

A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints

Nicolas Boumal

Inria & D.I., UMR 8548, Ecole Normale Supérieure, Paris, France

`nicolas.boumal@inria.fr`

June 2, 2015

Abstract

We propose a new algorithm to solve optimization problems of the form $\min f(X)$ for a smooth function f under the constraints that X is positive semidefinite and the diagonal blocks of X are small identity matrices. Such problems often arise as the result of relaxing a rank constraint (lifting). In particular, many estimation tasks involving phases, rotations, orthonormal bases or permutations fit in this framework, and so do certain relaxations of combinatorial problems such as Max-Cut. The proposed algorithm exploits the facts that (1) such formulations admit low-rank solutions, and (2) their rank-restricted versions are smooth optimization problems on a Riemannian manifold. Combining insights from both the Riemannian and the convex geometries of the problem, we characterize when second-order critical points of the smooth problem reveal KKT points of the semidefinite problem. We compare against state of the art, mature software and find that, on certain interesting problem instances, what we call the staircase method is orders of magnitude faster, is more accurate and scales better. Code is available.

1 Introduction

This paper considers the generic problem of estimating matrices $Y_1, \dots, Y_m \in \mathbb{R}^{d \times p}$, $p \geq d$, with orthonormal rows, that is, such that $Y_i Y_i^\top = I_d$ (identity of size d) for all i . We further focus on problems where only *relative* information is available, that is, information about $Y_i Y_j^\top$ for some of the pairs (i, j) is available, but there is no information about individual Y_i 's. As will be detailed below, particular cases of this come up in a number of applications. For example, when $p = d = 1$, the variables Y_i reduce to $\{\pm 1\}$ and the products $Y_i Y_j^\top$ indicate whether Y_i and Y_j have the same sign or not, allowing to model certain combinatorial problems. When $d = 1, p > 1$, the variables reduce to unit-norm vectors, and the products correspond to inner products between them. Finally, when $d = p > 1$, the matrices are orthogonal, and the products $Y_i Y_j^\top = Y_i Y_j^{-1}$ represent relative orthogonal transformations, such as rotations, reflections and permutations.

For ease of notation, we stack the orthonormal matrices on top of each other to form $Y \in \mathbb{R}^{n \times p}$, $n = md \geq p$. Then, $X = YY^\top$ is a block matrix whose block $X_{ij} \in \mathbb{R}^{d \times d}$ corresponds to the relative product $Y_i Y_j^\top$. Define the (transposed) *Stiefel manifold* as

$$\text{St}(d, p) = \{Z \in \mathbb{R}^{d \times p} : ZZ^\top = I_d\},$$

and the set of Y 's obtained by stacking as

$$\begin{aligned} \text{St}(d, p)^m &= \left\{ Y \in \mathbb{R}^{md \times p} : Y^\top = \begin{pmatrix} Y_1^\top & Y_2^\top & \dots & Y_m^\top \end{pmatrix} \text{ and } Y_1, \dots, Y_m \in \text{St}(d, p) \right\} \\ &= \left\{ Y \in \mathbb{R}^{n \times p} : (YY^\top)_{ii} = I_d \text{ for } i = 1 \dots m \right\}. \end{aligned} \quad (1)$$

This paper is concerned with solving optimization problems of the form

$$\min_{Y \in \mathbb{R}^{n \times p}} g(Y) = f(YY^\top), \quad \text{subject to } Y \in \text{St}(d, p)^m, \quad (\text{RP}_p)$$

with twice continuously differentiable cost $f: \mathbb{S}^{n \times n} \rightarrow \mathbb{R}$ defined over the symmetric matrices. Here, g encodes for example the negative likelihood of Y with respect to available data. The restriction that $g(Y)$ be only a function of YY^\top encodes the property that only relative information is available, through $(YY^\top)_{ij} = Y_i Y_j^\top$. This induces invariance of the cost under right-action of the orthogonal group. Indeed, $g(YQ) = g(Y)$ for any orthogonal matrix Q of size p . This also implies that solutions of (RP_p) are only defined up to this group action.

Problem (RP_p) is computationally hard. In particular, for $d = p = 1$ and linear f , it covers the NP-hard Max-Cut problem [37]. Following that and other previous work [13, 67, 9], we consider a relaxation through the following observation. For all $Y \in \text{St}(d, p)^m$, the matrix $X = YY^\top$ is positive semidefinite, its diagonal blocks X_{ii} are identity matrices I_d , and it has rank at most p . Conversely, any matrix X with those properties can be factored as YY^\top with $Y \in \text{St}(d, p)^m$. In other words, problem (RP_p) is equivalent to optimizing f over the convex set

$$\mathcal{C} = \{X \in \mathbb{S}^{n \times n} : X \succeq 0 \text{ and } X_{ii} = I_d \text{ for } i = 1 \dots m\}, \quad (2)$$

with the additional constraint $\text{rank}(X) \leq p$. As often, the rank constraint is the culprit. Indeed, continuing with the Max-Cut example (linear f), optimization over \mathcal{C} without the rank

constraint is a *semidefinite program*, which can be solved to arbitrary precision in polynomial time [77].

This is motivation to study the relaxation obtained by ignoring the rank constraint (for linear f , it is also the dual of the dual of (RP_p)):

$$\min_{X \in \mathbb{R}^{n \times n}} f(X), \quad \text{subject to } X \in \mathcal{C}. \quad (\text{P})$$

The optimal cost of (P) is a lowerbound on that of (RP_p) . Furthermore, if (P) admits a solution $X = YY^\top$ with $Y \in \mathbb{R}^{n \times p}$, that is, a solution of rank at most p , then Y is a solution of (RP_p) . When that is not the case, a higher-rank solution X may still be projected to a (hopefully good) initial guess for the nonconvex problem (RP_p) . See [52, 9] for a discussion of approximation results related to these projections. The price to pay is that \mathcal{C} is much higher dimensional than $\text{St}(d, p)^m$: this is called a *lift* [13].

For linear f , solutions of (P) can of course be computed using standard SDP solvers, such as interior point methods (IPM). Unfortunately, as demonstrated in Section 5, IPM’s do not scale well. The main reason for it is that, as the name suggests, IPM’s iterate inside the interior of the search space \mathcal{C} . The latter is formed by full-rank, dense matrices of size n : this quickly becomes unmanageable.

The full-rank operations seem even more wasteful considering that, still for linear f , problem (P) always admits a solution of rank at most

$$p^* = \frac{\sqrt{1 + 4md(d+1)} - 1}{2} < (d+1)\sqrt{m} \ll n. \quad (3)$$

Indeed, this follows a general result of Barvinok [12] and Pataki [56] regarding extreme points of *spectrahedra*¹, that is, intersections of the positive semidefinite cone with an affine subspace—the geometry of \mathcal{C} is discussed in Section 2.2. This prompted Burer and Monteiro [19, 20] to propose SDPLR, a generic SDP solver which exploits the low-rank phenomenon. Applying SDPLR to our problem amounts to computing a local minimizer Y of (RP_p) for some small p , using classical nonlinear optimization algorithms and penalizing for the constraints in a Lagrangian way. Then, p is increased as needed until YY^\top can be certified as a solution to the SDP.

SDPLR is powerful and generic, and the theory accompanying the algorithm brings great insight into the problem. But it also has some downsides we want to improve on in the context of (P) . First, it is not an easy matter to guarantee convergence to (even local) optimizers in the nonlinear subproblems. Furthermore, since constraints are enforced by penalization, they are not accurately satisfied by the returned solution. Finally, we would like to allow for nonlinear f . Nevertheless, Section 5 shows SDPLR improves significantly upon IPM’s.

Journée et al. [42] build upon SDPLR, observing that certain SDP’s harbor an elegant Riemannian geometry that can be put to good algorithmic use. In particular, they cover what here corresponds to the case $d = 1$ and observe that, as remains true for $d > 1$, (RP_p) is an optimization problem on a smooth space: $\text{St}(d, p)^m$ is a *Riemannian manifold*—this geometry is detailed in Section 2.1. Allowing for smooth nonlinear f , they apply essentially the SDPLR machinery, replacing the nonlinear programming algorithms for (RP_p) by Riemannian

¹The name spectrahedron for the search space of a semidefinite program echoes the name polyhedron for the search space of a linear program.

optimization algorithms [4]. These algorithms exploit the smooth structure of the nonconvex search space, resulting in constraint satisfaction up to numerical accuracy, as well as notable speedups.

As a further refinement, Journée et al. [42] address the invariance of f under orthogonal group action. Instead of optimizing $f(YY^\top)$ over $\text{St}(d, p)^m$, they optimize over the quotient space $\text{St}(d, p)_*^m / \sim$, where \sim is an equivalence relation defined over $\text{St}(d, p)_*^m$ (the full-rank elements of $\text{St}(d, p)^m$) by $Y \sim \tilde{Y} \Leftrightarrow YY^\top = \tilde{Y}\tilde{Y}^\top$. The advantage is that this quotient space, which is still a smooth Riemannian manifold, is now one-to-one with the rank- p matrices in \mathcal{C} . Unfortunately, the geometry breaks down at rank-deficient Y 's (to see this, notice that equivalence classes of different rank have different dimension; see also Figure 1). The breakdown is problematic since, as will become clear, it is desirable to converge to rank-deficient Y 's. Furthermore, that paper too asks for computation of local optimizers of the subproblems, which, on Riemannian manifolds too, is a difficult task.

In both [20] and [42], one of the keys to practical efficiency is (well-justified) optimism: (RP_p) is first solved for small values of p , and p is increased only as needed. In both papers, it is observed that, in practice, it often suffices to reach p just above the rank of the target solution of (P), which may be quite small; but there is no theory to confirm this. We do not prove such a strong result either, but we give some nontrivial bounds on “how high one must lift”, refining some results of [20].

1.1 Contribution

In this paper, we describe the Riemannian geometry of $\text{St}(d, p)^m$ in order to frame (RP_p) as a Riemannian optimization problem. We use existing algorithms [4] and the Manopt toolbox [17] to compute critical points of (RP_p) , that is, points where the (Riemannian) gradient of the cost function vanishes. In practice, those algorithms tend to converge to second-order critical points, that is, points where the (Riemannian) Hessian is also positive semidefinite, because all other critical points are unstable fixed points of the iteration.

For $p > d$, $\text{St}(d, p)^m$ is a connected², compact and smooth space. Since we further assume sufficient smoothness in f too, this makes for a nice problem with no delicate limit cases to handle. Furthermore, Riemannian optimization algorithms iterate on the manifold directly: all iterates satisfy constraints up to numerical accuracy.

We then turn our attention to computing Karush-Kuhn-Tucker (KKT) points for (P). These are points that satisfy first-order necessary optimality conditions. If f is convex, the conditions are also sufficient. Our goal is to compute KKT points via the computation of second-order critical points of (RP_p) , which is lower-dimensional. A key property that makes this possible is the availability of an explicit dual matrix $S(X)$ (21) which intervenes in both sets of conditions.

Using this dual matrix, we show that rank-deficient second-order critical points Y reveal KKT points $X = YY^\top$. Furthermore, when a computed second-order critical point is full rank, it is shown how to use it as a warm-start for the computation of a second-order critical point of (RP_p) with a larger value of p . It is guaranteed that if p is allowed to grow up to n , then all second-order critical points reveal KKT points, so that the procedure terminates.

²For $p = d$, $\text{St}(d, p)^m$ has 2^m disconnected components, because the orthogonal group has two components: matrices with determinant $+1$ and -1 . This is a strong incentive to relax at least to $p = d + 1$.

This is formalized in Algorithm 1, which we call the *Riemannian Staircase*, as it lifts (RP_p) to (P) step by step, instead of all at once.

The above points rest extensively on work discussed earlier in this introduction [19, 20, 42], and improve upon those along the lines announced in the same. In particular, we do not require the computation of local optimizers of (RP_p) , and we avoid the geometry breakdown tied to the quotient approach in [42]. We also stress that the latter reference only covers $d = 1$, and SDPLR only covers linear f .

We further take particular interest in understanding how large p may grow in the staircase algorithm. We view this part as our principal theoretical contribution. This investigation calls for inspection of the convex geometry of \mathcal{C} , with particular attention to its faces and their dimension. To this effect, we use results by Pataki [56] to describe the face of \mathcal{C} which contains a given X in its relative interior, and we quote the lower-bound on the dimension of that face as a function of $\text{rank}(X)$. We further argue that this bound is almost always tight, and we give an essentially tight upper bound on the dimension of a face, generalizing a result of Laurent and Poljak [44] to $d > 1$.

Using this facial description of \mathcal{C} , we establish that for strongly concave f , for $p > p^*$ (3), all second-order critical points of (RP_p) reveal KKT points. Also, for concave f , we show the same for $p > \frac{d+1}{d+3}n$, and argue that $p > p^*$ is sufficient under an additional condition we believe to be mild. Hence, *for linear f , above a certain threshold for p , all second-order critical points of (RP_p) are global optimizers.* The condition is stronger than the one proposed in [20], and the statement is about second-order critical points, rather than about local optimizers of (RP_p) . There are no similar results for convex f , as then solutions can have any rank.

We close the paper with numerical experiments showing the efficiency of the staircase algorithm to solve (P) on certain synchronization problems involving rotations and permutations, as compared to IPM's and SDPLR.

We note that, up to a linear change of variable, problem (P) also encompasses constraints of the form $X_{ii} = B_i$ where each B_i is positive definite. We assume all diagonal blocks have identical size d as this simplifies exposition, but the proposed method can easily accommodate inhomogeneous sizes, and many of the developments go through for complex matrices as well.

1.2 Applications

Problem (RP_p) and its relaxation (P) appear in numerous applications. Many of those belong to the class of *synchronization problems*, which consist in estimating group elements from measurements of pairwise ratios. Further applications are also described, e.g., in [66, 52, 9].

Combinatorial problems can be modeled in (RP_p) with $d = p = 1$. A seminal example is **Max-Cut**: the problem of clustering a graph in two classes, so as to maximize the weight of the edges joining the two classes. The cost f is linear, proportional to the graph Laplacian. Its relaxation to (P) is the subject of a deep analysis by Goemans and Williamson [37], which paved the way to the type of lifts considered here. See [31, eq. (3)] for a recent application of Max-Cut to genomics. The same setup, but with different linear costs, appears in the **stochastic block model** [1], in **community detection** [27], in **maximum a posteriori (MAP) inference in Markov random fields** with binary variables and pairwise interactions [36] and in **robust PCA** [49, Alg. 1]. All of these study the effects of the relaxation

on the final outcome, mostly under random data models. Their linear cost matrices are often structured (sparse or low-rank), which is easily exploited here.

Spherical embeddings is the general problem of estimating points on a sphere in \mathbb{R}^p , and appears notably in machine learning for classification [82] and in the fundamental problem of **packing spheres on a sphere** [26]. It is modeled by (RP_p) with $d = 1, p > 1$. The same setup also models **correlation matrix completion and approximation** [38]. In the latter, an algorithm to solve (P) is proposed, which inspired [42], which inspired this work.

Synchronization of rotations is the problem of estimating m rotation matrices (orthogonal matrices with determinant 1, to exclude reflections), based on pairwise relative rotation measurements. It is modeled in (RP_p) with $d = p > 1$ (often, 2 or 3) and comes up in **structure from motion** [7], **pose graph estimation** [21], **global registration** [24], the **generalized Procrustes problem** [73] and **simultaneous localization and mapping (SLAM)** [22]. It serves in global camera path estimation [18], scan alignment [15, 79], and sensor network localization and the molecule problem [29, 30]. In many of these problems, translations must be estimated as well, and it has often been shown that rotations and translations are best estimated separately [22, Fig. 1]. Here, (RP_p) can easily accommodate the determinant constraint: it comes down to picking one of the connected components of $\text{St}(d, p)^m$, as in [16]. The relaxation (P) ignores this, though; see [63] for relaxations which explicitly model this difference (at additional computational cost). The problem of estimating orthogonal matrices appears notably in the **noncommutative little Grothendieck problem** [52, 9]. In the latter, the relaxation (P) with linear f is called **Orthogonal-Cut**, and its effect on (RP_p) is analyzed. The same relaxation with a nonsmooth cost, for robust estimation, is proposed and analyzed in [79]. See also [8] for another robust formulation of the same problem, based on low-rank + sparse modeling.

