P}_1$. ∎

In fact, since $\Phi'(\mathbf{X})$ depends polynomially on the entries of $\mathbf{X}$, we can state a little more.

Proposition 2.7. *If there exists a single tensor $\mathbf{X}_0 \in \mathbb{R}^{n_1 \times \cdots \times n_d}$ satisfying (2.5), then almost all $\mathbf{X} \in \mathbb{R}^{n_1 \times \cdots \times n_d}$ satisfy (2.5). In particular, (2.5) holds for almost all $\mathbf{X} \in \mathcal{S}$ with respect to the standard Lebesgue measure, and Conjecture 2.3 is true for $n_1 \times \cdots \times n_d$ tensors.*

Proof. The genericity claim follows from a standard logic. Set $r := \dim(\mathcal{P})$ for brevity and assume $\operatorname{rank}(\Phi'(\mathbf{X}_0)) = r$. Since, by (2.3), the rank of $\Phi'(\mathbf{X})$ cannot be larger than r, it suffices to prove that it takes at least this value for all $\mathbf{X} \in \mathbb{R}^{n_1 \times \cdots \times n_d}$. This property can be encoded as $p(\mathbf{X}) \neq 0$, where p is a polynomial in the entries of $\mathbf{X}$ (for instance, one may use the sum of squares of all $r \times r$ minors of a matrix representation of $\Phi'(\mathbf{X})$). Since p is a polynomial, it holds that either $p(\mathbf{X}) \neq 0$ for almost all $\mathbf{X} \in \mathbb{R}^{n_1 \times \cdots \times n_d}$, or $p \equiv 0$. By assumption, the latter is not the case.

To prove that the property (2.5) is generic with respect to the Lebesgue measure on $\mathcal{S}$ as well, we could apply the same argument by using real-analytic charts of $\mathcal{S}$ (the set of zeros of a nonzero real-analytic function is of measure zero). However, it also follows from what we have already proved. Namely, by (2.1), $\mathrm{rank}(\Phi'(t\mathbf{X})) = \mathrm{rank}(\Phi'(\mathbf{X}))$ for all $t \neq 0$. This implies that the set of $\mathbf{X} \in \mathcal{S}$ with $\mathrm{rank}(\Phi'(\mathbf{X})) \leq \dim(\mathcal{P}_1) = \dim(\mathcal{P}) - 1$ must be of measure zero, since otherwise there is a set of positive volume with $\mathrm{rank}(\Phi'(\mathbf{X})) \leq \dim(\mathcal{P}) - 1$, which, in light of the previous considerations, is not possible. ∎

3. Proof of Theorem 2.4 (Conjecture 2.3 for $n \times \cdots \times n$ tensors). By Proposition 2.7, Theorem 2.4 is proven via the construction of an $n \times \cdots \times n$ tensor $\mathbf{X}_0$ for each $d \geq 3$ and $n \geq 2$ such that $\Phi'(\mathbf{X}_0)$ has rank $dn(n+1)/2 - d + 1$. By the definition of Φ, we have

$$\Phi'(\mathbf{X})[\mathbf{H}] = \left(M_{\mathbf{X}}^{(1)}(M_{\mathbf{H}}^{(1)})^T + M_{\mathbf{H}}^{(1)}(M_{\mathbf{X}}^{(1)})^T, \ldots, M_{\mathbf{X}}^{(d)}(M_{\mathbf{H}}^{(d)})^T + M_{\mathbf{H}}^{(d)}(M_{\mathbf{X}}^{(d)})^T \right). \tag{3.1}$$

We first discuss the case $d = 3$, $n = 2$ separately, and then give a general construction that is valid for all other cases.

Case $d = 3$, $n = 2$. Consider a general $2 \times 2 \times 2$ tensor $\mathbf{X}$ with its first matricization

$$M_{\mathbf{X}}^{(1)} = \mathbf{X} = \left(\begin{array}{cc|cc} a & c & e & g \\ b & d & f & h \end{array} \right).$$

A matrix representation of $\Phi'(\mathbf{X}) : \mathbb{R}^{2\times2\times2} \to \mathbb{R}^{2\times2}_{\mathrm{sym}} \times \mathbb{R}^{2\times2}_{\mathrm{sym}} \times \mathbb{R}^{2\times2}_{\mathrm{sym}}$ can be directly computed from (3.1):

$$\Phi'(\mathbf{X}) = \left(\begin{array}{cccccccc} 2a & 0 & 2c & 0 & 2e & 0 & 2g & 0 \\ b & a & d & c & f & e & h & g \\ 0 & 2b & 0 & 2d & 0 & 2f & 0 & 2h \\ \hline 2a & 2b & 0 & 0 & 2e & 2f & 0 & 0 \\ c & d & a & b & g & h & e & f \\ 0 & 0 & 2c & 2d & 0 & 0 & 2g & 2h \\ \hline 2a & 2b & 2c & 2d & 0 & 0 & 0 & 0 \\ e & f & g & h & a & b & c & d \\ 0 & 0 & 0 & 0 & 2e & 2f & 2g & 2h \end{array} \right).$$

