



Completely positive factorization by a Riemannian smoothing method

Zhijian Lai¹ · Akiko Yoshise²

Received: 17 December 2021 / Accepted: 19 September 2022 / Published online: 9 October 2022
© The Author(s) 2022

Abstract

Copositive optimization is a special case of convex conic programming, and it consists of optimizing a linear function over the cone of all completely positive matrices under linear constraints. Copositive optimization provides powerful relaxations of NP-hard quadratic problems or combinatorial problems, but there are still many open problems regarding copositive or completely positive matrices. In this paper, we focus on one such problem; finding a completely positive (CP) factorization for a given completely positive matrix. We treat it as a nonsmooth Riemannian optimization problem, i.e., a minimization problem of a nonsmooth function over a Riemannian manifold. To solve this problem, we present a general smoothing framework for solving nonsmooth Riemannian optimization problems and show convergence to a stationary point of the original problem. An advantage is that we can implement it quickly with minimal effort by directly using the existing standard smooth Riemannian solvers, such as Manopt. Numerical experiments show the efficiency of our method especially for large-scale CP factorizations.

Keywords Completely positive factorization · Smoothing method · Nonsmooth Riemannian optimization problem

Mathematics Subject Classification 15A23 · 15B48 · 90C48 · 90C59

✉ Akiko Yoshise
yoshise@sk.tsukuba.ac.jp

¹ Graduate School of Systems and Information Engineering, University of Tsukuba, Tsukuba, Ibaraki 305-8573, Japan

² Faculty of Engineering, Information and Systems, University of Tsukuba, Tsukuba, Ibaraki 305-8573, Japan

1 Introduction

The space of $n \times n$ real symmetric matrices \mathcal{S}_n is endowed with the trace inner product $\langle A, B \rangle := \text{trace}(AB)$. A matrix $A \in \mathcal{S}_n$ is called *completely positive* if for some $r \in \mathbb{N}$ there exists an entrywise nonnegative matrix $B \in \mathbb{R}^{n \times r}$ such that $A = BB^T$, and we call B a *CP factorization* of A . We define \mathcal{CP}_n as the set of $n \times n$ completely positive matrices, equivalently characterized as

$$\mathcal{CP}_n := \{BB^T \in \mathcal{S}_n \mid B \text{ is a nonnegative matrix}\} = \text{conv}\{xx^T \mid x \in \mathbb{R}_+^n\},$$

where $\text{conv}(S)$ denotes the convex hull of a given set S . We denote the set of $n \times n$ copositive matrices by $\mathcal{COP}_n := \{A \in \mathcal{S}_n \mid x^T Ax \geq 0 \text{ for all } x \in \mathbb{R}_+^n\}$. It is known that \mathcal{COP}_n and \mathcal{CP}_n are duals of each other under the trace inner product; moreover, both \mathcal{CP}_n and \mathcal{COP}_n are proper convex cones [1, Section2.2]. For any positive integer n , we have the following inclusion relationship among other important cones in conic optimization:

$$\mathcal{CP}_n \subseteq \mathcal{S}_n^+ \cap \mathcal{N}_n \subseteq \mathcal{S}_n^+ \subseteq \mathcal{S}_n^+ + \mathcal{N}_n \subseteq \mathcal{COP}_n,$$

where \mathcal{S}_n^+ is the cone of $n \times n$ symmetric positive semidefinite matrices and \mathcal{N}_n is the cone of $n \times n$ symmetric nonnegative matrices. See the monograph [1] for a comprehensive description of \mathcal{CP}_n and \mathcal{COP}_n .

Conic optimization is a subfield of convex optimization that studies minimization of linear functions over proper cones. Here, if the proper cone is \mathcal{CP}_n or its dual cone \mathcal{COP}_n , we call the conic optimization problem a *copositive programming* problem. Copositive programming is closely related to many nonconvex, NP-hard quadratic and combinatorial optimizations [2]. For example, consider the so-called standard quadratic optimization problem,

$$\min\{x^T Mx \mid \mathbf{e}^T x = 1, x \in \mathbb{R}_+^n\}, \tag{1}$$

where $M \in \mathcal{S}_n$ is possibly not positive semidefinite and \mathbf{e} is the all-ones vector. Bomze et al. [3] showed that the following completely positive reformulation,

$$\min\{\langle M, X \rangle \mid \langle E, X \rangle = 1, X \in \mathcal{CP}_n\},$$

where E is the all-ones matrix, is equivalent to (1). Burer [4] reported a more general result, where any quadratic problem with binary and continuous variables can be rewritten as a linear program over \mathcal{CP}_n . As an application to combinatorial problems, consider the problem of computing the independence number $\alpha(G)$ of a graph G with n nodes. De Klerk and Pasechnik [5] showed that

$$\alpha(G) = \max\{\langle E, X \rangle \mid \langle A + I, X \rangle = 1, X \in \mathcal{CP}_n\},$$

where A is the adjacency matrix of G . For surveys on applications of copositive programming, see [2, 6–9].

The difficulty of the above problems lies entirely in the completely positive conic constraint. Note that because neither \mathcal{COP}_n nor \mathcal{CP}_n is self-dual, the primal-dual interior

point method for conic optimization does not work as is. Besides this, there are many open problems related to completely positive cones. One is checking membership in \mathcal{CP}_n , which was shown to be NP-hard by [10]. Computing or estimating the cp-rank, as defined later in (3), is also an open problem. We refer the reader to [9, 11] for a detailed discussion of those unresolved issues.

In this paper, we focus on finding a CP factorization for a given $A \in \mathcal{CP}_n$, i.e., the *CP factorization problem*:

$$\text{Find } B \in \mathbb{R}^{n \times r} \text{ s.t. } A = BB^T \text{ and } B \geq 0, \tag{CPfact}$$

which seems to be closely related to the membership problem $A \in \mathcal{CP}_n$. Sometimes, a matrix is shown to be completely positive through duality, or rather, $\langle A, X \rangle \geq 0$ for all $X \in \mathcal{COP}_n$, but in this case, a CP factorization will not necessarily be obtained.

1.1 Related work on CP factorization

Various methods of solving CP factorization problems have been studied. Jarre and Schmallowsky [12] stated a criterion for complete positivity, based on the augmented primal dual method to solve a particular second-order cone problem. Dickinson and Dür [13] dealt with complete positivity of matrices that possess a specific sparsity pattern and proposed a method for finding CP factorizations of these special matrices that can be performed in linear time. Nie [14] formulated the CP factorization problem as an \mathcal{A} -truncated K -moment problem, for which the author developed an algorithm that solves a series of semidefinite optimization problems. Sponsel and Dür [15] considered the problem of projecting a matrix onto \mathcal{CP}_n and \mathcal{COP}_n by using polyhedral approximations of these cones. With the help of these projections, they devised a method to compute a CP factorization for any matrix in the interior of \mathcal{CP}_n . Bomze [16] showed how to construct a CP factorization of an $n \times n$ matrix based on a given CP factorization of an $(n - 1) \times (n - 1)$ principal submatrix. Dutour Sikirić et al. [17] developed a simplex-like method for a rational CP factorization that works if the input matrix allows a rational CP factorization.

In 2020, Groetzner and Dür [18] applied the alternating projection method to the CP factorization problem by posing it as an equivalent feasibility problem (see (FeasCP)). Shortly afterwards, Chen et al. [19] reformulated the split feasibility problem as a difference-of-convex optimization problem and solved (FeasCP) as a specific application. In fact, we will solve this equivalent feasibility problem (FeasCP) by other means in this paper. In 2021, Boğ and Nguyen [20] proposed a projected gradient method with relaxation and inertia parameters for the CP factorization problem, aimed at solving

$$\min_X \{ \|A - XX^T\|^2 \mid X \in \mathbb{R}_+^{n \times r} \cap \mathcal{B}(0, \sqrt{\text{trace}(A)}) \}, \tag{2}$$

where $\mathcal{B}(0, \varepsilon) := \{X \in \mathbb{R}^{n \times r} \mid \|X\| \leq \varepsilon\}$ is the closed ball centered at 0. The authors argued that its optimal value is zero if and only if $A \in \mathcal{CP}_n$.

1.2 Our contributions and organization of the paper

Inspired by the idea of Groetzner and Dür [18], wherein CPfact is shown to be equivalent to a feasibility problem called (FeasCP), we treat the problem (FeasCP) as a non-smooth Riemannian optimization problem and solve it through a general Riemannian smoothing method. Our contributions are summarized as follows:

1. Although it is not explicitly stated in [18], (FeasCP) is actually a Riemannian optimization formulation. We propose a new Riemannian optimization technique and apply it to the problem.
2. In particular, we present a general framework of Riemannian smoothing for the nonsmooth Riemannian optimization problem and show convergence to a stationary point of the original problem.
3. We apply the general framework of Riemannian smoothing to CP factorization. Numerical experiments clarify that our method is competitive with other efficient CP factorization methods, especially for large-scale matrices.

In Sect. 2, we review the process to reconstruct CPfact into another feasibility problem; in particular, we take a different approach to this problem from those in other studies. In Sect. 3, we describe the general framework of smoothing methods for Riemannian optimization. To apply it to the CP factorization problem, we employ a smoothing function named LogSumExp. Section 4 is a collection of numerical experiments for CP factorization. As a meaningful supplement, in Sect. 5, we conduct further experiments (FSV problem and robust low-rank matrix completion) to explore the numerical performance of various sub-algorithms and smoothing functions on different applications.

2 Preliminaries

2.1 cp-rank and cp-plus-rank

First, let us recall some basic properties of completely positive matrices. Generally, many CP factorizations of a given A may exist, and they may vary in their numbers of columns. This gives rise to the following definitions: the cp-rank of $A \in \mathcal{S}_n$, denoted by $\text{cp}(A)$, is defined as

$$\text{cp}(A) := \min\{r \in \mathbb{N} \mid A = BB^T, B \in \mathbb{R}^{n \times r}, B \geq 0\}, \quad (3)$$

where $\text{cp}(A) = \infty$ if $A \notin \mathcal{CP}_n$. Similarly, we can define the cp-plus-rank as

$$\text{cp}^+(A) := \min\{r \in \mathbb{N} \mid A = BB^T, B \in \mathbb{R}^{n \times r}, B > 0\}.$$

Immediately, for all $A \in \mathcal{S}_n$, we have

$$\text{rank}(A) \leq \text{cp}(A) \leq \text{cp}^+(A). \quad (4)$$

Every CP factorization B of A is of the same rank as A since $\text{rank}(XX^\top) = \text{rank}(X)$ holds for any matrix X . The first inequality of (4) comes from the fact that for any CP factorization B ,

$$\text{rank}(A) = \text{rank}(B) \leq \text{the number of columns of } B.$$

The second is trivial by definition.

Note that computing or estimating the cp-rank of any given $A \in \mathcal{CP}_n$ is still an open problem. The following result gives a tight upper bound of the cp-rank for $A \in \mathcal{CP}_n$ in terms of the order n .

Theorem 2.1 (Bomze, Dickinson, and Still [21, Theorem 4.1]) *For all $A \in \mathcal{CP}_n$, we have*

$$\text{cp}(A) \leq \text{cp}_n := \begin{cases} n & \text{for } n \in \{2, 3, 4\} \\ \frac{1}{2}n(n+1) - 4 & \text{for } n \geq 5. \end{cases} \tag{5}$$

The following result is useful for distinguishing completely positive matrices in either the interior or on the boundary of \mathcal{CP}_n .

Theorem 2.2 (Dickinson [22, Theorem 3.8]) *We have*

$$\begin{aligned} \text{int}(\mathcal{CP}_n) &= \{A \in \mathcal{S}_n \mid \text{rank}(A) = n, \text{cp}^+(A) < \infty\} \\ &= \{A \in \mathcal{S}_n \mid \text{rank}(A) = n, A = BB^\top, B \in \mathbb{R}^{n \times r}, B \geq 0, \\ &\quad b_j > 0 \text{ for at least one column } b_j \text{ of } B\}. \end{aligned}$$

2.2 CP factorization as a feasibility problem

Groetzner and Dür [18] reformulated the CP factorization problem as an equivalent feasibility problem containing an orthogonality constraint.

