A RIEMANNIAN TRUST REGION METHOD FOR THE CANONICAL TENSOR RANK APPROXIMATION PROBLEM*

PAUL BREIDING[†] AND NICK VANNIEUWENHOVEN[‡]

Abstract. The canonical tensor rank approximation problem (TAP) consists of approximating a real-valued tensor by one of low canonical rank, which is a challenging non-linear, non-convex, constrained optimization problem, where the constraint set forms a non-smooth semi-algebraic set. We introduce a Riemannian Gauss–Newton method with trust region for solving small-scale, dense TAPs. The novelty of our approach is threefold. First, we parametrize the constraint set as the Cartesian product of Segre manifolds, hereby formulating the TAP as a Riemannian optimization problem, and we argue why this parametrization is theoretically a good choice. Second, an original ST-HOSVD-based retraction operator is proposed. Third, we introduce a hot restart mechanism that efficiently detects when the optimization process is tending to an ill-conditioned tensor rank decomposition and which often yields a quick escape path from such spurious decompositions. Numerical experiments show improvements of up to three orders of magnitude in terms of the expected time to compute a successful solution over existing state-of-the-art methods.

Key words. CP decomposition, Riemannian optimization, trust region method, Gauss–Newton method, ST-HOSVD retraction, hot restarts

AMS subject classifications. 15A69, 53B21, 53B20, 65K10, 90C53, 14P10, 65Y20, 65F35

1. Introduction. A *simple* or *rank-1 tensor* is the tensor product of vectors. Identifying tensors in coordinates with the d -arrays representing them with respect to some basis, the tensor product is given explicitly by the *Segre map*, namely

$$(1.1) \quad \otimes : \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \times \cdots \times \mathbb{R}^{n_d} \rightarrow \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_d} \\ (\mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^d) \mapsto \left[a_{i_1}^{(1)} a_{i_2}^{(2)} \cdots a_{i_d}^{(d)} \right]_{i_1, i_2, \dots, i_d=1}^{n_1, n_2, \dots, n_d},$$

where $\mathbf{a}^k = [a_i^{(k)}]_{i=1}^{n_k}$. The image of this map comprises the simple tensors. The *tensor rank decomposition* or *canonical polyadic decomposition* (CPD) that was proposed by Hitchcock [35] expresses $\mathfrak{A} \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_d}$ as a linear combination of simple tensors:

$$(CPD) \quad \mathfrak{A} = \sum_{i=1}^r \mathbf{a}_i^1 \otimes \mathbf{a}_i^2 \otimes \cdots \otimes \mathbf{a}_i^d, \quad \mathbf{a}_i^k \in \mathbb{R}^{n_k}.$$

The smallest r such that a decomposition in (CPD) is possible, is called the (canonical) *rank* of $\mathfrak{A}$ and we denote it by $\text{rank}(\mathfrak{A})$.

The set of all tensors of rank bounded by r is then given by

$$\sigma_r := \{\mathfrak{A} \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_d} \mid \text{rank}(\mathfrak{A}) \leq r\}.$$

We say that σ_r is *non-defective* if on a dense subset of σ_r , in the Euclidean topology, a tensor only admits finitely many CPDs. *Throughout this paper, we assume that σ_r is non-defective.* This is not a serious limitation. Provided that $r < \frac{n_1 n_2 \cdots n_d}{n_1 + n_2 + \cdots + n_d - d + 1}$, an even stronger property namely holds [18, 19]: with few exceptions, there exists a Euclidean-dense subset U of σ_r such that every $\mathfrak{A} \in U$ is *r-identifiable*, meaning that the summands in the factorization in (CPD) are even *uniquely determined*.

The CPD arises naturally in a myriad of applications; see, for example, those listed in [57]. An extensive class where potentially high-rank CPDs are computed from dense tensors originates in machine learning and algebraic statistics, where the parameters of certain latent structure models admit a CPD [4, 6, 28, 56]. Such models were recently surveyed in a unified tensor-based framework [6], including exchangeable single topic models, naïve Bayes models, hidden Markov models, and Gaussian mixture models.

*Submitted to the editors.

Funding: The first author was partially supported by DFG research grant BU 1371/2-2. The second author was supported by a Postdoctoral Fellowship of the Research Foundation–Flanders (FWO).

[†]Max-Planck-Institute for Mathematics in the Sciences Leipzig. (breiding@mis.mpg.de)

[‡]KU Leuven, Department of Computer Science. (nick.vannieuwenhoven@cs.kuleuven.be)

Mostly, the tensors that one works with are only approximations of a theoretical low-rank tensor, as a consequence of measurement and representation errors or due to numerical computations. This is why the *tensor rank approximation problem* (TAP), which consists of approximating a given tensor $\mathfrak{B} \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ by a tensor $\mathfrak{A}$ of low canonical rank, i.e.,

$$(\text{TAP}) \quad \min_{\mathfrak{A} \in \sigma_r} \frac{1}{2} \|\mathfrak{A} - \mathfrak{B}\|_F^2,$$

is the usual optimization problem that has to be solved. There are two standard assumptions that we make throughout this paper. First, we seek only one solution of this problem, rather than all solutions. Second, for some inputs $\mathfrak{B}$ this optimization problem is *ill-posed* in the sense that only an *infimum* exists. Unfortunately, the set of all such $\mathfrak{B}$ can have positive Lebesgue measure; see de Silva and Lim [22].

The setting that we are specifically interested in consists of *dense* input tensors $\mathfrak{B}$ that can be well approximated by a tensor of small rank r whose *condition number* [10] (see section 2) is at most moderately large. In particular, we only consider the case of strictly subgeneric ranks $r < \frac{n_1 \dots n_d}{n_1 + \dots + n_d - d + 1}$.

REMARK 1.1 (Tucker compression). *It is known that if a tensor $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ of rank r admits an orthogonal Tucker decomposition [65]*

$$\mathfrak{B} = (Q_1, \dots, Q_d) \cdot \mathfrak{G} := \sum_{i_1=1}^{r_1} \dots \sum_{i_d=1}^{r_d} s_{i_1, \dots, i_d} \mathbf{q}_{i_1}^1 \otimes \dots \otimes \mathbf{q}_{i_d}^d$$

with core tensor $\mathfrak{G} \in \mathbb{R}^{r_1 \times \dots \times r_d}$ and orthonormal factor matrices $Q_k = [\mathbf{q}_i^k]_{i=1}^{r_k} \in \mathbb{R}^{n_k \times r_k}$, then one of $\mathfrak{B}$'s CPDs is $\mathfrak{B} = \sum_{i=1}^r (Q_1 \mathbf{s}_i^1) \otimes \dots \otimes (Q_d \mathbf{s}_i^d)$ where $\mathfrak{G} = \sum_{i=1}^r \mathbf{s}_i^1 \otimes \dots \otimes \mathbf{s}_i^d$ is a CPD of the core tensor $\mathfrak{G}$. The algorithm developed in this paper also applies when an orthogonal Tucker decomposition of $\mathfrak{B}$ is provided as input. In this case it suffices to compute a CPD of the (dense) core tensor $\mathfrak{G}$. This is useful when $r_i \ll n_i$ as it decreases the computational complexity dramatically; see Appendix A.

In practice, if $\mathfrak{B}$ can be well-approximated by a tensor $\mathfrak{A}$ of rank r , then a standard strategy for solving TAPs first constructs an orthogonal Tucker decomposition of multilinear rank componentwise bounded by $(r, \dots, r)$, and then solves the TAP for the core tensor. Since efficient algorithms for Tucker compression exist [14, 27, 47, 54], the key variable determining the computational cost for solving TAPs is typically the size of the approximation rank r rather than the dimensions n_i .

Several algorithms were proposed in the literature for solving (TAP). The state of the art can broadly be divided into two classes: alternating least squares (ALS) methods and general numerical optimization methods. While in 2008 the review article [40] considered ALS (see [15, 16, 32]) as “the ‘workhorse’ algorithm for CP,” we believe that nowadays quasi-Newton methods, such as those in [2, 24, 33, 48, 49, 51, 59, 64], have supplanted ALS-like strategies. For example, the default algorithm in the Tensorlab v3 software package [70] for finding an approximate CPDs is `nls_gnd1`, which is a Gauss–Newton (GN) method with trust region. Furthermore, experiments in [2, 51, 59, 64] have demonstrated that modern implementations of GN methods as described in section 3 outperform standard ALS methods on all but the simplest of problems in terms of execution time and iteration count. For this reason, we do not consider ALS in the rest of this paper.

Instead, we adopt the *Riemannian optimization* framework [1] to solve (TAP). The set σ_r is known to be semi-algebraic [9, 22], entailing that it is locally diffeomorphic to a Euclidean space at most, but not all, points.¹ Riemannian optimization methods [1] exploit such a local smooth structure on the domain for effectively solving optimization problems. Such methods were already proposed for other tensor decompositions, like Tucker, hierarchical Tucker, and tensor train decompositions;

¹More precisely, if σ_r is of dimension d , then by “most points” we mean the locus $U \subset \sigma_r$ of d -dimensional smooth points of σ_r : let $\sigma_r = \bigcup_{i=1}^k \mathcal{M}_i^{d_i}$ be a Nash stratification [9, Proposition 2.9.10], where $\mathcal{M}_i^{d_i}$ is a d_i -dimensional Nash submanifold, and let $S = \{i \mid d_i = d\}$. Then, $U := \bigcup_{j \in S} \mathcal{M}_j^d$, so that U is locally (Nash) diffeomorphic to $\mathbb{R}^d$.

see respectively [38, 55], [20, 66], and [36, 41, 60]. Hitherto, a Riemannian optimization method is lacking for the TAP. This paper develops such a method.

1.1. Contributions. The main contribution of this paper is an efficient Riemannian trust region (RTR) method for finding well-conditioned [10, 67] solutions of (TAP). The numerical experiments in section 7 show that this method can outperform state-of-the-art classic optimization methods on the key performance criterion, namely the expected time to success, by up to three orders of magnitude.

The design of our method was guided by a geometrical analysis of the TAP; see subsections 3.3 and 5.2. We argue that none of the state-of-the-art general optimization methods that were adapted to the TAP exploit the essential geometry of the optimization domain (see section 3). This analysis yields three main insights: first, it suggests that the parametrization of the domain of the optimization problem should be *homothetical*; second, parameterizing σ_r via the manifold of rank-one tensors is probably among the best choices and allows us to formulate the TAP as a Riemannian optimization problem; and, third, certain CPDs are ill-behaved for (Riemannian) Gauss–Newton methods, hence we argue that they should be avoided, e.g., via *hot restarts* (see section 5). This last point motivates the name for our method: Riemannian Gauss–Newton method with hot restarts, or RGN-HR for short.

Unfortunately, the main strength of RGN-HR is also its significant limitation: by design it can find only well-conditioned optimizers of (TAP). In particular, our method can only find *isolated* solutions. We believe that this is not a severe limitation in practice because in many applications only (locally) unique rank-one summands are of interest. We say that non-isolated solutions are *ill-posed* or, equivalently, that their condition number is unbounded; see Proposition 5.4.

1.2. Outline. The rest of the paper is structured as follows. In the next section, we recall the condition number of the CPD. In section 3, we analyze the geometry of the TAP, argue that this structure has not been exploited hitherto, and explain how it can be exploited in a Riemannian optimization method. Section 4 describes the outline of RGN-HR. The strategy for dealing with ill-conditioned Hessian approximations is described in greater detail in section 5. Riemannian optimization methods require the specification of a retraction operator. We explain our choice, a product ST-HOSVD retraction, in section 6. Numerical experiments demonstrating the efficacy of RGN-HR are featured in section 7, and section 8 illustrates the application of our method on a tensor originating in fluorescence spectroscopy. The final section presents our conclusions and outlook.

1.3. Notation. For ease of reading, vectors are typeset in lower-case boldface letters ($\mathbf{x}$); matrices in upper-case letters (A); tensors in upper-case fraktur letters ($\mathfrak{A}, \mathfrak{B}$); and varieties and manifolds in upper-case calligraphic letters ($\mathcal{S}, \mathcal{M}$).

The *vectorization* of a tensor $\mathfrak{A} = [a_{i_1, \dots, i_d}] \in \mathbb{R}^{n_1 \times \dots \times n_d}$ is defined in the usual way as $\text{vec}(\mathfrak{A}) := [a_{1, \dots, 1} \ a_{2, 1, \dots, 1} \ \dots \ a_{n_1, \dots, n_{d-1}, n_d-1} \ a_{n_1, \dots, n_d}]^T$. A real vector space $\mathbb{R}^n$ is endowed with the Euclidean inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T \mathbf{y}$. This induces the Euclidean norm $\|\mathbf{x}\| := \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$. The corresponding *spectral norm* of a matrix $M \in \mathbb{R}^{m \times n}$ is denoted by $\|M\|_2 := \max_{\mathbf{x} \in \mathbb{R}^n, \|\mathbf{x}\|=1} \|M\mathbf{x}\|$. The n -dimensional unit sphere in $\mathbb{R}^{n+1}$ is denoted by $\mathbb{S}^n$. We make $\mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ a normed vector space by defining the norm of a tensor $\mathfrak{A}$ to be the *Frobenius norm* $\|\mathfrak{A}\|_F := \|\text{vec}(\mathfrak{A})\|$.

The Moore–Penrose *pseudoinverse* of M [61, chapter 3, section 1] is denoted by $M^\dagger$. We denote by $\varsigma_{\max}(M)$, $\varsigma_{\min}(M)$, and $\varsigma_k(M)$, the largest, smallest, and k th largest singular value, respectively, of the linear operator or matrix M .

The set of rank-1 tensors in $\mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ is a smooth manifold called the *Segre manifold*² [42, section 4.3.5], which we denote by $\mathcal{S}_{n_1, \dots, n_d}$. Moreover, we define

$$\Sigma = \sum_{k=1}^d (n_k - 1) \quad \text{and} \quad \Pi = \prod_{k=1}^d n_k,$$

²It is standard to call the projective variety of which $\mathcal{S}$ is the cone the *Segre variety*. To avoid confusion and for brevity, we prefer “Segre manifold” to “affine cone over the Segre variety.”

such that $\dim(\mathcal{S}_{n_1, \dots, n_d}) = \Sigma + 1$ and $\dim(\mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}) = \Pi$. We often abbreviate $\mathcal{S} := \mathcal{S}_{n_1, \dots, n_d}$ when the tuple of integers $(n_1, \dots, n_d)$ is clear from the context.

Let $\mathcal{M}$ be a manifold and $x \in \mathcal{M}$. The *tangent space* to $\mathcal{M}$ at x is denoted by $T_x \mathcal{M}$, and the *tangent bundle* of $\mathcal{M}$ is denoted by $\mathcal{T}\mathcal{M}$; see Lee [43].

2. The condition number of CPDs. We recall the condition number of the CPD from [10]. It plays a pivotal role in the analyses in sections 3 and 5 which led to the main insights for improving the state-of-the-art optimization methods.

The *Segre manifold* in $\mathbb{R}^{n_1 \times \dots \times n_d}$ is denoted by $\mathcal{S} := \mathcal{S}_{n_1, \dots, n_d}$. It can be obtained as the image of the Segre map, as defined in (1.1), after removing the zero tensor:

$$\mathcal{S} = \{\mathbf{a}^1 \otimes \dots \otimes \mathbf{a}^d \mid \mathbf{a}^k \in \mathbb{R}^{n_k}\} \setminus \{0\}.$$

The set of tensors of rank at most r is then defined as the image of the *addition map*

$$(2.1) \quad \Phi : \mathcal{S}^{\times r} := \mathcal{S} \times \dots \times \mathcal{S} \rightarrow \mathbb{R}^{n_1 \times \dots \times n_d}, (\mathfrak{B}_1, \dots, \mathfrak{B}_r) \mapsto \mathfrak{B}_1 + \dots + \mathfrak{B}_r.$$

In section 1, we denoted it by σ_r , but for emphasizing the dependence on $\mathcal{S}$ we will write $\sigma_r(\mathcal{S}) := \Phi(\mathcal{S}^{\times r})$. The derivative of Φ at $\mathbf{p} = (\mathfrak{B}_1, \dots, \mathfrak{B}_r) \in \mathcal{S}^{\times r}$ is

$$d_{\mathbf{p}}\Phi : T_{\mathfrak{B}_1}\mathcal{S} \times \dots \times T_{\mathfrak{B}_r}\mathcal{S} \rightarrow T_{\Phi(\mathbf{p})}\mathbb{R}^{\Pi}, (\dot{\mathbf{p}}_1, \dots, \dot{\mathbf{p}}_r) \mapsto \dot{\mathbf{p}}_1 + \dots + \dot{\mathbf{p}}_r.$$

The (local) *condition number* [13, 34] of the CPD at $\mathbf{p}$ is

$$\kappa(\mathbf{p}, \mathfrak{A}) := \lim_{\epsilon \rightarrow 0} \max_{\mathfrak{A}' \in (B_{\epsilon}(\mathfrak{A}) \cap \sigma_r(\mathcal{S}))} \frac{\|\Phi_{\mathbf{p}}^{-1}(\mathfrak{A}) - \Phi_{\mathbf{p}}^{-1}(\mathfrak{A}')\|_F}{\|\mathfrak{A} - \mathfrak{A}'\|_F},$$

where $B_{\epsilon}(\mathfrak{A})$ is an ϵ -ball centered at $\mathfrak{A}$, if there is a local inverse function $\Phi_{\mathbf{p}}^{-1}$ of Φ at $\mathbf{p}$; otherwise, the local condition number is defined as $\kappa(\mathbf{p}, \mathfrak{A}) := \infty$. The norm in the numerator of the above definition is the Euclidean norm on $(\mathbb{R}^{n_1 \times \dots \times n_d})^{\times r} \simeq \mathbb{R}^{\Pi \times r}$; that is, if $\mathbf{p}' = (\mathfrak{B}'_1, \dots, \mathfrak{B}'_r) \in \mathcal{S}^{\times r}$, then

$$\|\mathbf{p} - \mathbf{p}'\|_F := \|\text{vec}(\mathfrak{B}_1 - \mathfrak{B}'_1) \quad \dots \quad \text{vec}(\mathfrak{B}_r - \mathfrak{B}'_r)\|_F.$$

Note that $\kappa(\mathbf{p}, \mathfrak{A})$ is completely determined by the choice of $\mathbf{p}$. We therefore set $\kappa(\mathbf{p}) := \kappa(\mathbf{p}, \mathfrak{A})$ and call it the *condition number of the decomposition* $\mathbf{p}$. In [10] we showed that it is the inverse of the least singular value of $d_{\mathbf{p}}\Phi$:

$$(2.2) \quad \kappa(\mathbf{p}) = \frac{1}{\varsigma_n(d_{\mathbf{p}}\Phi)}, \quad \text{where } n := r \cdot \dim \mathcal{S} = r(\Sigma + 1).$$

By the assumption that σ_r is non-defective we have $\dim \mathcal{S}^{\times r} = \dim \sigma_r(\mathcal{S})$. Hence, the condition number is finite on a dense subset of $\mathcal{S}^{\times r}$; this follows from [10, proof of Theorem 1.1]. If $\mathfrak{A} \in \sigma_r$ is r -identifiable then all $\mathbf{p} \in \Phi^{-1}(\mathfrak{A})$ admit the same condition number [10]. In this case it makes sense to speak about the *condition number of* $\mathfrak{A}$.