We know that $\mathrm{rank}(\Phi'(\mathbf{X})) \leq 7$ for all $\mathbf{X}$. Choosing the particular tensor $\mathbf{X}_0$ given by the matricization

$$M_{\mathbf{X}_0}^{(1)} = \left(\begin{array}{cc|cc} 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{array} \right),$$

we have

$$\Phi'(\mathbf{X}_0) = \left(\begin{array}{cccccccc} 2 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 & 0 & 2 \\ \hline 2 & 2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 & 0 & 2 \\ \hline 2 & 2 & 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 \end{array} \right), \qquad Z = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 & 2 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 & 2 \\ 2 & 2 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 \end{pmatrix}. \tag{3.2}$$

The 7×7 matrix Z is obtained by omitting rows 1, 4 and column 4 from $\Phi'(\mathbf{X}_0)$. One calculates that $\det(Z) = 16$, which implies that $\mathrm{rank}\,\Phi'(\mathbf{X}_0) = 7 = dn(n+1)/2 - d + 1$.

General construction. Let $\mathbf{X}_0$ be the $n \times \cdots \times n$ tensor of order d that has all entries zeros except for

$$\mathbf{X}_0(k, \dots, k) = 1, \quad k = 1, \dots, n,$$

and

$$\begin{aligned} \mathbf{X}_0(k, 1, 1, \dots, 1) &= 1, \quad k = 1, \dots, n, \\ \mathbf{X}_0(1, k, 1, \dots, 1) &= 1, \quad k = 1, \dots, n, \\ &\vdots \\ \mathbf{X}_0(1, \dots, 1, 1, k) &= 1, \quad k = 1, \dots, n. \end{aligned}$$

In other words, $\mathbf{X}_0$ has ones on its diagonal and in all fibers intersecting with $(1, \dots, 1)$. For example, when $d = 3$ and $n = 4$, this results in a matricization

$$M_{\mathbf{X}_0}^{(1)} = \left(\begin{array}{cccc|cccc|cccc|cccc} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array}\right). \tag{3.3}$$

Note that $\mathbf{X}_0$ is supersymmetric, and hence all matricizations are essentially the same.

Theorem 3.1. *For all d, n satisfying $d = 3, n \geq 3$ or $d \geq 4, n \geq 2$, the tensor $\mathbf{X}_0$ constructed above satisfies* (2.5).

The rest of this section is concerned with the proof of Theorem 3.1. It is worth mentioning that for $d = 3$, $n = 2$ the corresponding tensor $\mathbf{X}_0$ does not satisfy (2.5) because $\operatorname{rank}(\Phi'(\mathbf{X}_0)) = 5 < 7 = \dim(\mathcal{P})$ in this case.

3.1. Theorem 3.1 for $d = 3$, $n = 3$. This case is treated explicitly, in analogy to the case $d = 3$, $n = 2$ discussed above. As the involved matrices are quite large, we refrain from displaying them and only describe the construction. We first construct the 18×27 matrix representation of $\Phi'(\mathbf{X}_0)$ analogous to (3.2). After omitting rows $1, 7$ and columns $4, 6, 7, 8, 9, 11, 13, 15, 17, 21, 23$, the resulting 16×16 matrix has determinant 384. Thus, $\operatorname{rank}(\Phi'(\mathbf{X}_0)) = 16 = dn(n+1)/2 - d + 1$. A MATLAB package reproducing this and other results from this paper is available online at http://anchp.epfl.ch/software/misc.

3.2. Theorem 3.1 for remaining cases. To prove Theorem 3.1 for the remaining cases, we show (2.6) by constructing elements $\mathbf{H}$ from the tangent space $T_{\mathcal{S}}(\mathbf{X}_0)$ such that $\Phi'(\mathbf{X}_0)$ applied to these elements yields a basis of the space $\mathcal{P}_0$. Using (3.1), it is easy to verify that such a basis is obtained from the following two lemmas by sweeping through all combinations of j, k, and ℓ as indicated. In particular, note that the diagonal vectors appearing in Lemma 3.2(i) form a basis of the subspace of all vectors in $\mathbb{R}^n$ whose entries sum up to zero.