The common lines problem in Cryo-EM is an important biomedical imaging instance of (RP_p) , where orthonormal matrices are to be estimated with $d = 2, p = 3$ [80].

Phase synchronization and recovery can be modeled with $p = d = 2$ (as phases are rotations in \mathbb{R}^2). It is sometimes attractive to model phases as unit-modulus complex numbers instead, as is done in [66] for phase synchronization, with the same SDP relaxation. This can be used for **clock synchronization**. See [10] for a study of the tightness of this SDP, and [28] for an application to **ranking**. The **Phase-Cut** algorithm for phase recovery uses the same SDP [78], with a different linear cost. While not explicitly treated, many of the results in this paper extend to the complex case.

1.3 Related work

Problem (RP_p) is an instance of optimization on manifolds [4]. Optimization over orthonormal matrices is also studied in, e.g., [34, 81]. Being equivalent to (P) with a rank constraint, (RP_p) also falls within the scope of optimization over matrices with bounded rank [64, 50], where the latter is also an extension of [42], in the same way that this work is. The particular case of optimization over bounded-rank positive semidefinite matrices with linear constraints was already addressed in [72]. The same without positive semidefiniteness constraint is studied

recently in [46], also with a discussion of global optimality of second-order critical points. With a linear cost f , problem (RP_p) (which then has a quadratic cost g) is a subclass of quadratically constrained quadratic programming (QCQP). QCQP's and their SDP relaxations have been extensively studied, notably in [53, 68], with particular attention to approximation ratios. For (P) , these approximation ratios can be found in [9].

In part owing to the success of (P) with linear f in adequately solving a myriad of hard problems, there has been strong interest in developing fast, large-scale SDP solvers. The present paper is one example of such a solver, restricted to the class of problems (P) . SDPLR is a more generic such solver [19, 20]. See also [76] for a review, and [32] for a recent low-complexity example with precise convergence results, but which does not handle constraints.

Much of this paper is concerned with characterizing the rank of solutions of (P) , especially with respect to how large p must be allowed to grow in (RP_p) to solve (P) . There is also considerable value in determining under what conditions (P) admits solutions of the desired rank for a specific application, that is: when is the relaxation tight? This question is partially answered in [10] for the closely related phase synchronization problem, under a *stochastic* model for the data. See [1] for a proof in the stochastic block model, and [6] for a study of phase transitions in random convex programs. There also exist *deterministic* tightness results, typically relying on special structure in a graph underlying the problem data. See for example [69, 63, 61]. See also Appendix A for a deterministic proof of tightness in the case of single-cycle synchronization of rotations. The proof rests on the availability of a closed-form expression for the dual matrix S (21), and for the solution to be certified. With the same ingredients, it is easy to show, for example, that (P) is tight for Max-Cut when the graph is bipartite.

Semidefinite relaxations in the form of (P) with additional constraints have also appeared in the literature. In particular, this occurs in estimation of rotations, with explicit care for the determinant constraints: Saunderson et al. [63] explicitly constrain off-diagonal blocks to belong to the convex hull of the rotation group; this is not necessary for the orthogonal group—see Proposition 2.1. Similarly, for synchronization of permutations in joint shape matching, off-diagonal blocks are restricted to be doubly stochastic [25, 40]. Finally, in recent work, Bandeira et al. [11] study a more powerful class of synchronization problems with additional linear constraints of various forms. An example with an additional *nonlinear* constraint appears in [79], which imposes an upperbound on the spectral norm of X . All of these are motivation to generalize the framework studied here, in future work.

We mention in passing that the MaxBet and MaxDiff problems [74] do not fall within the scope of this paper. Indeed, although they also involve estimating orthonormal matrices as in (RP_p) , their cost function has a different type of invariance, which would also lead to a different type of relaxation.

1.4 Notation

The size parameters obey $1 \leq d \leq p \leq n = md$. Throughout this paper, matrices $A \in \mathbb{R}^{n \times n}$ are thought of as block matrices with blocks of size $d \times d$. Subscript indexing such as A_{ij} refers to the block on the i th row and j th column of blocks, $1 \leq i, j \leq m$. For $Z \in \mathbb{R}^{n \times p}$, Z_i refers to the i th slice of size $d \times p$, $1 \leq i \leq m$. The Kronecker product is written \otimes and vec vectorizes a matrix by stacking its columns on top of each other. A real number a is rounded down as $\lfloor a \rfloor$. The operator norm $\|A\|_{\text{op}} = \sigma_{\max}(A)$ is the largest singular value of a matrix,

and its Frobenius norm $\|A\|_F$ is the ℓ_2 -norm of $\text{vec}(A)$. $\mathbb{S}^{n \times n}$ is the set of symmetric matrices of size n , and $A \succeq 0$ means $A \in \mathbb{S}^{n \times n}$ is positive semidefinite. $\text{sym}(A) = (A + A^\top)/2$ extracts the symmetric part of a matrix. $O(d)$ is the group of orthogonal matrices of size d . $\ker \mathcal{L}$ denotes the null-space, or kernel, of a linear operator.

2 Geometry

Both search spaces of (RP_p) and (P) enjoy rich geometry, which leads to efficient analytical and numerical tools for the study of these optimization problems. The former is a smooth Riemannian manifold, while the latter is a compact convex set. Figure 1 depicts the two.

2.1 Smooth geometry of the rank-restricted search space

Endow $\mathbb{R}^{n \times p}$ with the classical Euclidean metric $\langle U_1, U_2 \rangle = \text{trace}(U_1^\top U_2)$, corresponding to the Frobenius norm: $\|U\|_F^2 = \langle U, U \rangle$. We view the search space of (RP_p) as a submanifold of $\mathbb{R}^{n \times p}$ and endow it with the Riemannian submanifold geometry [4]. First, define a linear operator $\text{symblockdiag}: \mathbb{R}^{n \times n} \rightarrow \mathbb{S}^{n \times n}$ which symmetrizes diagonal blocks and zeroes out all other blocks:

$$\text{symblockdiag}(M)_{ij} = \begin{cases} \frac{M_{ii} + M_{ii}^\top}{2} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

This allows for a simple definition of the manifold via an equality constraint as

$$\text{St}(d, p)^m = \left\{ Y \in \mathbb{R}^{n \times p} : \text{symblockdiag}(YY^\top) = I_n \right\}. \quad (5)$$

The set is non-empty if $p \geq d$. It is connected if $p > d$. Counting dimensions yields $\dim \text{St}(d, p)^m = np - md(d+1)/2$. The tangent space to $\text{St}(d, p)^m$ at Y is a subspace of $\mathbb{R}^{n \times p}$ obtained by differentiating the equality constraint:

$$\text{T}_Y \text{St}(d, p)^m = \left\{ \dot{Y} \in \mathbb{R}^{n \times p} : \text{symblockdiag}(\dot{Y}Y^\top + Y\dot{Y}^\top) = 0 \right\}. \quad (6)$$

Among the tangent vectors are all vectors of the form $Y\Omega$, for $\Omega \in \mathbb{R}^{p \times p}$ skew-symmetric: these correspond to “vertical directions”, in the sense that following them does not affect the product YY^\top (at first order). Each tangent space is equipped with a restriction of the metric $\langle \cdot, \cdot \rangle$, thus making $\text{St}(d, p)^m$ a Riemannian submanifold of $\mathbb{R}^{n \times p}$. The orthogonal projector from the embedding space $\mathbb{R}^{n \times p}$ to the tangent space at Y is

$$\text{Proj}_Y(Z) = Z - \text{symblockdiag}(ZY^\top)Y. \quad (7)$$

The total computational cost of a projection is thus $\mathcal{O}(m \cdot d^2 p) = \mathcal{O}(ndp)$ flops.

The optimization problem (RP_p) involves a function $g(Y) = f(YY^\top)$ defined over $\mathbb{R}^{n \times p}$. Denote its classical, Euclidean gradient at Y as $\nabla g(Y)$. The Riemannian gradient of g at Y , $\text{grad } g(Y)$, is defined as the unique tangent vector at Y such that, for all tangent \dot{Y} , $\langle \text{grad } g(Y), \dot{Y} \rangle = \langle \nabla g(Y), \dot{Y} \rangle$. Naturally, this is given by the projection of the classical gradient to the tangent space [4, eq. (3.37)]:

$$\text{grad } g(Y) = \text{Proj}_Y(\nabla g(Y)) = 2 \text{Proj}_Y(\nabla f(YY^\top)Y), \quad (8)$$




Figure 1: (Left) For $m = 3, d = 1$, the set \mathcal{C} (2) contains all positive semidefinite matrices of the form $X = [1, a, b; a, 1, c; b, c, 1]$. It is here represented in coordinates (a, b, c) . The interior of the shell contains the rank-3 matrices; the four extreme points (black dots) are the rank-1 matrices; and the remainder of the boundary is the (smooth) set of rank-2 matrices. That smooth geometry breaks down at the rank-1 matrices. Note that for $d > 1$, extreme points of rank d are no longer isolated. (Right) For $m = 3, d = 1, p = 2$, the set $\text{St}(d, p)^m$ (1) parameterizes the matrices of rank at most 2 in \mathcal{C} (redundantly). In this case, $\text{St}(d, p)^m$ corresponds to a product of three circles in 2D: $Y = [\cos \alpha_1, \sin \alpha_1; \cos \alpha_2, \sin \alpha_2; \cos \alpha_3, \sin \alpha_3] \in \text{St}(d, p)^m$. One of these degrees of freedom is redundant, because the factorization of $X \in \mathcal{C}$ as YY^\top is not unique. The figure represents the remaining degrees of freedom after fixing $\alpha_1 = 0$. Notice how accepting the redundancy in the parameterization allows for a smooth representation of the nonsmooth set of bounded rank matrices in \mathcal{C} . Color codes for $\|X\|_F = \|YY^\top\|_F$.

where $\nabla f(X)$ is the classical gradient of f , a symmetric³ matrix of size n , and we used (10). Furthermore, denote by $\nabla^2 g(Y)$ the classical Hessian of g at Y . This is a symmetric operator on $\mathbb{R}^{n \times p}$. The Riemannian Hessian of g at Y is a symmetric operator on the tangent space at Y obtained as the projection of the derivative of the gradient vector field [4, eq. (5.15)]:

$$\begin{aligned} \text{Hess } g(Y)[\dot{Y}] &= \text{Proj}_Y \left(D(Y \mapsto \text{Proj}_Y(\nabla g(Y)))(Y)[\dot{Y}] \right) \\ &= \text{Proj}_Y \left(\nabla^2 g(Y)[\dot{Y}] - \text{symblockdiag} \left(\nabla g(Y) Y^\top \right) \dot{Y} \right), \end{aligned} \quad (9)$$

where D denotes a classical directional derivative and we used $\text{Proj}_Y \circ \text{Proj}_Y = \text{Proj}_Y$. For future reference, we note these expressions of the derivatives of g in terms of those of f :

$$\nabla g(Y) = 2\nabla f(X)Y, \quad \text{and} \quad (10)$$

$$\nabla^2 g(Y)[\dot{Y}] = 2 \left(\nabla^2 f(X)[\dot{X}]Y + \nabla f(X)\dot{Y} \right), \quad \text{with } \dot{X} = \dot{Y}Y^\top + Y\dot{Y}^\top. \quad (11)$$

Optimization algorithms on Riemannian manifolds typically are iterative. As such, they require a means of moving away from a point Y along a prescribed tangent direction \dot{Y} , to reach a new point on the manifold: the next iterate. Since $Y + \dot{Y}$ does not, in general, belong to the manifold, extra operations are required. *Retractions* achieve exactly this [4, § 4.1]. One possible retraction for (1) is as follows. For each $d \times p$ “slice” i in $\{1, \dots, m\}$,

$$\left(\text{Retraction}_Y(\dot{Y}) \right)_i = U_i V_i^\top, \quad \text{with} \quad Y_i + \dot{Y}_i = U_i \Sigma_i V_i^\top, \quad (12)$$

where $U_i \Sigma_i V_i^\top$ is a thin singular value decomposition of the i th slice $Y_i + \dot{Y}_i$. This retraction projects each slice of $Y + \dot{Y}$ to the closest orthonormal matrix. Consequently, this is even a second-order retraction [2]. The total cost of computing a retraction is $\mathcal{O}(m \cdot (p^2 d + d^3)) = \mathcal{O}(np^2)$ flops.

2.2 Convex geometry of the full search space

The optimization problem (P) is defined over the compact convex set

$$\mathcal{C} = \{X \in \mathbb{S}^{n \times n} : X \succeq 0 \text{ and } X_{ii} = I_d \text{ for } i = 1 \dots m\}.$$

For $d = 1$, this is the *elliptope*, or set of correlation matrices [44]. Often, we hope to recover matrices X in \mathcal{C} such that X has rank d , or such that off-diagonal blocks X_{ij} are orthogonal. The following proposition shows that these two considerations are equivalent, and that the convex relaxation leading to \mathcal{C} is tight in that sense (the tightest relaxation would consider the convex hull of rank- d matrices in \mathcal{C} , but this is difficult to handle⁴).

Proposition 2.1. *For all $X \in \mathcal{C}$, all blocks X_{ij} are in the convex hull of $\text{O}(d)$, that is, $\sigma_{\max}(X_{ij}) \leq 1$. Furthermore, $\text{rank}(X) = d$ if and only if $X_{ij} \in \text{O}(d)$ for all i, j .*

³ $\nabla f(X)$ is symmetric because f is formally defined over the symmetric matrices. If the gradient of f over the square matrices is not symmetric, $\nabla f(X)$ is obtained by extracting its symmetric part.

⁴Let $\tilde{\mathcal{C}}$ be the convex hull of rank- d matrices in \mathcal{C} . The extreme points of $\tilde{\mathcal{C}}$ are these matrices [59, Cor. 18.3.1]. Thus, optimizing a linear cost function over $\tilde{\mathcal{C}}$ solves (RP_d) , which is NP-hard. Hence, there probably does exist an efficient representation of $\tilde{\mathcal{C}}$.

Proof. Since X is positive semidefinite, for all $i \neq j$, the submatrix formed by the blocks X_{ii}, X_{ij}, X_{ji} and X_{jj} is positive semidefinite. By Schur, this holds if and only if $X_{ij}^\top X_{ij} \preceq I_d$, which in turn happens if and only if all singular values of X_{ij} are at most 1. The set of such matrices is the convex hull of all orthogonal matrices of size d [62].

Consider $Y \in \text{St}(d, p)^m$ such that $\text{rank}(X) = p$ and $X = YY^\top$. Clearly, if $p = d$, $X_{ij} = Y_i Y_j^\top$ is orthogonal, since all Y_i 's are orthogonal. Conversely, if X_{ij} is orthogonal, then the rows of Y_i and Y_j span the same subspace. Indeed, $Y_i Y_j^\top Y_j Y_i^\top = I_d$. Multiply by Y_i on the right. Now notice that $Y_k^\top Y_k$ is an orthogonal projector onto the subspace spanned by the rows of Y_k . Since Y_i remains unaffected by such a projection first on the subspace of Y_j then again on the subspace of Y_i , they must span the same subspace. Hence, fixing $j = 1$, for each i , there exists $Q_i \in O(d)$ such that $Y_i = Q_i Y_1$. Finally, $Y = \text{diag}(I_d, Q_2, \dots, Q_m)(\mathbb{1}_{m \times 1} \otimes Y_1)$ (where \otimes is the Kronecker product), which confirms that Y and X have rank d . Notice that this proof further shows that X has rank d if and only if there is a spanning tree of edges (i, j) on an m -nodes graph such that the X_{ij} 's are orthogonal (in which case they are all orthogonal). \square

The set \mathcal{C} may be decomposed into *faces* of various dimensions.