Recall that the condition number describes the first-order behavior of the local inverse of Φ at $\mathbf{p}$, if it exists. It thus measures the *sensitivity* of $\mathbf{p}$ with respect to perturbations of $\mathfrak{A} = \Phi(\mathbf{p})$ within $\sigma_r(\mathcal{S})$. A large condition number means that a small perturbation $\mathfrak{A}' \approx \mathfrak{A}$ can result in a strongly perturbed CPD $\mathbf{p}' = (\mathfrak{B}'_1, \dots, \mathfrak{B}'_r) \in \mathcal{S}^{\times r}$. This is because we have the asymptotically sharp bound $\|\mathbf{p}' - \mathbf{p}\|_F \lesssim \kappa(\mathbf{p}') \cdot \|\mathfrak{A}' - \mathfrak{A}\|_F$. What about small condition numbers? Recall from [10] that the relative error between $\mathfrak{B}'_i$ in $\mathbf{p}'$ and $\mathfrak{B}_i$ in $\mathbf{p} = (\mathfrak{B}_1, \dots, \mathfrak{B}_r) \in \mathcal{S}^{\times r}$ is asymptotically bounded by

$$\frac{\|\mathfrak{B}'_i - \mathfrak{B}_i\|_F}{\|\mathfrak{B}'_i\|_F} \lesssim \kappa_i^{\text{rel}}(\mathbf{p}') \frac{\|\mathfrak{A}' - \mathfrak{A}\|_F}{\|\mathfrak{A}'\|_F}, \quad \text{where } \kappa_i^{\text{rel}}(\mathbf{p}') := \kappa(\mathbf{p}') \frac{\|\mathfrak{A}'\|_F}{\|\mathfrak{B}'_i\|_F}.$$

If $\kappa(\mathbf{p}') \|\mathfrak{A}' - \mathfrak{A}\|_F > \|\mathfrak{B}'_i\|_F$, then this “upper bound” is larger than 1; hence, interpreting $\mathfrak{B}'_i$ in an application is nonsensical without further analysis. Note that a CPD with a small condition number $\kappa(\mathbf{p}')$ can contain rank-1 terms $\mathfrak{B}'_i$ that are well-conditioned, i.e., $\|\mathfrak{B}'_i\|_F \approx \|\mathfrak{A}'\|_F$, and others that are ill-conditioned, i.e., $\|\mathfrak{B}'_i\|_F \ll \|\mathfrak{A}'\|_F$.

For the above reasons, in applications, where one naturally faces approximate data, a small condition number is required to ensure that at least some rank-1 tensors in the computed CPD $\mathbf{p}'$ are close to the corresponding terms in the true CPD $\mathbf{p}$. Following [10], we say that the TAP is *ill-posed* at the CPD $\mathbf{p} \in \mathcal{S}^{\times r}$ if $\kappa(\mathbf{p}) = \infty$. In this paper, we deal exclusively with CPDs that are not ill-posed solutions of (TAP). That is, we always assume that a (local) optimizer $\mathbf{p}$ of (TAP) has $\kappa(\mathbf{p}) < \infty$. To the best of our knowledge, RGN-HR is the first method that exploits this property.

3. Analysis of the TAP. In this section we explain how the approach for solving (TAP) in this article differs from existing methods. First, we summarize the state of the art, then recall the Riemannian optimization framework from [1] and propose our Riemannian formulation of the TAP, next consider the geometry behind the TAP to motivate theoretically why we think that our formulation is preferable, and finally argue that the local convergence rate estimates of the Riemannian formulation are superior to those of state-of-the-art optimization methods in certain cases.

3.1. State of the art. All of the optimization methods that we know of, i.e., [2, 15, 16, 23, 24, 32, 33, 44, 48–51, 59, 64], choose a standard parameterization of tensors of rank bounded by r , called *factor matrices*. That is, a tensor $\mathfrak{A}$ is represented as

$$(3.1) \quad \mathfrak{A} := \llbracket A_1, A_2, \dots, A_d \rrbracket := \sum_{i=1}^r \mathbf{a}_i^1 \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d,$$

where $A_k := [\mathbf{a}_i^k]_{i=1}^r \in \mathbb{R}^{n_k \times r}$. The TAP is then formulated as the following unconstrained optimization problem over $D := \mathbb{R}^{n_1 \times r} \times \dots \times \mathbb{R}^{n_d \times r}$,

$$(3.2) \quad \min_{(A_1, \dots, A_d) \in D} \frac{1}{2} \|\llbracket A_1, \dots, A_d \rrbracket - \mathfrak{B}\|_F^2,$$

which can be handled with traditional methods. In the remainder, we restrict our discussion to the widespread, effective GN methods for this problem, such as those proposed in [2, 33, 48, 49, 51, 59, 64].

A key step in GN methods is (approximately) solving the normal equations $(J^T J)\mathbf{p} = -J^T \mathbf{r}$, where $\mathbf{r} = \text{vec}(\llbracket A_1, A_2, \dots, A_d \rrbracket - \mathfrak{B})$ is the residual vector and the Jacobian matrix $J \in \mathbb{R}^{\Pi \times r(\Sigma+d)}$ has the following well-known block structure:

$$(3.3) \quad J = \begin{bmatrix} [I_{n_1} \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d]_{i=1}^r & \dots & [\mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^{d-1} \otimes I_{n_d}]_{i=1}^r \end{bmatrix}.$$

The matrix $J^T J$ is the GN approximation of the Hessian of the objective function in (3.2). Given the (approximate) solution $\mathbf{p}$ of the normal equations, the next iterate is then determined either by a line or plane search, or via the trust region framework.

A complication is that the Gram matrix $J^T J$ is never of full rank for geometrical reasons, so that the normal equations are ill-posed; see [67, section 2] and [subsection 3.3](#) below. The aforementioned methods differ in how they resolve this problem; for example, it can be solved by adding a regularization term, via a truncated conjugate gradient (CG) method, or by computing the Moore–Penrose pseudoinverse.

REMARK 3.1. *For reasons of numerical accuracy it is in general not recommended to solve the normal equations; rather one usually solves the equivalent least-squares problem $J\mathbf{p} = -\mathbf{r}$. However, in the case of the TAP, both the Jacobian matrix J and its Gram matrix $J^T J$ are very structured. The usual QR-factorization in a least-squares solver destroys this structure, resulting in significantly slower solves and excessive memory consumption. For this reason, the references [2, 33, 48, 49, 51, 59, 64] prefer solving the normal equations rather than the least-squares problem.*

3.2. A Riemannian formulation. Classic optimization methods can be generalized to the framework of *Riemannian optimization* [1] for solving optimization problems where the constraint set $\mathcal{M}$ is a *Riemannian manifold* [43, chapter 13]. A particular instance of a Riemannian optimization problem that is of relevance to this paper is the least-squares optimization problem

$$(3.4) \quad \min_{x \in \mathcal{M}} \frac{1}{2} \|F(x)\|^2,$$

where $F : \mathcal{M} \rightarrow \mathbb{R}^M$ is a smooth objective function with $M \geq \dim \mathcal{M}$. Riemannian optimization methods are defined with respect to a *Riemannian metric* [43, chapter 13] and a *retraction operator* R on $\mathcal{M}$; see [1, section 4.1], and [section 6](#) below. Since we will assume that $\mathcal{M}$ is embedded in $\mathbb{R}^N$, we take the metric induced by this ambient Euclidean space. The classic GN method generalizes to functions on manifolds; the outline of this Riemannian Gauss–Newton (RGN) method for solving (3.4) from [1, Algorithm 14] is recalled as Algorithm 3.1.

Algorithm 3.1: Riemannian Gauss-Newton method

Data: Riemannian manifold $\mathcal{M}$; retraction R on $\mathcal{M}$; function $F : \mathcal{M} \rightarrow \mathbb{R}^M$.
Input: A starting point $x_0 \in \mathcal{M}$.
Output: A sequence of iterates x_k .
1 **for** $k = 0, 1, 2, \dots$ **do**
2 Solve the GN equation $((d_{x_k} F)^* \circ (d_{x_k} F)) \eta_k = -(d_{x_k} F)^*(F(x_k))$ for the unknown $\eta_k \in T_{x_k} \mathcal{M}$;
3 Set $x_{k+1} \leftarrow R_{x_k}(\eta_k)$;
4 **end**

One might hope that the TAP could be formulated as in (3.4) with domain $\sigma_r(\mathcal{S})$; alas, this set is *not* a manifold. However, the product $\mathcal{S}^{\times r}$ is a manifold. We propose parameterizing $\sigma_r(\mathcal{S})$ via the addition map Φ , so that we can formulate the TAP as

$$(TAP^*) \quad \min_{\mathbf{p} \in \mathcal{S}^{\times r}} \frac{1}{2} \|\Phi(\mathbf{p}) - \mathfrak{B}\|_F^2.$$

Assuming that (TAP) is well posed, it is clear that $\mathbf{p} \in \mathcal{S}^{\times r}$ is a global minimizer of (TAP*) if and only if $\Phi(\mathbf{p})$ is a global minimizer of (TAP). It is also evident that every $\mathbf{p}$ such that $\Phi(\mathbf{p})$ is a local minimizer of (TAP) is a local minimizer of (TAP*). However, while we have no explicit examples, we expect based on considerations about r -identifiability that the converse is false at least for some $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$; that is, the formulation (TAP*) may introduce additional local minima.

3.3. Analysis. The formulation as Riemannian optimization problem in (TAP*) distinguishes the proposed approach from the state-of-the-art methods described in subsection 3.1 which solve (3.2) instead. In this subsection, we interpret both approaches as special cases of a family of possible Riemannian formulations of the TAP. We argue that in this family *the formulation in (TAP*) is theoretically a good choice when a RGN method is utilized for solving it*, the condition number of the obtained solution is small, and orthonormal bases are selected for the tangent spaces.

The parameterizations of $\sigma_r(\mathcal{S})$ via factor matrices in $\mathbb{R}^{n_1 \times r} \times \dots \times \mathbb{R}^{n_d \times r}$ and via rank-1 tensors in $\mathcal{S}^{\times r}$ are both special cases of the following scenario. Let $\mathcal{E}$ be a Riemannian manifold with $\dim \mathcal{E} \geq \dim \mathcal{S}^{\times r}$, and let $\Psi : \mathcal{E} \rightarrow \sigma_r(\mathcal{S})$ be a smooth, surjective parametrization of $\sigma_r(\mathcal{S})$ with corresponding optimization problem

$$(3.5) \quad \min_{x \in \mathcal{E}} \frac{1}{2} \|\Psi(x) - \mathfrak{B}\|_F^2.$$

We furthermore assume that there is a smooth map $\pi : \mathcal{E} \rightarrow \mathcal{S}^{\times r}$ so that the diagram

$$\begin{array}{ccc} \mathcal{E} & \xrightarrow{\pi} & \mathcal{S}^{\times r} \\ & \searrow \Psi & \downarrow \Phi \\ & & \sigma_r(\mathcal{S}) \end{array}$$

commutes, where Φ is the addition map from (2.1). We believe that a meaningful parametrization of $\sigma_r(\mathcal{S})$ should supply a method π for obtaining the collection of rank-1 summands in (CPD), as they are often of interest in applications. This explains why the assumption on the existence of π is justified.

In the case of factor matrices, π is the columnwise Khatri–Rao product $\odot$, i.e., $\pi(A_1, \dots, A_d) = A_1 \odot \dots \odot A_d := [\mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d]_{i=1}^r$, and $\Psi = \llbracket \cdot \rrbracket$ is the map from (3.1). On the other hand, in our approach, π is the identity map and $\Psi = \Phi$. Both the factor matrices and the parametrization over $\mathcal{S}^{\times r}$ may serve as guiding examples, but in the following discussion ($\mathcal{E}, \Psi$) can be any suitable parametrization of $\sigma_r(\mathcal{S})$.

Let $F(x) = \Psi(x) - \mathfrak{B}$. After fixing a choice of orthonormal bases for $T_x \mathcal{E}$ and $T_{\Psi(x)} \mathbb{R}^{n_1 \times \dots \times n_d}$, the RGN method applied to optimization problem (3.5) solves, in coordinates, either the normal equations or a least squares problem, i.e., either

$$(J^T J) \eta = -J^T (\Psi(x) - \mathfrak{B}) \quad \text{or} \quad J \eta = -(\Psi(x) - \mathfrak{B}),$$

respectively, where J is the matrix of $d_x F = d_x \Psi$ with respect to the chosen bases. The sensitivity of both computational problems depends critically on the *effective condition number* of J , namely $\kappa_2(J) := \|J^\dagger\|_2 \|J\|_2$; see, e.g., [29, section 5.3.8].

We now show that some parameterizations are inherently worse than others in the sense that $\kappa_2(J)$ will be large. To this end, we require the mild technical assumption that π is a surjective *submersion*: $d_x\pi$ is of maximal rank $n := \dim \mathcal{S}^{\times r}$ for all $x \in \mathcal{E}$. One verifies that the identity and Segre maps are surjective submersions.

Let $x \in \mathcal{E}$ and $\mathbf{p} := \pi(x) \in \mathcal{S}^{\times r}$. Choose any orthonormal basis for $T_{\mathbf{p}}\mathcal{S}^{\times r}$, and let P , respectively Q , be the matrix of $d_x\pi$, respectively $d_{\mathbf{p}}\Phi$, with respect to the chosen bases on the three tangent spaces. Let $m := \dim \mathcal{E} \geq \dim \mathcal{S}^{\times r} = n$, then we have $J \in \mathbb{R}^{\Pi \times m}$, $P \in \mathbb{R}^{n \times m}$, and $Q \in \mathbb{R}^{\Pi \times n}$. From the chain rule and the commuting diagram above it follows that $J = QP$. As in (2.2), let $\kappa(\mathbf{p})$ denote the condition number of the CPD at $\mathbf{p}$. Recall that the singular values of a linear operator are the singular values of any of its matrix representations with respect to orthonormal bases on the domain and image. Hence, $\kappa(\mathbf{p}) = (\varsigma_n(d_{\mathbf{p}}\Phi))^{-1} = (\varsigma_n(Q))^{-1}$. As explained before, we assume that $\kappa(\mathbf{p}) < \infty$, so that it follows that Q is injective and thus has a left inverse $Q^\dagger$. Moreover, P has full row rank because we assumed $d_x\pi$ is a submersion, so that P has a right inverse $P^\dagger$. We find $JP^\dagger = Q$ with pseudoinverse $PJ^\dagger = Q^\dagger$. On the other hand, we also have $Q^\dagger J = P$ and $J^\dagger Q = P^\dagger$. From the submultiplicativity of the spectral norm, we then derive $\|Q\| \leq \|J\|\|P^\dagger\|$, $\|Q^\dagger\| \leq \|P\|\|J^\dagger\|$, $\|P\| \leq \|J\|\|Q^\dagger\|$, and $\|P^\dagger\| \leq \|J^\dagger\|\|Q\|$. Combining the first and second bounds, and the third and fourth bounds yields, respectively,

$$(3.6) \quad \kappa_2(Q) \leq \kappa_2(J)\kappa_2(P) \quad \text{and} \quad \kappa_2(P) \leq \kappa_2(J)\kappa_2(Q).$$

We can write these condition numbers in terms of the singular values of the derivative operators, namely

$$(3.7) \quad \kappa_2(P) = \frac{\varsigma_1(d_x\pi)}{\varsigma_n(d_x\pi)}, \quad \kappa_2(Q) = \frac{\varsigma_1(d_{\mathbf{p}}\Phi)}{\varsigma_n(d_{\mathbf{p}}\Phi)} \quad \text{and} \quad \kappa_2(J) = \frac{\varsigma_1(d_x\Psi)}{\varsigma_n(d_x\Psi)},$$

which are intrinsic expressions that do not depend on the choice of bases. For proceeding, we need the following lemma, which is proved in [Appendix B](#).