Lemma 3.2 (diagonal entries). *For any $1 \leq j \leq d$ and $2 \leq k \leq n$ there exists $\mathbf{H} \in T_{\mathcal{S}}(\mathbf{X}_0)$ such that*

(i) *the diagonal of $M_{\mathbf{X}_0}^{(j)}(M_{\mathbf{H}}^{(j)})^T$ is the vector $(1,0,\dots,0,-1,0,\dots,0)$ with -1 at position k;*

(ii) *the diagonal of $M_{\mathbf{X}_0}^{(i)}(M_{\mathbf{H}}^{(i)})^T$ is zero for $i \neq j$.*

Lemma 3.3 (off-diagonal entries). *For any* $1 \leq j \leq d$ *and* $1 \leq \ell < k \leq n$ *there exists* $\mathbf{H} \in T_{\mathcal{S}}(\mathbf{X}_0)$ *such that*

(i) *the only nonzero of* $M^{(j)}_{\mathbf{X}_0}(M^{(j)}_{\mathbf{H}})^T$ *is at position* (k, ℓ)*;*

(ii) $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T = 0$ *for all* $i \neq j$*.*

In the following proofs, it is not explicitly stated, but is easy to check, that $\mathbf{H} \in T_{\mathcal{S}}(\mathbf{X}_0)$, that is, $\langle \mathbf{X}_0, \mathbf{H} \rangle = 0$ with respect to the Frobenius inner product. Note that $\langle \mathbf{X}, \mathbf{H} \rangle = \operatorname{tr}(M^{(j)}_{\mathbf{X}}(M^{(j)}_{\mathbf{H}})^T)$ for any $j = 1, \ldots, d$. Therefore, by (3.1), showing $\Phi'(\mathbf{X}_0)[\mathbf{H}] \in \mathcal{P}_0$ for some $\mathbf{H}$ actually automatically implies $\mathbf{H} \in T_{\mathcal{S}}(\mathbf{X}_0)$, since $\operatorname{tr}(M^{(j)}_{\mathbf{X}_0}(M^{(j)}_{\mathbf{H}})^T) = 0$.

Proof of Lemma 3.2. Because $\mathbf{X}_0$ is supersymmetric, it suffices to treat $j = 1$. For fixed $2 \leq k \leq n$ consider $\mathbf{H}$ with all entries zeros except for an entry 1 at position $(1, \ldots, 1)$ and an entry -1 at position $(k, 1, \ldots, 1)$. The pth diagonal entry of $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T$ is the Frobenius inner product of the slices $\mathbf{X}_0(\cdots, p, \cdots)$ and $\mathbf{H}(\cdots, p, \cdots)$ with p at position i.

When $i \neq j = 1$, the slice $\mathbf{H}(\cdots, p, \cdots)$ contains nonzero entries only when $p = 1$, namely an entry 1 at $(1, \ldots, [1], \ldots, 1)$ and an entry -1 at $(k, 1, \ldots, [1], \ldots, 1)$. (To simplify notation, we use square brackets to indicate the fixed index of slices.) Since the entries of $\mathbf{X}_0$ are 1 at both of these positions, the slices are orthogonal for $p = 1$ as well. In turn, we have proved that $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T$ has a zero diagonal when $i \neq 1$.

When $i = j = 1$, the slices $\mathbf{H}(p, \cdots)$ are nonzero only if $p = 1$ or $p = k$. In both cases they contain a single nonzero entry at $([1], 1, \ldots, 1)$, respectively $([k], 1, \ldots, 1)$, which will be multiplied by a 1 at the corresponding position of $\mathbf{X}_0(p, \cdots)$ when forming the inner product. Hence, the diagonal entries of $M^{(1)}_{\mathbf{X}_0}(M^{(1)}_{\mathbf{H}})^T$ are as asserted. ■

Proof of Lemma 3.3. Again, it suffices to consider $j = 1$. Three cases will be distinguished.

Case $d \geq 4$*,* $n \geq 2$*.* This case is simpler than the case $d = 3$ and therefore we treat it first. Given $1 \leq \ell < k \leq n$, we consider the tensor $\mathbf{H}$ that contains only zeros except for a nonzero at position $(\ell, k, \ldots, k)$.

The (p, q) entry of $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T$ is the Frobenius inner product of the slices $\mathbf{X}_0(\cdots, p, \cdots)$ and $\mathbf{H}(\cdots, q, \cdots)$ with p, q at positions i. When $i \neq 1$, this slice of $\mathbf{H}$ is nonzero only if $q = k$, in which case the nonzero entry is at $(\ell, k, \ldots, [k], \ldots, k)$. The nonzero entries of $\mathbf{X}_0$ are at multi-indices, where either all indices are the same or contain $d - 1$ indices equal to one. Since $1 \leq \ell < k$ and $d \geq 4$, it follows that $\mathbf{X}_0(\ell, k, \ldots, k, p, k, \ldots, k) = 0$ for any p, and hence none of the slices $\mathbf{X}_0(\cdots, p, \cdots)$ has a nonzero matching the one at $(\ell, k, \ldots, [k], \ldots, k)$. We conclude that $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T = 0$ for $i \neq 1$.