Given $A \in \mathcal{CP}_n$, we can easily get another CP factorization \widehat{B} with r' columns for every integer $r' \geq r$, if we also have a CP factorization B with r columns. The simplest way to construct such an $n \times r'$ matrix \widehat{B} is to append $k := r' - r$ zero columns to B , i.e., $\widehat{B} := [B, 0_{n \times k}] \geq 0$. Another way is called *column replication*, i.e.,

$$\widehat{B} := [b_1, \dots, b_{n-1}, \underbrace{\frac{1}{\sqrt{m}}b_n, \dots, \frac{1}{\sqrt{m}}b_n}_{m := r' - n + 1 \text{ columns}}], \tag{6}$$

where b_i denotes the i -th column of B . It is easy to see that $\widehat{B}\widehat{B}^\top = BB^\top = A$. The next lemma is easily derived from the previous discussion, and it implies that there always exists an $n \times \text{cp}_n$ CP factorization for any $A \in \mathcal{CP}_n$. Recall that definition of cp_n is given in (5).

Lemma 2.3 *Suppose that $A \in \mathcal{S}_n$, $r \in \mathbb{N}$. Then $r \geq \text{cp}(A)$ if and only if A has a CP factorization B with r columns.*

Let $\mathcal{O}(r)$ denote the orthogonal group of order r , i.e., the set of $r \times r$ orthogonal matrices. The following lemma is essential to our study. Note that many authors have proved the existence of such an orthogonal matrix X (see, e.g., [23, Lemma 2.1] and [18, Lemma 2.6]).

Lemma 2.4 *Let $B, C \in \mathbb{R}^{n \times r}$. $BB^T = CC^T$ if and only if there exists $X \in \mathcal{O}(r)$ with $BX = C$.*

The next proposition puts the previous two lemmas together.

Proposition 2.5 *Let $A \in \mathcal{CP}_n$, $r \geq \text{cp}(A)$, $A = \bar{B}\bar{B}^T$, where $\bar{B} \in \mathbb{R}^{n \times r}$ may possibly be not nonnegative. Then there exists an orthogonal matrix $X \in \mathcal{O}(r)$ such that $\bar{B}X \geq 0$ and $A = (\bar{B}X)(\bar{B}X)^T$.*

This proposition tells us that one can find an orthogonal matrix X which can turn a “bad” factorization \bar{B} into a “good” factorization $\bar{B}X$. Let $r \geq \text{cp}(A)$ and $\bar{B} \in \mathbb{R}^{n \times r}$ be an arbitrary (possibly not nonnegative) initial factorization $A = \bar{B}\bar{B}^T$. The task of finding a CP factorization of A can then be formulated as the following feasibility problem,

$$\text{Find } X \text{ s.t. } \bar{B}X \geq 0 \text{ and } X \in \mathcal{O}(r). \tag{FeasCP}$$

We should notice that the condition $r \geq \text{cp}(A)$ is necessary; otherwise, (FeasCP) has no solution even if $A \in \mathcal{CP}_n$. Regardless of the exact $\text{cp}(A)$ which is often unknown, one can use $r = \text{cp}_n$ in (5). Note that finding an initial matrix \bar{B} is not difficult. Since a completely positive matrix is necessarily positive semidefinite, one can use Cholesky decomposition or spectral decomposition and then extend it to r columns by using (6). The following corollary shows that the feasibility of (FeasCP) is precisely a criterion for complete positivity.

Corollary 2.6 *Set $r \geq \text{cp}(A)$, $\bar{B} \in \mathbb{R}^{n \times r}$ an arbitrary initial factorization of A . Then $A \in \mathcal{CP}_n$ if and only if (FeasCP) is feasible. In this case, for any feasible solution X , $\bar{B}X$ is a CP factorization of A .*

In this study, solving (FeasCP) is the key to finding a CP factorization, but it is still a hard problem because $\mathcal{O}(r)$ is nonconvex.

2.3 Approaches to solving (FeasCP)

Groetzner and Dür [18] applied the so-called alternating projections method to (FeasCP). They defined the polyhedral cone, $\mathcal{P} := \{X \in \mathbb{R}^{r \times r} : \bar{B}X \geq 0\}$, and rewrote (FeasCP) as

Find X s.t. $X \in \mathcal{P} \cap \mathcal{O}(r)$.

The alternating projections method is as follows: choose a starting point $X_0 \in \mathcal{O}(r)$; then compute $P_0 = \text{proj}_{\mathcal{P}}(X_0)$ and $X_1 = \text{proj}_{\mathcal{O}(r)}(P_0)$, and iterate this process. Computing the projection onto \mathcal{P} amounts to solving a second-order cone problem (SOCP), while computing the projection onto $\mathcal{O}(r)$ amounts to a singular value decomposition. Note that we need to solve an SOCP alternately at every iteration, which is still expensive in practice. A modified version without convergence involves calculating an approximation of $\text{proj}_{\mathcal{P}}(X_k)$ by using the Moore-Penrose inverse of \bar{B} ; for details, see [18, Algorithm 2].

Our way is to use the optimization form. Here, we denote by $\max(\cdot)$ (resp. $\min(\cdot)$) the *max-function* (resp. *min-function*) that selects the largest (resp. smallest) entry of a vector or matrix. Notice that $-\min(\cdot) = \max(-\cdot)$. We associate (FeasCP) with the following optimization problem:

$$\max_{X \in \mathcal{O}(r)} \{ \min(\bar{B}X) \}.$$

For consistency of notation, we turn the maximization problem into a minimization problem:

$$\min_{X \in \mathcal{O}(r)} \{ \max(-\bar{B}X) \}. \tag{OptCP}$$

The feasible set $\mathcal{O}(r)$ is known to be compact [24, Observation 2.1.7]. In accordance with the extreme value theorem [25, Theorem 4.16], OptCP attains the global minimum, say t . Summarizing these observations together with Corollary 2.6 yields the following proposition.

Proposition 2.7 *Set $r \geq \text{cp}(A)$, and let $\bar{B} \in \mathbb{R}^{n \times r}$ be an arbitrary initial factorization of A . Then the following statements are equivalent:*

1. $A \in \mathcal{CP}_n$.
2. (FeasCP) is feasible.
3. In OptCP, there exists a feasible solution X such that $\max(-\bar{B}X) \leq 0$; alternatively, $\min(\bar{B}X) \geq 0$.
4. In OptCP, the global minimum $t \leq 0$.

3 Riemannian smoothing method

The problem of minimizing a real-valued function over a Riemannian manifold \mathcal{M} , which is called Riemannian optimization, has been actively studied during the last few decades. In particular, the Stiefel manifold,

$$\text{St}(n, p) = \{X \in \mathbb{R}^{n \times p} \mid X^T X = I\},$$

(when $n = p$, it reduces to the orthogonal group) is an important case and is our main interest here. We treat the CP factorization problem, i.e., OptCP as a problem of minimizing a nonsmooth function over a Riemannian manifold, for which variants of subgradient methods [26], proximal gradient methods [27], and the alternating direction method of multipliers (ADMM) [28] have been studied.

Smoothing methods [29], which use a parameterized smoothing function to approximate the objective function, are effective on a class of nonsmooth optimizations in Euclidean space. Recently, Zhang, Chen and Ma [30] extended a smoothing steepest descent method to the case of Riemannian submanifolds in \mathbb{R}^n . This is not the first time that smoothing methods have been studied on manifolds. Liu and Boumal [31] extended the augmented Lagrangian method and exact penalty method to the Riemannian case. The latter leads to a nonsmooth Riemannian optimization problem to which they applied smoothing techniques. Cambier and Absil [32] dealt with the problem of robust low-rank matrix completion by solving a Riemannian optimization problem, wherein they applied a smoothing conjugate gradient method.

In this section, we propose a general Riemannian smoothing method and apply it to the CP factorization problem.

3.1 Notation and terminology of Riemannian optimization

Let us briefly review some concepts in Riemannian optimization, following the notation of [33]. Throughout this paper, \mathcal{M} will refer to a complete Riemannian submanifold of Euclidean space \mathbb{R}^n . Thus, \mathcal{M} is endowed with a Riemannian metric induced from the Euclidean inner product, i.e., $\langle \xi, \eta \rangle_x := \xi^T \eta$ for any $\xi, \eta \in T_x \mathcal{M}$, where $T_x \mathcal{M} \subseteq \mathbb{R}^n$ is the tangent space to \mathcal{M} at x . The Riemannian metric induces the usual Euclidean norm $\|\xi\|_x := \|\xi\| = \sqrt{\langle \xi, \xi \rangle_x}$ for $\xi \in T_x \mathcal{M}$. The tangent bundle $T\mathcal{M} := \bigsqcup_{x \in \mathcal{M}} T_x \mathcal{M}$ is a disjoint union of the tangent spaces of \mathcal{M} . Let $f : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function on \mathcal{M} . The Riemannian gradient of f is a vector field $\text{grad}f$ on \mathcal{M} that is uniquely defined by the identities: for all $(x, v) \in T\mathcal{M}$,

$$Df(x)[v] = \langle v, \text{grad}f(x) \rangle_x$$

where $Df(x) : T_x \mathcal{M} \rightarrow T_{f(x)} \mathbb{R} \cong \mathbb{R}$ is the differential of f at $x \in \mathcal{M}$. Since \mathcal{M} is an embedded submanifold of \mathbb{R}^n , we have a simpler statement for f that is also well defined on the whole \mathbb{R}^n :

$$\text{grad}f(x) = \text{Proj}_x(\nabla f(x)),$$

where $\nabla f(x)$ is the usual gradient in \mathbb{R}^n and Proj_x denotes the orthogonal projector from \mathbb{R}^n to $T_x \mathcal{M}$. For a subset $D \subseteq \mathbb{R}^n$, the function $h \in C^1(D)$ is smooth, i.e., continuously differentiable on D . Given a point $x \in \mathbb{R}^n$ and $\delta > 0$, $\mathcal{B}(x, \delta)$ denotes a closed ball of radius δ centered at x . \mathbb{R}_{++} denotes the set of positive real numbers. We use subscript notation x_i to select the i th entry of a vector and superscript notation x^k to designate an element in a sequence $\{x^k\}$.

3.2 Ingredients

Now let us consider the nonsmooth Riemannian optimization problem (NROP):

$$\min_{x \in \mathcal{M}} f(x), \tag{NROP}$$

where $\mathcal{M} \subseteq \mathbb{R}^n$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a *proper lower semi-continuous* function (maybe nonsmooth or even non-Lipschitzian) on \mathbb{R}^n . For convenience, the term smooth Riemannian optimization problem (SROP) refers to NROP when $f(\cdot)$ is continuously differentiable on \mathbb{R}^n . To avoid confusion in this case, we use g instead of f ,

$$\min_{x \in \mathcal{M}} g(x). \tag{SROP}$$

Throughout this subsection, we will refer to many of the concepts in [30].

First, let us review the usual concepts and properties related to generalized subdifferentials in \mathbb{R}^n . For a proper lower semi-continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the *Fréchet subdifferential* and the *limiting subdifferential of f at $x \in \mathbb{R}^n$* are defined as

$$\begin{aligned} \hat{\partial}f(x) &:= \{ \nabla h(x) \mid \exists \delta > 0 \text{ such that } h \in C^1(\mathcal{B}(x, \delta)) \text{ and} \\ &\quad f - h \text{ attains a local minimum at } x \text{ on } \mathbb{R}^n \}, \\ \partial f(x) &:= \{ \lim_{\ell \rightarrow \infty} v^\ell \mid v^\ell \in \hat{\partial}f(x^\ell), (x^\ell, f(x^\ell)) \rightarrow (x, f(x)) \}. \end{aligned}$$

The definition of $\hat{\partial}f(x)$ above is not the standard one: the standard definition follows [34, 8.3 Definition]. But these definitions are equivalent by [34, 8.5 Proposition]. For locally Lipschitz functions, the *Clarke subdifferential at $x \in \mathbb{R}^n$* , $\partial^\circ f(x)$, is the convex hull of the limiting subdifferential. Their relationship is as follows:

$$\hat{\partial}f(x) \subseteq \partial f(x) \subseteq \partial^\circ f(x).$$

Notice that if f is convex, $\partial f(x)$ and $\partial^\circ f(x)$ coincide with the classical subdifferential in convex analysis [34, 8.12 Proposition].