LEMMA 3.2. *The spectral norm of $d_{\mathbf{p}}\Phi$ satisfies $1 \leq \|d_{\mathbf{p}}\Phi\| \leq \sqrt{r}$.*

Since $\|d_{\mathbf{p}}\Phi\| = \varsigma_1(d_{\mathbf{p}}\Phi)$ and $\varsigma_n(d_{\mathbf{p}}\Phi) = \kappa(\mathbf{p})^{-1}$, we derive from plugging (3.7) into (3.6), and using [Lemma 3.2](#) that

$$(3.8) \quad \max \left\{ \kappa(\mathbf{p}) \cdot \frac{\varsigma_n(d_x\pi)}{\varsigma_1(d_x\pi)}, \frac{1}{\sqrt{r}} \cdot \left(\kappa(\mathbf{p}) \frac{\varsigma_n(d_x\pi)}{\varsigma_1(d_x\pi)} \right)^{-1} \right\} \leq \kappa_2(J).$$

Note that the condition number $\kappa(\mathbf{p})$ should be interpreted as a constant here, because it is an inherent property of the decomposition $\mathbf{p}$ that does not depend on the parametrization we choose for $\mathcal{S}^{\times r}$.

Based on the above lower bound, a natural idea is constructing a parameterization $\mathcal{E}$ with maps Ψ and π in such a way that the lower bound in (3.8) is minimized. However, we do not know how it should be constructed, if it exists at all. Therefore, we settle for a good choice instead. Recall from the previous section that decompositions $\mathbf{p}$ with large condition number are often not of interest in applications, because the components cannot be interpreted rigorously due to their sensitivity to noise. We thus focus on those cases with a small $\kappa(\mathbf{p})$, in which case it is clear that a small lower bound in (3.8) is obtained by choosing a parametrization π with likewise small $\kappa_2(P)$. The minimal $\kappa_2(P) = 1$ and corresponds to a *homothetical* parametrization $\pi : \mathcal{E} \rightarrow \mathcal{S}^{\times r}$; that is, $\varsigma_1(d_x\pi) = \varsigma_n(d_x\pi)$ for all $x \in \mathcal{E}$. Based on this analysis, we propose the choice $\pi = \text{id}$, yielding the lower bound $\kappa(\mathbf{p}) \leq \kappa_2(J)$.

An example that may lead to a large lower bound in (3.8) is the parameterization via factor matrices with π equal to the r -fold Segre map. In this case, the parameter space is $\mathcal{E} = \mathbb{R}^{n_1 \times r} \times \dots \times \mathbb{R}^{n_d \times r}$. Assume that $(A_1, \dots, A_d) \in \mathcal{E}$. Then, $d_x\pi$ with respect to the standard basis on $T_x\mathcal{E} \simeq \mathbb{R}^{n_1 r + \dots + n_d r}$ and the standard basis on $T_{\pi(x)}(\mathbb{R}^\Pi)^{\times r} \simeq \mathbb{R}^{r\Pi}$ is

$$T := \text{diag}(T_1, \dots, T_r), \quad \text{where } T_i = [I \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d \quad \dots \quad \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^{d-1} \otimes I]$$

and $A_k = [\mathbf{a}_i^k]_{i=1}^r$. The matrix T_i is what [\[67\]](#) calls a *Terracini matrix* corresponding to a rank-1 tensor. Its spectrum was completely characterized in [\[67, corollary 18\]](#) in

the *norm-balanced case*, i.e., the case whereby $\alpha_i := \|\mathbf{a}_i^1\| = \dots = \|\mathbf{a}_i^d\|$. It follows from these results that $\varsigma_1(T_i) = \sqrt{d}\alpha_i^{d-1}$ and the least nonzero singular value is α_i^{d-1} . From the block diagonality of T , it follows that $\frac{\varsigma_1(d_x \pi)}{\varsigma_n(d_x \pi)} = \sqrt{d} \left(\frac{\max_i \alpha_i}{\min_i \alpha_i} \right)^{d-1}$, which is large for all tensors with large differences in norm between the rank-1 tensors, yielding the lower bound $\frac{\sqrt{d}}{\sqrt{r} \kappa(\mathbf{p})} \left(\frac{\max_i \alpha_i}{\min_i \alpha_i} \right)^{d-1} \leq \kappa_2(J)$.

3.4. Convergence. We showed in [11] that the multiplicative constant in the local linear rate of convergence of a RGN method to a non-exact local optimizer $\mathbf{p}^*$, i.e., $f(\mathbf{p}^*) \neq 0$, is a multiple of the square of the condition number $\kappa(\mathbf{p}^*)$ in (2.2).

Likewise, it follows from the proof³ of Theorem 7 in [25] that the multiplicative constant in the local linear rate of convergence of GN methods to a non-exact solution x^* of the objective function $f(x) = \frac{1}{2}\|g(x)\|^2$ whose Jacobian J_g has constant rank on an open neighborhood containing x^* is a multiple of $\|J_g^\dagger\|_2^2$. Consequently, the local convergence behavior of the GN methods mentioned in subsection 3.1 applied to the TAP is governed by $\|J^\dagger\|_2^2$ with J as in (3.3). In [67] it was argued that $\kappa(\mathbf{p}) := \|J^\dagger\|_2$ is naturally the condition number of the TAP in (3.2) at the CPD $\mathbf{p}$.

As we have shown at the end of the previous subsection, $\|J^\dagger\|_2$ with J as in (3.3) will be large at $\mathbf{p} = (\mathfrak{B}_1, \dots, \mathfrak{B}_r)$ if there are large differences in the scaling factors $\alpha_i := \|\mathfrak{B}_i\|_F$ of the rank-1 terms in $\mathbf{p}$. The convergence of GN methods will, at least theoretically, deteriorate with increasing $\frac{\max_i \alpha_i}{\min_i \alpha_i}$. However, since we proved in [10, Proposition 5.1] that the condition number in (2.2) is *independent of the scaling of the rank-1 terms* appearing in the CPD $\mathbf{p}$, i.e., $\kappa((t_1 \mathfrak{B}_1, \dots, t_r \mathfrak{B}_r)) = \kappa(\mathbf{p})$ for all $t_i \in \mathbb{R} \setminus \{0\}$, it follows that, theoretically, the local rate of convergence of RGN methods for (TAP*) is not affected by large differences in scaling. *These methods might thus significantly outperform the GN methods for (3.2) for CPDs with large differences in scale.* This behavior is confirmed experimentally in section 7.

4. The proposed Riemannian trust region method. The considerations of the preceding section and because we are not aware of any evident homothetical parametrization of $\mathcal{S}^{\times r}$ other than the trivial one, i.e., $\pi = \text{id}$, motivated us to avoid optimizing over a flat parameter space. Instead we will optimize over the Riemannian manifold $\mathcal{S}^{\times r}$. To this end, let us write

$$f(\mathbf{p}) = \frac{1}{2} \|F(\mathbf{p})\|^2 \quad \text{with } F : \mathcal{S}^{\times r} \rightarrow \mathbb{R}^\Pi, \mathbf{p} \mapsto \text{vec}(\Phi(\mathbf{p}) - \mathfrak{B}),$$

where $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ is the tensor that we want to approximate. In this notation, (TAP*) is written as $\min_{\mathbf{p} \in \mathcal{S}^{\times r}} f(\mathbf{p})$. We propose a RGN method with *trust region* [1, section 7] to solve this non-linear least squares problem. The key innovation is the addition of a scheme we call *hot restarts* for effectively dealing with ill-conditioned Hessian approximations; see section 5.

Recall that in RGN methods with trust region the trust region subproblem (TRS) consists of (approximately) solving

$$(4.1) \quad \min_{\mathbf{p} \in B_\Delta} m_{\mathbf{p}}(\mathbf{p}) \quad \text{with } m_{\mathbf{p}}(\mathbf{p}) := f(\mathbf{p}) + \langle d_{\mathbf{p}} f, \mathbf{p} \rangle + \frac{1}{2} \langle \mathbf{p}, ((d_{\mathbf{p}} \Phi)^* \circ d_{\mathbf{p}} \Phi)(\mathbf{p}) \rangle,$$

where the trust region $B_\Delta := \{\mathbf{p} \in T_{\mathbf{p}} \mathcal{S}^{\times r} \mid \|\mathbf{p}\| \leq \Delta\}$ with *trust region radius* $\Delta > 0$. The high-level outline of RGN-HR is stated as Algorithm 4.1. Aside from the addition of step 2, this is a conventional RGN with trust region method; see [1, chapter 7]. The reader familiar with RTR methods can skim the next subsections, where we specialize the necessary components to the TAP. The interesting constituents, namely the hot restarts technique, hidden in step 7, and the novel retraction operator in step 8 are described in sections 5 and 6 respectively.

4.1. Choice of parameterization. We can choose *any* convenient representation of points on $\mathcal{S}^{\times r}$. An efficient data-sparse choice consists of representing a rank-1 tensor $p_i = \alpha_i \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d$, where $\mathbf{a}_i^k \in \mathbb{S}^{n_k-1}$, as $(\alpha_i, \mathbf{a}_i^1, \dots, \mathbf{a}_i^d)$. However, representing an element $\mathbf{v} \in \mathbb{S}^n \subset \mathbb{R}^{n+1}$ using n parameters is not convenient from a

³Specifically the penultimate step in the penultimate formulae on page 208.

Algorithm 4.1: RGN method with trust region for the TAP

Data: Retraction R on $\mathcal{S}^{\times r}$.
Input: The tensor to approximate $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$.
Output: A sequence of iterates $\mathbf{p}^{(k)} \in \mathcal{S}^{\times r}$.
1 Choose random initial points $p_i \in \mathcal{S} \subset \mathbb{R}^\Pi$;
2 Solve the least-squares problem $[p_i]_{i=1}^r \mathbf{x} = \text{vec}(\mathfrak{B})$ for $\mathbf{x}$, and set $p_i \leftarrow x_i p_i$;
3 Let $\mathbf{p}^{(1)} \leftarrow (p_1, \dots, p_r)$, and set $k \leftarrow 0$;
4 Choose a trust region radius $\Delta > 0$;
5 **while** the method has not converged **do**
6 Compute the gradient and GN Hessian approximation;
7 Solve the TRS (4.1) for the search direction $\mathbf{p}_k \in B_\Delta \subset T_{\mathbf{p}} \mathcal{S}^{\times r}$;
8 Compute the tentative next iterate $\mathbf{p}^{(k+1)} \leftarrow R_{\mathbf{p}^{(k)}}(\mathbf{p}_k)$;
9 Compute the trustworthiness ρ_k ;
10 Accept or reject the next iterate based on the trustworthiness ρ_k . If the step was
11 accepted, increment k ;
12 Update the trust region radius Δ .
12 **end**

programming perspective. For this reason, we prefer *norm-balanced* representatives where $p_i = \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d$ is represented as the tuple $\mathbf{p}_i = (\mathbf{a}_i^1, \dots, \mathbf{a}_i^d)$ with $\mathbf{a}_i^k \in \mathbb{R}^{n_k}$ and $\|\mathbf{a}_i^1\| = \dots = \|\mathbf{a}_i^d\|$. This choice incurs a minor cost of $d - 1$ parameters relative to minimal parameterizations. The CPD $\mathbf{p} \in \mathcal{S}^{\times r}$ is then represented sparsely using $r(\Sigma + d)$ parameters as $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_r)$.

We can also choose a convenient basis for the tangent space of $\mathcal{S}^{\times r}$. Suppose that $p_i = \alpha_i \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d$ are the rank-1 tensors, where $\mathbf{a}_i^k \in \mathbb{S}^{n_k-1}$. There is a specific orthonormal basis of $T_{p_i} \mathcal{S}$ that enables computationally efficient implementations of the basic operations, which are discussed in [Appendix A](#). For $k = 1$, let $U_{i,1} := I_{n_1}$ be the standard basis. For $k = 2, \dots, d$, we choose an orthonormal basis $U_{i,k}$ of $T_{\mathbf{a}_i^k} \mathbb{S}^{n_k-1}$ as the first $n_k - 1$ columns of the Q factor of a rank-revealing QR-decomposition $(I_{n_k} - \mathbf{a}_i^k (\mathbf{a}_i^k)^T)P = QR$ with $Q^T Q = I_{n_k}$, R upper triangular, and P a permutation matrix. Then the columns of

$$(4.2) \quad T_{p_i} := [U_{i,1} \otimes \mathbf{a}_i^2 \otimes \dots \otimes \mathbf{a}_i^d \quad \dots \quad \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^{d-1} \otimes U_{i,d}] = [T_{i,1} \quad \dots \quad T_{i,d}]$$

form an orthonormal basis of $T_{p_i} \mathcal{S}$. This $\Pi \times (\Sigma + 1)$ matrix is *never computed explicitly*, but rather we exploit its tensor structure. It suffices to store the matrices $U_{i,k} \in \mathbb{R}^{n_k \times (n_k-1)}$ for $k = 2, \dots, d$ and $i = 1, \dots, r$.

4.2. Random starting points. We choose the initial points $p_i \in \mathcal{S}$ in step 1 randomly by sampling the elements of $\mathbf{a}_i^k$, $i = 1, \dots, r$ and $k = 1, \dots, d$, i.i.d. from a standard normal distribution. Then, $p_i := \mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d \in \mathbb{R}^{n_1 \times \dots \times n_d}$, which is represented sparsely with norm-balanced representatives. After the above random initialization, we solve the linear least squares problem

$$(4.3) \quad \min_{x_1, \dots, x_r \in \mathbb{R}} \left\| \sum_{i=1}^r x_i p_i - \mathfrak{B} \right\|_F$$

to determine the optimal linear combination of the initial p_i 's. This least-squares problem admits a unique solution with probability 1, because the p_i 's are generically linearly independent; see, e.g., [19, Corollary 4.5]. All of the standard methods for solving (4.3) can be employed; in [Appendix A](#) we describe how this operation is implemented efficiently without computing the p_i 's explicitly. Letting $x_1^*, \dots, x_r^*$ denote the optimal coefficients, we take $x_i^* p_i \in \mathcal{S}$ rather than p_i as “optimally scaled” rank-1 tensors in the initial rank- r approximation of $\mathfrak{B}$.

The main motivation for this extra step is that it can be implemented more efficiently than one solve of the TRS. We observed that this modification generally reduces the number of iterations by a small amount. For the same reason, we also employed this modification in the hot restarts strategy discussed in [section 5](#).

4.3. The trust region scheme. The trust region radius Δ is updated according to the standard scheme [46, chapter 4]. Define minimum and maximum radii as

$$\Delta_{\min}(\mathbf{p}) := 10^{-1} \cdot \sqrt{\frac{d}{r} \cdot \sum_{i=1}^r \|\mathbf{a}_i^1\|_F^2} \quad \text{and} \quad \Delta_{\max} := \frac{1}{2} \|\mathfrak{B}\|_F,$$

respectively. As initial radius in step 4 we take $\Delta = \min\{\Delta_{\min}(\mathbf{p}^{(1)}), \Delta_{\max}\}$.

In step 9, the trustworthiness of the model at $\mathbf{p}^{(k)} \in \mathcal{S}^{\times r}$ is evaluated as

$$\rho_k = \frac{f(\mathbf{p}^{(k)}) - f(\mathbf{p}^{(k+1)})}{m_{\mathbf{p}^{(k)}}(0) - m_{\mathbf{p}^{(k)}}(\mathbf{p}_k)}.$$

where $\mathbf{p}_k$ is the solution of the TRS in step 7, and $\mathbf{p}^{(k+1)}$ is computed in step 8. If the trustworthiness exceeds 0.2, then $\mathbf{p}^{(k+1)}$ is accepted as the next iterate in step 10; otherwise, the iterate is rejected.

In step 11, the trust region radius is updated as follows. If the trustworthiness $\rho_k > 0.6$, then the trust region is enlarged by taking $\Delta = \min\{2\|\mathbf{p}_k\|, \Delta_{\max}\}$, where $\mathbf{p}_k$ is the solution of the TRS in step 7. Otherwise, the radius is adjusted as

$$\Delta = \min\left\{\left(\frac{1}{3} + \frac{2}{3} \cdot (1 + e^{-14 \cdot (\rho_k - \frac{1}{3})})^{-1}\right) \Delta, \Delta_{\max}\right\},$$

which is the strategy in Tensorlab's `cpd_nls` method [70] with different constants. The effect of the scaling factor $\frac{1}{3} + \frac{2}{3}(1 + e^{-14(\rho_k - \frac{1}{3})})^{-1}$ is shown below:

When $\rho_k < 0.2$, the trust region radius is strongly reduced by multiplying with approximately $\frac{1}{3}$, between $0.2 \leq \rho_k \leq 0.47$ the reduction factor increases approximately linearly with ρ_k , and finally for $0.47 \leq \rho_k \leq 0.6$ the radius is kept nearly constant.