When $i = 1$, we note that the slice $\mathbf{H}(q, \cdots)$ has a nonzero entry only if $q = \ell$, namely at $([\ell], k, \ldots, k)$. Since $k > 1$, a slice $\mathbf{X}_0(p, \cdots)$ has a nonzero at the same place only if $p = k$. Hence the only nonzero entry of $M^{(1)}_{\mathbf{X}_0}(M^{(1)}_{\mathbf{H}})^T$ is at $(p, q) = (k, \ell)$.

Case $d = 3$*,* $2 \leq \ell < k \leq n$*.* In this case, we again consider the tensor $\mathbf{H}$ which contains only zeros except at entry (ℓ, k, k).

For $i = 2$, the (p, q) entry of $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T$ is the Frobenius inner product of the slices $\mathbf{X}_0(\cdot, p, \cdot)$ and $\mathbf{H}(\cdot, q, \cdot)$. This slice of $\mathbf{H}$ is nonzero only if $q = k$, in which case the nonzero entry is at $(\ell, [k], k)$. The slice $\mathbf{X}_0(\cdot, p, \cdot)$, on the other hand, has possible nonzeros at $(p, [p], p)$

and $(\alpha, [p], \beta)$, where α or β is equal to one, or both are if $p = 1$. Since $2 \le \ell < k$, it is not possible that $\alpha = \ell, \beta = k$ and, hence, $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T = 0$. The argument for $i = 3$ is analogous.

Considering $i = 1$, we note that the slice $\mathbf{H}(q, \cdot, \cdot)$ is zero unless $q = \ell$, in which case it has a single nonzero entry at $([\ell], k, k)$. Since $k > \ell \ge 1$, the slice $\mathbf{X}_0(p, \cdot, \cdot)$ has a nonzero at the same place only if $p = k$. The only nonzero entry of $M^{(1)}_{\mathbf{X}_0}(M^{(1)}_{\mathbf{H}})^T$ is therefore at $(p, q) = (k, \ell)$.

Case $d = 3$, $n \ge 4$, $1 = \ell < k \le n$. Let us first assume $2 < k < n$. Then we consider the tensor $\mathbf{H}$ which contains only zero entries, except that the slice $\mathbf{H}(1, \cdot, \cdot)$ contains the submatrix

$$\mathbf{H}([1], k-1:k+1, k-1:k+1) = \begin{pmatrix} 0 & -\frac{1}{2} & \frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} \\ \frac{1}{2} & -\frac{1}{2} & 0 \end{pmatrix}.$$

For example, when $n = 4$, $k = 3$, the resulting tensor has the matricization

$$M^{(1)}_{\mathbf{H}} = \left(\begin{array}{cccc|cccc|cccc|cccc} 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & \frac{1}{2} & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & \frac{1}{2} & -\frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right). \tag{3.4}$$

A comparison with (3.3) immediately shows that $M^{(1)}_{\mathbf{X}_0}(M^{(1)}_{\mathbf{H}})^T$ has a single nonzero entry at $(k, 1)$. One also easily checks that $M^{(i)}_{\mathbf{X}_0}(M^{(i)}_{\mathbf{H}})^T = 0$ for $i = 2, 3$. For instance, for $i = 3$ this can be seen from the fact that the "frontal" slices of $\mathbf{X}_0$ depicted in (3.3) are pairwise orthogonal in the Frobenius inner product to the frontal slices of $\mathbf{H}$ given in (3.4). This reasoning remains valid for larger n, as $\mathbf{H}$ will only have additional zero blocks.

When $k = 2$, one considers the submatrix

$$\mathbf{H}([1], 2:4, 2:4) = \begin{pmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix},$$

whereas for $k = n$ one chooses

$$\mathbf{H}([1], n-2:n, n-2:n) = \begin{pmatrix} 0 & \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 0 & -\frac{1}{2} \\ -\frac{1}{2} & -\frac{1}{2} & 1 \end{pmatrix}.$$

This completes the proof of Lemma 3.3. ■

4. Numerical evidence for unequal mode sizes. The construction of the tensor $\mathbf{X}_0$ from section 3 does not extend to cases where the mode sizes $n_1, \dots, n_d$ are not identical. At this point, we are not aware of a construction that admits rigorous verification of the condition of Proposition 2.7 for *general* $n_1, \dots, n_d$. For a *specific* choice of $n_1, \dots, n_d$, it is certainly possible to extend the technique from section 3.1. In the following, we use a simpler approach to provide strong numerical evidence for $d = 3$ and a range of (small) mode sizes.