Example 1 (Bagirov, Karmitsa, and Mäkelä [35, Theorem 3.23]) From a result on the pointwise max-function in convex analysis, we have

$$\partial \max(x) = \text{conv}\{e_i \mid i \in \mathcal{I}(x)\},$$

where e_i 's are the standard bases of \mathbb{R}^n and $\mathcal{I}(x) = \{i \mid x_i = \max(x)\}$.

Next, we extend our discussion to include generalized subdifferentials of a nonsmooth function on submanifolds \mathcal{M} . The *Riemannian Fréchet subdifferential* and the *Riemannian limiting subdifferential of f at $x \in \mathcal{M}$* (see, e.g., [30, Definition 3.1]) are defined as

$$\begin{aligned} \hat{\partial}_{\mathcal{R}}f(x) &:= \{ \text{grad}h(x) \mid \exists \delta > 0 \text{ such that } h \in C^1(\mathcal{B}(x, \delta)) \text{ and } f - h \text{ attains a local minimum at } x \text{ on } \mathcal{M} \}, \\ \partial_{\mathcal{R}}f(x) &:= \{ \lim_{\ell \rightarrow \infty} v^\ell \mid v^\ell \in \hat{\partial}_{\mathcal{R}}f(x^\ell), (x^\ell, f(x^\ell)) \rightarrow (x, f(x)) \}. \end{aligned}$$

If $\mathcal{M} = \mathbb{R}^n$, the above definitions coincide with the usual Fréchet and limiting subdifferentials in \mathbb{R}^n . Moreover, it follows directly that, for all $x \in \mathcal{M}$, one has $\hat{\partial}_{\mathcal{R}}f(x) \subseteq \partial_{\mathcal{R}}f(x)$. According to [30, Proposition 3.2], if x is a local minimizer of f on \mathcal{M} , then $0 \in \hat{\partial}_{\mathcal{R}}f(x)$. Thus, we call a point $x \in \mathcal{M}$ a *Riemannian limiting stationary point of NROP* if

$$0 \in \partial_{\mathcal{R}}f(x). \tag{7}$$

In this paper, we will treat it as a necessary condition for a local solution of NROP to exist.

The smoothing function is the most important tool of the smoothing method.

Definition 3.1 (Zhang and Chen [36, Definition 3.1]) A function $\tilde{f}(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}_{++} \rightarrow \mathbb{R}$ is called a *smoothing function* of $f : \mathbb{R}^n \rightarrow \mathbb{R}$, if $\tilde{f}(\cdot, \mu)$ is continuously differentiable in \mathbb{R}^n for any $\mu > 0$,

$$\lim_{z \rightarrow x, \mu \downarrow 0} \tilde{f}(z, \mu) = f(x)$$

and there exist a constant $\kappa > 0$ and a function $\omega : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ such that

$$|\tilde{f}(x, \mu) - f(x)| \leq \kappa \omega(\mu) \quad \text{with} \quad \lim_{\mu \downarrow 0} \omega(\mu) = 0. \tag{8}$$

Example 2 (Chen, Wets, and Zhang [37, Lemma 4.4]) The LogSumExp function, $\text{lse}(x, \mu) : \mathbb{R}^n \times \mathbb{R}_{++} \rightarrow \mathbb{R}$, given by

$$\text{lse}(x, \mu) := \mu \log(\sum_{i=1}^n \exp(x_i/\mu)),$$

is the smoothing function of $\max(x)$ because we can see that

(i) $\text{lse}(\cdot, \mu)$ is smooth on \mathbb{R}^n for any $\mu > 0$. Its gradient $\nabla_x \text{lse}(x, \mu)$ is given by $\sigma(\cdot, \mu) : \mathbb{R}^n \rightarrow \Delta^{n-1}$,

$$\nabla_x \text{lse}(x, \mu) = \sigma(x, \mu) := \frac{1}{\sum_{\ell=1}^n \exp(x_{\ell}/\mu)} [\exp(x_1/\mu), \dots, \exp(x_n/\mu)]^T, \tag{9}$$

where $\Delta^{n-1} := \{x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 1, x_i \geq 0\}$ is the unit simplex.

(ii) For all $x \in \mathbb{R}^n$ and $\mu > 0$, we have $\max(x) < \text{lse}(x, \mu) \leq \max(x) + \mu \log(n)$. Then, the constant $\kappa = \log(n)$ and $\omega(\mu) = \mu$. The above inequalities imply that $\lim_{z \rightarrow x, \mu \downarrow 0} \text{lse}(z, \mu) = \max(x)$.

Gradient sub-consistency or consistency is crucial to showing that any limit point of the Riemannian smoothing method is also a limiting stationary point of NROP.

Definition 3.2 (Zhang, Chen and Ma [30, Definition 3.4 & 3.9]) A smoothing function \tilde{f} of f is said to satisfy *gradient sub-consistency on \mathbb{R}^n* if, for any $x \in \mathbb{R}^n$,

$$G_{\tilde{f}}(x) \subseteq \partial f(x), \tag{10}$$

where the *subdifferential of f associated with \tilde{f} at $x \in \mathbb{R}^n$* is given by

$$G_{\tilde{f}}(x) := \{u \in \mathbb{R}^n \mid \nabla_x \tilde{f}(z_k, \mu_k) \rightarrow u \text{ for some } z_k \rightarrow x, \mu_k \downarrow 0\}.$$

Similarly, \tilde{f} is said to satisfy *Riemannian gradient sub-consistency on \mathcal{M}* if, for any $x \in \mathcal{M}$,

$$G_{\tilde{f}, \mathcal{R}}(x) \subseteq \partial_{\mathcal{R}} f(x), \tag{11}$$

where the *Riemannian subdifferential of f associated with \tilde{f} at $x \in \mathcal{M}$* is given by

$$G_{\tilde{f}, \mathcal{R}}(x) = \{v \in \mathbb{R}^n \mid \text{grad} \tilde{f}(z_k, \mu_k) \rightarrow v \text{ for some } z_k \in \mathcal{M}, z_k \rightarrow x, \mu_k \downarrow 0\}.$$

If one substitutes the inclusion with equality in (10), then \tilde{f} satisfies *gradient consistency* on \mathbb{R}^n , and similarly in (11) for \mathcal{M} . Thanks to the following useful proposition from [30], we can induce gradient sub-consistency on \mathcal{M} from that on \mathbb{R}^n if f is locally Lipschitz.

Proposition 3.3 (Zhang, Chen and Ma [30, Proposition 3.10]) *Let f be a locally Lipschitz function and \tilde{f} a smoothing function of f . For \tilde{f} , if gradient sub-consistency holds on \mathbb{R}^n , then Riemannian gradient sub-consistency holds on \mathcal{M} as well.*

The next example illustrates Riemannian gradient sub-consistency on \mathcal{M} for $\text{lse}(x, \mu)$ in Example 2, since any convex function is locally Lipschitz continuous.

Example 3 (Chen, Wets, and Zhang [37, Lemma 4.4]) The smoothing function $\text{lse}(x, \mu)$ of $\max(x)$ satisfies gradient consistency on \mathbb{R}^n . That is, for any $x \in \mathbb{R}^n$,

$$\partial \max(x) = G_{\text{lse}}(x) = \left\{ \lim_{x^k \rightarrow x, \mu_k \downarrow 0} \sigma(x^k, \mu_k) \right\}.$$

Note that the original assertion of [37, Lemma 4.4] is gradient consistency in the Clarke sense, i.e., $\partial^\circ \max(x) = G_{\text{lse}}(x)$.

3.3 Riemannian smoothing method

Motivated by the previous papers [30–32] on smoothing methods and Riemannian manifolds, we propose a general Riemannian smoothing method. Algorithm 1 is the basic framework of this general method.

Algorithm 1: Basic Riemannian smoothing method for (NROP)

Initialization: Given $\theta \in (0, 1)$, $\mu_0 > 0$, and $x_{-1} \in \mathcal{M}$, select a smoothing function \tilde{f} and a Riemannian algorithm (called sub-algorithm here) for (SROP).

for $k = 0, 1, 2, \dots$

Solve

$$x_k = \arg \min_{x \in \mathcal{M}} \tilde{f}(x, \mu_k)$$

approximately by using the chosen sub-algorithm, starting at x^{k-1}

if final convergence test is satisfied

stop with approximate solution x_k

end if

 Set $\mu_{k+1} = \theta\mu_k$

end for

Now let us describe the convergence properties of the basic method. First, let us assume that the function $\tilde{f}(x, \mu_k)$ has a minimizer on \mathcal{M} for each value of μ_k .

Theorem 3.4 *Suppose that each x^k is an exact global minimizer of (12) in Algorithm 1. Then every limit point x^* of the sequence $\{x^k\}$ is a global minimizer of the problem NROP.*

Proof Let \bar{x} be a global solution of NROP, that is,

$$f(\bar{x}) \leq f(x) \quad \text{for all } x \in \mathcal{M}.$$

From the Definition 3.1 of the smoothing function, there exist a constant $\kappa > 0$ and a function $\omega : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ such that, for all $x \in \mathcal{M}$,

$$-\kappa\omega(\mu) \leq \tilde{f}(x, \mu) - f(x) \leq \kappa\omega(\mu) \tag{13}$$

with $\lim_{\mu \downarrow 0} \omega(\mu) = 0$. Substituting x^k and combining with the global solution \bar{x} , we have that

$$\tilde{f}(x^k, \mu_k) \geq f(x^k) - \kappa\omega(\mu_k) \geq f(\bar{x}) - \kappa\omega(\mu_k).$$

By rearranging this expression, we obtain

$$-\kappa\omega(\mu_k) \leq \tilde{f}(x^k, \mu_k) - f(\bar{x}). \tag{14}$$

Since x^k minimizes $\tilde{f}(x, \mu_k)$ on \mathcal{M} for each μ_k , we have that $\tilde{f}(x^k, \mu_k) \leq \tilde{f}(\bar{x}, \mu_k)$, which leads to

$$\tilde{f}(x^k, \mu_k) - f(\bar{x}) \leq \tilde{f}(\bar{x}, \mu_k) - f(\bar{x}) \leq \kappa\omega(\mu_k). \tag{15}$$

The second inequality above follows from (13). Combining (14) and (15), we obtain

$$|\tilde{f}(x^k, \mu_k) - f(\bar{x})| \leq \kappa\omega(\mu_k). \tag{16}$$

Now, suppose that x^* is a limit point of $\{x^k\}$, so that there is an infinite subsequence \mathcal{K} such that $\lim_{k \in \mathcal{K}} x^k = x^*$. Note that $x^* \in \mathcal{M}$ because \mathcal{M} is complete. By taking the limit as $k \rightarrow \infty, k \in \mathcal{K}$, on both sides of (16), again by the definition of the smoothing function, we obtain

$$|f(x^*) - f(\bar{x})| = \lim_{k \in \mathcal{K}} |\tilde{f}(x^k, \mu_k) - f(\bar{x})| \leq \lim_{k \in \mathcal{K}} \kappa \omega(\mu_k) = 0.$$

Thus, it follows that $f(x^*) = f(\bar{x})$. Since $x^* \in \mathcal{M}$ is a point whose objective value is equal to that of the global solution \bar{x} , we conclude that x^* , too, is a global solution. □

This strong result requires us to find a global minimizer of each subproblem, which, however, cannot always be done. The next result concerns the convergence properties of the sequence $\tilde{f}(x^k, \mu_k)$ under the condition that \tilde{f} has the following additional property:

$$0 < \mu_2 < \mu_1 \implies \tilde{f}(x, \mu_2) < \tilde{f}(x, \mu_1) \text{ for all } x \in \mathbb{R}^n. \tag{17}$$

Example 4 The above property holds for $\text{lse}(x, \mu)$ in Example 2; i.e., we have $\text{lse}(x, \mu_2) < \text{lse}(x, \mu_1)$ on \mathbb{R}^n , provided that $0 < \mu_2 < \mu_1$.