4.4. Gradient and Hessian approximation in coordinates. For solving the TRS in step 7, we need explicit expressions for the gradient and approximate Riemannian Hessian of the objective function. The derivative of f at $\mathbf{p} \in \mathcal{S}^{\times r}$ is

$$\mathrm{d}_{\mathbf{p}} f = \frac{1}{2} \mathrm{d}_{\mathbf{p}} \langle F(\mathbf{p}), F(\mathbf{p}) \rangle = \langle \mathrm{d}_{\mathbf{p}} F, F(\mathbf{p}) \rangle = \langle \mathrm{d}_{\mathbf{p}} \Phi, F(\mathbf{p}) \rangle.$$

Let $\mathbf{p} = (p_1, \dots, p_r)$ and T_{p_i} be defined as in (4.2). With respect to the standard basis on $\mathbb{R}^{\Pi}$ and the orthonormal basis given by the columns of

$$(4.4) \quad B := \text{diag}(T_{p_1}, \dots, T_{p_r})$$

on the domain, the derivative $\mathrm{d}_{\mathbf{p}} f$ is represented by the *gradient* vector

$$\nabla_{\mathbf{p}} f := T_{\mathbf{p}}^T F(\mathbf{p}) = T_{\mathbf{p}}^T \cdot \text{vec}(\Phi(\mathbf{p}) - \mathfrak{B}),$$

where $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ and $T_{\mathbf{p}}$ is the “Jacobian” matrix of $\mathrm{d}_{\mathbf{p}} \Phi$, i.e.,

$$(4.5) \quad T_{\mathbf{p}} := [T_{p_1} \quad \dots \quad T_{p_r}]$$

with T_{p_i} as in (4.2). Appendix A.2 gives an efficient algorithm for computing $\nabla_{\mathbf{p}} f$.

We approximate the Riemannian Hessian by the GN approximation, namely $(\mathrm{d}_{\mathbf{p}} F)^*(\mathrm{d}_{\mathbf{p}} F)$; see, e.g., [1, section 8.4]. It follows from the above computations that its matrix with respect to the basis B on both the domain and image is $T_{\mathbf{p}}^T T_{\mathbf{p}}$. Efficiently constructing it is covered in Appendix A.2.

Let $\mathbf{t} = B\mathbf{x} \in T_{\mathbf{p}}S^{\times r}$, and define

$$\mathbf{r}_{\mathbf{p}} := \text{vec}(\Phi(\mathbf{p}) - \mathfrak{B}), \quad \mathbf{g}_{\mathbf{p}} := T_{\mathbf{p}}^T \mathbf{r}_{\mathbf{p}}, \quad \text{and} \quad H_{\mathbf{p}} := T_{\mathbf{p}}^T T_{\mathbf{p}}.$$

Then, we find in coordinates with respect to the orthonormal basis B that

$$\begin{aligned} \langle d_{\mathbf{p}}f, \mathbf{t} \rangle &= \langle BT_{\mathbf{p}}^T \mathbf{r}_{\mathbf{p}}, B\mathbf{x} \rangle = \mathbf{g}_{\mathbf{p}}^T \mathbf{x}, \quad \text{and} \\ \langle \mathbf{t}, ((d_{\mathbf{p}}\Phi)^* \circ d_{\mathbf{p}}\Phi)(\mathbf{t}) \rangle &= \langle B\mathbf{x}, BT_{\mathbf{p}}^T T_{\mathbf{p}} B^T B\mathbf{x} \rangle = \mathbf{x}^T H_{\mathbf{p}} \mathbf{x}. \end{aligned}$$

4.5. Solving the TRS. The key step in RTR methods is solving the TRS in (4.1) for obtaining a suitable tangent direction along which the retraction can proceed. Many standard strategies exist for computing or approximating the solution of the TRS [46]. We choose the cheap dogleg heuristic [46, section 4.1].

Note that $H_{\mathbf{p}}$ is symmetric positive semi-definite. Then the *Newton direction* $\mathbf{p}_{\mathbf{N}}$ and the *Cauchy point* are given respectively by

$$(4.6) \quad H_{\mathbf{p}} \mathbf{p}_{\mathbf{N}} := -\mathbf{g}_{\mathbf{p}} \quad \text{and} \quad \mathbf{p}_{\mathbf{C}} = -\frac{\mathbf{g}_{\mathbf{p}}^T H_{\mathbf{p}} \mathbf{g}_{\mathbf{p}}}{\mathbf{g}_{\mathbf{p}}^T \mathbf{g}_{\mathbf{p}}} \mathbf{g}_{\mathbf{p}}.$$

If $\mathbf{p}_{\mathbf{N}}$ is outside of the trust region B_{Δ} , but $\mathbf{p}_{\mathbf{C}} \in B_{\Delta}$, then the dogleg step is the intersection with the boundary of the ball B_{Δ} of vector pointing from $\mathbf{p}_{\mathbf{C}}$ to $\mathbf{p}_{\mathbf{N}}$:

$$\mathbf{p}_{\mathbf{I}} := \mathbf{p}_{\mathbf{C}} + (\tau - 1)(\mathbf{p}_{\mathbf{N}} - \mathbf{p}_{\mathbf{C}}),$$

where $1 \leq \tau \leq 2$ is the unique solution such that $\|\mathbf{p}_{\mathbf{I}}\|^2 = \Delta^2$. Then, the dogleg step approximates the optimal TRS solution of (4.1) by the next rule:

$$(4.7) \quad \hat{\mathbf{p}} = \begin{cases} \mathbf{p}_{\mathbf{N}} & \text{if } \|\mathbf{p}_{\mathbf{N}}\| \leq \Delta, \\ \mathbf{p}_{\mathbf{C}} & \text{if } \|\mathbf{p}_{\mathbf{N}}\| > \Delta \text{ and } \|\mathbf{p}_{\mathbf{C}}\| \geq \Delta, \\ \mathbf{p}_{\mathbf{I}} & \text{otherwise.} \end{cases}$$

As mentioned in Remark 3.1, computing $\mathbf{p}_{\mathbf{N}}$ via the normal equations is not numerically stable; however, a stable least-squares solve via a QR -decomposition is computationally not attractive. We choose solving the normal equations as this requires much fewer operations due to $H_{\mathbf{p}}$'s structure; see Appendix A.

4.6. Stopping criterion. The RGN-HR method is halted based on a multi-component stopping criterion involving 5 parameters; the first 4 are standard: an absolute tolerance τ_f on the objective function value, a relative tolerance $\tau_{\Delta f}$ on the improvement of the objective function value, a relative tolerance on the step size $\tau_{\Delta x}$, and a maximum number of iterations $k_{\max}$. The last parameter $r_{\max}$ specifies the maximum number of restarts (see section 5).

Specifically, the method halts in the k th iteration if either

$$f(\mathbf{p}^{(k)}) \leq \tau_f, \quad \frac{|f(\mathbf{p}^{(k-1)}) - f(\mathbf{p}^{(k)})|}{f(\mathbf{p}^{(1)})} \leq \tau_{\Delta f}, \quad \text{or} \quad \frac{\|\mathbf{p}_k\|}{\left(\sum_{k=1}^d \|A_k\|_F^2\right)^{\frac{1}{2}}} \leq \tau_{\Delta x},$$

where $A_k = [\mathbf{a}_i^k]_{i=1}^r$ are the factor matrices corresponding to the norm-balanced representatives. If neither of these conditions is satisfied after $k_{\max}$ iterations or after $r_{\max}$ restarts, then the method halts.

5. The hot restarts strategy. The main motivation for adding the trust region scheme to (quasi-)Newton methods consists of obtaining a globally convergent method that still has a local superlinear rate of convergence. In principle, the trust region mechanism can handle ill-conditioned Hessian approximations $H_{\mathbf{p}} = T_{\mathbf{p}}^T T_{\mathbf{p}}$, where $T_{\mathbf{p}}$ is as in (4.5), without special considerations. However, we observed in practice that the RGN with trust region method progressed very slowly near ill-conditioned $H_{\mathbf{p}}$'s.

In subsection 5.2, we investigate the geometry of decompositions $\mathbf{p}$ inducing singular Hessian approximations $H_{\mathbf{p}}$. The main conclusion is that *these decompositions should be avoided altogether*. We believe that one should generally escape ill-conditioned Hessian approximations as quickly as possible. Several techniques for this were considered before in the context of the CPD, notably by extrapolating previous

Algorithm 5.1: Hot restarts

```

1 Let  $(\mathbf{a}_i^1, \dots, \mathbf{a}_i^d)$  be a norm-balanced representative of  $p_i$ ;
2  $t \leftarrow 1$ ;
3  $\hat{\alpha} \leftarrow \min\{\frac{1}{4}, 10 \cdot \frac{\|\mathfrak{B} - \Phi(\mathbf{p})\|_F}{\|\mathfrak{B}\|_F}\}$ ;
4 Compute the Cholesky decomposition  $H_{\mathbf{p}} = L_{\mathbf{p}} L_{\mathbf{p}}^T$ ;
5 while Cholesky decomposition failed or  $\min_i (L_{\mathbf{p}})_{ii} < 10^{-5}$  do
6    $\alpha \leftarrow t \cdot \hat{\alpha}$ ;
7   for  $i = 1, \dots, r$  do
8     Let the elements of  $\mathbf{n}_k \in \mathbb{R}^{n_k}$  be sampled i.i.d. from  $N(0, 1)$ ;
9      $\mathbf{p}'_i \leftarrow ((1 - \alpha)\mathbf{a}_i^1 + \alpha \frac{\|\mathbf{a}_i^1\|}{\|\mathbf{n}_1\|} \mathbf{n}_1) \otimes \dots \otimes ((1 - \alpha)\mathbf{a}_i^d + \alpha \frac{\|\mathbf{a}_i^d\|}{\|\mathbf{n}_d\|} \mathbf{n}_d)$ ;
10  end
11 Solve  $\min_{\mathbf{x}} \|\mathfrak{B} - \sum_{i=1}^r x_i \mathbf{p}'_i\|_F$ ;
12  $\mathbf{p} \leftarrow (x_1 \mathbf{p}'_1, x_2 \mathbf{p}'_2, \dots, x_r \mathbf{p}'_r)$ ;
13  $\Delta \leftarrow \min\{\Delta_{\min}(\mathbf{p}), \Delta_{\max}\}$ ;
14  $t \leftarrow t + 1$ ;
15 Compute the Cholesky decomposition  $H_{\mathbf{p}} = L_{\mathbf{p}} L_{\mathbf{p}}^T$ ;
16 end

```

search directions [7, 17, 23, 53], using stochastic gradient-descent algorithms, using the true Hessian, or exploiting third-order information [5, 45]. Notwithstanding extensive experimentation,⁴ we found that the following Monte Carlo approach typically resulted in the best results: randomly sample nearby points of $\mathbf{p}$ on $\mathcal{S}^{\times r}$ until one is found with a reasonably well-conditioned Hessian approximation. The specific details of this *hot restarts* scheme are stated in Algorithm 5.1, which is executed before computing the dogleg step in (4.7).

REMARK 5.1 (Global convergence). *Proving global convergence for any input $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ is impossible. Assume that a hot restart is triggered whenever the CPD $\mathbf{p}^{(k)} \in \mathcal{S}^{\times r}$ in the execution of Algorithm 4.1 has $\kappa(\mathbf{p}^{(k)}) \geq \tau > 1$. Let $\mathcal{G}_{\leq \tau} := \{\mathbf{p} \in \mathcal{S}^{\times r} \mid \kappa(\mathbf{p}) \leq \tau\}$ be the locus of decompositions whose condition number is below the jumping threshold. Since Algorithm 4.1 produces a sequence of points $\mathbf{p}^{(k)} \in \mathcal{G}_{\leq \tau}$, it can at most converge to local minimizers of (TAP*) that are in $\mathcal{G}_{\leq \tau}$. Hence, we can at most prove global convergence for inputs $\mathfrak{B} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ whose best rank- r approximation $\mathfrak{B}^*$ exists and is such that $\Phi^{-1}(\mathfrak{B}^*)$ has at least one CPD contained in $\mathcal{G}_{\leq \tau}$. A major technical obstacle then still remains because the hot restarts procedure will introduce a stochastic aspect, further complicating any proof of global convergence. A full investigation is outside of the scope of this paper.*

REMARK 5.2 (Local convergence). *Local convergence is proved as follows. Let τ be as above. Taking $\frac{\kappa}{\tau} < \alpha < 1$ in Theorem 1 of [11] establishes conditions for linear or quadratic local convergence, including attraction radii and multiplicative constants. If $\frac{\kappa}{\tau} \geq 1$, then there is no proof of convergence, as expected.*

5.1. Recognizing ill-conditioned decompositions. For determining if a decomposition $\mathbf{p}$ is ill-conditioned, one could compute or approximate the least singular value of $H_{\mathbf{p}}$. For computational efficiency, however, we prefer relying on properties of the Cholesky decomposition $H_{\mathbf{p}} = L_{\mathbf{p}} L_{\mathbf{p}}^T$, where $L_{\mathbf{p}}$ is lower triangular. Recall that $H_{\mathbf{p}} = T_{\mathbf{p}}^T T_{\mathbf{p}}$ and that $T_{\mathbf{p}}$ is the matrix of $d_{\mathbf{p}} \Phi$. By Lemma 3.2, $\|d_{\mathbf{p}} \Phi\|_2 \geq 1$ and so $\varsigma_{\max}(H_{\mathbf{p}}) \geq 1$. Combining this with the fact that $\varsigma_k(H_{\mathbf{p}}) = (\varsigma_k(L_{\mathbf{p}}))^2$ and that one has $\varsigma_{\min}(A) \leq |\lambda_{\min}(A)|$, where $\lambda_{\min}$ is the smallest eigenvalue in magnitude, we find

$$\kappa_2(H_{\mathbf{p}}) = \frac{\varsigma_{\max}(H_{\mathbf{p}})}{\varsigma_{\min}(H_{\mathbf{p}})} \geq (\varsigma_{\min}(L_{\mathbf{p}}))^{-2} \geq \max_i |(L_{\mathbf{p}})_{i,i}|^{-2}.$$

In the last step we used that the diagonal entries of a triangular matrix are its eigenvalues. Moreover, a breakdown will occur while computing the Cholesky decomposition if $H_{\mathbf{p}}$ is very ill-conditioned [34, p. 200]. For these reasons, we apply a hot restart

⁴We performed preliminary experiments with regularization, using the true Hessian, switching to steepest descent or conjugate gradients near ill-conditioned decompositions, and a simple randomization procedure wherein a multiple of a random vector is added to the dogleg direction in which the weight increases as the decomposition becomes more ill-conditioned.

if a breakdown occurred or if one of the diagonal elements of $L_{\mathbf{p}}$ is less than 10^{-5} , so $\kappa_2(H_{\mathbf{p}}) \geq 10^{10}$; see line 5 of Algorithm 5.1. After executing Algorithm 5.1, the Cholesky decomposition is reused for efficiently solving $L_{\mathbf{p}}L_{\mathbf{p}}^T\mathbf{p}_N = -\mathbf{g}_{\mathbf{p}}$ in the dogleg step via backsubstitutions.

The constants appearing in line 3 of Algorithm 5.1 were empirically chosen. We choose $\hat{\alpha}$ in function of the relative residual, for the following reason. If the relative residual $\tau = \frac{\|\mathcal{B} - \Phi(\mathbf{p})\|_F}{\|\mathcal{B}\|_F}$ is small, one hopes to be close to a global optimizer, so spending computational resources to search an acceptable decomposition in a small radius about $\mathbf{p}$ is justified. However, if ill-conditioning is already encountered when $\tau \gtrsim 10^{-1}$, then our experiments suggested that aggressively restarting, i.e., Monte Carlo sampling with a large $\hat{\alpha}$, was the most cost-effective strategy. As long as a satisfyingly well-conditioned decomposition is not sampled, α is gradually increased.

5.2. Analysis. In this subsection, we argue why CPDs with a singular Hessian approximation $H_{\mathbf{p}}$ are troublesome for the RGN method. The arguments below are not intended to be a proof, but rather they provide an intuition and offer an outline of a potential proof strategy. A complete proof is beyond the scope of this work.

From the properties of Gram matrices, we have the following connection between $\kappa(\mathbf{p})$ from (2.2) and $H_{\mathbf{p}}$:

$$(5.1) \quad \kappa(\mathbf{p}) = \begin{cases} \infty & \text{if } H_{\mathbf{p}} \text{ is singular,} \\ \|T_{\mathbf{p}}^\dagger\|_2 = \|H_{\mathbf{p}}^{-1}\|_2^{\frac{1}{2}} & \text{otherwise.} \end{cases}$$

Assume that some computed local optimizer $\mathbf{p}^*$ of (TAP*) is well-conditioned, and hence, by (5.1), the associated $H_{\mathbf{p}^*}$ is reasonably well-conditioned. However, during the execution of the RTR method, we may encounter points $\mathbf{p}$, for which $H_{\mathbf{p}}$ is almost singular. We sketch an informal argument why escaping the neighborhood of such points may take many iterations in a RGN method, with or without trust region.

Following [10, equation (2.3)], the locus of ill-posed CPDs is defined as follows:

$$\mathcal{I}_r := \{\mathbf{q} \in \mathcal{S}^{\times r} \mid \kappa(\mathbf{q}) = \infty\} = \{\mathbf{q} \in \mathcal{S}^{\times r} \mid H_{\mathbf{q}} = T_{\mathbf{q}}^T T_{\mathbf{q}} \text{ is singular}\};$$

the second equality is by (5.1). The following lemma, proved in Appendix B, implies that a randomly chosen CPD $\mathbf{p}$ lies on $\mathcal{I}_r$ with probability zero, but also that there is a nonzero probability of coming close to $\mathcal{I}_r$; we do not know how large it is.