Table 1

Computed values of $\sigma_{N-d+1}\big(\Phi'(\mathbf{X})\big)/\sigma_1\big(\Phi'(\mathbf{X})\big)$ for random $n_1 \times n_2 \times n_3$ tensors $\mathbf{X}$.

$n_1 = 2$	$n_3 = 2$	$n_3 = 3$	$n_3 = 4$	$n_3 = 5$
$n_2 = 2$	0.034	0.027	0.016	6×10^{-17}
$n_2 = 3$	0.035	0.046	0.013	0.011
$n_2 = 4$	0.012	0.028	0.032	0.033
$n_2 = 5$	7×10^{-17}	0.030	0.034	0.028

$n_1 = 3$	$n_3 = 3$	$n_3 = 4$	$n_3 = 5$	$n_3 = 6$
$n_2 = 3$	0.060	0.071	0.053	0.045
$n_2 = 4$	0.060	0.061	0.053	0.040
$n_2 = 5$	0.054	0.055	0.060	0.051
$n_2 = 6$	0.034	0.046	0.052	0.050

$n_1 = 4$	$n_3 = 4$	$n_3 = 5$	$n_3 = 6$	$n_3 = 7$
$n_2 = 4$	0.059	0.055	0.053	0.048
$n_2 = 5$	0.059	0.067	0.059	0.057
$n_2 = 6$	0.043	0.059	0.059	0.057
$n_2 = 7$	0.055	0.057	0.057	0.053

$n_1 = 5$	$n_3 = 5$	$n_3 = 6$	$n_3 = 7$	$n_3 = 8$
$n_2 = 5$	0.066	0.062	0.059	0.046
$n_2 = 6$	0.061	0.062	0.056	0.053
$n_2 = 7$	0.057	0.054	0.055	0.056
$n_2 = 8$	0.047	0.054	0.056	0.054

We used MATLAB with Tensor Toolbox [2] to construct the matrix representation of $\Phi'(\mathbf{X})$ with the range restricted to $\mathbb{R}^{n_1\times n_1}_{\text{sym}} \times \cdots \times \mathbb{R}^{n_d\times n_d}_{\text{sym}}$; see (3.2) for an example. This matrix has size $N \times (n_1 \cdots n_d)$ with $N = \frac{1}{2}\sum_{j=1}^{d} n_j(n_j+1)$, and, by Proposition 2.7, Conjecture 2.3 holds if $\Phi'(\mathbf{X})$ has rank $N-d+1$ for some $\mathbf{X}$. To verify this condition numerically we have computed $\sigma_{N-d+1}\big(\Phi'(\mathbf{X})\big)/\sigma_1\big(\Phi'(\mathbf{X})\big)$ for random $\mathbf{X}$, where $\sigma_k(\cdot)$ denotes the kth largest singular value of a matrix. If this ratio is sufficiently larger than 10^{-16} in double precision, then Conjecture 2.3 is likely to hold, because singular values are perfectly well conditioned [5]. Table 1 displays the results obtained for tensors constructed by typing

```
rand('seed',0); X = rand(n1,n2,n3);
```

in MATLAB. The condition of Proposition 2.7 is confirmed for all mode sizes tested, with the notable exceptions $(n_1, n_2, n_3) = (2, 2, 5)$ and $(n_1, n_2, n_3) = (2, 5, 2)$, for which the compatibility condition $n_j \le n_j^c$ of Conjecture 2.3 is not satisfied.

5. A fast iterative method for assigning higher-order singular values.

In [8], an alternating projection method for prescribing higher-order singular values was proposed. This method was observed to converge linearly to a feasible set of singular values, but no theoretical analysis was provided. Based on the results of the present paper, we develop a variant of the Newton method for which local quadratic convergence can be proven.

5.1. Newton method.

Given a collection of symmetric, positive definite matrices $(G^{(1)}, \ldots, G^{(d)})$, we aim at finding a tensor $\mathbf{X}$ such that

$$\Phi(\mathbf{X}) = (G^{(1)}, \ldots, G^{(d)}). \tag{5.1}$$

Our method and the analysis simplify if we do not impose a normalization on $\|\mathbf{X}\|_F$. We only assume that the right-hand side of (5.1) is contained in the linear space $\mathcal{P}$ from (2.2), that is, the traces of all $G^{(j)}$ are equal. For the rest of this section, we restrict the codomain of Φ to $\mathcal{P}$, that is, $\Phi : \mathbb{R}^{n_1\times\cdots\times n_d} \to \mathcal{P}$.