Note that under the equality,

$$\sum_{i=1}^n \exp(x_i/\mu) = \exp\{\text{lse}(x, \mu)/\mu\},$$

the i -th component of $\sigma(x, \mu)$ can be rewritten as

$$\sigma_i(x, \mu) = \exp\{(x_i - \text{lse}(x, \mu))/\mu\}.$$

For any fixed $x \in \mathbb{R}^n$, consider the derivative of a real function $\mu \rightarrow \text{lse}(x, \cdot) : \mathbb{R}_{++} \rightarrow \mathbb{R}$. Then we have

$$\begin{aligned} \nabla_{\mu} \text{lse}(x, \mu) &= \text{lse}/\mu - \frac{\sum_{i=1}^n x_i \exp(x_i/\mu)}{\mu \exp(\text{lse}/\mu)} = (\text{lse} - \sum_{i=1}^n x_i \exp\{(x_i - \text{lse})/\mu\})/\mu \\ &= (\text{lse} - \sum_{i=1}^n x_i \sigma_i)/\mu \leq 0, \end{aligned}$$

where “lse, σ ” are shorthand for $\text{lse}(x, \mu)$ and $\sigma(x, \mu)$. For the last inequality above, we observe that $\sigma \in \Delta^{n-1}$; hence, the term $\sum_{i=1}^n x_i \sigma_i$ is a convex combination of all entries of x , which implies that $\sum_{i=1}^n x_i \sigma_i \leq \max(x) < \text{lse}$. This completes the proofs of our claims.

In [32], the authors considered a special case of Algorithm 1, wherein the smoothing function $\tilde{f}(x, \mu) = \sqrt{\mu^2 + x^2}$ of $|x|$ also satisfies (17) and a Riemannian conjugate gradient method is used for (12).

Theorem 3.5 *Suppose that $f^* := \inf_{x \in \mathcal{M}} f(x)$ exists and the smoothing function \tilde{f} has property (17). Let $f^k := \tilde{f}(x^k, \mu_k)$. Then the sequence $\{f^k\}$ generated by Algorithm 1 is strictly decreasing and bounded below by f^* ; hence,*

$$\lim_{k \rightarrow \infty} |f^k - f^{k-1}| = 0.$$

Proof For each $k \geq 1$, x_k is obtained by approximately solving

$$\min_{x \in \mathcal{M}} \tilde{f}(x, \mu_k),$$

starting at x^{k-1} . Then at least, we have

$$\tilde{f}(x^{k-1}, \mu_k) \geq \tilde{f}(x^k, \mu_k) = f^k.$$

Since $\mu_k = \theta \mu_{k-1} < \mu_{k-1}$, property (17) ensures

$$f^{k-1} = \tilde{f}(x^{k-1}, \mu_{k-1}) > \tilde{f}(x^{k-1}, \mu_k).$$

The claim that sequence $\{f^k\}$ is strictly decreasing follows from these two inequalities.

Suppose that, for all $\mu > 0$ and for all $x \in \mathbb{R}^n$,

$$\tilde{f}(x, \mu) \geq f(x). \tag{18}$$

Then for each k ,

$$f^k = \tilde{f}(x^k, \mu_k) \geq f(x^k) \geq \inf_{x \in \mathcal{M}} f(x) = f^*,$$

which proves our claims.

Now, we show (18) is true if the smoothing function has property (17). Fix any $x \in \mathbb{R}^n$; (17) implies that $\tilde{f}(x, \cdot)$ is strictly decreasing as $\mu \rightarrow 0$. Actually, $\tilde{f}(x, \cdot)$ is monotonically increasing on $\mu > 0$. On the other hand, from the definition of the smoothing function, we have that

$$\lim_{\mu \downarrow 0} \tilde{f}(x, \mu) = f(x).$$

Hence, we have $\inf_{\mu > 0} \tilde{f}(x, \mu) = f(x)$, as claimed. □

Note that the above weak result does not ensure that $\{f^k\} \rightarrow f^*$. Next, for better convergence (compared with Theorem 3.5) and an effortless implementation (compared with Theorem 3.4), we propose an enhanced Riemannian smoothing method: Algorithm 2. This is closer to the version in [30], where the authors use the Riemannian steepest descent method for solving the smoothed problem (19).

Algorithm 2: Enhanced Riemannian smoothing method for (NROP)

Initialization: Given $\theta \in (0, 1)$, $\mu_0 > 0$ and a nonnegative sequence $\{\delta_k\}$ with $\delta_k \rightarrow 0$, and $x_{-1} \in \mathcal{M}$, select a smoothing function \tilde{f} and a Riemannian algorithm (called sub-algorithm here) for (SROP).

for $k = 0, 1, 2, \dots$

Solve

$$x_k = \arg \min_{x \in \mathcal{M}} \tilde{f}(x, \mu_k)$$

approximately by using the chosen sub-algorithm, starting at x^{k-1} , such that

$$\|\text{grad } \tilde{f}(x^k, \mu_k)\| < \delta_k$$

if final convergence test is satisfied

stop with approximate solution x_k

end if

Set $\mu_{k+1} = \theta\mu_k$

end for

The following result is adapted from [30, Proposition 4.2 & Theorem 4.3]. Readers are encouraged to refer to [30] for a discussion on the stationary point associated with \tilde{f} on \mathcal{M} .

Theorem 3.6 *In Algorithm 2, suppose that the chosen sub-algorithm has the following general convergence property for SROP:*

$$\liminf_{\ell \rightarrow \infty} \|\text{grad}g(x^\ell)\| = 0. \tag{21}$$

Moreover, suppose that, for all μ_k , the function $\tilde{f}(\cdot, \mu_k)$ satisfies the convergence assumptions of the sub-algorithm needed for g above and \tilde{f} satisfies the Riemannian gradient sub-consistency on \mathcal{M} . Then

1. *For each k , there exists an x^k satisfying (20); hence, Algorithm 2 is well defined.*
2. *Every limit point x^* of the sequence $\{x^k\}$ generated by Algorithm 2 is a Riemannian limiting stationary point of NROP (see (7)).*

Proof Fix any μ_k . By (21), we have $\liminf_{\ell \rightarrow \infty} \|\text{grad}\tilde{f}(x^\ell, \mu_k)\| = 0$. Hence, there is a convergent subsequence of $\|\text{grad}\tilde{f}(x^\ell, \mu_k)\|$ whose limit is 0. This means that, for any $\epsilon > 0$, there exists an integer ℓ_ϵ such that $\|\text{grad}\tilde{f}(x^{\ell_\epsilon}, \mu_k)\| < \epsilon$. If $\epsilon = \delta_k$, we get $x^k = x^{\ell_\epsilon}$. Thus, statement (1) holds.

Next, suppose that x^* is a limit point of $\{x^k\}$ generated by Algorithm 2, so that there is an infinite subsequence \mathcal{K} such that $\lim_{k \in \mathcal{K}} x^k = x^*$. From (1), we have

$$\lim_{k \in \mathcal{K}} \|\text{grad}\tilde{f}(x^k, \mu_k)\| \leq \lim_{k \in \mathcal{K}} \delta_k = 0,$$

and we find that $\text{grad}\tilde{f}(x^k, \mu_k) \rightarrow 0$ for $k \in \mathcal{K}$, $x^k \in \mathcal{M}$, $x^k \rightarrow x^*$, $\mu_k \downarrow 0$. Hence,

Table 1 List of smoothing functions of the absolute value function $|x|$ with κ and $\omega(\mu)$ in (8)

	κ	$\omega(\mu)$
$\tilde{f}_1(x, \mu) = \begin{cases} x & \text{if } x > \frac{\mu}{2} \\ \frac{x^2}{\mu} + \frac{\mu}{4} & \text{if } x \leq \frac{\mu}{2} \end{cases}$	$\frac{1}{4}$	μ
$\tilde{f}_2(x, \mu) = \sqrt{\mu^2 + x^2}$	1	μ
$\tilde{f}_3(x, \mu) = 2\mu \log(1 + e^{\frac{x}{\mu}}) - x$	$2 \log(2)$	μ
$\tilde{f}_4(x, \mu) = x \tanh(\frac{x}{\mu})$, where $\tanh(z)$ is the hyperbolic tangent function.	1	μ
$\tilde{f}_5(x, \mu) = x \operatorname{erf}(\frac{x}{\mu})$, where $\operatorname{erf}(z) := \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$ is the Gauss error function.	$\frac{2}{e\sqrt{\pi}}$	μ

$$0 \in G_{\tilde{f}, \mathcal{R}}(x^*) \subseteq \partial_{\mathcal{R}} f(x^*).$$

□

Now let us consider the selection strategy of the nonnegative sequence $\{\delta_k\}$ with $\delta_k \rightarrow 0$. In [30], when $\mu_{k+1} = \theta \mu_k$ shrinks, the authors set

$$\delta_{k+1} := \rho \delta_k \tag{22}$$

with an initial value of δ_0 and constant $\rho \in (0, 1)$. In the spirit of the usual smoothing methods described in [29], one can set

$$\delta_k := \gamma \mu_k \tag{23}$$

with a constant $\gamma > 0$. The latter is an adaptive rule, because μ_k determines subproblem (19) and its stopping criterion at the same time. The merits and drawbacks of the two rules require more discussion, but the latter seems to be more reasonable.

We conclude this section by discussing the connections with [32] and [30]. Our work is based on an efficient unification of them. [32] focused on a specific algorithm and did not discuss the underlying generalities, whereas we studied a general framework for Riemannian smoothing. Recall that the “smoothing function” is the core tool of the smoothing method. In addition to what are required by its definition (see Definition 3.1), it needs to have the following “additional properties” (AP) in order for the algorithms to converge:

(AP1) Approximate from above, i.e., (17). (Needed in Algorithm 1)

(AP2) (Riemannian) gradient sub-consistency, i.e., Definition 3.2. (Needed in Algorithm 2)

We find that not all smoothing functions satisfy (AP1) and for some functions it is hard to prove whether (AP2) holds. For example, all the functions in Table 1 are smoothing functions of $|x|$, but only the first three meet (AP1); the last two do not. In [29], the authors showed that the first one in Table 1, $\tilde{f}_1(x, \mu)$, has property (AP2).

The others remain to be verified, but doing so will not be a trivial exercise. To a certain extent, Algorithm 1 as well as Theorem 3.5 guarantee a fundamental convergence result even if one has difficulty in showing whether one's smoothing function satisfies (AP2). Therefore, it makes sense to consider Algorithms 1 and 2 together for the sake of the completeness of the general framework.

Algorithm 2 expands on the results of [30]. It allows us to use any standard method of SROP, not just steepest descent, to solve the smoothed problem (19). Various standard Riemannian algorithms for SROP, such as the Riemannian conjugate gradient method [38] (which often performs better than Riemannian steepest descent), the Riemannian Newton method [39, Chapter 6], and the Riemannian trust region method [39, Chapter 7], have extended the concepts and techniques used in Euclidean space to Riemannian manifolds. As shown by Theorem 3.6, no matter what kind of sub-algorithm is implemented for (19), it does not affect the final convergence as long as the chosen sub-algorithm has property (21). On the other hand, we advocate that the sub-algorithm should be viewed as a “Black Box” and the user should not have to care about the code details of the sub-algorithm at all. We can directly use an existing solver, e.g., Manopt [40], which includes the standard Riemannian algorithms mentioned above. Hence, we can choose the *most suitable* sub-algorithm for the application and quickly implement it with minimal effort.