Definition 2.1 (faces, §18 in [59]). *A face of \mathcal{C} is a convex subset \mathcal{F} of \mathcal{C} such that every (closed) line segment in \mathcal{C} with a relative interior point in \mathcal{F} has both endpoints in \mathcal{F} . The empty set and \mathcal{C} itself are faces of \mathcal{C} .*

By [59, Thm. 18.2], the collection of relative interiors⁵ of the non-empty faces forms a partition of \mathcal{C} . That is, each $X \in \mathcal{C}$ is in the relative interior of exactly one face of \mathcal{C} , called \mathcal{F}_X . Furthermore, all faces of \mathcal{C} are exposed [58, Cor. 1], that is, for every face \mathcal{F} , there exists a linear function f such that \mathcal{F} is the set of solutions of (P). Of particular interest are the zero-dimensional faces of \mathcal{C} (singletons), called its *extreme points*.

Definition 2.2 (Extreme and exposed points). *$X \in \mathcal{C}$ is an extreme point of \mathcal{C} if there does not exist $X', X'' \in \mathcal{C} \setminus \{X\}$ and $0 < \lambda < 1$ such that $X = \lambda X' + (1 - \lambda)X''$. X is an exposed point of \mathcal{C} if there exists C such that X is the unique maximizer of $\langle C, X \rangle$ in \mathcal{C} .*

In other words, X is extreme if it does not lie on an open line segment included in \mathcal{C} . Since \mathcal{C} is compact, it is the convex hull of its extreme points [59, Cor. 18.5.1]. Extreme points are of interest notably because they often arise as the solution of optimization problems. Specifically, if f is a concave function (in particular, if f is linear), then f attains its minimum on \mathcal{C} at one of its extreme points [59, Cor. 32.3.2].

Following the construction in the proof of [56, Thm. 2.1], given $Y \in \mathbb{R}^{n \times p}$ of full rank such that $X = YY^\top$ ($\text{rank}(X) = p$), we find that

$$\mathcal{F}_X = \left\{ \tilde{X} = Y(I_p + A)Y^\top : A \in \ker \mathcal{L}_X \text{ and } I_p + A \succeq 0 \right\}, \text{ with} \quad (13)$$

$$\mathcal{L}_X : \mathbb{S}^{p \times p} \rightarrow (\mathbb{S}^{d \times d})^m : A \mapsto \mathcal{L}_X(A) = (Y_1 A Y_1^\top, \dots, Y_m A Y_m^\top). \quad (14)$$

The dimension of \mathcal{F}_X is the dimension of the kernel of \mathcal{L}_X . The rank-nullity theorem gives a lowerbound (see also Theorem 3.16 for an upperbound):

$$\dim \mathcal{F}_X = \frac{p(p+1)}{2} - \text{rank } \mathcal{L}_X \geq \frac{p(p+1)}{2} - m \frac{d(d+1)}{2} \triangleq \Delta. \quad (15)$$

⁵The relative interior of a singleton is the singleton.

It follows that extreme points X (i.e., points such that $\dim \mathcal{F}_X = 0$) have small rank:

$$d \leq \text{rank}(X) \leq p^* := \left(\sqrt{1 + 4md(d+1)} - 1 \right) / 2. \quad (16)$$

Note that $\Delta \geq 0$ when $p \geq p^*$.

Remark 2.2. For linear f , (P) admits an extreme point as global optimizer, so that (RP_p) and (P) have the same optimal value as soon as $p \geq p^*$ (16). In other words: for linear f , (RP_p) is not NP-hard if $p \geq p^*$.

Not all feasible X 's with rank as in (16) are extreme. For example, setting $d = 1$ and $m \geq 3$ as in Figure 1, select two distinct, admissible matrices of rank 1, X_0 and X_1 . For all $0 < \lambda < 1$, the matrix $X_\lambda = \lambda X_1 + (1 - \lambda)X_0$, lying on the open line segment between X_0 and X_1 , is admissible and has rank 2. Thus, X_λ satisfies (16), but it is not an extreme point, by construction. Notwithstanding, the expectation that \mathcal{L}_X is generically of full rank suggests that almost all feasible X 's satisfying (16) should be extreme; an intuition that is supported by Figure 1. More generally, in Theorem B.1, we prove for $d = 1$ that $\dim \mathcal{F}_X = \Delta$ for almost all X of rank p .

Many applications look for solutions of rank d . All X 's of rank d are exposed (hence extreme), meaning they can all be recovered as unique solutions of (P).

Proposition 2.3. For all $X \in \mathcal{C}$, $\|X\|_F^2 \leq m^2 d$. Furthermore, $\|X\|_F^2 = m^2 d$ if and only if $\text{rank}(X) = d$. In particular, each $X \in \mathcal{C}$ of rank d is an exposed extreme point of \mathcal{C} .

Proof. Let $\sigma_1(X_{ij}) \geq \dots \geq \sigma_d(X_{ij})$ denote the singular values of X_{ij} . By Proposition 2.1, $\sigma_k(X_{ij}) \leq 1$ for all i, j, k . Hence,

$$\|X\|_F^2 = \sum_{i,j=1}^m \|X_{ij}\|_F^2 = \sum_{i,j=1}^m \sum_{k=1}^d \sigma_k^2(X_{ij}) \leq m^2 d. \quad (17)$$

The upperbound is attained if and only if $\sigma_k(X_{ij}) = 1$ for all i, j, k , thus, if and only if all X_{ij} 's are orthogonal. By Proposition 2.1, this is the case if and only if $\text{rank}(X) = d$. Now consider X has rank d . We show it is exposed (and hence extreme):

$$\max_{\hat{X} \in \mathcal{C}} \langle X, \hat{X} \rangle \leq \max_{\|\hat{X}\|_F^2 \leq m^2 d} \langle X, \hat{X} \rangle \leq \|X\|_F \cdot \max_{\|\hat{X}\|_F^2 \leq m^2 d} \|\hat{X}\|_F = m^2 d. \quad (18)$$

The second inequality follows by Cauchy-Schwarz, and equality is attained if and only if $\hat{X} = X$, which is in \mathcal{C} . Thus, both max problems admit X as unique solution, confirming that X is an exposed extreme point. \square

Remark 2.4. Proposition 2.3 is an extension of Theorem 1 in [48] to the case $d > 1$. We note that the proof used in the latter reference does not generalize to $d > 1$. Indeed, let $\lambda(X) \in \mathbb{R}^{md}$ denote the vector of (nonnegative) eigenvalues of X . It holds that $d \leq \|\lambda(X)\|_0 = \text{rank}(X) \leq md$, $\|\lambda(X)\|_1 = \text{trace}(X) = md$ and $\|\lambda(X)\|_2 = \|X\|_F \leq m\sqrt{d}$. It is well known that the 1-norm and the 2-norm of a vector coincide if and only if it has at most one nonzero entry. For $d = 1$, Malick [48] uses this to show that $\|X\|_F$ is maximal if and only if X has rank 1. For $d > 1$ though, there exist vectors $x \in \mathbb{R}^{md}$ with the prescribed 1-norm and maximal 2-norm, but with sparsity larger than d . For example, $m = d = 2$ and $x = [8, 2, 2, 0]/3$.

3 From second-order critical points Y to KKT points X

In this section, we show that rank-deficient second-order critical points Y of (RP_p) yield KKT points $X = YY^\top$ of (P) . Furthermore, when Y is second-order critical but X is not KKT, it is shown how to escape the saddle point by increasing p . If p increases all the way to n , then all second-order critical points reveal KKT points. The proofs parallel those in [42]. The main novelty is explicit bounds on p such that *all* second-order critical points of (RP_p) reveal KKT points. The proofs bring us to consider the facial structure of \mathcal{C} .

A key ingredient for all proofs in this section is the availability of an explicit matrix $S(X)$ (21) which is positive semidefinite if and only if X is KKT (Theorem 3.3). The formula for S is simply read off from the first-order optimality conditions of (RP_p) , owing to smoothness of the latter.

Lemma 3.1 (Necessary optimality conditions for (P)). *$X \in \mathcal{C}$ is called a KKT point for (P) if there exist a symmetric matrix $\hat{S} \in \mathbb{S}^{n \times n}$ and a symmetric, block-diagonal matrix $\hat{\Lambda} \in \mathbb{S}^{n \times n}$ (dual variables) such that*

$$\hat{S}X = 0, \quad \hat{S} = \nabla f(X) + \hat{\Lambda}, \quad \text{and} \quad \hat{S} \succeq 0.$$

If X is a local optimizer for (P) , then X is a KKT point. If f is convex, all KKT points are global optimizers.

Proof. Apply Theorems 3.25 and 3.34, and Example 3.36 in [60]. KKT conditions are necessary since Slater's condition holds: I_n is feasible for (P) and it is strictly positive definite. \square

Lemma 3.2 (Necessary optimality conditions for (RP_p)). *Let $Y \in \text{St}(d, p)^m$ and $X = YY^\top$. A critical point of (RP_p) satisfies $\text{grad } g(Y) = 0$, that is,*

$$\left(\nabla f(X) - \text{symblockdiag}(\nabla f(X)X) \right) Y = 0. \quad (19)$$

A second-order critical point is a critical point which satisfies $\text{Hess } g(Y) \succeq 0$, that is,

$$\text{for all } \dot{Y} \in T_Y \text{St}(d, p)^m, \quad \left\langle \dot{Y}, \nabla^2 g(Y)[\dot{Y}] - \text{symblockdiag}(\nabla g(Y)Y^\top) \dot{Y} \right\rangle \geq 0. \quad (20)$$

If Y is a local optimizer for (RP_p) , then it is a second-order critical point.

Proof. This is a direct generalization of the classical necessary optimality conditions for unconstrained optimization [14, Prop. 1.1.1] to the Riemannian setting, as per the formalism in [4, 83]. Use equations (7), (8) and (9) and the fact that Proj_Y is self-adjoint to obtain equations (19) and (20). \square

Lemmas 3.1 and 3.2 suggest the definition of an (as yet merely tentative) formula for the dual certificate \hat{S} , based on (19):

$$S = S(X) = \nabla f(X) - \text{symblockdiag}(\nabla f(X)X). \quad (21)$$

Indeed, for any critical point Y of (RP_p) , it holds (with $X = YY^\top$) that $SY = 0$, so that $SX = 0$. In that case, for $p = d$, S can be advantageously interpreted as a graph Laplacian

(up to a change of variable⁶), as is often the case for dual certificates of estimation problems on graphs [10, 33]. We now show that $S(X)$ is indeed the unique possible dual certificate for any feasible X . (This is similar to, but different from, Theorem 4 in [42]; a complex version appears in [10] for $d = 1$.)

Theorem 3.3 (S is the right certificate). *$X \in \mathcal{C}$ is a KKT point for (P) if and only if S (21) is positive semidefinite. If so, $\hat{S} = S$ is the unique dual certificate for Lemma 3.1.*

Proof. We show the if and only if parts of the first statement separately.

\Leftarrow : By construction, $\text{trace}(SX) = 0$. Since both S and X are positive semidefinite, this implies $SX = 0$. Apply Lemma 3.1 with $\hat{S} = S$ and $\hat{\Lambda} = -\text{symblockdiag}(\nabla f(X)X)$.

\Rightarrow : Since X is a KKT point for (P), there exist $\hat{S} \succeq 0$ and $\hat{\Lambda}$ symmetric, block-diagonal satisfying the conditions in Lemma 3.1. In particular, $\hat{S}X = 0$ and $\nabla f(X) = \hat{S} - \hat{\Lambda}$. Thus, $\nabla f(X)X = -\hat{\Lambda}X$ and $\text{symblockdiag}(\nabla f(X)X) = -\text{symblockdiag}(\hat{\Lambda}X) = -\hat{\Lambda}$.

Here, we used both the fact that $\hat{\Lambda}$ is symmetric, block-diagonal and the fact that $X_{ii} = I_d$. Consequently, $S = \nabla f(X) - \text{symblockdiag}(\nabla f(X)X) = \hat{S} - \hat{\Lambda} + \hat{\Lambda} = \hat{S} \succeq 0$.

The last point also shows there exists only one pair $(\hat{S}, \hat{\Lambda})$ certifying X is a KKT point. \square

Notice how the Riemannian structure underlying problem (P) made it possible to simply read off an analytical expression for a dual certificate from the necessary optimality conditions of (RP_p) . This smooth geometry also leads to uniqueness of the dual certificate (this is connected to nondegeneracy [5, Thm. 7]). Theorem 3.3 makes for an unusually comfortable situation and will be helpful throughout the paper.

All important points of (P) give rise to important points in (RP_p) for sufficiently large p , meaning we do not run the risk of missing features of the former by studying the latter.

Proposition 3.4 (No information loss). *Let $X \in \mathcal{C}$ and $p \geq \text{rank}(X)$. Consider any $Y \in \text{St}(d, p)^m$ such that $X = YY^\top$. If X is a local minimizer of (P), then Y is a local minimizer of (RP_p) . If X is a KKT point, then Y is a critical point.*

Proof. First claim: X is optimal in a neighborhood $\mathcal{U} \subseteq \mathcal{C}$, $X \in \text{int}\mathcal{U}$. Let U be the set of $\tilde{Y} \in \text{St}(d, p)^m$ such that $\tilde{Y}\tilde{Y}^\top \in \mathcal{U}$; U is a neighborhood of Y in $\text{St}(d, p)^m$ and Y is optimal in it. Second claim: X is KKT $\Rightarrow SX = 0 \Leftrightarrow SY = 0 \Leftrightarrow Y$ is critical. \square

For convex f , we can make the following statement regarding uniqueness of the solution.

Theorem 3.5. *Assume f is convex. If $X \in \mathcal{C}$ is an extreme point for (P) (which is true in particular if $\text{rank}(X) = d$), and $S \succeq 0$, and $\text{rank}(X) + \text{rank}(S) = n$ (strict complementarity), then X is the unique global optimizer of (P).*

⁶For Y a critical point of (RP_p) with $p = d$, we have $Y = (y_1^\top \dots y_m^\top)^\top$ with orthogonal Y_i 's, and $SY = 0$. Write $\nabla f(YY^\top) := -C$ for short. $SY = 0$ implies, after some algebra, that the $(CYY^\top)_{ii}$'s are symmetric. This can be used to see that $\hat{S} = \text{diag}(Y_1, \dots, Y_m)^\top S(YY^\top) \text{diag}(Y_1, \dots, Y_m)$ is a matrix with off-diagonal blocks equal to $-Y_i^\top C_{ij} Y_j$ and diagonal blocks equal to $\sum_{j \neq i} Y_i^\top C_{ij} Y_j$. Thus, \hat{S} is exactly the Laplacian of the graph with m nodes and edge “weights” given by the matrices $Y_i^\top C_{ij} Y_j$.

Proof. From Theorem 3.3, it is clear that X is a global optimizer. We prove by contradiction that it is unique. Let $X' \neq X$ be another global optimizer. Since (P) is a convex problem in this setting, f is constant over the whole (optimal) segment $t \mapsto X + t(X' - X)$ for $t \in [0, 1]$. Hence, the directional derivative of f at X along $\dot{X} = X' - X$ is zero:

$$\begin{aligned}
0 &= \langle \nabla f(X), \dot{X} \rangle \\
&= \langle S + \text{symblockdiag}(\nabla f(X)X), \dot{X} \rangle && \text{(definition of } S \text{ (21))} \\
&= \langle S, \dot{X} \rangle && \text{(diagonal blocks of } \dot{X} \text{ are zero)} \\
&= \langle S, X' \rangle && (SX = 0).
\end{aligned}$$

Since both S and X' are positive semidefinite, it ensues that $SX' = 0$. (Note that for linear f , this shows S is the dual certificate for all global optimizers of (P), not only for X .) Hence, $S\dot{X} = 0$.