If we are only interested in perturbing singular values, we compute the spectral decomposition of the Gram matrices for an initial tensor $\mathbf{X}_0$,

$$G^{(j)}_{\mathbf{X}_0} = U^{(j)}_{\mathbf{X}_0}\Lambda^{(j)}_{\mathbf{X}_0}(U^{(j)}_{\mathbf{X}_0})^T,$$

and set

$$G^{(j)} = U^{(j)}_{\mathbf{X}_0}\Lambda^{(j)}(U^{(j)}_{\mathbf{X}_0})^T$$

with

$$\Lambda^{(j)} = \Lambda^{(j)}_{\mathbf{X}_0} + O(\epsilon).$$

After a suitable normalization, $(G^{(1)}, \dots, G^{(d)}) \in \mathcal{P}$. In view of our results, (5.1) is likely to have at least one solution for sufficiently small ϵ.

Applying the Newton method to (5.1) requires solving an equation of the form

$$\Phi'(\mathbf{X}_n)[\mathbf{H}_n] = \Phi(\mathbf{X}_n) - (G^{(1)}, \dots, G^{(d)}). \tag{5.2}$$

Because $\mathcal{P}$ is linear, the right-hand side is contained in $\mathcal{P}$.

Suppose that $\Phi'(\mathbf{X}_n) : \mathbb{R}^{n_1\times\cdots\times n_d} \to \mathcal{P}$ satisfies the full rank condition (2.5) of Proposition 2.7. Then $\Phi'(\mathbf{X}_n)$ has full row rank, in other words, (5.2) is consistent. Following [3, section 4.4], we choose the solution of smallest norm,

$$\mathbf{H}_n = \Phi'(\mathbf{X}_n)^+\big(\Phi(\mathbf{X}_n) - (G^{(1)}, \dots, G^{(d)})\big), \tag{5.3}$$

where $\Phi'(\mathbf{X}_n)^+ : \mathcal{P} \to \mathbb{R}^{n_1\times\cdots\times n_d}$ denotes the Moore–Penrose pseudoinverse of $\Phi'(\mathbf{X}_n)$.[1] The next iterate is

$$\mathbf{X}_{n+1} = \mathbf{X}_n - \mathbf{H}_n. \tag{5.4}$$

5.2. Convergence analysis. In the following, we study the convergence of (5.3)–(5.4). For this purpose, we need to choose a norm on $\mathcal{P}$. Given $\tilde{G} = (\tilde{G}^{(1)}, \dots, \tilde{G}^{(d)}) \in \mathcal{P}$, we let $\|\tilde{G}\|$ be the norm obtained by taking the Euclidean norm of the vector containing all entries in the upper triangular parts of the symmetric matrices $\tilde{G}^{(j)}$. Based on this norm, we define

$$\gamma(\mathbf{X}) := 1/\sup\{\|\Phi'(\mathbf{X})^+\tilde{G}\|_F : \tilde{G} \in \mathcal{P},\ \|\tilde{G}\| = 1\} = \sigma_{\dim(\mathcal{P})}(\Phi'(\mathbf{X})), \tag{5.5}$$

which is positive for any $\mathbf{X}$ satisfying (2.5).

Theorem 5.1. *Assume that $\mathbf{X}_0$ satisfies (2.5) and*

$$\|\Phi(\mathbf{X}_0) - (G^{(1)}, \dots, G^{(d)})\| \le \frac{\gamma(\mathbf{X}_0)^2}{6\sqrt{d}} \tag{5.6}$$

for $(G^{(1)}, \dots, G^{(d)}) \in \mathcal{P}$. Then all iterates $\mathbf{X}_n$ defined by (5.3)–(5.4) satisfy (2.5) and converge to a solution of (5.1). Moreover,

$$\|\mathbf{X}_{n+1} - \mathbf{X}_n\|_F \le \frac{1}{2}\omega\|\mathbf{X}_n - \mathbf{X}_{n-1}\|_F^2$$

holds with $\omega = 6\sqrt{d}/\gamma(\mathbf{X}_0)$.

[1] Letting A denote the matrix representation of $\Phi'(\mathbf{X}_n)$ with the codomain restricted to $\mathcal{P}$, the matrix representation of $\Phi'(\mathbf{X}_n)^+$ is given by $A^T(AA^T)^{-1}$.