4 Numerical experiments on CP factorization

The numerical experiments in Sects. 4 and 5 were performed on a computer equipped with an Intel Core i7-10700 at 2.90GHz with 16GB of RAM using Matlab R2022a. Our Algorithm 2 is implemented in the Manopt framework [40] (version 7.0). The number of iterations to solve the smoothed problem (19) with the sub-algorithm is recorded in the total number of iterations. We refer readers to the supplementary material of this paper for the available codes.¹

In this section, we describe numerical experiments that we conducted on CP factorization in which we solved OptCP using Algorithm 2, where different Riemannian algorithms were employed as sub-algorithms and $\text{lse}(-BX, \mu)$ was used as the smoothing function. To be specific, we used three built-in Riemannian solvers of Manopt 7.0 — steepest descent (SD), conjugate gradient (CG), and trust regions (RTR), denoted by SM_SD, SM_CG and SM_RTR, respectively. We compared our algorithms with the following non-Riemannian numerical algorithms for CP factorization that were mentioned in Sect. 1.1. We followed the settings used by the authors in their papers.

- SpFeasDC_ls [19]: A difference-of-convex functions approach for solving the split feasibility problem, it can be applied to (FeasCP). The implementation

¹ Alternatively, <https://github.com/GALVINLAI/General-Riemannian-Smoothing-Method>.

details regarding the parameters we used are the same as in the numerical experiments reported in [19, Section 6.1].

- RIPG_mod [20]: This is a projected gradient method with relaxation and inertia parameters for solving (2). As shown in [20, Section 4.2], RIPG_mod is the best among the many strategies of choosing parameters.
- APM_mod [18]: A modified alternating projection method for CP factorization; it is described in Sect. 2.3.

We have shown that $\text{lse}(x, \mu)$ is a smoothing function of $\max(x)$ with gradient consistency. $\text{lse}(\cdot, \mu)$ of the matrix argument can be simply derived from entrywise operations. Then from the properties of compositions of smoothing functions [41, Proposition 1 (3)], we have that $\text{lse}(-\bar{B}X, \mu)$ is a smoothing function of $\max(-\bar{B}X)$ with gradient consistency. In practice, it is important to avoid numerical overflow and underflow when evaluating $\text{lse}(x, \mu)$. Overflow occurs when any x_i is large and underflow occurs when all x_i are small. To avoid these problems, we can shift each component x_i by $\max(x)$ and use the following formula:

$$\text{lse}(x, \mu) = \mu \log\left(\sum_{i=1}^n \exp((x_i - \max(x))/\mu)\right) + \max(x),$$

whose validity is easy to show.

The details of the experiments are as follows. If $A \in \mathcal{CP}_n$ was of full rank, for accuracy reasons, we obtained an initial \bar{B} by using Cholesky decomposition. Otherwise, \bar{B} was obtained by using spectral decomposition. Then we extended \bar{B} to r columns by column replication (6). We set $r = \text{cp}(A)$ if $\text{cp}(A)$ was known or r was sufficiently large. We used RandOrthMat.m [42] to generate a random starting point X^0 on the basis of the Gram-Schmidt process.

For our three algorithms, we set $\mu_0 = 100, \theta = 0.8$ and used an adaptive rule (23) of $\delta_k := \gamma \mu_k$ with $\gamma = 0.5$. Except for RIPG_mod, all the algorithms terminated successfully at iteration k , where $\min(\bar{B}X^k) \geq -10^{-15}$ was attained before the maximum number of iterations (5,000) was reached. In addition, SpFeasDC_Is failed when $\bar{L}_k > 10^{10}$. Regarding RIPG_mod, it terminated successfully when $\|A - X_k X_k^T\|^2 / \|A\|^2 < 10^{-15}$ was attained before at most 10,000 iterations for $n < 100$, and before at most 50,000 iterations in all other cases. In the tables of this section, we report the rounded success rate (Rate) over the total number of trials, although the definitions of ‘‘Rate’’ in the different experiments (described in Sections 4.1-4.4) vary slightly from one experiment to the other. We will describe them later.

4.1 Randomly generated instances

We examined the case of randomly generated matrices to see how the methods were affected by the order n or r . The instances were generated in the same way as in [18, Section 7.7]. We computed C by setting $C_{ij} := |B_{ij}|$ for all i, j , where B is a random $n \times 2n$ matrix based on the Matlab command randn, and we took $A = CC^T$ to be factorized. In Table 2, we set $r = 1.5n$ and $r = 3n$ for the values $n \in \{20, 30, 40, 100, 200, 400, 600, 800\}$. For each pair of n and r , we generated 50

Table 2 CP factorization of random completely positive matrices

Methods	SM_SD			SM_CG			SM_RTR			SpFeasDC_Is			RIPG_mod			APM_mod			
	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	
$n(r = 1.5n)$																			
20	1	0.0409	49	1	0.0394	41	1	0.0514	35	1	0.0027	24	1	0.0081	1229	0.32	0.3502	2318	
30	1	0.0549	58	1	0.0477	44	1	0.0690	36	1	0.0075	24	1	0.0231	1481	0.04	1.0075	2467	
40	1	0.0735	65	1	0.0606	46	1	0.0859	37	1	0.0216	46	1	0.0574	1990	0	-	-	
100	1	0.2312	104	1	0.1520	56	1	0.4061	45	1	0.2831	109	1	0.8169	4912	0	-	-	
200	1	1.0723	167	1	0.5485	69	1	1.9855	53	1	2.2504	212	1	5.2908	9616	0	-	-	
400	1	14.6	314	1	4.1453	86	1	22.1	69	1	36.9	636	1	90.6	17987	0	-	-	
600	1	50.6	474	1	14.7	105	1	46.4	80	1	140.1	882	1	344.7	26146	0	-	-	
800	1	133.3	643	1	30.0	109	1	93.5	89	1	413.3	1225	1	891.1	34022	0	-	-	
Methods																			
$n(r = 3n)$																			
20	1	0.0597	52	1	0.0551	42	1	0.0842	37	1	0.0057	15	1	0.0105	1062	0.30	0.7267	2198	
30	1	0.0793	59	1	0.0673	45	1	0.1161	39	1	0.0128	17	1	0.0336	1127	0	-	-	
40	1	0.1035	67	1	0.0882	48	1	0.1961	41	1	0.0256	19	1	0.0822	1460	0	-	-	
100	1	0.5632	103	1	0.3395	57	1	1.4128	50	1	0.8115	86	1	1.1909	4753	0	-	-	
200	1	4.5548	163	1	2.4116	68	1	14.3	65	1	8.1517	184	1	9.2248	9402	0	-	-	
400	1	46.5	296	1	19.2	89	1	77.1	80	1	124.3	453	1	156.6	17563	0	-	-	
600	1	209.0	446	1	65.7	99	1	294.5	76	1	981.8	795	1	616.7	25336	0	-	-	
800	1	609.7	628	1	160.4	114	1	650.7	83	1	4027.4	1070	1	1289.4	26820	0	-	-	

instances if $n \leq 100$ and 10 instances otherwise. For each instance, we initialized all the algorithms at the same random starting point X^0 and initial decomposition \bar{B} , except for RIPG_mod. Note that each instance A was assigned only one starting point.

Table 2 lists the average time in seconds (Time_s) and the average number of iterations (Iter_s) among the successful instances. For our three Riemannian algorithms, Iter_s contains the number of iterations of the sub-algorithm. Table 2 also lists the rounded success rate (Rate) over the total number (50 or 10) of instances for each pair of n and r . Boldface highlights the two best results in each row.

As shown in Table 2, except for APM_mod, each method had a success rate of 1 for all pairs of n and r . Our three algorithms outperformed the other methods on the large-scale matrices with $n \geq 100$. In particular, SM_CG with the conjugate-gradient method gave the best results.

4.2 A specifically structured instance

Let \mathbf{e}_n denote the all-ones vector in \mathbb{R}^n and consider the matrix [18, Example 7.1],

$$A_n = \begin{pmatrix} 0 & \mathbf{e}_{n-1}^\top \\ \mathbf{e}_{n-1} & I_{n-1} \end{pmatrix}^\top \begin{pmatrix} 0 & \mathbf{e}_{n-1}^\top \\ \mathbf{e}_{n-1} & I_{n-1} \end{pmatrix} \in \mathcal{CP}_n.$$

Theorem 2.2 shows that $A_n \in \text{int}(\mathcal{CP}_n)$ for every $n \geq 2$. By construction, it is obvious that $\text{cp}(A_n) = n$. We tried to factorize A_n for the values $n \in \{10, 20, 50, 75, 100, 150\}$ in Table 3. For each A_n , using $r = \text{cp}(A_n) = n$ and the same initial decomposition \bar{B} , we tested all the algorithms on the same 50 randomly generated starting points, except for RIPG_mod. Note that each instance was assigned 50 starting points.

Table 3 lists the average time in seconds (Time_s) and the average number of iterations (Iter_s) among the successful starting points. It also lists the rounded success rate (Rate) over the total number (50) of starting points for each n . Boldface highlights the two best results for each n .

We can see from Table 3 that the success rates of our three algorithms were always 1, whereas the success rates of the other methods decreased as n increased. Likewise, SM_CG with the conjugate-gradient method gave the best results.

4.3 An easy instance on the boundary of \mathcal{CP}_n

Consider the following matrix from [43, Example 2.7]:

$$A = \begin{pmatrix} 41 & 43 & 80 & 56 & 50 \\ 43 & 62 & 89 & 78 & 51 \\ 80 & 89 & 162 & 120 & 93 \\ 56 & 78 & 120 & 104 & 62 \\ 50 & 51 & 93 & 62 & 65 \end{pmatrix}.$$

Table 3 CP factorization of a family of specifically structured instances

Methods	SM_SD			SM_CG			SM_RTR			SpFeasDC_Is			RIPG_mod			APM_mod			
	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	Rate	Time _s	Iter _s	
$n(r = n)$																			
10	1	0.0399	71	1	0.0313	49	1	0.0424	45	1	0.0043	149	1	0.0074	2085	0.80	0.0174	616	
20	1	0.0486	85	1	0.0408	63	1	0.0637	55	0.98	0.0139	201	0.74	0.0212	3478	0.90	0.0591	864	
50	1	0.2599	295	1	0.1073	101	1	0.2104	76	0.98	0.3389	770	0	-	-	0.76	0.6948	1416	
75	1	0.3843	329	1	0.1923	135	1	0.4293	93	0.98	1.0706	1186	0	-	-	0.64	1.4809	1510	
100	1	0.7459	458	1	0.3289	168	1	0.9074	108	0.80	1.6653	1083	0	-	-	0.60	2.8150	1690	
150	1	1.8076	647	1	0.7837	241	1	2.6030	145	0.70	3.7652	1170	0	-	-	0.35	9.9930	2959	

The sufficient condition from [43, Theorem 2.5] ensures that this matrix is completely positive and $\text{cp}(A) = \text{rank}(A) = 3$. Theorem 2.2 tells us that $A \in \text{bd}(\mathcal{CP}_5)$, since $\text{rank}(A) \neq 5$.