LEMMA 5.3. *If $\sigma_r(\mathcal{S})$ is non-defective, then $\mathcal{I}_r$ is contained in a non-empty, strict subvariety of $\mathcal{S}^{\times r}$.*

Let $\mathbf{p}, \mathbf{q} \in \mathcal{S}^{\times r}$ with $\mathbf{q}$ close to $\mathbf{p}$, and assume that the Hessian approximation $H_{\mathbf{q}}$ is singular, i.e., $\mathbf{q} \in \mathcal{I}_r$. Let $T_{\mathbf{p}} = USV^T$ be the compact SVD of $T_{\mathbf{p}}$, where V is orthogonal, U has orthonormal columns, and $S = \text{diag}(\varsigma_1, \varsigma_2, \dots, \varsigma_{r(\Sigma+1)})$ with $\varsigma_1 \geq \varsigma_2 \geq \dots \geq \varsigma_{r(\Sigma+1)} \geq 0$. Let $n := r(\Sigma+1)$ and write $V = [\mathbf{v}_1 \ \dots \ \mathbf{v}_n]$. Assuming $\varsigma_n > 0$, or equivalently $\kappa(\mathbf{p}) < \infty$, by (4.6), the Newton direction is

$$(5.2) \quad \mathbf{p}_N = -H_{\mathbf{p}}^\dagger \mathbf{g}_{\mathbf{p}} = -VS^{-1}U^T \mathbf{r}_{\mathbf{p}} = -\sum_{i=1}^n \alpha_i \varsigma_i^{-1} \mathbf{v}_i,$$

where $\alpha_i := (U^T \mathbf{r}_{\mathbf{p}})_i$. Since $\mathbf{p}$ and $\mathbf{q}$ are close, $\varsigma_n^2 = \|H_{\mathbf{p}}^{-1}\|_2^{-1} \approx \|H_{\mathbf{q}}^{-1}\|_2^{-1} = 0$, and so $\mathbf{p}_N$ is pointing mostly in the direction of $\mathbf{v}_n$. As long as $\mathbf{v}_n$ is a descent direction of the objective $\frac{1}{2}\|F(\mathbf{p})\|^2$, the dogleg step is expected to yield an escape from $\mathbf{q}$.

Unfortunately, there are two classes of tensors, whose decompositions could be particularly troublesome when coming close to them during the RGN method: tensors with infinitely many decompositions and tensors whose *border ranks* [42, section 2.4] are not equal to their ranks.⁵

5.2.1. Tensors with infinitely many decompositions. Consider a decomposition $\mathbf{q}$ that is not isolated. To show this is troublesome, we need the next result.

⁵It is an interesting open question whether these decompositions could actually be attractive for the RGN method without trust region.

PROPOSITION 5.4. Let $\mathbf{q} = (q_1, \dots, q_r) \in \mathcal{S}^{\times r}$ and $\mathfrak{B} = \Phi(\mathbf{q})$. Let $\Phi^{-1}(\mathfrak{B})$ contain a positive-dimensional analytic submanifold $\mathcal{E}$ without boundary with $\mathbf{q} \in \mathcal{E}$.

1. We have $\mathcal{E} \subset \mathcal{I}_r$.
2. If $\mathbf{z} = (z_1, z_2, \dots, z_r) \in T_{\mathbf{q}}\mathcal{E} \subset T_{\mathbf{q}}\mathcal{S}^{\times r}$ and $\mathbf{z}_i$ are the local coordinates of $\mathbf{z}_i$ with respect to the basis of $T_{q_i}\mathcal{S}$ given by the columns of the matrix T_{q_i} from (4.2), then $\mathbf{z}^T := [\mathbf{z}_1^T \quad \mathbf{z}_2^T \quad \dots \quad \mathbf{z}_r^T] \in \ker T_{\mathbf{q}}$.

Proof. By construction, $\mathcal{E}$ is a smooth algebraic variety. From [19, Lemma 6.5] it follows that for every $\mathbf{p} = (p_1, \dots, p_r) \in \mathcal{E}$ the corresponding $\langle T_{p_1}\mathcal{S}, \dots, T_{p_r}\mathcal{S} \rangle$ is of dimension strictly less than $r \cdot \dim \mathcal{S} \leq \Pi$. This entails that $T_{\mathbf{p}}$ has linearly dependent columns, proving the first part. Furthermore, $0 = (d_{\mathbf{q}}\Phi)(\mathbf{z}) = z_1 + z_2 + \dots + z_r$, which written in coordinates is $T_{\mathbf{q}}\mathbf{z} = T_{q_1}\mathbf{z}_1 + T_{q_2}\mathbf{z}_2 + \dots + T_{q_r}\mathbf{z}_r = 0$, by (4.2) and (4.5). $\square$

Consequently, the kernel of $H_{\mathbf{q}} = T_{\mathbf{q}}^T T_{\mathbf{q}}$ contains at least the tangent space $T_{\mathbf{q}}\mathcal{E}$. The implication for nearby points $\mathbf{p} \approx \mathbf{q}$ is as follows. Let $V = [\mathbf{v}_1 \dots \mathbf{v}_n]$ be the matrix of right singular vectors of $T_{\mathbf{p}}$, and $e := \dim \mathcal{E}$. Then, because $\mathbf{p} \approx \mathbf{q}$ and since the kernel of $T_{\mathbf{q}}$ is at least e -dimensional, we expect the linear space $\langle \mathbf{v}_{n-e+1}, \dots, \mathbf{v}_n \rangle$ to be almost parallel to $T_{\mathbf{q}}\mathcal{E}$. Moreover, at least e singular values of $T_{\mathbf{p}}$ will be small, so that, by (5.2), the Newton-direction $\mathbf{p}_N$ is significantly parallel to $T_{\mathbf{p}}\mathcal{E}$. Retracting along $\mathbf{p}_N$ or the dogleg step is then expected to move $\mathbf{p}$ to $\mathbf{p}'$ along a path that could be mostly parallel to $\mathcal{E}$. The resulting point $\mathbf{p}'$ would then also lie close to $\mathcal{E}$, i.e., close to one of the other decompositions of $\Phi(\mathbf{q})$. Iterating this reasoning suggests that points close to $\mathcal{E}$ could be hard to escape with an RGN method.

5.2.2. Open boundary tensors. Let $\mathfrak{B}$ be an r -open boundary tensor: $\mathfrak{B}$ has border rank [42, section 2.4] equal to r , but $\text{rank}(\mathfrak{B}) > r$. This means that there exist sequences of tensors $\Phi(\mathbf{p}_i) \rightarrow \mathfrak{B}$ as $i \rightarrow \infty$, where $\mathbf{p}_i \in \mathcal{S}^{\times r}$. The rank- r TAP is ill-posed in this case, as only an infimum exists [22]. CPDs $\mathbf{p} \in \mathcal{S}^{\times r}$ with $\Phi(\mathbf{p})$ near $\mathfrak{B}$ are also troublesome for RGN methods, as we argue next.

Let $\mathbf{p}(t) = (p_1(t), \dots, p_r(t)) \in \mathcal{S}^{\times r}$ be a smooth curve with $\lim_{t \rightarrow 0} \Phi(\mathbf{p}(t)) = \mathfrak{B}$. The proof of [10, theorem 1.4] shows that scalar functions $\nu_1(t), \dots, \nu_r(t)$ exist so that

$$\mathbf{h}(t) = (\nu_1(t)p_1(t), \nu_2(t)p_2(t), \dots, \nu_r(t)p_r(t)) \in T_{\mathbf{p}(t)}\mathcal{S}^{\times r}$$

tends to a nonzero vector $\mathbf{h}_{\star} \in \ker T$, where $T := \lim_{t \rightarrow 0} T_{\mathbf{p}(t)}$ with $T_{\mathbf{p}(t)}$ as in (4.5). This limit T exists, as $T_{\mathbf{p}(t)} \in \text{St}(\Sigma + 1, \mathbb{R}^{\Pi})^{\times r}$ lives in the r -fold product of the Stiefel manifold of matrices in $\mathbb{R}^{\Pi \times (\Sigma + 1)}$ with orthonormal columns, which is closed.

Assume now that $\mathbf{p} := \mathbf{p}(\epsilon)$ and $\mathbf{h} := \mathbf{h}(\epsilon)$ for some specific $\epsilon \approx 0$. Whenever $\|\mathbf{h} - \mathbf{h}_{\star}\|$ is small, the Newton direction $\mathbf{p}_N$ in (5.2) is expected to have a significant component in the direction of $\mathbf{h}$. Note that the entries of $\mathbf{h}$ are constant multiples of the $p_i(\epsilon)$'s. Hence, adding any multiple of $\mathbf{h}$ to $\mathbf{p}$ yields an element of $\mathcal{S}^{\times r}$. Many retraction operators $R_x : T_x\mathcal{M} \rightarrow \mathcal{M}$ on a manifold $\mathcal{M}$, including both retractions discussed in section 6, are the identity when $x + \mathbf{h} \in \mathcal{M}$. In this case,

$$\mathbf{p}' := R_{\mathbf{p}}(\mathbf{h}) = \mathbf{p} + \mathbf{h} = ((1 + \nu_1(\epsilon))p_1(\epsilon), \dots, (1 + \nu_r(\epsilon))p_r(\epsilon)).$$

As a result, only the norms of the rank-1 terms in the CPD are altered by such a retraction along $\mathbf{h}$. Since the matrix $T_{\mathbf{p}}$ is invariant under this scaling [10, Proposition 4.4], we have $T_{\mathbf{p}} = T_{\mathbf{p}'}$. Hence, the same argument also applies at the next iterate $\mathbf{p}'$, suggesting again that CPDs $\mathbf{p}$ with $\Phi(\mathbf{p})$ close to an r -open boundary tensor $\mathfrak{B}$ could be hard to escape with RGN methods.

6. The product ST-HOSVD retraction. A characteristic ingredient of a Riemannian optimization method is an effective retraction operator. It appears in step 8 of Algorithm 4.1 for pushing forward the current iterate $\mathbf{p}^{(k)} \in \mathcal{S}^{\times r}$ along a tangent vector $\mathbf{p}_k \in T_{\mathbf{p}^{(k)}}\mathcal{S}^{\times r}$. The result is the (tentative) next iterate $\mathbf{p}' \in \mathcal{S}^{\times r}$.

A *retraction* is a map taking a tangent vector $\xi_p \in T_p\mathcal{M}$ to a manifold $\mathcal{M}$ at p to the manifold itself. It can be regarded as an operator that approximates the action of the exponential map to first order [1]. A formal definition is as follows [1, 3].

DEFINITION 6.1. Let $\mathcal{M}$ be a manifold. A retraction R is a map $T\mathcal{M} \rightarrow \mathcal{M}$ that satisfies all of the following properties for every $p \in \mathcal{M}$:

1. $R(p, 0_p) = p$;

2. there exists an open neighborhood $\mathcal{N} \subset \mathcal{TM}$ of $(p, 0_p)$ such that the restriction $R|_{\mathcal{N}}$ is well-defined and a smooth map;
3. local rigidity: $d_{0_x} R(x, \cdot) = \text{Id}_{T_x \mathcal{M}}$ for all $(x, 0_x) \in \mathcal{N}$.

We let $R_p(\cdot) = R(p, \cdot)$ denote the retraction R with foot at p .

Choosing an efficient retraction is critical for attaining good computational performance. Fortunately, we can exploit the following well-known result that on a product manifold, the product of individual retractions specifies a retraction.

LEMMA 6.2. *Let $\mathcal{M}_1, \dots, \mathcal{M}_r$ be manifolds. Let $R_i : \mathcal{TM}_i \rightarrow \mathcal{M}_i$ be retractions. Then the following is a retraction on $\mathcal{M}_1 \times \dots \times \mathcal{M}_r$:*

$$R((p_1, \xi_{p_1}), \dots, (p_r, \xi_{p_r})) = (R_1(p_1, \xi_{p_1}), \dots, R_r(p_r, \xi_{p_r})),$$

where $p_i \in \mathcal{M}_i$ and $\xi_{p_i} \in T_{p_i} \mathcal{M}_i$.

Because $\mathcal{S}^{\times r}$ is a product manifold, it thus suffices to find an efficient retraction for the Segre manifold $\mathcal{S}$. Since $\mathcal{S}$ coincides with the manifold of tensors of multilinear rank $(1, \dots, 1)$, we can apply the truncated higher-order singular value decomposition (T-HOSVD) retraction from [41, Proposition 2.3]. This retraction is defined as the rank- $(1, \dots, 1)$ T-HOSVD approximation of $\mathfrak{A} + \mathbf{p}$ [21], i.e., $\hat{R}_{\mathfrak{A}}(\mathbf{p}) := (\hat{\mathbf{q}}_1 \hat{\mathbf{q}}_1^T, \dots, \hat{\mathbf{q}}_d \hat{\mathbf{q}}_d^T) \cdot (\mathfrak{A} + \mathbf{p})$, where $\mathfrak{A} \in \mathcal{S} \subset \mathbb{R}^{n_1 \times \dots \times n_d}$, $\mathbf{p} \in T_{\mathfrak{A}} \mathcal{S} \subset \mathbb{R}^{n_1 \times \dots \times n_d}$, and $\hat{\mathbf{q}}_k \in \mathbb{R}^{n_k}$ is the dominant left singular vector of the *standard flattening* $(\mathfrak{A} + \mathbf{p})_{(k)}$ [21].

Instead of applying the above T-HOSVD retraction, we propose a retraction based on the sequentially truncated HOSVD (ST-HOSVD) [30, 69], as its computational complexity is lower [69]. For rank-1 tensors $\mathfrak{A} \in \mathcal{S} \subset \mathbb{R}^{n_1 \times \dots \times n_d}$, we define a map $\mathcal{TS} \rightarrow \mathcal{S}$ as follows

$$(6.1) \quad R_{\mathfrak{A}}(\mathbf{p}) := (\mathbf{q}_1 \mathbf{q}_1^T, \dots, \mathbf{q}_d \mathbf{q}_d^T) \cdot (\mathfrak{A} + \mathbf{p}),$$

where $\mathbf{p} \in T_{\mathfrak{A}} \mathcal{S} \subset \mathbb{R}^{n_1 \times \dots \times n_d}$, and the $\mathbf{q}_k$'s are the factors of the rank- $(1, \dots, 1)$ ST-HOSVD of $\mathfrak{A} + \mathbf{p}$. While it is not indicated in the notation, one should remember that the vector $\mathbf{q}_i$ depends on $\mathbf{q}_1, \dots, \mathbf{q}_{i-1}$ [69]. We prove the next result in Appendix B.

LEMMA 6.3. *The map in (6.1) defines a retraction $\mathcal{TS} \rightarrow \mathcal{S}$.*

Let $\mathbf{p} = (p_1, \dots, p_r) \in \mathcal{S}^{\times r}$ and let $\mathbf{t} \in T_{\mathbf{p}} \mathcal{S}^{\times r}$. The tangent vector $\mathbf{t}$ is given in coordinates by $\mathbf{x} \in \mathbb{R}^{r(\Sigma+1)}$ with respect to the orthonormal basis B in (4.4). Then, the product retraction on $\mathcal{S}^{\times r}$ is

$$R_{\mathbf{p}}(\mathbf{t}) := (R_{p_1}(T_{p_1} \mathbf{x}_1), \dots, R_{p_r}(T_{p_r} \mathbf{x}_r)),$$

where $\mathbf{x}^T = [\mathbf{x}_1^T \ \dots \ \mathbf{x}_r^T]$. An efficient implementation is given in Appendix A.1.

REMARK 6.4. *The above product ST-HOSVD retraction was proved valid only in a neighborhood $\mathcal{N} \subset \mathcal{TS}^{\times r}$ about $(\mathbf{p}, 0)$ by Lemma 6.2 and Lemma 6.3. Before applying the retraction $R_{\mathbf{p}}$ to $\mathbf{t}$, in theory, one should verify that $(\mathbf{p}, \mathbf{t}) \in \mathcal{N}$. In practice, this strict regime is not enforced; retractions generally are also applied outside of the region where the smoothness of $R|_{\mathcal{N}}$ can be proved. In our implementation, the retraction is applied to all inputs, regardless of whether $(\mathfrak{A}, \mathbf{t}) \in \mathcal{N}$.*

7. Numerical experiments. RGN-HR was implemented⁶ in Matlab, employing some functionality of Tensorlab v3.0 [70]. All numerical experiments were conducted in Matlab R2017a using 2 computational threads on a computer system comprising two Intel Xeon E5-2697 v3 CPU's with 14 cores each (all clocked at 2.6GHz) and 128GB of total main memory.

We compare RGN-HR with state-of-the-art nonlinear least-squares (NLS) solvers designed for the TAP from Tensorlab v3 [70], specifically the trust region method with dogleg steps, called `nls_gnd1`. It employs factor matrices as parameterization and uses the GN approximation of the Hessian of (3.2), i.e., $J^T J$ with J as in (3.3), as explained in subsection 3.1. As a result, the normal equations are ill-posed, and, hence, the Newton step $\mathbf{p}_N$ is not well defined. Therefore, Tensorlab (approximately) computes the minimum-norm solution of the least-squares problem $J \mathbf{p}_N = -\mathbf{r}$, where $\mathbf{r}$ is the

⁶Our implementation can be obtained from <https://arxiv.org/abs/1709.00033>.

residual. Both a *direct* (`LargeScale = false`) and *iterative* (`LargeScale = true`) solution of this least-squares problem is supported. The direct algorithm computes the Moore-Penrose pseudoinverse and applies it to $-\mathbf{r}$, while the iterative method is LSQR with block Jacobi preconditioner; see [59] and the user manual of Tensorlab [70] for details. We refer to them as GNDL and GNDL-PCG, respectively. Tensorlab offers two additional NLS solvers (`nls_lm` and `nls_cgs`) that we did not find competitive.