Proof. Let us first note that $\Phi'(\mathbf{Y})$ is linear in $\mathbf{Y}$ and (3.1) implies the bound

$$\|\Phi'(\mathbf{Y})[\mathbf{H}]\|^2 \leq \sum_{j=1}^{d} \|M_{\mathbf{Y}}^{(j)}(M_{\mathbf{H}}^{(j)})^T + M_{\mathbf{H}}^{(j)}(M_{\mathbf{Y}}^{(j)})^T\|_F^2 \leq 4d\|\mathbf{Y}\|_F^2\|\mathbf{H}\|_F^2.$$

Hence, the induced operator norm of $\Phi'(\mathbf{Y})$ satisfies $\|\Phi'(\mathbf{Y})\| \leq 2\sqrt{d}\|\mathbf{Y}\|_F$.

Now, let $F(\mathbf{X}) := \Phi(\mathbf{X}) - (G^{(1)}, \ldots, G^{(d)})$. By (5.6),

$$\|F'(\mathbf{X}_0)^+ F(\mathbf{X}_0)\|_F = \|\Phi'(\mathbf{X}_0)^+ F(\mathbf{X}_0)\|_F \leq \frac{\gamma(\mathbf{X}_0)}{6\sqrt{d}} =: \delta.$$

Set $\rho := 2\delta = \gamma(\mathbf{X}_0)/(3\sqrt{d})$. Then for every $\mathbf{X}$ such that $\|\mathbf{X} - \mathbf{X}_0\| \leq \rho$ it holds that

$$\gamma(\mathbf{X}) \geq \gamma(\mathbf{X}_0) - \|\Phi'(\mathbf{X}_0 - \mathbf{X})\| \geq \gamma(\mathbf{X}_0) - 2\sqrt{d}\rho = \gamma(\mathbf{X}_0)/3 > 0,$$

where the first inequality follows from Weyl's inequality for perturbed singular values. Hence, $\mathbf{X}$ satisfies (2.5) and

$$\begin{aligned}\|\Phi'(\mathbf{X})^+(\Phi'(\mathbf{Y}) - \Phi'(\mathbf{X}))[\mathbf{Y} - \mathbf{X}]\|_F &\leq \frac{3}{\gamma(\mathbf{X}_0)}\|(\Phi'(\mathbf{Y} - \mathbf{X}))[\mathbf{Y} - \mathbf{X}]\| \\ &\leq \frac{6\sqrt{d}}{\gamma(\mathbf{X}_0)}\|\mathbf{Y} - \mathbf{X}\|_F^2 = \omega\|\mathbf{Y} - \mathbf{X}\|_F^2.\end{aligned}$$

Because of $\delta\omega \leq 1$, all conditions of a Newton–Kantorovich-like theorem for underdetermined systems are satisfied, and the claim of the theorem follows from [3, Theorem 4.19]. ∎

5.3. Numerical examples. Numerically, we observed rather robust convergence of the Newton method (5.3)–(5.4). In the following, we give examples that are representative of our observations.

Example 5.2. *We created a random $10 \times 10 \times 10$ tensor $\mathbf{X}_0$, computed the Gram matrices, and perturbed every Gram matrix by a random perturbation of norm ε. If a perturbed matrix happens to be indefinite, it is shifted by the smallest eigenvalue to become positive semidefinite. All Gram matrices are normalized to have trace 1. The Newton method has been applied with starting tensor $\mathbf{X}_0$ to match the perturbed Gram matrices. Figure 1 shows the obtained results for $\varepsilon \in \{10^{-1}, 10^{-2}, 10^{-3}\}$. For $\varepsilon = 10^{-2}$ and $\varepsilon = 10^{-3}$, we clearly observe local quadratic convergence. For $\varepsilon = 10^{-1}$, one of the Gram matrices has a zero eigenvalue. Clearly, such perturbed Gram matrices cannot represent an interior point in the image of Φ, and therefore we cannot expect local quadratic convergence. It turns out that the Newton method still converges, but it deteriorates to linear convergence.*

Example 5.3. *We use the same tensor $\mathbf{X}_0$ as in Example 5.2 as a starting point, but now prescribe, rather arbitrarily, diagonal Gram matrices*

$$G^{(1)} = G^{(2)} = G^{(3)} = \frac{1}{10}\,\mathrm{diag}(1, 2, \ldots, 10). \tag{5.7}$$