We found that all the algorithms could easily factorize this matrix. However, our three algorithms returned a CP factorization B whose smallest entry was as large as possible. In fact, they also maximized the smallest entry in the $n \times r$ symmetric factorization of A , since OptCP is equivalent to

$$\max_{A=XX^T, X \in \mathbb{R}^{n \times r}} \{ \min(X) \}.$$

When we did not terminate as soon as $\min(\bar{B}X^k) \geq -10^{-15}$, for example, after 1000 iterations, our algorithms gave the following CP factorization whose the smallest entry is around $2.8573 \gg -10^{-15}$:

$$A = BB^T, \text{ where } B \approx \begin{pmatrix} 3.5771 & 4.4766 & \mathbf{2.8573} \\ 2.8574 & 3.0682 & 6.6650 \\ 8.3822 & 7.0001 & 6.5374 \\ 5.7515 & 2.8574 & 7.9219 \\ 2.8574 & 6.7741 & 3.3085 \end{pmatrix}.$$

4.4 A hard instance on the boundary of \mathcal{CP}_n

Next, we examined how well these methods worked on a hard matrix on the boundary of \mathcal{CP}_n . Consider the following matrix on the boundary taken from [44]:

$$A = \begin{pmatrix} 8 & 5 & 1 & 1 & 5 \\ 5 & 8 & 5 & 1 & 1 \\ 1 & 5 & 8 & 5 & 1 \\ 1 & 1 & 5 & 8 & 5 \\ 5 & 1 & 1 & 5 & 8 \end{pmatrix} \in \text{bd}(\mathcal{CP}_5).$$

Since $A \in \text{bd}(\mathcal{CP}_5)$ and A is of full rank, it follows from Theorem 2.2 that $\text{cp}^+(A) = \infty$; i.e., there is no strictly positive CP factorization for A . Hence, the global minimum of (OptCP), $t = 0$, is clear. None of the algorithms could decompose this matrix under our tolerance, 10^{-15} , in the stopping criteria. As was done in [18, Example7.3], we investigated slight perturbations of this matrix. Given

$$MM^T =: C \in \text{int}(\mathcal{CP}_5) \text{ with } M = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \end{pmatrix},$$

we factorized $A_\lambda := \lambda A + (1 - \lambda)C$ for different values of $\lambda \in [0, 1)$ using $r = 12 > \text{cp}_5 = 11$. Note that $A_\lambda \in \text{int}(\mathcal{CP}_5)$ provided $0 \leq \lambda < 1$ and A_λ approached the boundary as $\lambda \rightarrow 1$. We chose the largest $\lambda = 0.9999$. For each A_λ , we tested

Table 4 Success rate of CP factorization of A_λ for values of λ from 0.6 to 0.9999

λ	SM_RTR	SM_SD	RIPG_mod	SM_CG	SpFeasDC_ls	APM_mod
0.6	1	1	1	1	1	0.42
0.65	1	1	1	1	1	0.44
0.7	1	1	1	1	1	0.48
0.75	1	1	1	1	1	0.52
0.8	1	1	1	1	0.96	0.46
0.82	1	1	1	1	0.98	0.4
0.84	1	1	1	1	0.86	0.24
0.86	1	1	1	1	0.82	0.1
0.88	1	1	1	1	0.58	0.18
0.9	1	1	1	1	0.48	0.18
0.91	1	1	1	1	0.4	0.14
0.92	1	1	1	1	0.2	0.18
0.93	1	1	0.98	1	0.22	0.22
0.94	1	1	0.98	1	0.1	0.2
0.95	1	1	1	1	0.12	0.32
0.96	1	1	0.96	0.98	0.06	0.34
0.97	1	1	0.86	0.82	0.06	0.14
0.98	1	1	0.76	0.28	0.02	0
0.99	1	0.68	0.42	0	0	0
0.999	1	0	0.14	0	0	0
0.9999	1	0	0	0	0	0

all of the algorithms on 50 randomly generated starting points and computed the success rate over the total number of starting points.

Table 4 shows how the success rate of each algorithm changes as A_λ approaches the boundary. The table sorts the results from left to right according to overall performance. Except for SM_RTR, whose success rate was always 1, the success rates of all the other algorithms significantly decreased as λ increased to 0.9999. Surprisingly, the method of SM_CG, which performed well in the previous examples, seemed unable to handle instances close to the boundary.

5 Further numerical experiments: comparison with [30, 32]

As described at the end of Sect. 3, the algorithms in [30] and [32] are both special cases of our algorithm. In this section, we compare them to show whether it performs better when we use other sub-algorithms or other smoothing functions. We applied Algorithm 2 to two problems: finding a sparse vector (FSV) in a subspace and robust low-rank matrix completion, which are

problems implemented in [30] and [32], respectively. Since they both involve approximations to the ℓ_1 norm, we applied the smoothing functions listed in Table 1.

We used the *six* solvers built into Manopt 7.0, namely, steepest descent; Barzilai-Borwein (i.e., gradient-descent with BB step size); Conjugate gradient; trust regions; BFGS (a limited-memory version); ARC (i.e., adaptive regularization by cubics).

5.1 FSV problem

The FSV problem is to find the sparsest vector in an n -dimensional linear subspace $W \subseteq \mathbb{R}^m$; it has applications in robust subspace recovery, dictionary learning, and many other problems in machine learning and signal processing [45, 46]. Let $Q \in \mathbb{R}^{m \times n}$ denote a matrix whose columns form an orthonormal basis of W : this problem can be formulated as

$$\min_{x \in S^{n-1}} \|Qx\|_0,$$

where $S^{n-1} := \{x \in \mathbb{R}^n \mid \|x\| = 1\}$ is the sphere manifold, and $\|z\|_0$ counts the number of nonzero components of z . Because this discontinuous objective function is unwieldy, in the literature, one instead focuses on solving the ℓ_1 norm relaxation given below:

$$\min_{x \in S^{n-1}} \|Qx\|_1, \tag{24}$$

where $\|z\|_1 := \sum_i |z_i|$ is the ℓ_1 norm of the vector z .

Our synthetic problems of the ℓ_1 minimization model (24) were generated in the same way as in [30]: i.e., we chose $m \in \{4n, 6n, 8n, 10n\}$ for $n = 5$ and $m \in \{6n, 8n, 10n, 12n\}$ for $n = 10$. We defined a sparse vector $e_n := (1, \dots, 1, 0, \dots, 0)^T \in \mathbb{R}^m$, whose first n components are 1 and the remaining components are 0. Let the subspace W be the span of e_n and some $n - 1$ random vectors in \mathbb{R}^m . By `mgson.m` [47], we generated an orthonormal basis of W to form a matrix $Q \in \mathbb{R}^{m \times n}$. With this construction, the minimum value of $\|Qx\|_0$ should be equal to n . We chose the initial points by using the `M.rand()` tool of Manopt 7.0 that returns a random point on the manifold M and set $x_0 = \text{abs}(M.rand())$. The nonnegative initial point seemed to be better in the experiment. Regarding the the settings of our Algorithm 2, we chose the same smoothing function $\tilde{f}_1(x, \mu)$ in Table 1 and the same gradient tolerance strategy (22) as in [30]: $\mu_0 = 1, \theta = 0.5, \delta_0 = 0.1, \rho = 0.5$. We compared the numerical performances when using different sub-algorithms. Note that with the choice of the steepest-descent method, our Algorithm 2 is exactly the same as the one in [30].

For each (n, m) , we generated 50 pairs of random instances and random initial points. We claim that an algorithm successfully terminates if $\|Qx_k\|_0 = n$, where x_k is the k -th iteration. Here, when we count the number of nonzeros of Qx_k , we truncated the entries as

$$(Qx_k)_i = 0 \quad \text{if } |(Qx_k)_i| < \tau, \tag{25}$$

Table 5 Number of successes from 50 pairs of random instances and random initial points for the ℓ_1 minimization model (24) and $n = 5$

(n, m)	τ	Steepest-descent	Barzilai-Borwein	Conjugate-gradient	Trust-regions	BFGS	ARC
(5, 20)	10^{-5}	21	19	0	22	23	23
	10^{-6}	21	19	0	22	23	23
	10^{-7}	21	19	0	22	23	23
	10^{-8}	16	19	0	22	23	23
(5, 30)	10^{-5}	36	42	0	34	36	35
	10^{-6}	36	42	0	34	36	35
	10^{-7}	36	42	0	34	36	35
	10^{-8}	34	42	0	34	36	35
(5, 40)	10^{-5}	44	47	1	44	47	45
	10^{-6}	44	47	0	44	47	45
	10^{-7}	44	47	0	44	47	45
	10^{-8}	43	47	0	44	47	45
(5, 50)	10^{-5}	47	47	2	45	45	45
	10^{-6}	47	47	2	45	45	45
	10^{-7}	47	47	0	45	45	45
	10^{-8}	46	47	0	45	45	45
(5, 20)	10^{-9}	0	19	0	22	23	23
	10^{-10}	0	19	0	22	23	23
	10^{-11}	0	19	0	22	23	19
	10^{-12}	0	18	0	22	22	17
(5, 30)	10^{-9}	8	42	0	34	36	35
	10^{-10}	1	42	0	34	36	35
	10^{-11}	0	42	0	34	36	33
	10^{-12}	0	42	0	34	34	29
(5, 40)	10^{-9}	3	47	0	44	47	45
	10^{-10}	2	47	0	44	47	45
	10^{-11}	1	47	0	44	47	44
	10^{-12}	0	46	0	44	44	36
(5, 50)	10^{-9}	5	47	0	45	45	45
	10^{-10}	2	47	0	45	45	45
	10^{-11}	0	47	0	45	45	45
	10^{-12}	0	47	0	45	45	37

where $\tau > 0$ is a tolerance related to the precision of the solution, taking values from 10^{-5} to 10^{-12} . Tables 5 and 6 report the number of successful cases out of 50 cases. Boldface highlights the best result for each row.

As shown in Tables 5 and 6, surprisingly, the conjugate-gradient method, which performed best on the CP factorization problem in Sect. 4, performed

Table 6 Number of successes from 50 pairs of random instances and random initial points for the ℓ_1 minimization model (24) and $n = 10$

(n, m)	τ	Steepest-descent	Barzilai-Borwein	Conjugate-gradient	Trust-regions	BFGS	ARC
(10, 60)	10^{-5}	24	28	0	28	28	25
	10^{-6}	24	28	0	28	28	25
	10^{-7}	24	28	0	28	28	25
	10^{-8}	23	28	0	28	28	25
(10, 80)	10^{-5}	39	37	1	40	39	40
	10^{-6}	39	37	0	40	39	40
	10^{-7}	39	37	0	40	39	40
	10^{-8}	39	37	0	40	39	40
(10, 100)	10^{-5}	45	48	3	45	43	41
	10^{-6}	45	48	0	45	43	41
	10^{-7}	45	48	0	45	43	41
	10^{-8}	45	48	0	45	43	41
(10, 120)	10^{-5}	44	46	1	44	44	43
	10^{-6}	44	46	0	44	44	43
	10^{-7}	44	46	0	44	44	43
	10^{-8}	44	46	0	44	44	43
(10, 60)	10^{-9}	3	28	0	28	28	25
	10^{-10}	0	28	0	28	28	25
	10^{-11}	0	28	0	28	28	22
	10^{-12}	0	28	0	28	27	12
(10, 80)	10^{-9}	5	37	0	40	39	40
	10^{-10}	0	37	0	40	39	40
	10^{-11}	0	37	0	40	39	39
	10^{-12}	0	37	0	40	37	30
(10, 100)	10^{-9}	13	48	0	45	43	41
	10^{-10}	2	48	0	45	43	41
	10^{-11}	0	48	0	45	43	40
	10^{-12}	0	48	0	45	43	37
(10, 120)	10^{-9}	14	46	0	44	44	43
	10^{-10}	0	46	0	44	44	43
	10^{-11}	0	46	0	44	44	43
	10^{-12}	0	46	0	44	43	40

worst on the FSV problem. In fact, it was almost useless. Moreover, although the steepest-descent method employed in [30] was not bad at obtaining low-precision solutions with $\tau \in \{10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}\}$, it had difficulty obtaining high-precision solutions with $\tau \in \{10^{-9}, 10^{-10}, 10^{-11}, 10^{-12}\}$. The remaining four sub-algorithms easily obtained high-precision solutions, with the

Barzilai-Borwein method performing the best in most occasions. Combined with the results in Sect. 4, this shows that in practice, the choice of sub-algorithm in the Riemannian smoothing method (Algorithm 2) is highly problem-dependent. For the other smoothing functions in Table 1, we obtained similar results as in Tables 5 and 6.