Let $p = \text{rank}(X)$ and $Y \in \mathbb{R}^{n \times p}$ be a full-rank matrix such that $X = YY^\top$. Strict complementarity and $SX = 0$ imply the columns of Y form a basis for the kernel of S . Hence, $S\dot{X} = 0$ ensures that $\dot{X} = YAY^\top$ for some $A \in \mathbb{S}^{p \times p}$, $A \neq 0$, such that $\text{symblockdiag}(YAY^\top) = 0$. In other words, $X' = X + \dot{X} = Y(I_p + A)Y^\top$ is in the face \mathcal{F}_X . This is a contradiction, since $\mathcal{F}_X = \{X\}$. \square

Remark 3.6. *In general, the condition that X be an extreme point in the previous theorem cannot be removed. Indeed, if $f(X) \equiv 0$, then all admissible X 's are globally optimal and $S(X) \equiv 0 \succeq 0$. In particular, $X = I_n$ satisfies strict complementarity, but if $m > 1$, it is not extreme, and it is not a unique global optimizer. Likewise, any rank- d admissible X is extreme and globally optimal, but does not satisfy strict complementarity. Similar examples can be built with nonzero linear costs $f(X) = \langle C, X \rangle$, where the sparsity pattern of C corresponds to a disconnected graph.*

Conversely, it is not true in general that uniqueness of the global optimizer implies extremeness or strict complementarity. Simply consider $f(X) = \|X - X_0\|_F^2$ with $X_0 \in \mathcal{C}$: the global optimizer $X = X_0$ is unique, and $\nabla f(X_0) = 0$, so that $S(X_0) = 0$. For m large enough, X_0 can be chosen to be both not extreme and rank deficient. For an illustration of this in semidefinite programs, see the nice example after Theorem 10 in [5].

Continuing with convex cost functions, KKT points of (P) coincide with global optimizers. This and the fact that (RP_p) is a relaxation of (P) lead to the following summary regarding global optimality conditions. For (RP_p), these sufficient conditions are conclusive whenever the relaxation is tight.

Corollary 3.7 (Global optimality conditions). *Assume f is convex and let $Y \in \text{St}(d, p)^m$; $X = YY^\top$ is globally optimal for (P) if and only if $S \succeq 0$. If so, then Y is a global optimizer for the nonconvex problem (RP_p). Furthermore, if X is extreme (in particular, if $\text{rank}(X) = d$) and if $\text{rank}(X) + \text{rank}(S) = n$, then X is the unique global optimizer of (P) and Y is the unique global optimizer of (RP_p), up to orthogonal action YQ , $Q \in \text{O}(p)$.*

Returning to the general case of f not necessarily convex, we establish links between (second-order) critical points of (RP_p) and KKT points of (P). In doing so, it is useful to reformulate the second-order optimality condition (20) on (RP_p) in terms of S (21) and f . Let

$Y \in \text{St}(d, p)^m$ and $X = YY^\top$. Then, Y is a second-order critical point for (RP_p) if and only if it is critical and, for all $\dot{Y} \in \text{TySt}(d, p)^m$, with $\dot{X} = \dot{Y}Y^\top + Y\dot{Y}^\top$, it holds that (using (11)):

$$\begin{aligned} \langle \dot{Y}, \text{Hess } g(Y)[\dot{Y}] \rangle &= \langle \dot{Y}, \nabla^2 g(Y)[\dot{Y}] - \text{symblockdiag}(\nabla g(Y)Y^\top) \dot{Y} \rangle \\ &= 2\langle \dot{Y}, \nabla^2 f(X)[\dot{X}]Y + \nabla f(X)\dot{Y} - \text{symblockdiag}(\nabla f(X)X) \dot{Y} \rangle \\ &= 2\langle \dot{Y}, \nabla^2 f(X)[\dot{X}]Y + S\dot{Y} \rangle \\ &= \langle \dot{X}, \nabla^2 f(X)[\dot{X}] \rangle + 2\langle \dot{Y}, S\dot{Y} \rangle \geq 0. \end{aligned} \quad (22)$$

Since (RP_p) is essentially equivalent to (P) with the additional constraint $\text{rank}(X) \leq p$, assuming Y is optimal for (RP_p) , we expect X to be a KKT point at least if either of the following holds: (1) if Y is rank deficient, since then the extra constraint is not active, meaning it is “as if” we were solving (P) ; or (2) if $p = n$, since then the extra constraint is vacuous. The two following theorems show this still holds for second-order critical points Y .

Theorem 3.8. *If Y is a rank-deficient, second-order critical point for (RP_p) , then $X = YY^\top$ is a KKT point for (P) .*

Proof. By Theorem 3.3, we must show that S (21) is positive semidefinite. Since Y is rank deficient, there exists $z \in \mathbb{R}^p$ such that $z \neq 0$ and $Yz = 0$. Furthermore, for all $x \in \mathbb{R}^n$, the matrix $\dot{Y} = xz^\top$ is such that $Y\dot{Y}^\top = 0$. In particular, \dot{Y} is a tangent vector at Y (6). Since Y is second-order critical, inequality (22) holds, and here simplifies to:

$$\langle \dot{Y}, S\dot{Y} \rangle = \langle xz^\top, Sxz^\top \rangle = \|z\|^2 \cdot x^\top Sx \geq 0.$$

This holds for all $x \in \mathbb{R}^n$. Thus, S is positive semidefinite. \square

Theorem 3.9. *If Y is square ($p = n$) and it is a second-order critical point for (RP_n) , then $X = YY^\top$ is a KKT point for (P) . If Y is full-rank, it needs only be first-order critical for X to be a KKT point.*

Proof. If Y is rank deficient, then the result follows from Theorem 3.8. If Y is full rank, then it is invertible. Since Y is also a critical point, first-order optimality conditions (19) imply $SY = 0$, hence $S = 0$. This completes the proof, as per Theorem 3.3. \square

In particular, if f is convex and (P) has a unique solution of rank r , then all second-order critical points of (RP_p) have rank either r or p . Thus, if p is larger than the rank of a solution of (P) , we may hope that minimizing (RP_p) until we reach a second-order critical point will result in a rank-deficient Y , revealing a KKT point. Unfortunately, in general, we cannot guarantee rank deficiency beforehand. For those cases, the following theorem and corollary provide a means of escaping unsatisfactory critical points.

Theorem 3.10 (Escape direction (rank-deficient)). *Let Y be a rank-deficient critical point for (RP_p) such that $X = YY^\top$ is not a KKT point of (P) . Then, for all nonzero vectors $z \in \mathbb{R}^p$ and $u \in \mathbb{R}^n$ such that $Yz = 0$ and $u^\top Su < 0$ (21), $\dot{Y} = uz^\top \in \text{TySt}(d, p)^m$ is a descent direction for g from Y .*

Proof. By Theorem 3.3, u exists because X is not a KKT point. Since $Y\dot{Y}^\top = 0$, \dot{Y} is indeed tangent at Y (6). From (22), it follows that

$$\langle \dot{Y}, \text{Hess } g(Y)[\dot{Y}] \rangle = 2\langle \dot{Y}, S\dot{Y} \rangle = 2\|z\|^2 \cdot u^\top Su < 0.$$

As a result, this truncated Taylor expansion holds (we use the fact that the retraction (12) is second-order and assume $\|z\| = 1$):⁷

$$\phi(t) := g(\text{Retraction}_Y(t\dot{Y})) = g(Y) + (u^\top Su)t^2 + \mathcal{O}(t^4).$$

Since $u^\top Su < 0$, there exists $t_0 > 0$ such that $\phi(t) < \phi(0)$ for all t in $]0, t_0[$. \square

Corollary 3.11 (Escape direction (full-rank)). *Let Y be a full-rank critical point for (RP_p) such that $X = YY^\top$ is not a KKT point of (P). Let $Y_+ = (Y \ 0_{n \times (p_+ - p)}) \in \text{St}(d, p_+)^m$. Then, (a) $g(Y_+) = g(Y)$; (b) Y_+ is a critical point for (RP_{p_+}) ; and (c) for all $u \in \mathbb{R}^n$ such that $u^\top Su < 0$, $\dot{Y} = ue_{p+1}^\top \in T_{Y_+} \text{St}(d, p_+)^m$ is a descent direction for g from Y_+ , where $e_{p+1} \in \mathbb{R}^{p_+}$ is a zero vector except for its $(p+1)^{\text{st}}$ entry, equal to 1.*

Proof. Since $Y_+ Y_+^\top = YY^\top$, Y_+ is indeed feasible for (RP_{p_+}) and $g(Y_+) = g(Y)$. Since Y is critical, $SY = 0$, so $SY_+ = 0$: Y_+ is a critical point. The rest follows from Theorem 3.10. \square

Later in this section, we show how to escape full-rank points without increasing the rank, under additional assumptions—see Proposition 3.18.

An important question remains: for moderate p , how frequently should we expect to encounter second-order critical points Y that do not correspond to KKT points of (P)? We provide partial answers for concave f below. In particular, this covers the important case of linear costs. The stronger results do not include strictly convex functions, as for these (P) can have solutions of arbitrary rank.

The result below shows that Y can only be a critical point of (RP_p) if YY^\top is a critical point for (P) restricted to the face \mathcal{F}_X ; and similarly for second-order critical points. This brings a useful corollary.

Lemma 3.12. *Let $Y \in \text{St}(d, p)^m$ and $X = YY^\top \in \mathcal{C}$. If Y is a critical point (resp., a second-order critical point) of (RP_p) , then X is a critical point (resp., a second-order critical point) of $\min_{\tilde{X} \in \mathcal{F}_X} f(\tilde{X})$, where \mathcal{F}_X (13) is the face of \mathcal{C} which contains X in its relative interior.*

⁷ We remark that the residue is of order 4 rather than 3, since ϕ is an even function of t . Indeed, let $J = \text{diag}(1, \dots, 1, -1) \in \text{O}(p)$ and consider the SVD of $Y = U\Sigma V^\top$, such that the last column of V is $z/\|z\|$. Then, $Y - t\dot{Y} = (YV - t\dot{Y}V)V^\top = (Y + t\dot{Y})VJV^\top$, because the last column of YV is zero, and $\dot{Y}V$'s only nonzero column is the last one. Hence, $\text{Retraction}_Y(-t\dot{Y}) = \text{Retraction}_Y(t\dot{Y})VJV^\top$. These yield the same matrix X , thus the same value of g . If the residue is bounded by Lt^4 , then $t = \sqrt{-u^\top Su/2L}$ ensures $\phi(t) \leq \phi(0) - (u^\top Su)^2/4L$. In particular, for $f(X) = \langle C, X \rangle$, the fourth-order term follows and could be used to estimate L , to be used in a line-search for the escape:

$$\frac{1}{4}t^4 \cdot \left[\langle C, AXA \rangle + u^\top D(3 \text{symblockdiag}(CX) - 4C)u \right] + \mathcal{O}(t^6),$$

with $A = \text{symblockdiag}(uu^\top)$ and $D = \text{diag}(\|u_1\|^2, \dots, \|u_m\|^2) \otimes I_d$.

Proof. If Y is first-order critical, then $SY = 0$ (21). This implies that $\nabla f(X)$ is orthogonal to all directions $\dot{X} = Y\dot{Y}^\top + \dot{Y}Y^\top$ for $\dot{Y} \in \text{TySt}(d, p)^m$. Indeed, $\text{symblockdiag}(\dot{X}) = 0$, hence

$$\langle \nabla f(X), \dot{X} \rangle = \langle S, \dot{X} \rangle = 0.$$

(The subspace spanned by all such \dot{X} 's has dimension $np - m \frac{d(d+1)}{2} - \frac{p(p-1)}{2}$.) From (13), observe that \dot{X} is parallel to the face \mathcal{F}_X iff $\dot{X} = YAY^\top$ for $A \in \mathbb{S}^{p \times p}$ with $\text{symblockdiag}(YAY^\top) = 0$. Thus, $\nabla f(X)$ is, in particular, orthogonal to \mathcal{F}_X . This shows X is a first-order critical point for $\min_{\tilde{X} \in \mathcal{F}_X} f(\tilde{X})$, since X is in the relative interior of the face. Further consider (22) for second-order critical Y and all $\dot{X} = YAY^\top$. Since $\dot{Y} = \frac{1}{2}YA$ and $SY = 0$, it follows that

$$0 \leq \langle \dot{X}, \nabla^2 f(X)[\dot{X}] \rangle + 2\langle \dot{Y}, S\dot{Y} \rangle = \langle \dot{X}, \nabla^2 f(X)[\dot{X}] \rangle.$$

Thus, X is a second-order critical point for the face-restricted optimization problem. \square

The following corollary can be put in perspective with [20, Thm.3.4]. The latter states a similar result for (P) with general linear equality constraints, for linear f . Their result characterizes local optimizers of (RP_p), whereas the following result characterizes first- and second-order critical points (computationally more manageable objects).

Corollary 3.13. *Theorem 3.8 and Lemma 3.12 imply the following, for $Y \in \text{St}(d, p)^m$ and $X = YY^\top \in \mathcal{C}$.*

- *If f is linear and Y is critical, then f is constant over \mathcal{F}_X . If furthermore Y is second-order critical and $p > \lfloor p^* \rfloor$ (16), then either X is globally optimal for (P), or (Y has full rank and) the face \mathcal{F}_X has positive dimension (15) and is suboptimal.*
- *If f is strongly concave and Y is second-order critical, then X is an extreme point (since $\dot{X} = YAY^\top \neq 0 \Rightarrow \langle \dot{Y}, \text{Hess } g(Y)[\dot{Y}] \rangle = \langle \dot{X}, \nabla^2 f(X)[\dot{X}] \rangle < 0$, a contradiction). In particular, all second-order critical points of (RP_p) have rank at most $\lfloor p^* \rfloor$, and $p > \lfloor p^* \rfloor \Rightarrow X$ is KKT (since Y must be rank-deficient).*
- *If f is convex (resp., strictly convex) and Y is critical, then X is optimal (resp., the unique optimizer) of (P) restricted to \mathcal{F}_X .*

For linear f , Corollary 3.13 is not quite sufficient to determine how large p must be to exclude “bad” second-order critical points. Paraphrasing the comment following [20, Thm.3.4], the latter showed that, for linear f and $p > p^*$ (essentially), local optima of (RP_p) are global optima, with the caveat that positive-dimensional faces (over which f must be constant) may harbor non-global local optima. In the literature, this has sometimes been quoted as saying that local optima are global optima if f is not constant over any proper face of \mathcal{C} [see, e.g., 49, footnote 3], but there is no indication that this is a mild condition.⁸

We thus set out to further refine the implications of second-order criticality of Y . We do so by leveraging the tight relationship (22) between the Hessian of the cost on (RP_p) and the dual certificate S .

⁸In fact, we found in numerical experiments (not reported) that, for $d = 1$, $\Delta = 1$ (almost all faces have dimension 1, $p > p^*$) and a random linear cost $\langle C, X \rangle$, we could easily find a face \mathcal{F}_X of dimension 1 over which the cost is constant but not optimal.

Theorem 3.14. *Assume f is concave. Let $Y \in \text{St}(d, p)^m$ be a second-order critical point for (RP_p) . The matrix $X = YY^\top$ belongs to the relative interior of the face \mathcal{F}_X (13). If Y is rank-deficient, then $S \succeq 0$. If Y is full-rank, then S has at most*

$$\left\lfloor \frac{\dim \mathcal{F}_X - \Delta}{p} \right\rfloor \quad (23)$$

negative eigenvalues. Hence, if X is not a KKT point for (P), then it has rank p and $\dim \mathcal{F}_X \geq \Delta + p$.