GNDL is Tensorlab’s default method for computing dense small-scale CPDs, i.e., $r(\Sigma + d) \leq 100$, whereas GNDL-PCG is the default method for computing CPDs in the case $r(\Sigma + d) > 100$. It should be stressed that for fixed d , GNDL and RGN-HR both *have the same asymptotic computational complexity*. Hence, this is a comparison between methods appropriate for the same class of TAPs, i.e., dense, small-scale, low-residual problems. On the other hand, one may not expect that RGN-HR will be competitive with GNDL-PCG for medium and large-scale problems with $r(\Sigma + d)$ much larger than 1000. The reason is that GNDL-PCG was specifically designed for such problems by crudely approximating the Newton direction using an asymptotically faster algorithm. Indeed, the complexity per iteration of RGN-HR is roughly $r^3(\Sigma + d)^3$, while GNDL-PCG’s complexity is only $k_{\text{CG}} \cdot r^2(\Sigma + d)^2$. The exact Newton direction is recovered only with $k_{\text{CG}} = \mathcal{O}(r(\Sigma + d))$. Recall from the introduction that small-scale dense TAPs often result from an orthogonal Tucker compression of a medium or large-scale tensor; hence, solving small-scale TAPs is an important computational kernel.

For the above reason, we experiment with dense, small-scale CPDs. We compare RGN-HR with GNDL and GNDL-PCG for some order-3 tensors; namely, in [subsection 7.3](#) for tensors in $\mathbb{R}^{15 \times 15 \times 15}$ and in [subsection 7.4](#) for tensors in $\mathbb{R}^{13 \times 11 \times 9}$.

7.1. Experimental setup. Let $\mathfrak{A} \sim \mathcal{N}$ denote that the entries of $\mathfrak{A} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ are identically and independently distributed (i.i.d.) following a standard normal distribution. For the tensor models in [subsections 7.3](#) and [7.4](#) below, the performance of methods for solving TAPs is evaluated as follows:

1. Randomly sample an order- d rank- r decomposition $\mathbf{p} \in \mathcal{S}^{\times r}$ from one of the models, and let $\mathfrak{A} := \Phi(\mathbf{p})$;
2. create a perturbed tensor $\mathfrak{B} := \frac{\mathfrak{A}}{\|\mathfrak{A}\|} + 10^{-e} \frac{\mathfrak{C}}{\|\mathfrak{C}\|}$, where $\mathfrak{C} \sim \mathcal{N}$ and $e > 0$;
3. randomly sample factor matrices $\mathbf{q} = (A_1, \dots, A_d) \in \mathcal{E} := \mathbb{R}^{n_1 \times r} \times \dots \times \mathbb{R}^{n_d \times r}$ with $A_k \sim \mathcal{N}$; and
4. solve the TAP for $\mathfrak{B}$ from this random starting point $\mathbf{q}$ with each method.

Our main performance criterion is the *expected time to success* (ETS), which measures how much time must be spent on average to find a solution of the TAP with method M by trying multiple random starting points if necessary. Let t_{success} denote the average time (in seconds) a successful attempt takes to compute a solution of the TAP, and let p_{success} denote the fraction of successful attempts. Similarly, write t_{fail} for the average time (in seconds) needed for an unsuccessful attempt. Then,

$$\text{ETS}_M := \sum_{k=0}^{\infty} p_{\text{success}} p_{\text{fail}}^k (t_{\text{success}} + k \cdot t_{\text{fail}}) = \frac{p_{\text{fail}} t_{\text{fail}} + p_{\text{success}} t_{\text{success}}}{p_{\text{success}}}.$$

We call $\mathbf{p}' \in \mathcal{S}^{\times r}$ a *(local) solution* of the TAP if its residual is within 10% of the error level, i.e., if $\|\mathfrak{A} - \Phi(\mathbf{p}')\| \leq 1.1 \cdot 10^{-e}$, and its condition number is within a factor 50 of the condition number of the CPD $\mathbf{p}$, i.e., $\kappa(\mathbf{p}') \leq 50\kappa(\mathbf{p})$. In our comparisons, we present the *speedup* $\text{ETS}_M / \text{ETS}_{\text{RGN-HR}}$ of RGN-HR relative to method M.

The ETS was empirically estimated for each of the methods by estimating p_{success} , t_{fail} , and t_{success} via their corresponding sample statistics; we sample only one tensor from the model (in step 1), one perturbation (in step 2), and $k \in \mathbb{N}$ random starting points (in step 3). For the models in [subsections 7.3](#) and [7.4](#) we took $k = 25$ and $k = 50$, respectively.

7.2. Parameter choices. Both in RGN-HR and Tensorlab some parameters related to the stopping criterion should be selected. The `nls_gndl` method has the same 4 standard parameters as RGN-HR, described in [subsection 4.6](#), namely τ_f , $\tau_{\Delta f}$, $\tau_{\Delta x}$, and k_{max} . In addition, for RGN-HR we have a maximum number of restarts r_{max} , and for the `nls_gndl` with `LargeScale = true` we have a relative tolerance on the

improvement of the residual in the LSQR method τ_{CG} , and a maximum number of iterations of the LSQR method k_{CG} . In every experiment, $\tau_f = 0$, $\tau_{CG} = 10^{-6}$, $k_{CG} = 75$, and $r_{\max} = 500$ was selected. Furthermore, we selected $\tau_{\Delta f} = 10^{-2e}$ for RGN-HR, and $\tau_{\Delta f} = 10^{-2(e+1)}$ for the Tensorlab methods. We suggest this last distinction, as we generally observed that Tensorlab’s ETS would be worse if we took 10^{-2e} , as it negatively affected its probability of success.

7.3. Model 1. The first model is a family $\mathcal{F}_r(c, s) \subset \mathbb{R}^{15 \times r} \times \mathbb{R}^{15 \times r} \times \mathbb{R}^{15 \times r}$ with 2 parameters in which the factor matrices have correlated columns, controlled by $0 \leq c < 1$, and the norms of the rank-1 terms approximately satisfy an exponential increase from about 1 to about 10^s , where $s \in \mathbb{N}$. Let R_c be the upper triangular factor in the Cholesky decomposition $R_c^T R_c := c \mathbf{1} \mathbf{1}^T + (1 - c)I$, where $\mathbf{1} \in \mathbb{R}^r$ is the vector of ones. Then,

$$\mathcal{F}_r(c, s) := \{(A_1, A_2, A_3) \in \mathcal{E} \mid A_k = N_k R_c \text{diag}(10^{\frac{s}{3r}}, 10^{\frac{2s}{3r}}, \dots, 10^{\frac{rs}{3r}}), \\ \text{where } N_k \sim \mathcal{N}, k = 1, 2, 3\}.$$

We observed empirically that for fixed r, s , increasing c increases the condition number of the CPD. For $(r, c, s, e) \in \{15, 20, 25, 30\} \times \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 0.95\} \times \{1, 2, 3, 4\} \times \{3, 5, 7\}$, we sampled a random CPD from $\mathcal{F}_r(c, s)$, applied one random perturbation (as in step 2 of the experimental setup), generated $k = 25$ random starting points from whence each of the optimization methods starts, and record the time and whether the CPD was a solution. The ETS is estimated from these data. The tested methods were RGN-HR, a variant of RGN-HR called RGN-Reg described below, GNDL, and GNDL-PCG.⁷ We took $\tau_{\Delta x} = 10^{-12}$ and $k_{\max} = 1500$. The speedups of RGN-HR with respect to the competing methods are shown in Figure 7.1. A value α means that the corresponding method has an ETS that is α times the ETS of RGN-HR; $\alpha = \infty$ means that the competing method could not solve the TAP.

We do not include the results for the experiment’s parameter choices $e = 3$ and $c = 0.95$ because the condition number ranged between about 223 and 1186 for these configurations. At error level $e = 3$, this entails that the rank-1 terms in the computed CPD can only be guaranteed to capture between 0 and $0.65 \approx -\log_{10}(223 \cdot 10^{-3})$ significant correct digits of the true CPD. Hence, these “results” are meaningless.

RGN-Reg is a variation of RGN-HR without hot restarts. Ill-conditioned normal equations are resolved by applying Tikhonov regularization:

$$\left(H_{\mathbf{p}} + 10^{-10} \left(\frac{\|\Phi(\mathbf{p}) - \mathfrak{B}\|}{\|\mathfrak{B}\|} \right)^{\frac{3}{4}} \|H_{\mathbf{p}}\|_F I \right) \mathbf{p}_N = -\mathbf{g}_{\mathbf{p}},$$

where the notation is as in (4.6). The reason for including this method is to illustrate that the proposed hot restarts mechanism provides a superior way of handling ill-conditioned Hessian approximations. The top rows of the tables in Figure 7.1(a)–(c) show the speedup of RGN-HR with respect to RGN-Reg. In almost all configurations, the hot restarts of the former significantly outperformed the latter’s regularization.

The experiments show that RGN-HR outperforms both GNDL and GNDL-PCG in this small-scale model. Specifically, it outperforms GNDL, by at least 1 and up to 3 orders of magnitude in almost all configurations. This difference is partly explained by the way in which GNDL solves $J\mathbf{p}_N = -\mathbf{r}$, namely by explicitly applying the pseudoinverse of J . This is more expensive than solving a linear system via a Cholesky factorization.⁸ The other gains are attributed to the reduction of the number of iterations, which can be ascribed to only two key differences between RGN-HR and GNDL, namely hot restarts for avoiding ill-conditioned CPDs altogether; and the Riemannian formulation. Indeed, both methods use a trust region, employ dogleg steps, and feature similar stopping criteria and parameter selections. It is particularly noteworthy that RGN-HR’s performance greatly improves relative to GNDL-PCG (and

⁷In fact, we also tested Tensorlab’s `nls_lm` with both direct and iterative solves, i.e., with the `LargScale` option set alternatively to `false` and `true`. It was not competitive with the methods in Figures 7.1 and 7.2, successfully solving only 31%, respectively 58%, of the combinations.

⁸For square matrices computing an SVD is 20 times as expensive; see [26, p. 78 and p. 239–240].

(a) $e = 3$

(b) $e = 5$

(c) $e = 7$

Figure 7.1: Speedups in terms of the ETS (based on 25 samples) of RGN-HR with respect to RGN-Reg, and Tensorlab's GNDL and GNDL-PCG on tensors sampled from model $\mathcal{F}_r(c, s)$ with perturbation factors $e = 3, 5, 7$.

GNDL							GNDL-PCG						
s	r	5	7	9	11	13	r	5	7	9	11	13	
4	4	28	40	59	failed		6	67	35	inf	failed		
3	7	16	46	381	inf		7	21	23	24	inf		
2	4	17	31	168	452		13	18	22	27	24		
1	4	5	20	31	166		17	7		16	24		
0	3	4	10	16	58		7	8	6	4	10		

Figure 7.2: Speedups in terms of the ETS (based on 50 samples) of RGN-HR with respect to RGN-Reg, and Tensorlab’s GNDL and GNDL-PCG on tensors sampled from model $\mathcal{G}_r(s)$ with perturbation factors $e = 5$. The label “failed” indicates that none of the methods could solve the problem in 50 attempts.

GNDL) when s increases. The reason is that they require more iterations to converge, while RGN-HR’s number of iterations remains relatively stable. *The analysis in subsection 3.4 already predicted this outcome.*

For a very significant fraction of starting points, the state-of-the-art methods halt at extremely ill-conditioned CPDs $\mathbf{p}'$. The table below shows the fraction of cases where $\kappa(\mathbf{p}) > 10^e$ among those CPDs $\mathbf{p}'$ whose backward error is very small, namely $\|\Phi(\mathbf{p}') - \mathfrak{B}\|_F \leq 1.1 \cdot 10^{-e}$. Since the forward error is asymptotically bounded by $\kappa(\mathbf{p})\|\Phi(\mathbf{p}') - \mathfrak{B}\|_F \approx 1.1$, this means that such CPDs are *completely uninterpretable*: no correct significant digits are present in the individual rank-1 terms, unless their norms would be orders of magnitude larger than $\|\Phi(\mathbf{p}')\|_F$. Nowadays the backward error is still the dominant criterion to “determine” if a computed CPD is “a solution” of the TAP. The table below highlights just how dangerous this practice is:

e	RGN-HR	RGN-Reg	GNDL	GNDL-PCG
3	23.6%	42.3%	51.0%	50.8%
5	0.0%	10.9%	9.1%	9.3%
7	0.0%	0.0%	0.0%	0.0%

Specifically for tensors corrupted by large amounts of noise, i.e., small e , the situation is dramatic. For $e = 3$, about *half* of the CPDs returned by the state-of-the-art, default solver in Tensorlab v3.0 whose backward error is as good as one can hope, i.e., about 10^{-e} , are extremely ill-conditioned. The proposed RGN-HR method explicitly attempts to avoid this problem via hot restarts. The table shows that it is fairly successful compared to the other methods, which do not use hot restarts, halving the fraction of evidently spurious CPDs to less than $\frac{1}{4}$ for $e = 3$, and eliminating the issue for $e = 5, 7$. As in [10, 67] we caution that a CPD can only be interpreted meaningfully if both the backward error and the condition number are small.

7.4. Model 2. Another challenging family is $\mathcal{G}_r(s) \subset \mathbb{R}^{13 \times r} \times \mathbb{R}^{11 \times r} \times \mathbb{R}^{9 \times r}$, where $s \geq 0$ is a parameter that controls how close the tensor is to the variety of tensors of multilinear rank componentwise bounded by $(r_1, r_2, r_3) = (5, 5, 5)$. Specifically,

$$\mathcal{G}_r(s) := \{(A_1, A_2, A_3) \in \mathcal{E} \mid A_k = N_k(10^{\frac{2-s}{2}} I_r + X_k Y_k^T) \text{diag}(5^{\frac{0}{r-1}}, \dots, 5^{\frac{r-1}{r-1}}), \\ \text{where } X_k, Y_k \in \mathbb{R}^{r \times r_k} \text{ and } N_k, X_k, Y_k \sim \mathcal{N}, k = 1, 2, 3\}.$$

Empirical evidence shows that for fixed r , letting $s \rightarrow 0$ increases the CPD’s condition number; a theoretical argument was sketched in [10, section 6.1].

For each $(r, s) \in \{5, 7, 9, 11, 13\} \times \{0, 1, \dots, 4\}$, we sampled one CPD from $\mathcal{G}_r(s)$, applied one random perturbation with $e = 5$, generated $k = 50$ random starting points from whence each of the methods starts, and recorded the time and whether the CPD was a solution. From these data, the ETS is estimated. We tested RGN-HR, GNDL, and GNDL-PCG with $\tau_{\Delta x} = 10^{-15}$ and $k_{\max} = 7500$. The speedups are shown in Figure 7.2.

Broadly the same observations hold as with the previous model: RGN-HR is up to 2 orders of magnitude faster than GNDL, and up to 1 order of magnitude faster than GNDL-PCG. In all configurations, RGN-HR was the fastest method.

Figure 7.3: Randomly selected convergence graphs for the model $\mathcal{G}_7(2)$ with error level $e = 5$. The objective value $f(\mathbf{p}) = \frac{1}{2}\|\Phi(\mathbf{p}) - \mathfrak{B}\|_F^2$ is plotted versus the time spent for the proposed RGN-HR method, the variant RGN-Reg, and Tensorlab’s GNDL, GNDL-PCG, LM and LM-PCG methods.

Finally, in Figure 7.3 we show some convergence plots for this model with rank $r = 7$ and $s = 2$. It includes Tensorlab’s `nls_lm` method, a Levenberg–Marquardt method, both with direct solves and iterative solves of the least-squares problem $\mathbf{J}\mathbf{p}_N = -\mathbf{r}$; we refer to them as respectively LM and LM-PCG.

The effect of the hot restarts in RGN-HR can clearly be seen in Figure 7.3. Observe in particular that a hot restart is often triggered after a brief period of stagnation of the convergence. Contrast this with the prolonged periods of stagnation in all of the other methods, including RGN-Reg. We surmise that stagnation is often caused by ill-conditioned Hessian approximations. In the plots on the right we see that the progress of RGN-HR and RGN-Reg is almost identical up to 0.4 seconds. At that point, the former detects that the condition number is too high and applies a hot restart, while the latter steadily reduces the objective value. Nevertheless, in the end, RGN-HR converges faster.

8. An application. The proposed algorithm can also be applied efficiently to larger data sets appearing in real applications. Fluorescence spectroscopy is an imaging technique that can be employed to detect the concentration of certain chemical compounds, called *fluorophores*, in a diluted mixture. This inexpensive analysis technique is widely employed in the life sciences; for example, it is employed for the identification of dissolved organic material in natural and waste water [37] and in food chemistry [58], among others. The theoretical model underlying the *emission–excitation matrices* that are obtained from a fluorescence spectroscopy analysis of a diluted, inert mixture of several fluorophores is the tensor rank decomposition [8].