5.2 Robust low-rank matrix completion

Low-rank matrix completion consists of recovering a rank- r matrix M of size $m \times n$ from only a fraction of its entries with $r \ll \min(m, n)$. The situation in *robust* low-rank matrix completion is one where only a few observed entries, called *outliers*, are perturbed, i.e.,

$$M = M_0 + S,$$

where M_0 is the unperturbed original data matrix of rank r and S is a sparse matrix. This is a case of adding non-Gaussian noise for which the traditional ℓ_2 minimization model,

$$\min_{X \in \mathcal{M}_r} \|\mathcal{P}_\Omega(X - M)\|_2$$

is not well suited to recovery of M_0 . Here, $\mathcal{M}_r := \{X \in \mathbb{R}^{m \times n} \mid \text{rank}(X) = r\}$ is a fixed rank manifold, Ω denotes the set of indices of observed entries, and $\mathcal{P}_\Omega : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is the projection onto Ω , defined as

$$Z_{ij} \xrightarrow{\mathcal{P}_\Omega} \begin{cases} Z_{ij} & \text{if } (i, j) \in \Omega \\ 0 & \text{if } (i, j) \notin \Omega. \end{cases}$$

In [32], the authors try to solve

$$\min_{X \in \mathcal{M}_r} \|\mathcal{P}_\Omega(X - M)\|_1,$$

because the sparsity-inducing property of the ℓ_1 norm leads one to expect exact recovery when the noise consists of just a few outliers.

In all of the experiments, the problems were generated in the same way as in [32]. In particular, after picking the values of m, n, r , we generated the ground truth $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$ with independent and identically distributed (i.i.d.) Gaussian entries of zero mean and unit variance and $M := UV^\top$. We then sampled $k := \rho r(m + n - r)$ observed entries uniformly at random, where ρ is the oversampling factor. In our experiments, we set $\rho = 5$ throughout. We chose the initial points X_0 by using the rank- r truncated SVD decomposition of $\mathcal{P}_\Omega(M)$.

Regarding the setting of our Algorithm 2, we tested all combinations of the five smoothing functions in Table 1 and six sub-algorithms mentioned before (30 cases in total). We set $\mu_0 = 100$ and chose an aggressive value of $\theta = 0.05$ for reducing μ , as in [32]. The stopping criterion of the loop of the sub-algorithm was set to a maximum of 40 iterations or the gradient tolerance (23), whichever was reached first. We monitored

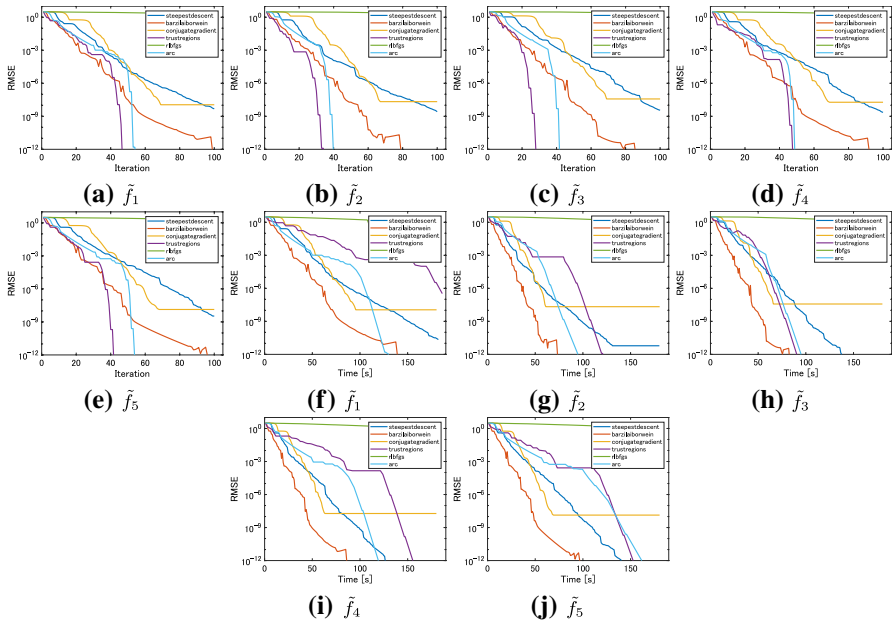


Fig. 1 Perfect low-rank matrix completion of a rank-10 5000×5000 matrix without any outliers using different smoothing functions in Table 1. **a–e** comprise the running iteration comparison; **f–j** comprise the time comparison

the iterations X_k through the root mean square error (RMSE), which is defined as the error on all the entries between X_k and the original low-rank matrix M_0 , i.e.,

$$\text{RMSE}(X_k, M_0) := \sqrt{\frac{\sum_{i=1}^m \sum_{j=1}^n (X_{k,ij} - M_{0,ij})^2}{mn}}.$$

5.2.1 Perfect low-rank matrix completion

As in [32], we first tested all the methods on a simple perfect matrix M (without any outliers) of size 5000×5000 and rank 10. The results are shown in Figure 1. We can see that the choice of smoothing function does not have much effect on numerical performance. In terms of the number of iterations ((a)–(e)), our Algorithm 2 inherits the convergence of its sub-algorithm at least Q-superlinearly when trust regions or ARC are used. But the single iteration cost of trust regions and ARC is high; they are not efficient in terms of time. Specifically, the conjugate-gradient method employed in [32] stagnates at lower precision. Overall, Barzilai-Borwein performed best in terms of time and accuracy.

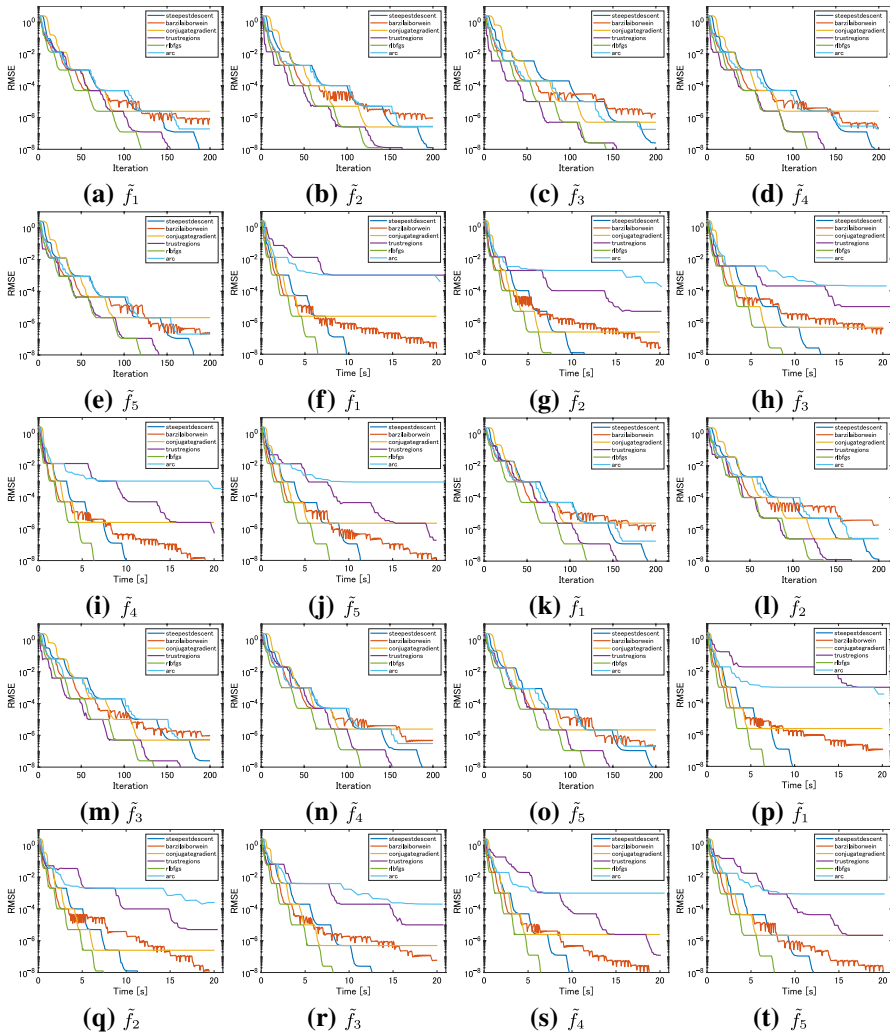


Fig. 2 Low-rank matrix completion with outliers for two rank-10 500×500 matrices by using different smoothing functions in Table 1. **a–j** corresponds to one matrix with outliers created by using $\mu_N = \sigma_N = 0.1$, while **k–t** corresponds to the other with outliers created by using $\mu_N = \sigma_N = 1$. **a–e** and **k–o** comprise the running iteration comparison; **f–j** and **p–t** comprise the time comparison

5.2.2 Low-rank matrix completion with outliers

Given a 500×500 matrix for which we observed the entries uniformly at random with an oversampling ρ of 5, we perturbed 5% of the observed entries by adding noise to them in order to create outliers. The added item was a random variable defined as $\mathcal{O} = \mathcal{S}_{\pm 1} \cdot \mathcal{N}(\mu_N, \sigma_N^2)$ where $\mathcal{S}_{\pm 1}$ is a random variable with equal probability of being equal to $+1$ or -1 , while $\mathcal{N}(\mu_N, \sigma_N^2)$ is a Gaussian random variable of mean μ_N and variance σ_N^2 .

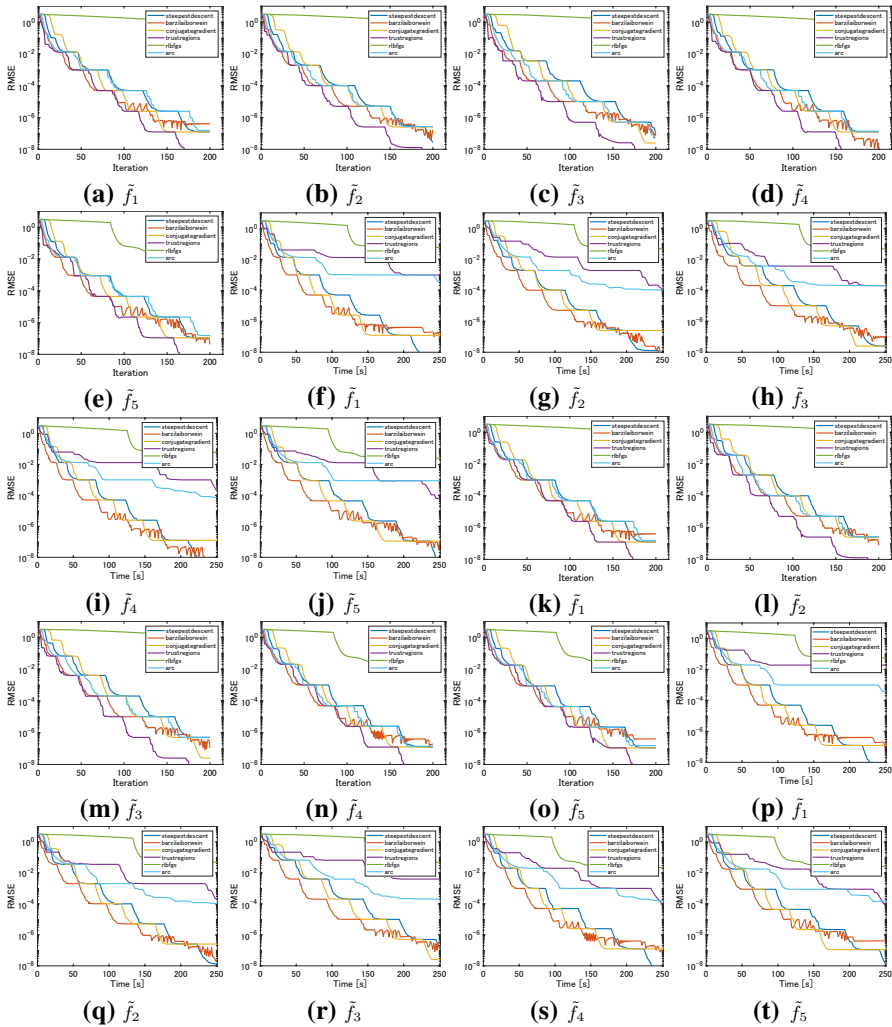


Fig. 3 Low-rank matrix completion with outliers for two rank-10 5000×5000 matrices by using different smoothing functions in Table 1. **a–j** corresponds to one matrix with outliers created by using $\mu_N = \sigma_N = 0.1$, while **k–t** corresponds to the other with outliers created by using $\mu_N = \sigma_N = 1$. **a–e** and **k–o** comprise the running iteration comparison; **f–j** and **p–t** comprise the time comparison

Figure 2 reports the results of two 500×500 instances with outliers generated using $\mu_N = \sigma_N = 0.1$ and $\mu_N = \sigma_N = 1$. Again, we can see that the choice of smoothing function does not have much effect. In most cases, BFGS and trust regions were better than the other methods in terms of number of iterations, and BFGS was the fastest. Furthermore, the conjugate-gradient method employed in [32] still stagnated at solutions with lower precision, approximately 10^{-6} , while steepest descent, BFGS, and trust regions always obtained solutions with at least 10^{-8} precision.