Proof. Given Theorems 3.3 and 3.8, we need only focus on full-rank Y 's. Since Y is second-order critical, inequality (22) holds. Since furthermore f is concave, $\langle \dot{X}, \nabla^2 f(X)[\dot{X}] \rangle \leq 0$ for all \dot{X} , so that $\langle \dot{Y}, S\dot{Y} \rangle \geq 0$ for all $\dot{Y} \in \text{T}_Y \text{St}(d, p)^m$. Let $k = \dim \text{St}(d, p)^m$ and $U \in \mathbb{R}^{np \times k}$, $U^\top U = I_k$ denote an orthonormal basis of the space spanned by the vectorized tangent vectors $\text{vec}(\dot{Y})$. Then, $U^\top(I_p \otimes S)U$ is positive semidefinite (for linear f , it inherits the spectrum of $\frac{1}{2}\text{Hess} g(Y)$). On the other hand, since Y is a critical point, $SY = 0$. Let $V \in \mathbb{R}^{np \times p^2}$, $V^\top V = I_{p^2}$ denote an orthonormal basis of the space spanned by the vectors $\text{vec}(YR)$ for $R \in \mathbb{R}^{p \times p}$. Clearly, $(I_p \otimes S)V = 0$. Let k' denote the dimension of the space spanned by the columns of both U and V , and let $W \in \mathbb{R}^{np \times k'}$, $W^\top W = I_{k'}$ be an orthonormal basis for this space. It follows that $M = W^\top(I_p \otimes S)W$ is positive semidefinite.

Let $\lambda_0 \leq \dots \leq \lambda_{n-1}$ denote the eigenvalues of S . Likewise, $\tilde{\lambda}_0 \leq \dots \leq \tilde{\lambda}_{np-1}$ denote the eigenvalues of $I_p \otimes S$. These are simply the eigenvalues of S , each repeated p times, thus: $\tilde{\lambda}_i = \lambda_{\lfloor i/p \rfloor}$. Let $\mu_0 \leq \dots \leq \mu_{k'-1}$ denote the eigenvalues of M . The Cauchy interlacing theorem states that, for all i ,

$$\tilde{\lambda}_i \leq \mu_i \leq \tilde{\lambda}_{i+np-k'}. \quad (24)$$

In particular, since $M \succeq 0$, we have $0 \leq \mu_0 \leq \lambda_{\lfloor (np-k')/p \rfloor}$. It remains to determine k' .

From Section 2.1, recall that $k = np - md(d+1)/2$. We now investigate how many new dimensions V adds to U . All matrices $R \in \mathbb{R}^{p \times p}$ admit a unique decomposition as $R = R_{\text{skew}} + R_{\ker \mathcal{L}} + R_{(\ker \mathcal{L})^\perp}$, where R_{skew} is skew-symmetric, $R_{\ker \mathcal{L}}$ is in the kernel of \mathcal{L}_X (14) and $R_{(\ker \mathcal{L})^\perp}$ is in the orthogonal complement of the latter in $\mathbb{S}^{p \times p}$. Clearly, YR_{skew} and $YR_{\ker \mathcal{L}}$ are tangent vectors, thus vectorized versions of these are already in the span of U . On the other hand, by definition, $YR_{(\ker \mathcal{L})^\perp}$ is not tangent at Y (if it is nonzero). This raises k' (the rank of W) to:

$$k' = k + p^2 - \frac{p(p-1)}{2} - \dim \mathcal{F}_X = np - m \frac{d(d+1)}{2} + \frac{p(p+1)}{2} - \dim \mathcal{F}_X. \quad (25)$$

Combine with $\lambda_{\lfloor (np-k')/p \rfloor} \geq 0$ and the definition of Δ (15) to conclude. \square

Theorem 3.14 is particularly meaningful for linear f , considering the intuition that for $p \geq p^*$, generically, $\dim \mathcal{F}_X = \Delta \geq 0$ (15) (Theorem B.1 gives a proof for $d = 1$). Thus, for such p , a second-order critical point is either globally optimal, or it maps to a face of abnormally high dimension, over which f must be constant and suboptimal. We could not produce an example of the latter. We summarize this in a corollary, followed by a question.

Corollary 3.15. *Assume f is linear and fix $p > p^*$, hence $\Delta > 0$. If $Y \in \text{St}(d, p)^m$ is a second-order critical point for (RP_p) but $X = YY^\top$ is not a global optimizer for (P), then $\text{rank}(X) = p$, $\dim \mathcal{F}_X \geq \Delta + p$ and f is constant over \mathcal{F}_X .*

Proof. Use Corollary 3.13, Theorem 3.14 and Theorem 3.8. \square

For linear $f(X) = \langle C, X \rangle$ and $p > p^*$, the question is the following: if C is sampled uniformly at random from the unit-norm symmetric matrices, what is the probability that f is constant over a face \mathcal{F}_X of dimension $\Delta + p$ or larger, with $\text{rank}(X) = p$? If it is zero, then almost surely all second-order critical points of (RP_p) are global optimizers. We do not answer this question here, but refer to Theorem B.1 to argue that there are few such faces.

Theorem 3.14 is motivation to investigate upper-bounds on the dimensions of faces of \mathcal{C} . The following result extends [44, Thm. 3.1(i)] to $d \geq 1$.

Theorem 3.16. *If $X \in \mathcal{C}$ has rank p , then the face \mathcal{F}_X (13) has dimension bounded as:*

$$\frac{p(p+1)}{2} - n \frac{d+1}{2} \leq \dim \mathcal{F}_X \leq \frac{p(p+1)}{2} - p \frac{d+1}{2}. \quad (26)$$

If p is an integer multiple of d , the upperbound is attained for some X .

Proof. Inequality (15) covers the lower bound. It remains to show that $\mathcal{L}_X(A) = 0$ imposes at least $p(d+1)/2$ linearly independent constraints on $A \in \mathbb{S}^{p \times p}$. Let $Y \in \text{St}(d, p)^m$ be such that $X = YY^\top$, and let $y_1, \dots, y_n \in \mathbb{R}^p$ denote the rows of Y , transposed. Greedily select p linearly independent rows of Y , in order, such that row i is picked iff it is linearly independent from rows y_1 to y_{i-1} . This is always possible since Y has full rank. Write $t = \{t_1 < \dots < t_p\}$ to denote the indices of selected rows. Write $s_k = \{((k-1)d+1), \dots, kd\}$ to denote the indices of rows in slice Y_k , and let $c_k = s_k \cap t$ be the indices of selected rows in that slice.

For $x_1, \dots, x_p \in \mathbb{R}^p$ linearly independent, the $p(p+1)/2$ symmetric matrices $x_i x_j^\top + x_j x_i^\top$ form a basis of $\mathbb{S}^{p \times p}$ —see for example [44, Lem. 2.1]. Defining $E_{ij} = y_i y_j^\top + y_j y_i^\top = E_{ji}$, this means $\mathcal{E}_t = \{E_{t_\ell, t_{\ell'}} : \ell, \ell' = 1 \dots p\}$ forms a basis of $\mathbb{S}^{p \times p}$. Similarly, since each slice Y_k has orthonormal rows, the matrices $\{E_{ij} : i, j \in s_k\}$ are linearly independent.

The constraint $\mathcal{L}_X(A) = 0$ means $\langle A, E_{ij} \rangle = 0$ for each k and for each $i, j \in s_k$. To establish the theorem, we need to extract a subset T of at least $p(d+1)/2$ of these $md(d+1)/2$ constraint matrices, and guarantee their linear independence. To this end, let

$$T = \{E_{ij} : k \in \{1, \dots, m\} \text{ and } i \in c_k, j \in s_k\}. \quad (27)$$

That is, for each slice k , T includes all constraints of that slice which involve at least one of the selected rows. For each slice k , there are $|c_k|d - \frac{|c_k|(|c_k|-1)}{2}$ such constraints—note the correction for double-counting the E_{ij} 's where both i and j are in c_k . Thus, using $|c_1| + \dots + |c_m| = p$, the cardinality of T is:

$$|T| = \sum_{k=1}^m |c_k|d - \frac{|c_k|(|c_k|-1)}{2} = p(d+1/2) - \frac{1}{2} \sum_{k=1}^m |c_k|^2. \quad (28)$$

We first show matrices in T are linearly independent. Then, we show $|T|$ is large enough.

Consider one $E_{ij} \in T$: $i = t_\ell$ for some ℓ (otherwise, permute i and j) and $i, j \in s_k$ for some k . By construction of t , we may expand y_j in terms of the rows selected in slices 1 to k , i.e., $y_j = \sum_{\ell'=1}^{\ell_k} \alpha_{j, \ell'} y_{t_{\ell'}}$, where $\ell_k = |c_1| + \dots + |c_k|$. As a result, E_{ij} expands in the basis \mathcal{E}_t as follows: $E_{ij} = \sum_{\ell'=1}^{\ell_k} \alpha_{j, \ell'} E_{t_\ell, t_{\ell'}}$. As noted before, E_{ij} 's contributed by a same slice k are linearly independent. Furthermore, they expand in only a subset of the basis, namely:

$\mathcal{E}_t^{(k)} = \{E_{t_\ell, t_{\ell'}} : \ell_{k-1} < \ell \leq \ell_k, \ell' \leq \ell_k\}$. For $k \neq k'$, $\mathcal{E}_t^{(k)}$ and $\mathcal{E}_t^{(k')}$ are disjoint. Hence, elements of T are linearly independent.

It remains to lowerbound (28). To this effect, use $|c_k| \leq d$ to obtain:

$$\sum_{k=1}^m |c_k|^2 \leq \max_{x \in \mathbb{R}^m: \|x\|_\infty \leq d, \|x\|_1 = p} \|x\|_2^2 = \left\lfloor \frac{p}{d} \right\rfloor d^2 + \left(p - \left\lfloor \frac{p}{d} \right\rfloor d\right)^2 \leq pd.$$

Indeed, the maximum is attained by making as many of the entries of x as large as possible—this can be verified using KKT conditions. In combination with (28), this confirms at least $p(d+1/2) - pd/2 = p(d+1)/2$ linearly independent constraints act on A , thus upperbounding $\dim \mathcal{F}_X$.

To conclude, we argue that the proposed upperbound is essentially tight. Indeed, build Y by repeating m times the d first rows of I_p , then by replacing its p first rows with I_p (to ensure Y is full-rank). If p/d is an integer, then exactly the p/d first slices each contribute $d(d+1)/2$ independent constraints, i.e., $\dim \mathcal{F}_{Y Y^\top} = p(p+1)/2 - p(d+1)/2$. \square

Theorems 3.3, 3.14 and 3.16 combined give a sufficient condition on p to ensure all second-order critical points of (RP_p) correspond to KKT points of (P) .

Corollary 3.17. *Assume f is concave and $p > \frac{d+1}{d+3}n$. If $Y \in \text{St}(d, p)^m$ is a second-order critical point for (RP_p) , then $X = Y Y^\top$ is a KKT point for (P) . If furthermore f is linear, then all second-order critical points of (RP_p) are global optimizers.*

Proof. Since $\text{rank}(X) \leq p$, we have $\dim \mathcal{F}_X - \Delta \leq (n-p)\frac{d+1}{2}$. Theorem 3.14 then gives $S \succeq 0$ (21) if $(n-p)(d+1) < 2p$, which is the case. Apply Theorem 3.3 to conclude. \square

In particular, for the Max-Cut SDP ($d=1$, f linear), this shows that computing a second-order critical point of (RP_p) with $p = \lfloor n/2 \rfloor + 1$ certainly solves (P) . This is an interesting and new result, but of course, in practice, it is desirable (and empirically sufficient) to take $p = \lfloor p^* \rfloor + 1$. In the unlikely event we would encounter a “bad” second-order critical point with such p , the following theorem provides an escape route (for concave f) which does not require increasing the rank. It proceeds by moving inside a face.

Proposition 3.18 (in-face rank reduction). *Let $Y \in \text{St}(d, p)^m$ have full-rank, $X = Y Y^\top$, and consider the symmetric operator \mathcal{H} on $\mathbb{S}^{p \times p}$ defined by $\mathcal{H}(A) = Y^\top \text{symblockdiag}(Y A Y^\top) Y$. \mathcal{H} is positive semidefinite and $\dim \ker \mathcal{H} = \dim \mathcal{F}_X$. If $A \in \ker \mathcal{H}$ is nonzero, then $X' = Y(I_p - A/\lambda_{\min}(A))Y^\top \in \mathcal{F}_X$ (on the boundary) and $\text{rank}(X') \leq p-1$.*

Proof. Recall the definitions of \mathcal{F}_X (13) and \mathcal{L}_X (14). All follows from $\mathcal{H} = \mathcal{L}_X^* \mathcal{L}_X$, where \mathcal{L}_X^* is the adjoint of \mathcal{L}_X . \square

The latter proposition suggests an explicit numerical method to compute A , by computing a minimal eigenvector of \mathcal{H} . Applying \mathcal{H} costs $\mathcal{O}(m(d^2 p + p^2 d))$ flops. Assuming $p = \lfloor p^* \rfloor + 1 = \Theta(d\sqrt{m})$ and that up to $p(p+1)/2$ applications are necessary, this brings the cost of computing A to $\mathcal{O}(d^2 n^3)$ flops.

Algorithm 1 Riemannian Staircase Algorithm

```
1: Input: Integers  $d < p_1 < p_2 < \dots < p_k \leq n$ ; an initial iterate  $Y_0 \in \text{St}(d, p_1)^m$ .
2: for  $i = 1 \dots k$  do
3:    $Y_i \leftarrow \text{RIEMANNIANOPTIMIZATION}(\text{St}(d, p_i)^m, g, Y_{i-1})$   $\triangleright$  Descent to 2nd order critical
4:   if  $i = k$  or  $\text{rank}(Y_i) < p_i$  then
5:     return  $Y_i$   $\triangleright$  Theorems 3.8 and 3.9
6:   else
7:      $Y_i \leftarrow (Y_i \ 0_{n \times (p_{i+1} - p_i)})$   $\triangleright$  Augment  $Y_i$  for the next rank
8:      $Z \leftarrow \text{ESCAPEDIRECTION}(\text{St}(d, p_{i+1})^m, g, Y_i)$   $\triangleright$  Corollary 3.11 + line-search
9:      $Y_i \leftarrow \text{Retraction}_{Y_i}(Z)$   $\triangleright$  Eq. (12)
10:  end if
11: end for
```

4 The Riemannian staircase algorithm

The above results suggest a simple algorithm to compute KKT points of (P): for some small value of $p \geq d + 1$, find a second-order critical point Y of (RP_p) . If Y is rank deficient, then Theorem 3.8 guarantees $X = YY^\top$ is KKT for (P). Otherwise, increase p and find a second-order critical point of (RP_{p+}) , possibly warm-starting as suggested by Corollary 3.11. Iterating this procedure, the worst-case scenario is when p increases all the way to n , in which case any second-order critical point of (RP_n) yields a KKT point of (P), as per Theorem 3.9. Specific results pertaining to classes of functions f limit how large p could grow. We call this the *Riemannian Staircase*, listed as Algorithm 1. Of course, the hope is that the algorithm returns for some small p , and in practice we find that it is often sufficient to take p just above the rank of a solution.

Algorithm 1 assumes availability of a procedure $\text{RIEMANNIANOPTIMIZATION}(\mathcal{M}, g, Y_0)$, which returns a second-order critical point of $g: \mathcal{M} \rightarrow \mathbb{R}$, with cost at most $g(Y_0)$. This assumption is discussed below.

Inside the *else*-block, the augmented Y_i (with additional columns of zeros) is (usually) a saddle point. Although the second-order procedure $\text{RIEMANNIANOPTIMIZATION}$ should be able to escape it, we make this step explicit via the procedure ESCAPEDIRECTION . The latter can be implemented using Corollary 3.11, which indicates how computing an eigenvector of S (21) associated to its smallest eigenvalue, combined with a line-search, allows to escape the saddle with strict cost decrease (unless that eigenvalue is nonnegative, in which case $Z = 0$ and Y_i is returned with $Y_i Y_i^\top$ being KKT).

For all sufficiently smooth f , taking $p_k = n$ guarantees Algorithm 1 returns Y such that YY^\top is a KKT point. For convex f , KKT points may have arbitrary rank, so that allowing large p_k seems necessary in general. For strongly concave f , it is sufficient to take $p_k = \lfloor p^* \rfloor + 1$ (Corollary 3.13); for concave (and linear) f , it is sufficient to take $p_k = \lfloor \frac{d+1}{d+3} n \rfloor + 1$ (Corollary 3.17), and it is expected that $p_k = \lfloor p^* \rfloor + 1$ should be sufficient (Corollary 3.15 and discussion).