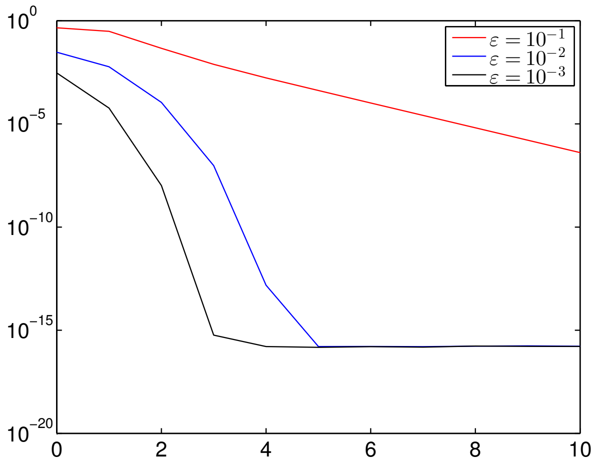

Figure 1. *Error* $\|\Phi(\mathbf{X}_n) - (G^{(1)}, \ldots, G^{(d)})\|$ *vs. iteration number* n *for Example* 5.2.

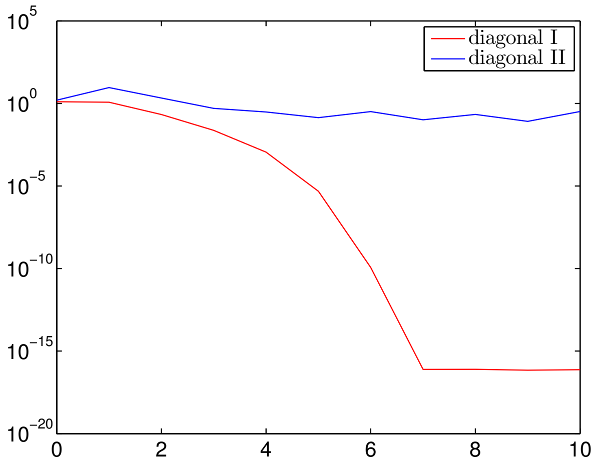

Figure 2. *Error* $\|\Phi(\mathbf{X}_n) - (G^{(1)}, \ldots, G^{(d)})\|$ *vs. iteration number* n *for Example* 5.3*, using the diagonal Gram matrices* (5.7) *(*`diagonal I`*) and* (5.8) *(*`diagonal II`*).*

The Newton method still converges quadratically; see Figure 2*. Having diagonal Gram matrices, the resulting tensor is an* HOSVD tensor *as defined in* [8]*. Despite the fact that all the Gram matrices actually are equal, the resulting tensor is* not *diagonal. In fact, a diagonal tensor never satisfies condition* (2.5)*.*

We now modify the last Gram matrix as follows:

$$G^{(1)} = G^{(2)} = \frac{1}{10}\,\mathrm{diag}(1, 2, \ldots, 10), \quad G^{(3)} = \mathrm{diag}(1, 0, \ldots, 0). \tag{5.8}$$

This is not a feasible configuration, and, not surprisingly, the Newton method does not converge for this example.

6. Conclusions and open problems. In this work, we have shown that the higher-order singular values can be moved in arbitrary directions for almost every tensor with identical mode sizes. Numerical evidence suggests that this property also holds for tensors with unequal, compatible mode sizes. While our results reveal insights into the independence of the higher-order singular values, a complete characterization of the set of feasible higher-order singular values remains an open problem. Also, it would be interesting to extend our results to other subspace-based tensor decompositions, such as the tensor train and hierarchical Tucker decompositions, for which one has to investigate more complicated systems of tensor matricizations.

REFERENCES

[1] M. Bachmayr, R. Schneider, and A. Uschmajew, *Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations*, Found. Comput. Math., 16 (2016), pp. 1423–1472.
[2] B. W. Bader, T. G. Kolda, et al., *MATLAB Tensor Toolbox Version* 2.6, http://www.sandia.gov/~tgkolda/TensorToolbox/, 2015.
[3] P. Deuflhard, *Newton Methods for Nonlinear Problems. Affine Invariance and Adaptive Algorithms*, Springer, Heidelberg, 2011.
[4] J. Dieudonné, *Treatise on Analysis. Vol.* III, Academic Press, New York, 1972.
[5] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, MD, 2013.
[6] L. Grasedyck, D. Kressner, and C. Tobler, *A literature survey of low-rank tensor approximation techniques*, GAMM-Mitt., 36 (2013), pp. 53–78.
[7] W. Hackbusch, *Tensor Spaces and Numerical Tensor Calculus*, Springer, Heidelberg, 2012.
[8] W. Hackbusch and A. Uschmajew, *On the interconnection between the higher-order singular values of real tensors*, Numer. Math., 135 (2017), pp. 875–894.
[9] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM Rev., 51 (2009), pp. 455–500, https://doi.org/10.1137/07070111X.
[10] L. De Lathauwer, B. De Moor, and J. Vandewalle, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1253–1278, https://doi.org/10.1137/S0895479896305696.