The tensor we consider was obtained from fluorescence spectroscopy measurements of five mixtures of three amino acids; in ideal circumstances the rank of the tensor would thus be 3.⁹ Detailed information about the acquisition of the data can be found in [12, 39]. The size of the tensor is $5 \times 201 \times 61$; the first factor corresponds to the mixtures, the second to the emission wavelengths (250–450 nm in steps of 1 nm), and the third to the excitation wavelengths (250–310 nm in steps of 1 nm).

⁹The data can be obtained at http://www.models.life.ku.dk/Amino_Acid_fluo.

Figure 8.1: Emission and excitation spectra of the three identified components in the mixtures of the amino acids using both the proposed algorithm (RGN) and Tensorlab’s `cpd_nls` algorithm (TL). We stress that all 12 spectra are plotted; the spectra recovered by the two methods are visually indistinguishable.

As explained in [Remark 1.1](#), it is common for large but low-rank tensors to start with a Tucker compression. We approximated $\mathfrak{B} \in \mathbb{R}^{5 \times 201 \times 61}$ by a rank-(5, 6, 6) ST-HOSVD [69] $(Q_1, Q_2, Q_3) \cdot \mathfrak{S} \approx \mathfrak{B}$ using Tensorlab’s `mlsvd` function. This took about 0.01 seconds, and results in a relative error of $1.236 \cdot 10^{-2}$. Both RGN-HR and GNDL-PCG were then applied to $\mathfrak{S}$ to compute a rank-3 approximation, taking about 0.08 and 0.06 seconds respectively. Let (M_1, M_2, M_3) and (M'_1, M'_2, M'_3) denote their respective factor matrices. Then, $(Q_1 M_1, Q_2 M_2, Q_3 M_3)$ and $(Q_1 M'_1, Q_2 M'_2, Q_3 M'_3)$ are the factor matrices of the corresponding rank-3 tensors $\mathfrak{B}_3$ and $\mathfrak{B}'_3$. The relative errors between $\mathfrak{B}_3$ and $\mathfrak{B}$ was approximately equal to $2.50493697 \cdot 10^{-2}$, and likewise for $\mathfrak{B}'_3$. Computing the best-3 approximation directly from $\mathfrak{B}$, without Tucker compression, yields an approximation error of about $2.505 \cdot 10^{-2}$ in both cases.

The columns of the second factor matrices of $\mathfrak{B}_3$ and $\mathfrak{B}'_3$ represent the emission spectra of the 3 identified components; they are plotted in the left graph in [Figure 8.1](#). The right graph in that figure visualizes the columns of the third factor matrices of $\mathfrak{B}_3$ and $\mathfrak{B}'_3$, which represent the excitation spectra. Note that the spectra recovered by RGN-HR are visually indistinguishable from those recovered by GNDL-PCG. For brevity, we do not plot the concentration profiles represented by the first factor matrices; their relative difference in Frobenius norm is approximately $2 \cdot 10^{-6}$.

The true concentrations are also provided in the data set. For both RGN-HR and GNDL-PCG, the columns of the first factor matrix correlate very well with the true data: correlation coefficients greater than 99.8% were obtained in all cases.

9. Conclusions. We proposed the first Riemannian optimization method for approximating a given tensor by one of low canonical rank. The theoretical analysis in [section 3](#) motivated why the proposed Riemannian formulation of the TAP should be preferred over the naive, overparameterized formulation involving factor matrices if a Riemannian Gauss–Newton method is employed. Specifically, our analysis predicts a great advantage of the proposed RGN method over state-of-the-art non-Riemannian GN methods when the CPD contains rank-1 terms whose norms are of different magnitudes. The numerical experiments in [section 7](#) confirmed this theory.

The second main innovation explored in this paper was to exploit information about the condition number in the optimization method. It was argued in [subsection 5.2](#) that certain ill-conditioned CPDs are hard to escape with the RGN process. We proposed hot restarts for escaping such regions of ill-conditioned CPDs. The numerical experiments revealed significant speedups attributed to these hot restarts.

Based on the numerical experiments, we believe that RGN-HR can be a suitable alternative for classic GN methods with direct solves of the TRS. Speedups between 3 and 2000 with respect to the state-of-the-art GN method `nls_gnd1` in Tensorlab v3.0 [70] were observed. Our experiments suggest that RGN-HR can be competitively applied up to $r(\Sigma + d) \approx 1000$, which is one order of magnitude larger than Tensorlab’s default choice for `nls_gnd1`.

An obstacle for extending the RGN-HR method to large-scale TAPs in which

the least-squares problem in the TRS is only approximately solved, for example via LSQR, is how a good estimate of the condition number (2.2) should be computed inexpensively. As RGN-HR was very competitive for small-scale problems, we believe that this is a promising direction for future work. Another interesting problem is to prove the observations in the informal analysis of subsection 5.2 rigorously.

Acknowledgments. We thank B. Vandereycken for helpful feedback he shared with us at the 7th Workshop on Matrix Equations and Tensor Techniques in Bologna, Italy held in February 2017. We also thank two anonymous referees for their valuable feedback that improved this article.

REFERENCES

- [1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008.
- [2] E. ACAR, D. DUNLAVY, T. KOLDA, AND M. MØRUP, *Scalable tensor factorizations for incomplete data*, Chemometr. Intell. Lab., 106 (2011), pp. 41–56.
- [3] R. L. ADLER, J.-P. DEDIEU, J. Y. MARGULIES, M. MARTENS, AND M. SHUB, *Newton’s method on Riemannian manifolds and a geometric model for the human spine*, IMA J. Numer. Anal., 22 (2002), pp. 359–390.
- [4] E. S. ALLMAN, C. MATIAS, AND J. A. RHODES, *Identifiability of parameters in latent structure models with many observed variables*, Ann. Statist., 37 (2009), pp. 3099–3132.
- [5] A. ANANDKUMAR AND R. GE, *Efficient approaches for escaping higher order saddle points in non-convex optimization*, JMLR: Workshop and Conference Proceedings, 49 (2016), pp. 1–22.
- [6] 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.
- [7] C. A. ANDERSSON AND R. BRO, *The n-way toolbox for MATLAB*, Chemom. Intell. Lab. Syst., 52 (2000), pp. 1–4.
- [8] C. J. APPELLOF AND E. R. DAVIDSON, *Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents*, Analytic Chemistry, 53 (1981), pp. 2053–2056.
- [9] J. BOCHNAK, M. COSTE, AND M. ROY, *Real Algebraic Geometry*, Springer-Verlag, 1998.
- [10] P. BREIDING AND N. VANNIEUWENHOVEN, *The condition number of join decompositions*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 287–309.
- [11] P. BREIDING AND N. VANNIEUWENHOVEN, *Convergence analysis of Riemannian Gauss–Newton methods and its connection with the geometric condition number*, Appl. Math. Letters, 78 (2018), pp. 42–50.
- [12] R. BRO, *Multi-way Analysis in the Food Industry. Models, Algorithms, and Applications*, PhD thesis, University of Amsterdam, 1998.
- [13] P. BÜRGISSEER AND F. CUCKER, *Condition: The Geometry of Numerical Algorithms*, vol. 349 of Grundlehren der mathematischen Wissenschaften, Springer, Heidelberg, 2013.
- [14] C. F. CAIAFA AND A. CICHOCKI, *Generalizing the column-row matrix decomposition to multi-way arrays*, Linear Algebra Appl., 433 (2010), pp. 557–573.
- [15] J. CARROLL AND J.-J. CHANG, *Analysis of individual differences in multidimensional scaling via an n-way generalization of “Eckart–Young” decomposition*, Psychometrika, 35 (1970), pp. 283–319.
- [16] B. CHEN, S. HE, Z. LI, AND S. ZHANG, *Maximum block improvement and polynomial optimization*, SIAM J. Optim., 22 (2012), pp. 87–107.
- [17] Y. CHEN, D. HAN, AND L. QI, *New ALS methods with extrapolating search directions and optimal step size for complex-valued tensor decompositions*, IEEE Trans. Signal Process., 59 (2011), pp. 5888–5898.
- [18] L. CHIANTINI, G. OTTAVIANI, AND N. VANNIEUWENHOVEN, *An algorithm for generic and low-rank specific identifiability of complex tensors*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 1265–1287.
- [19] L. CHIANTINI, G. OTTAVIANI, AND N. VANNIEUWENHOVEN, *Effective criteria for specific identifiability of tensors and forms*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 656–681.
- [20] C. DA SILVA AND F. J. HERRMANN, *Optimization on the Hierarchical Tucker manifold—applications to tensor completion*, Linear Algebra Appl., 481 (2015), pp. 131–173.
- [21] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1253–1278.
- [22] V. DE SILVA AND L.-H. LIM, *Tensor rank and the ill-posedness of the best low-rank approximation problem*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 1084–1127.
- [23] H. DE STERCK, *A nonlinear GMRES optimization algorithm for canonical tensor decomposition*, SIAM J. Sci. Comput., 34 (2012), pp. A1351–A1379.
- [24] H. DE STERCK AND K. MILLER, *An adaptive algebraic multigrid algorithm for low-rank canonical tensor decomposition*, SIAM J. Sci. Comput., 35 (2013), pp. B1–B24.
- [25] J.-P. DEDIEU AND M.-H. KIM, *Newton’s method for analytic systems of equations with constant rank derivatives*, J. Complexity, 18 (2002), pp. 187–209.
- [26] J. W. DEMMEL, *Applied Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1997.
- [27] P. DRINEAS AND M. W. MAHONEY, *A randomized algorithm for a tensor-based generalization*

- of the singular value decomposition, *Linear Algebra Appl.*, 420 (2007), pp. 553–571.
- [28] L. D. GARCIA, M. STILLMAN, AND B. STURMFELS, *Algebraic geometry of Bayesian networks*, *J. Symbolic Comput.*, 39 (2005), pp. 331–355.
 - [29] G. H. GOLUB AND C. VAN LOAN, *Matrix Computations*, The John Hopkins University Press, 3 ed., 1996.
 - [30] W. HACKBUSCH, *Tensor Spaces and Numerical Tensor Calculus*, vol. 42 of Springer Series in Computational Mathematics, Springer-Verlag, 2012.
 - [31] J. HARRIS, *Algebraic Geometry, A First Course*, vol. 133 of Graduate Text in Mathematics, Springer-Verlag, 1992.
 - [32] 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.
 - [33] C. HAYASHI AND F. HAYASHI, *A new algorithm to solve PARAFAC-model*, *Behaviormetrika*, 11 (1982), pp. 49–60.
 - [34] N. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, PA, 1996.
 - [35] F. L. HITCHCOCK, *The expression of a tensor or a polyadic as a sum of products*, *J. Math. Phys.*, 6 (1927), pp. 164–189.
 - [36] S. HOLTZ, T. ROHWEDDER, AND R. SCHNEIDER, *On manifolds of tensors of fixed TT-rank*, *BIT Numer. Math.*, 120 (2012), pp. 701–731.
 - [37] N. HUDSON, A. BAKER, AND D. REYNOLDS, *Fluorescence analysis of dissolved organic matter in natural, waste and polluted waters—a review*, *River Research and Applications*, 23 (2007), pp. 631–649.
 - [38] M. ISHTEVA, P.-A. ABSIL, S. VAN HUFFEL, AND L. DE LATHAUWER, *Best low multilinear rank approximation of higher-order tensors, based on the Riemannian trust-region scheme*, *SIAM J. Matrix Anal. Appl.*, 32 (2011), pp. 115–135.
 - [39] H. A. L. KIERS, *A three-step algorithm for CANDECOMP/PARAFAC analysis of large data sets with multicollinearity*, *J. Chemometrics*, 12 (1998).
 - [40] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, *SIAM Rev.*, 51 (2009), pp. 455–500.
 - [41] D. KRESSNER, M. STEINLECHNER, AND B. VANDEREYCKEN, *Low-rank tensor completion by Riemannian optimization*, *BIT Numer. Math.*, 54 (2014), pp. 447–468.
 - [42] J. M. LANDSBERG, *Tensors: Geometry and Applications*, vol. 128 of Graduate Studies in Mathematics, AMS, Providence, Rhode Island, 2012.
 - [43] J. M. LEE, *Introduction to Smooth Manifolds*, vol. 218 of Graduate Texts in Mathematics, Springer, New York, USA, 2 ed., 2013.
 - [44] Z. LI, A. USCHMAJEW, AND S. ZHANG, *On convergence of the maximum block improvement method*, *SIAM J. Optim.*, 25 (2015), pp. 210–233.
 - [45] Y. NESTEROV AND B. T. POLYAK, *Cubic regularization of newton method and its global convergence*, *Math. Program., Ser. A*, 108 (2006), pp. 177–205.
 - [46] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer Series in Operation Research and Financial Engineering, Springer, 2 ed., 2006.
 - [47] I. V. OSELEDETS, D. V. SAVOSTIANOV, AND E. E. TYRTYSHNIKOV, *Tucker dimensionality reduction of three-dimensional arrays in linear time*, *SIAM J. Matrix Anal. Appl.*, 30 (2008), pp. 939–956.
 - [48] P. PAATERO, *A weighted non-negative least squares algorithm for three-way ‘PARAFAC’ factor analysis*, *Chemometr. Intell. Lab.*, 38 (1997), pp. 223–242.
 - [49] P. PAATERO, *The Multilinear Engine—a table-driven, least squares program for solving multilinear problems, including the n-way parallel factor analysis model*, *J. Comput. Graph. Stat.*, 8 (1999), pp. 854–888.
 - [50] A.-H. PHAN, P. TICHAVSKÝ, AND A. CICHOCKI, *Fast alternating LS algorithms for high order CANDECOMP/PARAFAC tensor factorizations*, *IEEE Trans. Signal Process.*, 61 (2013), pp. 4834–4846.
 - [51] A.-H. PHAN, P. TICHAVSKÝ, AND A. CICHOCKI, *Low complexity damped Gauss–Newton algorithms for CANDECOMP/PARAFAC*, *SIAM J. Matrix Anal. Appl.*, 34 (2013), pp. 126–147.
 - [52] Y. QI, P. COMON, AND L.-H. LIM, *Semialgebraic Geometry of Nonnegative Tensor Rank*, *SIAM J. Matrix Anal. Appl.*, 37 (2016), pp. 1556–1580.
 - [53] M. RAJUH, P. COMON, AND R. A. HARSHMAN, *Enhanced line search: a novel method to accelerate PARAFAC*, *SIAM J. Matrix Anal. Appl.*, 30 (2008), pp. 1128–1147.
 - [54] A. K. SAIBABA, *HOID: higher order interpolatory decomposition for tensors based on Tucker representation*, *SIAM J. Matrix Anal. Appl.*, 37 (2016), pp. 1223–1249.
 - [55] B. SAVAS AND L.-H. LIM, *Quasi-Newton methods on Grassmannians and multilinear approximations of tensors*, *SIAM J. Matrix Anal. Appl.*, 32 (2010), pp. 3352–3393.
 - [56] A. SHASHUA AND T. HAZAN, *Non-negative tensor factorization with applications to statistics and computer vision*, in *Proceedings of the 22nd international conference on Machine learning*, 2005.
 - [57] N. D. SIDIROPOULOS, L. DE LATHAUWER, X. FU, K. HUANG, E. E. PAPALEXAKIS, AND C. FALOUTSOS, *Tensor decomposition for signal processing and machine learning*, *IEEE Trans. Signal Proc.*, 65 (2017), pp. 3551–3582.
 - [58] A. SMILDE, R. BRO, AND P. GELADI, *Multi-way Analysis: Applications in the Chemical Sciences*, John Wiley & Sons, Hoboken, New Jersey, 2004.
 - [59] L. SORBER, M. VAN BAREL, AND L. DE LATHAUWER, *Optimization-based algorithms for tensor decompositions: canonical polyadic decomposition, decomposition in rank- $(l_r, l_r, 1)$ terms, and a new generalization*, *SIAM J. Optim.*, 23 (2013), pp. 695–720.
 - [60] M. STEINLECHNER, *Riemannian optimization for high-dimensional tensor completion*, *SIAM*

- J. Sci. Comput., 38 (2016), pp. S461–S484.
- [61] G. W. STEWART AND J.-G. SUN, *Matrix Perturbation Theory*, vol. 33 of Computer Science and Scientific Computing, Academic Press, 1990.
- [62] J. TENDEIRO, M. BENNANI DOSSE, AND J. TEN BERGE, *First and second order derivatives for CP and INDSCAL*, Chemom. Intell. Lab. Syst., 106 (2011), pp. 27–36.
- [63] G. TOMASI, *Practical and computational aspects in chemometric data analysis*, PhD thesis, The Royal Veterinary and Agricultural University, Frederiksberg, Denmark, 2006.
- [64] G. TOMASI AND R. BRO, *A comparison of algorithms for fitting the PARAFAC model*, Comput. Stat. Data An., 50 (2006), pp. 1700–1734.
- [65] L. R. TUCKER, *Some mathematical notes on three-mode factor analysis*, Psychometrika, 31 (1966), pp. 279–311.
- [66] A. USCHMAJEW AND B. VANDEREYCKEN, *The geometry of algorithms using hierarchical tensors*, Linear Algebra Appl., 439 (2013), pp. 133–166.
- [67] N. VANNIEUWENHOVEN, *Condition numbers for the tensor rank decomposition*, Linear Algebra Appl., 535 (2017), pp. 35–86.
- [68] N. VANNIEUWENHOVEN, K. MEERBERGEN, AND R. VANDEBRIL, *Computing the gradient in optimization algorithms for the CP decomposition in constant memory through tensor blocking*, SIAM J. Sci. Comput., 37 (2015), pp. C415–C438.
- [69] N. VANNIEUWENHOVEN, R. VANDEBRIL, AND K. MEERBERGEN, *A new truncation strategy for the higher-order singular value decomposition*, SIAM J. Sci. Comput., 34 (2012), pp. A1027–A1052.
- [70] N. VERVLIT, O. DEBALS, L. SORBER, M. VAN BAREL, AND L. DE LATHAUWER, *Tensorlab v3.0*, March 2016, <http://www.tensorlab.net>.