In the latter case, in the unlikely event that Algorithm 1 terminates with Y of size $n \times p_k$, $p_k \geq \lfloor p^* \rfloor + 1$, full-rank and second-order critical such that $X = YY^\top$ is *not* a KKT point of (P), it is possible to further optimize without increasing the rank. Indeed, since $\dim \mathcal{F}_X > 0$, Proposition 3.18 shows how to compute Y' such that $X' = Y'(Y')^\top$ is on the boundary of \mathcal{F}_X .

Since f is concave, $f(X') \leq f(X)$ (Lemma 3.12). Y' is critical and rank-deficient. If Y' is second-order critical, X' is KKT. Otherwise, Theorem 3.10 shows how to escape with a strict cost decrease. Iterating this procedure as needed, the cost decreases strictly (no cycling), and the rank p never exceeds p_k . We expect this procedure to terminate since (P) admits KKT points of rank at most $\lfloor p^* \rfloor$, but we do not prove this.

In practice, for the RIEMANNIANOPTIMIZATION procedure, we use the Riemannian trust-region method (RTR) [3], through the Manopt toolbox [17]. RTR is a descent method. It converges toward critical points regardless of the initial iterate (global convergence).⁹ Furthermore, the stable fixed points of RTR are local optimizers, thus making convergence to points which are not second-order critical unlikely (but not impossible). Should this happen, Theorem 3.10 shows how to escape. Admittedly, it is unclear how many times this might have to be repeated in the worst case.

Ideally, one would modify the RTR algorithm itself to ensure global convergence to second-order critical points, hopefully in a polynomial number of steps. To the best of our knowledge, algorithms with such properties have not yet been described in the Riemannian setting. Nevertheless, we are hopeful that this should be possible, in the light of recent work by Cartis et al. [23]. These authors indeed describe a modification of the classical trust-region method (for the unconstrained case) and guarantee polynomial-time convergence to approximate second-order critical points. Given how closely the classical convergence theory for the trust-region method has transferred to the Riemannian setting so far, it seems reasonable to expect this will too. Encouragingly, Sun et al. [71] achieved a strong result in this vein for dictionary learning with RTR on a sphere. This is an interesting research venue of its own.

RTR terminates once the norm of $\text{grad } g$ drops below a certain threshold. Thus, the returned Y is not exactly a critical point, and as a result it is not, in general, exactly rank deficient even when it should be. Numerically, we declare rank-deficiency when the condition number of $Y^\top Y$ exceeds some large threshold (say, 10^{10}). In future work, it would be interesting to relate *approximate* rank-deficient, second-order critical points Y to approximate KKT points YY^\top , i.e., to lowerbound the smallest eigenvalue of $S(YY^\top)$ given bounds on the gradient norm at Y , on the smallest Hessian eigenvalue at Y , and on the condition number of Y .

If a solution of rank q is sought but the obtained solution $Y_p \in \text{St}(d, p)^m$ has rank $p > q$, one heuristic is to project Y_p to $\text{St}(d, q)^m$ with any reasonable algorithm (call it Y_q)—for example, compute the thin SVD of $Y_p = U\Sigma V^\top$, retain only the first q columns of $U\Sigma$ and orthonormalize each $d \times q$ slice. Then, run RIEMANNIANOPTIMIZATION($\text{St}(d, q)^m, g, Y_q$). This typically returns a local optimizer of the hard problem. The hope is that the detour via the higher dimensional relaxation may help avoid many bad local traps.

⁹If the local optimizers of g were isolated, we could also guarantee local convergence at a quadratic rate, but $g(Y) = g(YQ)$ for all orthogonal Q , so this is never the case. In practice though, we do observe a characteristically superlinear convergence.

5 Special case: linear cost function

In the important special case where f is a linear function $f(X) = \langle C, X \rangle$ for some data matrix $C \in \mathbb{S}^{n \times n}$, the convex problem (P) is an SDP. As per Remark 2.2, it is equivalent to (RP_p) as soon as $p \geq p^*$. Remarkably, for any $X \in \mathcal{C}$, we obtain a lower-bound on the optimal value of the SDP, following an idea from Burer and Monteiro [20, § 6.1].

Proposition 5.1 (bounds on the SDP value). *Let $f(X) = \langle C, X \rangle$ be linear and let f^* denote the optimal value of (P). Then, for all $X \in \mathcal{C}$,*

$$f(X) + n \cdot \lambda_{\min}(S(X)) \leq f^* \leq f(X).$$

Proof. The dual of (P) is the following SDP:

$$\max \text{trace}(C - \tilde{S}), \text{ s.t. } C - \tilde{S} \text{ is symmetric, block-diagonal, and } \tilde{S} \succeq 0.$$

The matrix $\tilde{S} = S - \lambda_{\min}(S) \cdot I_n$ for $S = S(X)$ (21) is admissible. The result follows by strong duality, owing to Slater's condition. \square

Algorithm 1 solves the SDP by optimizing g in (RP_p), whose differentials are:

$$g(Y) = \text{trace}(Y^\top CY), \quad \nabla g(Y) = 2CY, \quad \nabla^2 g(Y)[\dot{Y}] = 2C\dot{Y}.$$

As an illustrative example, we here apply Algorithm 1 and competing SDP solvers to random instances of the orthogonal synchronization problem [9]. In this setting, one wishes to estimate m orthogonal matrices Q_1, \dots, Q_m , based on noisy measurements of the relative transformations $Q_i Q_j^\top$. See the introduction for applications.

In this benchmark, for increasing values of m , target matrices of size $d = 3$ are generated uniformly at random. The measurements of relative rotations are $H_{ij} = Q_i Q_j^\top + \sigma N_{ij}$ ($i < j$), where $\sigma = 0.3$ is the noise level and the N_{ij} 's are independent random noise matrices with i.i.d. normal entries. We also set $H_{ji} = H_{ij}^\top$ and $H_{ii} = I_d$. To estimate the Q_i 's from the H_{ij} 's, we set $C = -H/(nm)$ and solve (P). If the solution has rank d , this is equivalent to solving the maximum likelihood problem:

$$\min_{Q_1, \dots, Q_m \in \text{O}(d)} \sum_{i,j} \|H_{ij} - Q_i Q_j^\top\|_{\text{F}}^2.$$

Remarkably, for all instances generated, (P) admits a rank d solution, thus revealing the true maximum likelihood estimator: a hard quantity to compute, in general. This serendipitous phenomenon is partly explained in [10].

Figure 2 shows how much time it takes various solvers to find this solution of rank d (they all do). Algorithm 1 runs RTR once on (RP_p) with $p = d + 1$, with a random initial guess, and returns with an optimal rank d solution. We compare against interior point methods SeDuMi [70], SDPT3 [75] and Mosek [51] (the latter two via CVX [41]) as well as against SDPLR [19] with and without forcing the search rank to $d + 1$ (the forced version is labeled SDPLR*). We also depict how much time it takes to simply compute the top d eigenvectors of H , which, after projection, reveal an (empirically) equally good estimator for this problem, but with weaker guarantees [66, 9].




Figure 2: All methods solve (P) on problems from Section 5. The proposed staircase algorithm is the only one to return a solution which satisfies the constraints up to machine precision. It also returns the solution Y which is numerically closest to be of rank d . Its computational cost seems to grow at the same rate as that of merely computing d dominant eigenvectors of the data matrix (EIG), thus outperforming interior point methods as well as SDPLR.

6 Special case: Pseudo-Huber loss cost function

In the previous section, orthogonal synchronization is considered with Gaussian noise on the relative measurements. Maximum likelihood estimation then naturally leads to the minimization of a quadratic cost in Y , which simplifies to a linear cost in X .

When the relative measurements H_{ij} include outliers, least-squares are not expected to perform well. As an alternative, Wang and Singer [79] minimize a sum of *unsquared* errors, that is, they estimate the orthogonal matrices Q_i as the minimizers of $\sum_{i,j} \|H_{ij} - Q_i Q_j^\top\|_F$. A convex relaxation akin to the one from the previous section leads to solving (P) with the *least unsquared deviations* cost $f(X) = \sum_{i,j} \|H_{ij} - X_{ij}\|_F$ (LUD). This is similar in spirit to the convex relaxation for robust subspace estimation presented in [45]. The authors show that rounding the solutions of this convex program yields a good estimator, even if up to a (random) half of the data is random. In that regime, if the non-outliers are noiseless, solving the convex program achieves perfect recovery with high probability. They solve the problem with an alternating direction augmented Lagrangian method (ADM).

Tools in this paper do not directly apply to the LUD cost, because it is nonsmooth.¹⁰ Unfortunately, in our experiments we also found that smoothing the LUD cost typically leads to higher rank solutions, at a significant computational premium. We formalize this observation in the following theorem. The assumptions on H are not restrictive: they require just the slightest inconsistency in the measurements.

¹⁰Recent work on nonsmooth optimization on manifolds [43] may prove useful in this regard.

Theorem 6.1 (smoothing the LUD cost suppresses rank d solutions). *Let $\ell: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be an increasing function (with $\ell'(x) > 0$ if $x > 0$) such that $f: \mathbb{S}^{n \times n} \rightarrow \mathbb{R}$ defined by $f(X) = \sum_{i,j} \ell(\|X_{ij} - H_{ij}\|_F)$ is twice continuously differentiable, where each $H_{ij} = H_{ji}^\top$ verifies $\|H_{ij}\|_{\text{op}} \leq 1$ (which includes orthogonal matrices), and $H_{ii} = I_d$. If H is not a rank- d matrix in \mathcal{C} , then all KKT points of (P) have rank strictly larger than d . (Otherwise, $X = H$ is the unique KKT point.)*

Proof. The gradient of f with respect to X_{ij} is given by $\nabla f(X)_{ij} = w_{ij}(X_{ij} - H_{ij})$, with $w_{ij} = w_{ji} = \ell'(\|X_{ij} - H_{ij}\|_F) / \|X_{ij} - H_{ij}\|_F > 0$ if $X_{ij} \neq H_{ij}$, and $w_{ij} = 0$ otherwise. This is well defined by assumption. For contradiction, assume X is a KKT point of (P) and $\text{rank}(X) = d$. By Theorem 3.3, S (21) is positive semidefinite. In particular, its diagonal blocks are positive semidefinite:

$$\begin{aligned} S_{ii} &= \nabla f(X)_{ii} - \text{sym} \left(\sum_j \nabla f(X)_{ij} X_{ji} \right) \\ &= \sum_{j \neq i} w_{ij} \text{sym} \left(H_{ij} X_{ij}^\top - X_{ij} X_{ij}^\top \right) \\ &= \sum_{j \neq i} w_{ij} \text{sym} \left(H_{ij} X_{ij}^\top - I_d \right) \succeq 0. \end{aligned}$$

The last equality follows from the fact that, since $\text{rank}(X) = d$, each X_{ij} is orthogonal (Proposition 2.1). Since $\|H_{ij}\|_{\text{op}} \leq 1$, each term $\text{sym}(H_{ij} X_{ij}^\top - I_d)$ is *negative* semidefinite. Thus, the S_{ii} 's are zero (simultaneously positive and negative semidefinite), showing that $S = 0$ (by Schur's complement). Hence, the off-diagonal blocks are zero too: $S_{ij} = \nabla f(X)_{ij} = w_{ij}(X_{ij} - H_{ij}) = 0$, implying $X = H$: a contradiction. \square

Remark 6.2. *The latter theorem applies in particular for $\ell(x) = x^2$. Thus, the convex problem (P) with $f(X) = \|X - H\|_F^2$ is not expected to admit rank d solutions in the presence of even the smallest noise. This is in sharp contrast with the linear case, $f(X) = -\text{trace}(HX)$, even though these two costs differ only by a constant over the rank- d feasible X 's. The key difference is that the linear cost is also concave. Generally speaking, this suggests that rank- d solutions can be obtained reliably only if either the cost is nonsmooth, or if it is concave.*

In view of these results, we take interest in minimizing the related smoothed cost:

$$g(Y) = f(YY^\top) = \sum_{i,j} \ell_\varepsilon(\|H_{ij}Y_j - Y_i\|_F), \quad \ell_\varepsilon(x) = \sqrt{x^2 + \varepsilon^2} - \varepsilon \rightarrow |x| \text{ as } \varepsilon \rightarrow 0. \quad (29)$$

It is important to note that this function f , although it bears much resemblance with the convex LUD cost (in fact, they coincide when $\text{rank}(X) = d$ and $\varepsilon = 0$), is *strongly concave* in X . Indeed, $\|H_{ij}Y_j - Y_i\|_F^2$ is affine in X_{ij} , so that f is a sum of square roots of affine functions of X , and the terms under the square roots are larger than $\varepsilon^2 > 0$ and smaller than $\varepsilon^2 + (\|H_{ij}\|_F + \sqrt{d})^2$. The following marks the dependence in X more explicitly:

$$\begin{aligned} f(X) &= \sum_{i,j} \sqrt{\|H_{ij}\|_F^2 + \|I_d\|_F^2 - 2\langle H_{ij}, X_{ij} \rangle + \varepsilon^2} - \varepsilon, \\ \nabla f(X)_{ij} &= \frac{-1}{\sqrt{\|H_{ij}\|_F^2 + \|I_d\|_F^2 - 2\langle H_{ij}, X_{ij} \rangle + \varepsilon^2}} H_{ij}. \end{aligned}$$

The only difference with a smoothed LUD cost is the term $\|I_d\|_F^2$ which appears instead of $\|X_{ij}\|_F^2$. Considering that the aim is for the X_{ij} 's to be orthogonal, which maximizes their norm by Proposition 2.3, refraining from minimizing $\|X_{ij}\|_F^2$ appears as a good start.

We may still compute a KKT point for (P), but there is no guarantee that such a point will be even a local minimizer anymore. On the bright side, Corollary 3.13 states that, by strong concavity of f , all KKT points of (P) are extreme points (thus they have rank at most $\lfloor p^* \rfloor$ (16)) and, for $p > p^*$, all second-order critical points of (RP_p) reveal KKT points of (P). The numerical experiment below shows that, empirically, even for $\varepsilon > 0$, the proposed algorithm typically converges to a rank- d KKT point of excellent quality. Furthermore, as ε is decreased, the quality of the found KKT point increases (with warm-starting).

The general observation that good performance can be achieved in practice using a concave rather than a convex cost function is also found, for example, in [47]. The price to pay, as often, is a lack of generic purpose global optimality guarantees.

We now use the proposed robust formulation of orthogonal synchronization to situations where the sought matrices are in fact permutations¹¹ (without modifying the algorithms). Synchronization of permutations notably arises in image association problems in computer vision [40, 55].

Let the Q_i 's be permutations to estimate and let the H_{ij} 's be measurements of the relative permutations $Q_i Q_j^\top$. A subset of the measurements of a given size is selected uniformly at random and replaced by uniformly random permutations (outliers). The other measurements are correct. If perfect recovery of the Q_i 's is achieved, then the permutations are recovered.

Figure 3 exhibits the perfect recovery phenomenon hinted by Wang and Singer [79]. We say “hinted” as the chosen scenario does not exactly fit the assumptions of these authors. Even in the face of many outliers, the true permutations are recovered, showing the applicability of the proposed methods to permutation estimation.

In practice, we minimize f for some starting value $\varepsilon = 1$, then re-solve for decreasing values down to $\varepsilon = 10^{-3}$, warm-starting each new solve with the previous solution. The staircase method starts with a search rank $p = d+1$. For up to 80% outliers, RTR converges to a rank- d , second-order critical point of g ($\sigma_{d+1}(Y) \approx 10^{-10}$, $\|\text{grad } g(Y)\| \leq 10^{-6}$ and $\lambda_{\min}(\text{Hess } g(Y)) \geq -10^{-10}$, after scaling) without the need to increase p , thus rapidly identifying a KKT point of (P) which appears to be a global optimizer. We compare with ADM [79] optimizing the LUD cost, but not with the IPM's, as they rapidly run out of memory.