Next, we ran the same experiment on larger 5000×5000 matrices, with 5% outliers. Figure 3 illustrates the results of these experiments, with $\mu_N = \sigma_N = 0.1$ and $\mu_N = \sigma_N = 1$. In most cases, trust regions still outperformed the other methods in terms of number of iterations, while BFGS performed poorly. Barzilai-Borwein and the conjugate-gradient method were almost as good in terms of time.

6 Concluding remarks

We examined the problem of finding a CP factorization of a given completely positive matrix and treated it as a nonsmooth Riemannian optimization problem. To this end, we studied a general framework of Riemannian smoothing for Riemannian optimization. The numerical experiments clarified that our method can compete with other efficient CP factorization methods, in particular on large-scale matrices.

Let us summarize the relation of our approach to the existing CP factorization methods. Groetzner and Dür [18] and Chen et al. [19] proposed different methods to solve (FeasCP). Boş and Nguyen [20] tried to solve another model (2). However, the methods they used do not belong to the Riemannian optimization techniques, but are rather Euclidean ones, since they treated the set $\mathcal{O}(r) := \{X \in \mathbb{R}^{r \times r} : X^T X = I\}$ as a usual constraint in Euclidean space. By comparison, we recognize the existence of manifolds, namely, the Stiefel manifold $\mathcal{M} = \mathcal{O}(r)$, and use optimization techniques specific to them. This change in perspective suggests the possibility of using the rich variety of Riemannian optimization techniques. As the experiments in Sect. 4 show, our Riemannian approach is faster and more reliable than the Euclidean methods.

In the future, we plan to extend Algorithm 2 to the case of general manifolds and, particularly, to quotient manifolds. This application is believed to be possible, although effort should be put into establishing analogous convergence results. In fact, convergence has been verified in a built-in example in Manopt 7.0 [40]: `robust_pca.m` computes a robust version of PCA on data and optimizes a nonsmooth function over a Grassmann manifold. The nonsmooth term consists of the l_2 norm, which is not squared, for robustness. In `robust_pca.m`, Riemannian smoothing with a pseudo-Huber loss function is used in place of the l_2 norm.

As in the other numerical methods, there is no guarantee that Algorithm 2 will find a CP factorization for every $A \in \mathcal{CP}_n$. It follows from Proposition 2.7 that $A \in \mathcal{CP}_n$ if and only if the global minimum of OptCP, say t , is such that $t \leq 0$. Since our methods only converge to a stationary point, Algorithm 2 provides us with a local minimizer at best. We are looking forward to finding a global minimizer of OptCP in our future work.

Supplementary information The online version contains supplementary material available at (<https://doi.org/10.1007/s10589-022-00417-4>).

Acknowledgements This work was supported by the Research Institute for Mathematical Sciences, an International Joint Usage/Research Center, at Kyoto University, JSPS KAKENHI Grant, number (B)19H02373, and JST SPRING Grant, number JPMJSP2124. The authors would like to sincerely thank

the anonymous reviewers and the coordinating editor for their thoughtful and valuable comments which have significantly improved the paper.

Data availability The data that support the findings of this study are available from the corresponding author upon request.

Declarations

Conflict of interest All authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Berman, A., Shaked-Monderer, N.: *Completely Positive Matrices*. World Scientific, Singapore (2003)
2. Dür, M., Rendl, F.: Conic optimization: a survey with special focus on copositive optimization and binary quadratic problems. *EURO J. Comput. Optim.* **9**, 100021 (2021)
3. Bomze, I.M., Dür, M., De Klerk, E., Roos, C., Quist, A.J., Terlaky, T.: On copositive programming and standard quadratic optimization problems. *J. Glob. Optim.* **18**, 301–320 (2000)
4. Burer, S.: On the copositive representation of binary and continuous nonconvex quadratic programs. *Math. Program.* **120**, 479–495 (2009)
5. De Klerk, E., Pasechnik, D.V.: Approximation of the stability number of a graph via copositive programming. *SIAM J. Optim.* **12**, 875–892 (2002)
6. Bomze, I.M.: Copositive optimization—recent developments and applications. *Eur. J. Oper. Res.* **216**, 509–520 (2012)
7. Bomze, I.M., Schachinger, W., Uchida, G.: Think co(mpletely) positive! Matrix properties, examples and a clustered bibliography on copositive optimization. *J. Glob. Optim.* **52**, 423–445 (2012)
8. Burer, S.: A gentle, geometric introduction to copositive optimization. *Math. Program.* **151**, 89–116 (2015)
9. Dür, M.: Copositive programming—a survey. In: Diehl, M., Glineur, F., Jarlebring, E., Michiels, W. (eds.) *Recent advances in optimization and its applications in engineering*, pp. 3–20. Springer, Berlin, Heidelberg (2010)
10. Dickinson, P.J., Gijben, L.: On the computational complexity of membership problems for the completely positive cone and its dual. *Comput. Optim. Appl.* **57**, 403–415 (2014)
11. Berman, A., Dür, M., Shaked-Monderer, N.: Open problems in the theory of completely positive and copositive matrices. *Electron. J. Linear Algebra.* **29**, 46–58 (2015)
12. Jarre, F., Schmallowsky, K.: On the computation of C^* certificates. *J. Glob. Optim.* **45**, 281–296 (2009)
13. Dickinson, P.J., Dür, M.: Linear-time complete positivity detection and decomposition of sparse matrices. *SIAM J. Matrix Anal. Appl.* **33**, 701–720 (2012)
14. Nie, J.: The A -truncated K -moment problem. *Found. Comput. Math.* **14**, 1243–1276 (2014)
15. Sponsel, J., Dür, M.: Factorization and cutting planes for completely positive matrices by copositive projection. *Math. Program.* **143**, 211–229 (2014)
16. Bomze, I.M.: Building a completely positive factorization. *Cent. Eur. J. Oper. Res.* **26**, 287–305 (2018)

17. Doutor Sikirić, M., Schürmann, A., Vallentin, F.: A simplex algorithm for rational cp-factorization. *Math. Program.* **187**, 25–45 (2020)
18. Groetzner, P., Dür, M.: A factorization method for completely positive matrices. *Linear Algebra Appl.* **591**, 1–24 (2020)
19. Chen, C., Pong, T.K., Tan, L., Zeng, L.: A difference-of-convex approach for split feasibility with applications to matrix factorizations and outlier detection. *J. Glob. Optim.* **78**, 107–136 (2020)
20. Boş, R.I., Nguyen, D.-K.: Factorization of completely positive matrices using iterative projected gradient steps. *Numer. Linear Algebra Appl.* **28**, e2391 (2021)
21. Bomze, I.M., Dickinson, P.J., Still, G.: The structure of completely positive matrices according to their cp-rank and cp-plus-rank. *Linear Algebra Appl.* **482**, 191–206 (2015)
22. Dickinson, P.J.: An improved characterization of the interior of the completely positive cone. *Electron. J. Linear Algebra.* **20**, 723–729 (2010)
23. Burer, S., Monteiro, R.D.C.: Local minima and convergence in low-rank semidefinite programming. *Math. Program.* **103**, 427–444 (2005)
24. Horn, R.A., Johnson, C.R.: *Matrix Analysis*. Cambridge University Press, Cambridge (2012)
25. Rudin, W.: *Principles of Mathematical Analysis*. McGraw-Hill, New York (1964)
26. Borckmans, P.B., Selvan, S.E., Boumal, N., Absil, P.-A.: A Riemannian subgradient algorithm for economic dispatch with valve-point effect. *J. Comput. Appl. Math.* **255**, 848–866 (2014)
27. De Carvalho Bento, G., da Cruz Neto, J.X., Oliveira, P.R.: A new approach to the proximal point method: convergence on general Riemannian manifolds. *J. Optim. Theory Appl.* **168**, 743–755 (2016)
28. Kovnatsky, A., Glashoff, K., Bronstein, M.M.: MADMM: a generic algorithm for non-smooth optimization on manifolds. In: *European Conference on Computer Vision*, pp. 680–696. Springer, Cham (2016)
29. Chen, X.: Smoothing methods for nonsmooth, nonconvex minimization. *Math. Program.* **134**, 71–99 (2012)
30. Zhang, C., Chen, X., Ma, S.: A Riemannian smoothing steepest descent method for non-Lipschitz optimization on submanifolds. *arXiv preprint arXiv:2104.04199*. (2021)
31. Liu, C., Boumal, N.: Simple algorithms for optimization on Riemannian manifolds with constraints. *Appl. Math. Optim.* **82**, 949–981 (2020)
32. Cambier, L., Absil, P.-A.: Robust low-rank matrix completion by Riemannian optimization. *SIAM J. Sci. Comput.* **38**, S440–S460 (2016)
33. Boumal, N.: *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press (2022)
34. Rockafellar, R.T., Wets, R.J.B.: *Variational Analysis*. Springer Science & Business Media, Berlin (2009)
35. Bagirov, A., Karmitsa, N., Mäkelä, M.M.: *Introduction to Nonsmooth Optimization: Theory, Practice and Software*. Springer International Publishing, Berlin (2014)
36. Zhang, C., Chen, X.: A smoothing active set method for linearly constrained non-Lipschitz nonconvex optimization. *SIAM J. Optim.* **30**, 1–30 (2020)
37. Chen, X., Wets, R.J.B., Zhang, Y.: Stochastic variational inequalities: residual minimization smoothing sample average approximations. *SIAM J. Optim.* **22**, 649–673 (2012)
38. Sato, H., Iwai, T.: A new, globally convergent Riemannian conjugate gradient method. *Optim.* **64**, 1011–1031 (2015)
39. Absil, P.-A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton (2009)
40. Boumal, N., Mishra, B., Absil, P.-A., Sepulchre, R.: Manopt, a Matlab toolbox for optimization on manifolds. *J. Mach. Learn. Res.* **15**, 1455–1459 (2014)
41. Bian, W., Chen, X.: Neural network for nonsmooth, nonconvex constrained minimization via smooth approximation. *IEEE Trans. Neural Netw. Learn. Syst.* **25**, 545–556 (2013)
42. Shilon O.: Randorthmat. MATLAB Central File Exchange. <https://www.mathworks.com/matlabcentral/fileexchange/11783-randorthmat> (2022). Accessed 7 April 2022
43. So, W., Xu, C.: A simple sufficient condition for complete positivity. *Oper. Matrices.* **9**, 233–239 (2015)
44. Dür, M., Still, G.: Interior points of the completely positive cone. *Electron. J. Linear Algebra.* **17**, 48–53 (2008)

45. Qu, Q., Sun, J., Wright, J.: Finding a sparse vector in a subspace: Linear sparsity using alternating directions. *Adv. Neural. Inf. Process. Syst.* **27** (2014)
46. Qu, Q., Zhu, Z., Li, X., Tsakiris, M.C., Wright, J., Vidal, R.: Finding the sparsest vectors in a subspace: Theory, algorithms, and applications. arXiv preprint [arXiv:2001.06970](https://arxiv.org/abs/2001.06970). (2020)
47. Chen, M.: Gram-Schmidt orthogonalization. MATLAB Central File Exchange. <https://www.mathworks.com/matlabcentral/fileexchange/55881-gram-schmidt-orthogonalization> (2022). Accessed 24 June 2022

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.