Appendix A. Implementation details. This appendix describes efficient formulations of the critical steps in Algorithm 4.1. It is shown below that the following number of operations per iteration (lines 6–11) are obtained:

line 6:	$\mathcal{O}(dr\Pi + r\Sigma + d^2r^2\Sigma^2)$
line 7:	$\mathcal{O}(r^3(\Sigma + 1)^3 + (r + 1)\Pi)$
line 8:	$\mathcal{O}(4r\Sigma + r2^{d+1})$
lines 9–11:	$\mathcal{O}(\Pi + r\Sigma)$

A.1. Retraction. As observed in [41, section 3.3], the T-HOSVD retraction is computed efficiently by exploiting the structure of the tangent vectors. Essentially the same observations are valid for the ST-HOSVD retraction with foot at the point $p_i = \alpha_i \mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^d$, where $\mathbf{a}_i^k \in \mathbb{S}^{n_k-1}$. For completeness, we explain this below.

The TRS (4.1) is solved in local coordinates with respect to the basis B in (4.4), yielding $\hat{\mathbf{p}}$ as solution. Write $\hat{\mathbf{p}}^T = [\mathbf{x}_1^T \cdots \mathbf{x}_r^T]$ with $\mathbf{x}_i \in \mathbb{R}^{\Sigma+1}$, and then the components of the tangent vector $\mathbf{t} = (\mathbf{t}_1, \dots, \mathbf{t}_r) \in \mathbb{T}_{\mathbf{p}}\mathcal{S}^{\times r}$ are

$$\mathbf{t}_i = T_{p_i}\mathbf{x}_i = \sum_{k=1}^d T_{i,k}\mathbf{x}_i^k \in \mathbb{T}_{p_i}\mathcal{S}.$$

Since $U_{i,1} = I$, we find by the multilinearity of the tensor product that

$$p_i + \mathbf{t}_i = (\mathbf{x}_i^1 + \alpha_i \mathbf{a}_i^1) \otimes \mathbf{a}_i^2 \otimes \cdots \otimes \mathbf{a}_i^d + \mathbf{a}_i^1 \otimes (U_{i,2}\mathbf{x}_i^2) \otimes \mathbf{a}_i^3 \otimes \cdots \otimes \mathbf{a}_i^d + \cdots + \mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^{d-1} \otimes (U_{i,d}\mathbf{x}_i^d).$$

Let $\beta_{i,k}\mathbf{q}_i^k = U_{i,k}\mathbf{x}_i^k$ with $\|\mathbf{q}_i^k\| = 1$ for $k = 2, \dots, d$. Also define $\beta_{i,1}\mathbf{q}_i^1 = \mathbf{x}_i^1 + \alpha_i \mathbf{a}_i^1$ with $\|\mathbf{q}_i^1\| = 1$. Then, we may write $p_i + \mathbf{t}_i$ as the orthogonal Tucker decomposition

$$(A.1) \quad p_i + \mathbf{t}_i = (\mathbf{q}_i^1, Q_2, \dots, Q_d) \cdot \mathfrak{S}, \text{ where } Q_k = [\mathbf{a}_i^k \quad \mathbf{q}_i^k] \text{ for } k = 2, \dots, d,$$

and where the order- d tensor $\mathfrak{S} \in \mathbb{R}^{1 \times 2 \times \cdots \times 2}$ is given by

$$s_{1,i_2,\dots,i_d} = \begin{cases} \beta_{i,1} & \text{if } i_2 = \cdots = i_d = 1, \\ \beta_{i,k} & \text{if } i_k = 2 \text{ and } i_\ell = 1 \text{ for } 2 \leq \ell \neq k \leq d, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $Q_k^T Q_k = I$, by definition of $U_{i,k}$ in (4.2). To complete the derivation, we need the following result about ST-HOSVD's, which is proved in Appendix B.

LEMMA A.1. *Let $\mathfrak{A} \in \mathbb{R}^{n_1 \times \cdots \times n_d}$ admit an orthogonal Tucker decomposition $\mathfrak{A} = (Q_1, \dots, Q_d) \cdot \mathfrak{B}$, where $Q_k \in \mathbb{R}^{n_k \times m_k}$ has orthonormal columns and $\mathfrak{B} \in \mathbb{R}^{m_1 \times \cdots \times m_d}$. Let $(Z_1, \dots, Z_d) \cdot \mathfrak{C}$ be a rank- $(r_1, \dots, r_d)$ ST-HOSVD approximation of $\mathfrak{B}$ corresponding to the processing order π . Then, $(Q_1 Z_1, \dots, Q_d Z_d) \cdot \mathfrak{C}$ is a rank- $(r_1, \dots, r_d)$ ST-HOSVD approximation of $\mathfrak{A}$ corresponding to π .*

The rank- $(1, \dots, 1)$ ST-HOSVD approximation of $p_i + \mathbf{t}_i = (\mathbf{q}_i^1, Q_2, \dots, Q_d) \cdot \mathfrak{S}$ in (A.1) can thus be computed efficiently from the ST-HOSVD approximation $(1, \mathbf{r}_i^2, \dots, \mathbf{r}_i^d) \cdot \lambda$ of $\mathfrak{S}$; it is namely $(\mathbf{q}_i^1, Q_2 \mathbf{r}_i^2, \dots, Q_d \mathbf{r}_i^d) \cdot \lambda$.

Computing the retraction as above requires the following number of operations per rank-1 tensor $p_i \in \mathcal{S}$: $\mathcal{O}(2\Sigma)$ operations for constructing $\mathbf{q}_i^1$ and $Q_2, \dots, Q_d$; $\mathcal{O}(2^{d+1})$ operations for computing the ST-HOSVD approximation of $\mathfrak{S}$; and $\mathcal{O}(2\Sigma)$ operations for recovering the ST-HOSVD approximation of $p_i + \mathbf{t}_i$ from the approximation of $\mathfrak{S}$. This results in $\mathcal{O}(r(4\Sigma + 2^{d+1}))$ operations for one product ST-HOSVD retraction.

A.2. Gradient and Hessian approximation. Recall that the set of tensors of rank bounded by r is the image of the map $\llbracket \cdot \rrbracket$ in (3.1). Its Jacobian is J as in (3.3). Let $\mathbf{p} = (\mathbf{a}_i^1 \otimes \dots \otimes \mathbf{a}_i^d)_{i=1}^r$ and write $D := \text{diag}(U_{1,1}, \dots, U_{r,1}, \dots, U_{1,d}, \dots, U_{r,d})$, where the $U_{i,j}$ are defined as in subsection 4.1. Comparing with (4.5), we note $T_{\mathbf{p}} = JD$.

For $\mathbf{x} \in \mathbb{R}^\Pi$, the operation $J^T \mathbf{x}$ is an important computational kernel called the *CP gradient* [50]. Computationally efficient implementations were proposed in [50, 68]. Since $T_{\mathbf{p}}^T \mathbf{g}_{\mathbf{p}}$ is equivalent to $D^T (J^T \mathbf{g}_{\mathbf{p}})$, one can employ any of the efficient methods for computing CP gradients; afterwards multiply the result by the block diagonal matrix D . In our Matlab implementation, the CP gradient is computed via left-to-right and right-to-left (RTL) contractions [50], [68, section 2.3]. This scheme requires $\mathcal{O}(dr\Pi + r\Sigma)$ operations.

Efficient algorithms for computing $J^T J$ were investigated in the literature; e.g., [48, 51, 59, 62, 63]. One can then efficiently compute $H_{\mathbf{p}}$ as $H_{\mathbf{p}} = T_{\mathbf{p}}^T T_{\mathbf{p}} = D^T (J^T J) D$. In our implementation we chose the method from [59] for constructing $J^T J$ efficiently. The computational complexity of this algorithm is about $\mathcal{O}((dr\Sigma)^2)$.

A.3. Optimal coefficients. The solution $\mathbf{x}^*$ of least-squares problem (4.3) is obtained by observing that it is equivalent to

$$\mathbf{x}^* = [\mathbf{a}_1^1 \otimes \dots \otimes \mathbf{a}_1^d \quad \dots \quad \mathbf{a}_r^1 \otimes \dots \otimes \mathbf{a}_r^d]^{\dagger} \text{vec}(\mathfrak{B}) =: (A_1 \odot \dots \odot A_d)^{\dagger} \mathbf{b},$$

where $\mathbf{b} = \text{vec}(\mathfrak{B}) \in \mathbb{R}^\Pi$ is the vectorization of $\mathfrak{B}$, $A_k = [\mathbf{a}_i^k]_{i=1}^r$ are the factor matrices and $\odot$ is the columnwise Khatri–Rao product. Recall that the optimal coefficients are computed either in step 2 of Algorithm 4.1 after randomly initializing the rank-1 tensors, or in step 11 of Algorithm 5.1 after randomly perturbing them. As a result, in both cases $A_1 \odot \dots \odot A_d$ has linearly independent columns with probability 1.¹⁰ Then, it is well-known [40] that

$$\mathbf{x}^* = (A_1 \odot \dots \odot A_d)^{\dagger} \mathbf{b} = ((A_1^T A_1) \otimes \dots \otimes (A_d^T A_d))^{-1} ((A_1 \odot \dots \odot A_d)^T \mathbf{b}),$$

where $\otimes$ is the element-wise or Hadamard product. The rightmost matrix-vector product can be interpreted as r simultaneous tensor-to-vector contractions, which we compute with RTL contractions. Thereafter, the linear system is constructed as the formula suggests and solved via a Cholesky factorization. The optimal coefficients can thus be computed in $\mathcal{O}(r\Pi + r^2\Sigma + r^3)$ operations.

Appendix B. Proofs of the lemma's.

Proof of Lemma 3.2. We have $d_{\mathbf{p}}\Phi(U_1 \mathbf{x}_1, \dots, U_r \mathbf{x}_r) = U_1 \mathbf{x}_1 + \dots + U_r \mathbf{x}_r$ from [10], where $\mathbf{x}_i \in \mathbb{R}^{\Sigma+1}$, and $U_i \in \mathbb{R}^{\Pi \times (\Sigma+1)}$ contains an orthonormal basis of $T_{p_i} \mathcal{S}$. Then, using the triangle inequality we obtain

$$\|d_{\mathbf{p}}\Phi\|_2 = \max_{\|U_1 \mathbf{x}_1\|^2 + \dots + \|U_r \mathbf{x}_r\|^2 = 1} \|U_1 \mathbf{x}_1 + \dots + U_r \mathbf{x}_r\| \leq \max_{c_1^2 + \dots + c_r^2 = 1} (c_1 + \dots + c_r),$$

where we set $c_i := \|U_i \mathbf{x}_i\| = \|\mathbf{x}_i\| \in \mathbb{R}^r$ in the last step. Since we have the inequality $\max_{\|\mathbf{c}\|=1} \|\mathbf{c}\|_1 \leq \sqrt{r} \|\mathbf{c}\| = \sqrt{r}$. This proves the upper bound.

To prove the lower bound, take any $\mathbf{x}_1$ of unit norm. Then,

$$\|d_{\mathbf{p}}\Phi(U_1 \mathbf{x}_1, U_2 0, \dots, U_r 0)\| = \|U_1 \mathbf{x}_1\| = 1,$$

which implies that $\|d_{\mathbf{p}}\Phi\|_2 \geq 1$. The proof is concluded. $\square$

¹⁰One of our assumptions is that $r < \frac{\Pi}{\Sigma+1}$ is strictly subgeneric, so that in particular $r < \Pi$.

Proof of Lemma 5.3. Let $\mathcal{S}_{\mathbb{C}}$ denote the complexification of $\mathcal{S}$, and let $\sigma_r(\mathcal{S}_{\mathbb{C}})$ denote the r -secant variety of $\mathcal{S}_{\mathbb{C}}$ [31]. Recall that the real points of $\mathcal{S}_{\mathbb{C}}$ are dense in the Zariski topology [52, section 5], that $\sigma_r(\mathcal{S})$ (and so $\sigma_r(\mathcal{S}_{\mathbb{C}})$ as well) is non-defective by assumption, and, hence, that $\text{rank } T_{\mathbf{p}} = \dim T_{\Phi(\mathbf{p})} \sigma_r(\mathcal{S}_{\mathbb{C}}) = \dim(\mathcal{S}_{\mathbb{C}})^{\times r}$ for all $\mathbf{p} \in \mathcal{S}^{\times r}$ in a Zariski-dense set by Terracini's lemma [42]. Then, applying Sard's theorem [31] to $\Phi_{\mathbb{C}} : (\mathcal{S}_{\mathbb{C}})^{\times r} \rightarrow \mathcal{S}_{\mathbb{C}}, (p_1, \dots, p_r) \mapsto p_1 + \dots + p_r$ shows that $\mathcal{I}_r$ is contained in the complex subvariety of critical points of $\Phi_{\mathbb{C}}$. $\square$

Proof of Lemma 6.3. The proof is a verbatim copy of the proof of [41, Proposition 2.3], where the necessary quasi-best approximation property is [30, Theorem 10.5]. $\square$

Proof of Lemma A.1. Without loss of generality we assume that $\mathbf{p} = [1 \ 2 \ \dots \ d]$.

Let $(U_1, \dots, U_d) \cdot \mathfrak{S}$ be the ST-HOSVD approximation of $\mathfrak{A}$. We prove by induction that $U_i = Q_i Z_i$. It suffices to show that the bases can be chosen to be the same because they completely determine the core tensor by the relationship $\mathfrak{S} = (U_1, \dots, U_d)^T \cdot \mathfrak{A}$.

Let USV^T be a compact SVD of $\mathfrak{A}_{(1)}$ and let $\tilde{U}\tilde{S}\tilde{V}^T$ be an SVD of $\mathfrak{B}_{(1)}$. Then,

$$USV^T = \mathfrak{A}_{(1)} = Q_1 \mathfrak{B}_{(1)} (Q_2 \otimes \dots \otimes Q_d)^T = (Q_1 \tilde{U}) \tilde{S} (\tilde{V}^T (Q_2 \otimes \dots \otimes Q_d)^T).$$

Both the leftmost and rightmost expressions specify an SVD of $\mathfrak{A}_{(1)}$. Since it is essentially unique, $\tilde{U}$ can be chosen so that $U = Q_1 \tilde{U}$. By definition of the ST-HOSVD in [69], $U = [U_1 \ X]$ and $\tilde{U} = [Z_1 \ Y]$, so that $U_1 = Q_1 Z_1$, proving the base case.

Assume now that $U_\ell = Q_\ell Z_\ell$ for all $\ell = 1, \dots, k-1$, then we prove that it holds for k as well. Define

$$\mathfrak{A}_{(k)}^{(k-1)} := \mathfrak{A}_{(k)}(U_1 \otimes \dots \otimes U_{k-1} \otimes I \otimes \dots \otimes I) \text{ and}$$

$$\mathfrak{B}_{(k)}^{(k-1)} := \mathfrak{B}_{(k)}(Z_1 \otimes \dots \otimes Z_{k-1} \otimes I \otimes \dots \otimes I).$$

Since $\mathfrak{A}_{(k)} = Q_k \mathfrak{B}_{(k)} (Q_1 \otimes \dots \otimes Q_{k-1} \otimes Q_{k+1} \otimes \dots \otimes Q_d)^T$, it follows that

$$\begin{aligned} \mathfrak{A}_{(k)}^{(k-1)} &= Q_k \mathfrak{B}_{(k)} (Q_1^T U_1 \otimes \dots \otimes Q_{k-1}^T U_{k-1} \otimes Q_{k+1}^T \otimes \dots \otimes Q_d^T) \\ &= Q_k \mathfrak{B}_{(k)} (Z_1 \otimes \dots \otimes Z_{k-1} \otimes Q_{k+1}^T \otimes \dots \otimes Q_d^T) \\ &= Q_k \mathfrak{B}_{(k)}^{(k-1)} (I \otimes \dots \otimes I \otimes Q_{k+1} \otimes \dots \otimes Q_d)^T. \end{aligned}$$

Let USV^T be the compact SVD of $\mathfrak{A}_{(k)}^{(k-1)}$ and let $\tilde{U}\tilde{S}\tilde{V}^T$ be the compact SVD of $\mathfrak{B}_{(k)}^{(k-1)}$. Then,

$$USV^T = \mathfrak{A}_{(k)}^{(k-1)} = (Q_k \tilde{U}) \tilde{S} (\tilde{V}^T (I \otimes \dots \otimes I \otimes Q_{k+1} \otimes \dots \otimes Q_d)^T).$$

The leftmost and rightmost expressions are both compact SVDs, hence $\tilde{U}$ can be chosen so that $U = Q_k \tilde{U}$. From the definition of the ST-HOSVD in [69] it again follows that $U_k = Q_k Z_k$, which concludes the proof. $\square$