Another idea for robust synchronization would be to use the solver for linear cost f in an iteratively reweighted least-squares scheme (IRLS), similarly to the successful approach in [45], aiming to compute the LUD solution. Preliminary experiments seem to indicate this works well. The resulting algorithm would bear some resemblance with the so-called Weiszfeld algorithm [39].

¹¹Permutation matrices are binary matrices with exactly one 1 on each row and column. They are orthogonal.




Figure 3: Synchronization of $m = 100$ permutations of size $d = 6$. As explained in Section 6, for each pair of permutations we are given a relative permutation measurement. Some fraction of those are exact: this varies on the horizontal axis. The other measurements (selected uniformly at random) are uniformly random. When the mean squared error (vertical axis) is close to zero (say, below 10^{-6}), the estimation is essentially perfect and we get back the true permutations. Remarkably, the staircase method with the pseudo-Huber loss cost (29) can accommodate up to 80% of outliers and still (empirically) achieve perfect recovery. It is faster and appears more resilient than ADM [79], but unfortunately, without access to the ground truth, we cannot claim we found a global optimum because (29) is concave. ADM, on the other hand, comes with guarantees as it solves a convex problem.

7 Conclusions and perspectives

We proposed a novel algorithm to compute KKT points for optimization problems over a class of spectrahedra that come up in relaxations of various problems involving orthonormal matrices. Our approach consists in exploiting the smooth geometry of bounded-rank subsets of those spectrahedra, to reduce the problem to Riemannian optimization. This effectively allows one to control how much lifting (dimension increase) is involved in the relaxation. An investigation of both the convex and the Riemannian geometries of the total and the bounded-rank problem showed that, under certain conditions, it is only necessary to inspect a low-dimensional portion of the boundary of the spectrahedron. Numerical experiments confirm the usefulness of this observation.

The present work triggers a number of questions for future investigation.

- Which spectrahedra are such that their elements of bounded rank form a smooth manifold? (Journée et al. [42] cover a number of such sets.) When the search space is of such form with additional constraints, can those be accommodated efficiently? This would be useful to address the SDP's in, e.g., [40, 63, 25, 11].
- What is the computational complexity of obtaining a second-order critical point of a sufficiently smooth function on a Riemannian manifold, up to a given accuracy? This might be answered by following work in [23, 71]. When Y is only approximately second-order critical and rank deficient, is YY^\top approximately KKT, in a certain sense? For linear f , Proposition 5.1 offers a positive answer.
- RTR naturally allows for preconditioning. Could this be useful for specific applications?
- For nonconvex f , the set of KKT points includes the local optimizers of (P), as well as a number of uninteresting points. All KKT points give rise to critical points, but not necessarily second-order critical points. Can this be used to improve guarantees? Could we compute second-order KKT points instead, thus possibly excluding even more spurious points? A starting point might be [60, Thm. 3.45] and [65].
- Assuming linear f , if (P) admits a unique solution of rank r (see, e.g., [10]), is it sufficient to explore (RP_p) with $p = r + 1$? Under the noise model of Section 5, this appears to be true with high probability. Perhaps, this could be investigated via the expected size of the attraction basin of the global optimizers, similarly to [71] in the context of dictionary learning.
- Finally, regarding Corollary 3.15 and its attached question: for a random linear cost function and $p > p^*$, what is the probability that (RP_p) admits second-order critical points which are not global optimizers?

Acknowledgments

The author thanks P.-A. Absil, A. d'Aspremont, A. Bandeira, X. Cheng, B. Gerencsér, Y. Khoo, B. Mishra, A. Singer and B. Vandereycken for fruitful discussions. Parts of this research were conducted while N.B. was a research fellow with the FNRS in Belgium, and while generously supported by a Research in Paris grant and the Chaire Havas “Chaire Economie et gestion des nouvelles données”.

References

- [1] E. Abbe, A.S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *arXiv preprint arXiv:1405.3267*, 2014.
- [2] P.-A. Absil and J. Malick. Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012. doi:[10.1137/100802529](https://doi.org/10.1137/100802529).
- [3] P.-A. Absil, C. G. Baker, and K. A. Gallivan. Trust-region methods on Riemannian manifolds. *Foundations of Computational Mathematics*, 7(3):303–330, 2007. doi:[10.1007/s10208-005-0179-9](https://doi.org/10.1007/s10208-005-0179-9).
- [4] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. ISBN 978-0-691-13298-3.
- [5] F. Alizadeh, J.-P. Haeberly, and M.L. Overton. Complementarity and nondegeneracy in semidefinite programming. *Mathematical Programming*, 77(1):111–128, 1997. doi:[10.1007/BF02614432](https://doi.org/10.1007/BF02614432).
- [6] D. Amelunxen, M. Lotz, M.B. McCoy, and J.A. Tropp. Living on the edge: phase transitions in convex programs with random data. *Information and Inference*, 2014. doi:[10.1093/imaiai/iau005](https://doi.org/10.1093/imaiai/iau005).
- [7] M. Arie-Nachimson, S.Z. Kovalsky, I. Kemelmacher-Shlizerman, A. Singer, and R. Basri. Global motion estimation from point matches. In *Proceedings of the 2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, 3DIMPVT '12, pages 81–88, Washington, DC, USA, 2012. IEEE Computer Society. doi:[10.1109/3DIMPVT.2012.46](https://doi.org/10.1109/3DIMPVT.2012.46).
- [8] F. Arrigoni, L. Magri, B. Rossi, P. Fragneto, and A. Fusiello. Robust absolute rotation estimation via low-rank and sparse matrix decomposition. In *3D Vision (3DV), 2014 2nd International Conference on*, volume 1, pages 491–498, 2014. doi:[10.1109/3DV.2014.48](https://doi.org/10.1109/3DV.2014.48).
- [9] A.S. Bandeira, C. Kennedy, and A. Singer. Approximating the little Grothendieck problem over the orthogonal and unitary groups. *arXiv preprint arXiv:1308.5207*, 2013.
- [10] A.S. Bandeira, N. Boumal, and A. Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *arXiv preprint arXiv:1411.3272*, 2014.
- [11] A.S. Bandeira, Y. Chen, and A. Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *arXiv preprint arXiv:1505.03840*, 2015.
- [12] A.I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(1):189–202, 1995. doi:[10.1007/BF02574037](https://doi.org/10.1007/BF02574037).
- [13] A. Beck. Quadratic matrix programming. *SIAM Journal on Optimization*, 17(4):1224–1238, 2007. doi:[10.1137/05064816X](https://doi.org/10.1137/05064816X).
- [14] D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1995.
- [15] F. Bonarrigo and A. Signoroni. An enhanced ‘optimization-on-a-manifold’ framework for global registration of 3D range data. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on*, pages 350–357, 2011. doi:[10.1109/3DIMPVT.2011.51](https://doi.org/10.1109/3DIMPVT.2011.51).
- [16] N. Boumal, A. Singer, and P.-A. Absil. Robust estimation of rotations from relative measurements by maximum likelihood. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 1156–1161, Dec 2013. doi:[10.1109/CDC.2013.6760038](https://doi.org/10.1109/CDC.2013.6760038).

- [17] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014. URL <http://www.manopt.org>.
- [18] G. Bourmaud, R. Mégret, A. Giremus, and Y. Berthoumieu. Global motion estimation from relative measurements in the presence of outliers. In D. Cremers, I. Reid, H. Saito, and M.-H. Yang, editors, *Computer Vision – ACCV 2014*, volume 9007 of *Lecture Notes in Computer Science*, pages 366–381. Springer International Publishing, 2015. doi:[10.1007/978-3-319-16814-2_24](https://doi.org/10.1007/978-3-319-16814-2_24).
- [19] S. Burer and R.D.C. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003. doi:[10.1007/s10107-002-0352-8](https://doi.org/10.1007/s10107-002-0352-8).
- [20] S. Burer and R.D.C. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, 2005.
- [21] G. Calafiore, L. Carlone, and F. Dellaert. Pose graph optimization in the complex domain: Lagrangian duality, conditions for zero duality gap, and optimal solutions. *arXiv preprint arXiv:1505.03437*, 2015.
- [22] L. Carlone, R. Tron, K. Daniilidis, and F. Dellaert. Initialization techniques for 3D SLAM: a survey on rotation estimation and its use in pose graph optimization. 2015. URL <http://www.cc.gatech.edu/~dellaert/pub/Carlone15icra.pdf>.
- [23] C. Cartis, N.I.M. Gould, and P.L. Toint. Complexity bounds for second-order optimality in unconstrained optimization. *Journal of Complexity*, 28(1):93–108, 2012. doi:[10.1016/j.jco.2011.06.001](https://doi.org/10.1016/j.jco.2011.06.001).
- [24] K.N. Chaudhury, Y. Khoo, A. Singer, and D. Cowburn. Global registration of multiple point clouds using semidefinite programming. *Arxiv preprint arXiv:1306.5226*, 2013.
- [25] Y. Chen, L.J. Guibas, and Q.X. Huang. Near-optimal joint object matching via convex relaxation. *arXiv preprint arXiv:1402.1473*, 2014.
- [26] H. Cohn and A. Kumar. Universally optimal distribution of points on spheres. *Journal of the American Mathematical Society*, 20(1):99–148, 2007.
- [27] M. Cucuringu. Synchronization over \mathbb{Z}_2 and community detection in bipartite multiplex networks. Accessed on the author’s personal home page, 2013.
- [28] M. Cucuringu. Sync-Rank: Robust ranking, constrained ranking and rank aggregation via eigenvector and semidefinite programming synchronization. *arXiv preprint arXiv:1504.01070*, 2015.
- [29] M. Cucuringu, Y. Lipman, and A. Singer. Sensor network localization by eigenvector synchronization over the Euclidean group. *ACM Transactions on Sensor Networks*, 8(3):19:1–19:42, 2012.
- [30] M. Cucuringu, A. Singer, and D. Cowburn. Eigenvector synchronization, graph rigidity and the molecule problem. *Information and Inference, a Journal of the IMA*, 1(1):21–67, 2012. doi:[10.1093/imaiai/ias002](https://doi.org/10.1093/imaiai/ias002).
- [31] S. Das and H. Vikalo. SDhaP: haplotype assembly for diploids and polyploids via semidefinite programming. *BMC genomics*, 16(1):260, 2015. doi:[10.1186/s12864-015-1408-5](https://doi.org/10.1186/s12864-015-1408-5).
- [32] C. De Sa, K. Olukotun, and C. Ré. Global convergence of stochastic gradient descent for some nonconvex matrix problems. *arXiv preprint arXiv:1411.1134*, 2014.
- [33] L. Demanet and V. Jugnon. Convex recovery from interferometric measurements. *arXiv preprint arXiv:1307.6864*, 2013.

- [34] A. Edelman, T.A. Arias, and S.T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.
- [35] F. Fogel, I. Waldspurger, and A. d’Aspremont. Phase retrieval for imaging problems. *arXiv preprint arXiv:1304.7735*, 2013.
- [36] R. Frostig, S. Wang, P.S. Liang, and C.D. Manning. Simple map inference via low-rank relaxations. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3077–3085. Curran Associates, Inc., 2014.
- [37] M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995. doi:[10.1145/227683.227684](https://doi.org/10.1145/227683.227684).
- [38] I. Grubišić and R. Pietersz. Efficient rank reduction of correlation matrices. *Linear algebra and its applications*, 422(2):629–653, 2007. doi:[10.1016/j.laa.2006.11.024](https://doi.org/10.1016/j.laa.2006.11.024).
- [39] R. Hartley, K. Aftab, and J. Trumpf. L1 rotation averaging using the Weiszfeld algorithm. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3041–3048. IEEE, 2011.
- [40] Q.X. Huang and L. Guibas. Consistent shape maps via semidefinite programming. In *Computer Graphics Forum*, volume 32, pages 177–186. Wiley Online Library, 2013.
- [41] CVX Research Inc. CVX: Matlab software for disciplined convex programming, version 2.0. <http://cvxr.com/cvx>, August 2012.
- [42] M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM Journal on Optimization*, 20(5):2327–2351, 2010. doi:[10.1137/080731359](https://doi.org/10.1137/080731359).
- [43] A. Kovnatsky, K. Glashoff, and M.M. Bronstein. MADMM: a generic algorithm for non-smooth optimization on manifolds. *arXiv preprint arXiv:1505.07676*, 2015.
- [44] M. Laurent and S. Poljak. On the facial structure of the set of correlation matrices. *SIAM Journal on Matrix Analysis and Applications*, 17(3):530–547, 1996. doi:[10.1137/0617031](https://doi.org/10.1137/0617031).
- [45] G. Lerman, M.B. McCoy, J.A. Tropp, and T. Zhang. Robust computation of linear models by convex relaxation. *Foundations of Computational Mathematics*, pages 1–48, 2014. doi:[10.1007/s10208-014-9221-0](https://doi.org/10.1007/s10208-014-9221-0).
- [46] X. Liu, H. Wang, X. Chen, and Y. Yuan. On the global optimality for linear constrained rank minimization problem. Optimization Online, 2015. URL http://www.optimization-online.org/DB_HTML/2014/11/4635.html.
- [47] R. Luss and M. Teboulle. Conditional gradient algorithms for rank-one matrix approximations with a sparsity constraint. *SIAM Review*, 55(1):65–98, 2013. doi:[10.1137/110839072](https://doi.org/10.1137/110839072).
- [48] J. Malick. The spherical constraint in Boolean quadratic programs. *Journal of Global Optimization*, 39(4):609–622, 2007. doi:[10.1007/s10898-007-9161-1](https://doi.org/10.1007/s10898-007-9161-1).
- [49] M. McCoy and J.A. Tropp. Two proposals for robust PCA using semidefinite programming. *Electronic Journal of Statistics*, 5:1123–1160, 2011. doi:[10.1214/11-EJS636](https://doi.org/10.1214/11-EJS636).
- [50] B. Mishra, G. Meyer, F. Bach, and R. Sepulchre. Low-rank optimization with trace norm penalty. *SIAM Journal on Optimization*, 23(4):2124–2149, 2013. doi:[10.1137/110859646](https://doi.org/10.1137/110859646).

- [51] ApS MOSEK. The MOSEK optimization tools version 7. <http://www.mosek.com/>, 2014.
- [52] A. Naor, O. Regev, and T. Vidick. Efficient rounding for the noncommutative Grothendieck inequality. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 71–80. ACM, 2013.
- [53] A. Nemirovski. Sums of random symmetric matrices and quadratic optimization under orthogonality constraints. *Mathematical Programming*, 109(2–3):283–317, 2007. doi:[10.1007/s10107-006-0033-0](https://doi.org/10.1007/s10107-006-0033-0).
- [54] S. Noschese, L. Pasquini, and L. Reichel. Tridiagonal Toeplitz matrices: properties and novel applications. *Numerical Linear Algebra with Applications*, 20(2):302–326, 2013. doi:[10.1002/nla.1811](https://doi.org/10.1002/nla.1811).
- [55] D. Pachauri, R. Kondor, G. Sargur, and V. Singh. Permutation diffusion maps (PDM) with application to the image association problem in computer vision. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 541–549. Curran Associates, Inc., 2014.
- [56] G. Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of operations research*, 23(2):339–358, 1998. doi:[10.1287/moor.23.2.339](https://doi.org/10.1287/moor.23.2.339).
- [57] J.R. Peters, D. Borra, B.E. Paden, and F. Bullo. Sensor network localization on the group of 3D displacements. 2014. URL http://engineering.ucsb.edu/~jrpeters/sensor_network_localization.pdf.
- [58] M. Ramana and A.J. Goldman. Some geometric results in semidefinite programming. *Journal of Global Optimization*, 7(1):33–50, 1995. doi:[10.1007/BF01100204](https://doi.org/10.1007/BF01100204).
- [59] R.T. Rockafellar. *Convex analysis*. Princeton University Press, Princeton, NJ, 1970.
- [60] A.P. Ruszczyński. *Nonlinear optimization*. Princeton University Press, Princeton, NJ, 2006.
- [61] G. Sagnol. A class of semidefinite programs with rank-one solutions. *Linear Algebra and its Applications*, 435(6):1446–1463, 2011. doi:[10.1016/j.laa.2011.03.027](https://doi.org/10.1016/j.laa.2011.03.027).
- [62] J. Saunderson, P.A. Parrilo, and A.S. Willsky. Semidefinite descriptions of the convex hull of rotation matrices. *arXiv preprint arXiv:1403.4914*, 2014.
- [63] J. Saunderson, P.A. Parrilo, and A.S. Willsky. Semidefinite relaxations for optimization problems over rotation matrices. In *Proc. 53rd IEEE Conference on Decision and Control*, 2014.
- [64] R. Schneider and A. Uschmajew. Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality. *SIAM Journal on Optimization*, 25(1):622–646, 2015. doi:[10.1137/140957822](https://doi.org/10.1137/140957822).
- [65] A. Shapiro. First and second order analysis of nonlinear semidefinite programs. *Mathematical Programming*, 77(1):301–320, 1997. doi:[10.1007/BF02614439](https://doi.org/10.1007/BF02614439).
- [66] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011. doi:[10.1016/j.acha.2010.02.001](https://doi.org/10.1016/j.acha.2010.02.001).
- [67] A. Singer and Y. Shkolnisky. Three-Dimensional Structure Determination from Common Lines in Cryo-EM by Eigenvectors and Semidefinite Programming. *SIAM Journal on Imaging Sciences*, 4(2):543–572, 2011. doi:[10.1137/090767777](https://doi.org/10.1137/090767777).

- [68] A.M.C. So. Improved approximation bound for quadratic optimization problems with orthogonality constraints. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1201–1209. Society for Industrial and Applied Mathematics, 2009.
- [69] S. Sojoudi and J. Lavaei. Exactness of semidefinite relaxations for nonlinear optimization problems with underlying graph structure. *SIAM Journal on Optimization*, 24(4):1746–1778, 2014. doi:[10.1137/130915261](https://doi.org/10.1137/130915261).
- [70] J.F. Sturm. Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11(1–4):625–653, 1999. doi:[10.1080/10556789908805766](https://doi.org/10.1080/10556789908805766).
- [71] J. Sun, Q. Qu, and J. Wright. Complete dictionary recovery over the sphere. *arXiv preprint arXiv:1504.06785*, 2015.
- [72] P. Tarazaga and M.W. Trosset. An optimization problem on subsets of the symmetric positive-semidefinite matrices. *Journal of Optimization Theory and Applications*, 79(3):513–524, 1993.
- [73] J.M.F. ten Berge. Orthogonal Procrustes rotation for two or more matrices. *Psychometrika*, 42(2):267–276, 1977. doi:[10.1007/BF02294053](https://doi.org/10.1007/BF02294053).
- [74] J.M.F. ten Berge. Generalized approaches to the MaxBet problem and the MaxDiff problem, with applications to canonical correlations. *Psychometrika*, 53(4):487–494, 1988. doi:[10.1007/BF02294402](https://doi.org/10.1007/BF02294402).
- [75] K.C. Toh, M.J. Todd, and R.H. Tütüncü. SDPT3—a MATLAB software package for semidefinite programming. *Optimization Methods and Software*, 11(1–4):545–581, 1999. doi:[10.1080/10556789908805762](https://doi.org/10.1080/10556789908805762).
- [76] S. Tu and J. Wang. Practical first order methods for large scale semidefinite programming. 2014. URL <http://www.cs.berkeley.edu/~stephentu/writeups/first-order-sdp.pdf>.
- [77] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996.
- [78] I. Waldspurger, A. d’Aspremont, and S. Mallat. Phase recovery, MaxCut and complex semidefinite programming. *Mathematical Programming*, 149(1–2):47–81, 2015. doi:[10.1007/s10107-013-0738-9](https://doi.org/10.1007/s10107-013-0738-9).
- [79] L. Wang and A. Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference*, 2(2):145–193, 2013. doi:[10.1093/imaiai/iat005](https://doi.org/10.1093/imaiai/iat005).
- [80] L. Wang, A. Singer, and Z. Wen. Orientation determination of cryo-em images using least unsquared deviations. *SIAM Journal on Imaging Sciences*, 6(4):2450–2483, 2013. doi:[10.1137/130916436](https://doi.org/10.1137/130916436).
- [81] Z. Wen and W. Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 142(1–2):397–434, 2013. doi:[10.1007/s10107-012-0584-1](https://doi.org/10.1007/s10107-012-0584-1).
- [82] R.C. Wilson and E.R. Hancock. Spherical embedding and classification. In E.R. Hancock, R.C. Wilson, T. Windeatt, I. Ulusoy, and F. Escolano, editors, *Structural, Syntactic, and Statistical Pattern Recognition*, volume 6218 of *Lecture Notes in Computer Science*, pages 589–599. Springer Berlin Heidelberg, 2010. doi:[10.1007/978-3-642-14980-1_58](https://doi.org/10.1007/978-3-642-14980-1_58).
- [83] W.H. Yang, L.-H. Zhang, and R. Song. Optimality conditions for the nonlinear programming problems on Riemannian manifolds. Optimization Online, 2012.

A Tightness for synchronization on a cycle

Consider synchronization on a cycle: the goal is to estimate orthogonal matrices $R_1, \dots, R_m \in O(d)$ based on one cycle of measurements: $H_{1,2} \approx R_1 R_2^\top, H_{2,3} \approx R_2 R_3^\top, \dots, H_{m,1} \approx R_m R_1^\top$. This is achieved by solving (RP_p) with $p = d$ and

$$g(Y) = \sum_{i=1}^m \|Y_i Y_{i+1}^\top - H_{i,i+1}\|_F^2 = \sum_{i=1}^m \langle Y_i Y_{i+1}^\top - H_{i,i+1} \rangle + \text{constant},$$

with the indexing convention that $Y_{m+1} \equiv Y_1$ and $H_{m,m+1} \equiv H_{m,1}$. Under a permissive condition on the measurements, Peters et al. [57] exhibit an explicit formula for the solution (they restrict their attention to rotation matrices, that is, orthogonal matrices of determinant +1). We show that under that same condition, the corresponding SDP relaxation (P) with

$$f(X) = \langle C, X \rangle, \quad C = - \begin{pmatrix} 0 & H_{1,2} & & & H_{m,1}^\top \\ H_{1,2}^\top & 0 & H_{2,3} & & \\ & H_{2,3}^\top & \ddots & \ddots & \\ & & \ddots & \ddots & H_{m-1,m}^\top \\ H_{m,1} & & & H_{m-1,m}^\top & 0 \end{pmatrix} \quad (30)$$

is tight: there exists a unique solution of rank d which reveals the global optimum.

The proof rests on two key ingredients: (a) we have an explicit formula for the solution X to certify, and (b) we have an explicit formula for the dual certificate $S(X)$ to check. It seems reasonable to expect that the result should carry over to connected graphs whose cycles have disjoint edges.

Theorem A.1. *Let $H_{1,2}, H_{2,3}, \dots, H_{m,1} \in O(d)$ represent orthogonal measurements on a cycle ($m \geq 3$) and define their product $P = H_{1,2} \cdot H_{2,3} \cdots H_{m,1} \in O(d)$. If -1 is not an eigenvalue of P , then the semidefinite program (P) with cost (30) admits a unique solution of rank d .*

Proof. Part 1: guessing X . When the measurements are perfectly consistent, $P = I_d$, and it is easy to construct X : set $Y_m = I_d$ and $Y_i = H_{i,i+1} Y_{i+1}$ for $i = 1 \dots m-1$; then $X_{ij} = Y_i Y_j^\top$. This construction does not use $H_{m,1}$ but still achieves $g(Y) = 0$ owing to $P = I_d$, hence X is optimal. When the cycle is inconsistent, it is reasonable to guess that the least-squares criterion will attempt to spread the inconsistency evenly over each edge [57]. One m th of the error is represented by an offset $P^{1/m}$ —taking this principal matrix root requires P not to have negative eigenvalues. We build Y by incorporating part of the error at each step, appropriately aligned. First define this recurrence: $Q_m = H_{m,1}$ and $Q_i = H_{i,i+1} Q_{i+1}$ for $i = (m-1) \dots 1$ (note that $Q_1 = P$). Then, $Y_m = I_d$ and $Y_i = Q_i P^{-1/m} Q_i^\top H_{i,i+1} Y_{i+1}$ for $i = (m-1) \dots 1$. As previously, $X_{ij} = Y_i Y_j^\top$. It is not hard to check that $X_{ij} = Q_i P^{(i-j)/m} Q_j^\top$. Of course, X is admissible for (P) with rank d .

Part 2: certifying X . By Theorem 3.5, it is sufficient to verify that $S(X)$ (21) is positive semidefinite with rank $(m-1)d$. Let U be a $d \times d$ unitary matrix such that $D = U^* P U$ is diagonal— U always exists since P is normal—and let $V = \text{diag}(Q_1 U, \dots, Q_m U)$ be a block-

diagonal unitary matrix. We use V to operate a change of variables on C and X :

$$V^*CV = - \begin{pmatrix} 0 & I_d & & D^{-1} \\ I_d & 0 & I_d & \\ & I_d & \ddots & \ddots \\ & & \ddots & \ddots & I_d \\ D & & & I_d & 0 \end{pmatrix}, \quad (V^*XV)_{ij} = D^{(i-j)/m}. \quad (31)$$

We used that D is unitary. Indeed, since P is orthogonal without 1 as an eigenvalue, its eigenvalues are such that $D = \text{diag}(e^{i\theta_1}, \dots, e^{i\theta_d})$ for $\theta_1, \dots, \theta_d \in]-\pi, \pi[$. The spectrum of $S(X)$ is identical to that of $V^*S(X)V$, thus we study:

$$V^*S(X)V = \begin{pmatrix} D^{1/m} + D^{-1/m} & -I_d & & -D^{-1} \\ -I_d & D^{1/m} + D^{-1/m} & -I_d & \\ & -I_d & \ddots & \ddots \\ & & \ddots & \ddots & -I_d \\ -D & & & -I_d & D^{1/m} + D^{-1/m} \end{pmatrix}. \quad (32)$$

All blocks of (32) are diagonal, so that its rows and columns may be permuted (without affecting its spectrum) to make it block diagonal, with k th block of size m given by

$$A = \begin{pmatrix} 2\cos(\theta_k/m) & -1 & & -e^{-i\theta_k} \\ -1 & 2\cos(\theta_k/m) & -1 & \\ & -1 & \ddots & \ddots \\ & & \ddots & \ddots & -1 \\ -e^{i\theta_k} & & & -1 & 2\cos(\theta_k/m) \end{pmatrix} = \begin{pmatrix} T & u \\ u^* & c \end{pmatrix}, \quad (33)$$

with $T \in \mathbb{R}^{(m-1) \times (m-1)}$, $u \in \mathbb{C}^{m-1}$ and $c = 2\cos(\theta_k/m)$. It remains to show that A is positive semidefinite with rank $m-1$ for any $\theta_k \in]-\pi, \pi[$. Fortunately, T is tridiagonal and Toeplitz, so that its whole spectrum is known explicitly [54]: $\lambda_j(T) = 2(\cos(\theta_k/m) - \cos(j\pi/m))$, for $j = 1 \dots m-1$. These eigenvalues are all positive. By the Cauchy interlacing theorem,

$$\lambda_1(A) \leq \lambda_1(T) \leq \lambda_2(A) \leq \lambda_2(T) \leq \dots \leq \lambda_{m-1}(T) \leq \lambda_m(A). \quad (34)$$

In particular, $\lambda_2(A), \dots, \lambda_m(A) > 0$. Since the vector $[e^{i(1/m)\theta_k}, e^{i(2/m)\theta_k}, \dots, e^{i(m/m)\theta_k}]^*$ is in the kernel of A , it must be that $\lambda_1(A) = 0$. This concludes the proof. \square

In general, the condition on the eigenvalues of P is necessary. Indeed, for $d = 1$ and $m = 3$, choose the measurements such that $P = -1$ (for example, $+1, +1$ and -1) and verify that none of the 4 admissible rank-1 matrices are optimal. For the frequent case where the measurements are rotation matrices (that is, orthogonal with determinant $+1$), the condition is not too restrictive: even if they were distributed uniformly at random, P would satisfy the condition almost surely.

B Generic face dimension

For $d = 1$, the following theorem shows almost all faces of \mathcal{C} have minimal dimension, as per the bound (15) [56, 12]. Key parts of the proof are due to Xiuyuan Cheng and Balázs Gerencsér.

Theorem B.1 (Generically, faces have minimal dimension). *For $d = 1$, if $Y \in \text{St}(d, p)^m$ is selected uniformly at random, then, almost surely, $\dim \mathcal{F}_{Y Y^\top} = \max(0, \Delta)$, with $\Delta = \frac{p(p+1)}{2} - m \frac{d(d+1)}{2}$ (13).*

We first provide a useful lemma.

Lemma B.2. *Let v_1, \dots, v_n be statistically independent random vectors in a vector space V of dimension k . If for all i and for all subspaces $U \subset V$ with $\dim U < k$, $\Pr[v_i \in U] = 0$, then, almost surely, $\dim \text{span}\{v_1, \dots, v_n\} = \min(n, k)$ (which is maximal).*

Proof. The proof is by recurrence. Define $U_t = \text{span}\{v_1, \dots, v_t\}$ for $t \in \{0, \dots, n\}$. Clearly, $\dim U_0 = 0$. Assume $\dim U_t = \min(t, k)$ almost surely (a.s.). If $t \geq k$, then $\dim U_{t+1} = \dim U_t = k$ (a.s.). Otherwise, since U_t is statistically independent from v_{t+1} , by assumption, $\dim U_{t+1} = t + 1$ (a.s.). Thus, for all t , $\dim U_t = \min(t, k)$ (a.s.). \square

Proof. Proof of Theorem B.1. Let $y_1, \dots, y_n \in \mathbb{R}^p$ denote the columns of Y^\top . A matrix $\dot{X} = Y A Y^\top$ with $A \in \mathbb{S}^{p \times p}$ is parallel to the face \mathcal{F}_X if $\langle Y A Y^\top, e_i e_i^\top \rangle = \langle A, y_i y_i^\top \rangle = 0$ for all i . We study the dimension s of the space spanned by the constraint matrices $A_i = y_i y_i^\top$, since $\dim \mathcal{F}_X = \frac{p(p+1)}{2} - s$. We do so with Y taken uniformly at random. Thus, the y_i 's are sampled independently from $\mathcal{N}(0, I_p)$, then scaled to unit norm. The dimension s does not depend on the scaling of the vectors y_i , so we may safely ignore it. We note in passing that the proof holds for more general distributions too.

We aim to apply Lemma B.2 with $V = \mathbb{S}^{p \times p}$, $k = \frac{p(p+1)}{2}$ and $v_i = A_i$. The v_i 's are i.i.d., hence we omit the subscripts. To verify the lemma's condition, let U be any proper subspace of $\mathbb{S}^{p \times p}$: there exists a symmetric matrix $X \neq 0$ in the orthogonal complement of U . It suffices to check that $\langle X, y y^\top \rangle = y^\top X y \neq 0$ (a.s.), with $y \sim \mathcal{N}(0, I_p)$. Diagonalize $X = Q D Q^\top$ with $Q \in \text{O}(p)$. Notice that $Q^\top y$ is distributed identically to y . Thus, defining $D = \text{diag}(\lambda_1, \dots, \lambda_p)$ and $y = (y^1, \dots, y^p)^\top$, it suffices to check that $\sum_{j=1}^p \lambda_j (y^j)^2 \neq 0$ (a.s.). This is indeed true, since the $(y^j)^2$ are independent: their (nontrivial) linear combination has a density which is a convolution of (scaled) χ_1^2 densities: this has no point mass at zero. \square

It seems reasonable to expect that this result remains valid for $d \geq 1$, but we are missing a proof.