

Hidden convexity, optimization, and algorithms on rotation matrices

Akshay Ramachandran¹, Kevin Shu², and Alex L. Wang^{1,3}

¹Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

²George Institute of Technology, Atlanta, GA

³Purdue University, West Lafayette, IN

May 1, 2024

Abstract

This paper studies hidden convexity properties associated with constrained optimization problems over the set of rotation matrices $\text{SO}(n)$. Such problems are nonconvex due to the constraint $X \in \text{SO}(n)$. Nonetheless, we show that certain linear images of $\text{SO}(n)$ are convex, opening up the possibility for convex optimization algorithms with provable guarantees for these problems. Our main technical contributions show that any two-dimensional image of $\text{SO}(n)$ is convex and that the projection of $\text{SO}(n)$ onto its strict upper triangular entries is convex. These results allow us to construct exact convex reformulations for constrained optimization problems over $\text{SO}(n)$ with a single constraint or with constraints defined by low-rank matrices. Both of these results are maximal in a formal sense.

1 Introduction

This paper studies a general class of optimization problems over rotations and orthogonal bases. This class of problems covers applications such as the point registration problem in computer graphics [24, 25], Wahba’s problem of satellite attitude determination [29], spacecraft orientation [23], and obstacle avoidance in robotics [7]. The main goal of this paper is show that in certain cases of interest, we can produce natural *convex relaxations* that exactly recover the optimal solutions for such problems.

Recall, $\text{O}(n)$ is the set of orthogonal matrices in \mathbb{R}^n , or more explicitly,

$$\text{O}(n) := \{X \in \mathbb{R}^{n \times n} : X^\top X = I\}.$$

On the other hand, $\text{SO}(n)$ is the set of (orientation-preserving) rotations on \mathbb{R}^n , defined as

$$\text{SO}(n) := \left\{ X \in \mathbb{R}^{n \times n} : \begin{array}{l} X^\top X = I \\ \det(X) = 1 \end{array} \right\}.$$

These sets are groups under matrix multiplication, referred to as the orthogonal and special orthogonal groups, respectively.

We consider optimization problems of the form

$$\sup_{X \in \text{SO}(n)} \{\langle A, X \rangle : \mathcal{B}(X) \in \mathcal{C}\}, \tag{1}$$




Figure 1: Consider the set $S \subseteq \mathbb{R}^2$ (white region with solid boundary) and its convex hull $\text{conv}(S)$. The projection of S onto the horizontal axis is convex. On the other hand, the maximum vertical direction achieved by a point $x \in S$ between the two vertical lines differs from the maximum vertical direction achieved by a point $x \in \text{conv}(S)$ between the two vertical lines.

and their $O(n)$ counterparts. Here, the objective is a *linear* function defined by $A \in \mathbb{R}^{n \times n}$, and the constraint is defined by a *linear* operator $\mathcal{B} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ and some *convex* set $\mathcal{C} \subseteq \mathbb{R}^m$. The notation $\langle A, B \rangle$ denotes the trace inner product $\text{tr}(A^\top B)$. We will also study feasibility variants of (1) where the goal is to identify an $X \in \text{SO}(n)$ or $X \in O(n)$ satisfying $\mathcal{B}(X) \in \mathcal{C}$, or to declare that no such X exists. We give additional motivation for these problems in Section 1.1, where we discuss constrained versions of Wahba’s problem and point registration [24, 29].

Problems of the form (1) are ostensibly *nonconvex* due to the constraint $X \in \text{SO}(n)$ or $X \in O(n)$. Nevertheless, we will show that certain families of such problems admit exact convex reformulations. To achieve this, our main technical contributions show that *the images of $\text{SO}(n)$ or $O(n)$ under certain linear maps are convex*. Such results—showing that certain transformations of nonconvex sets are convex—are often referred to as hidden convexity results and enable the application of convex optimization algorithms to nonconvex problems [26, 30].

To see how such a result might be useful in solving problems of the form (1), suppose that $\mathcal{L} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{1+m}$ is the linear map

$$\mathcal{L}(X) := \begin{pmatrix} \langle A, X \rangle \\ \mathcal{B}(X) \end{pmatrix}.$$

If the image of $\text{SO}(n)$ under \mathcal{L} is convex, then we would have that $\mathcal{L}(\text{SO}(n)) = \text{conv}(\mathcal{L}(\text{SO}(n))) = \mathcal{L}(\text{conv}(\text{SO}(n)))$. Here, $\text{conv}(\cdot)$ represents the convex hull of the argument. In this case, it would then follow that

$$\sup_{X \in \text{SO}(n)} \{ \langle A, X \rangle : \mathcal{B}(X) \in \mathcal{C} \} = \sup_{X \in \text{conv}(\text{SO}(n))} \{ \langle A, X \rangle : \mathcal{B}(X) \in \mathcal{C} \}.$$

In other words, convexity of the image $\mathcal{L}(\text{SO}(n))$ implies that the convex relaxation of (1) that simply replaces $\text{SO}(n)$ with $\text{conv}(\text{SO}(n))$ is exact. This exactness is in terms of objective value, however we will also see how to numerically recover an actual optimizer in $\text{SO}(n)$ or $O(n)$ in the settings we consider. Note that the convexity of $\mathcal{B}(\text{SO}(n))$ alone does not imply exactness (see for example Figure 1). Similar results can be derived for $O(n)$ and/or feasibility variants of (1) under corresponding convexity results.

The convex hulls of both $O(n)$ and $\text{SO}(n)$ can be described via linear matrix inequalities (LMIs). The first fact is well-known while the latter fact is due to Saunderson et al. [27]. For ease of reference, we collect both facts in the following proposition.

Proposition 1 (Classical/[27]). *The convex hull of $O(n)$ is equal to the operator norm ball and can be written as*

$$\text{conv}(O(n)) = \mathbb{B}_{\text{op}}(n) = \left\{ X \in \mathbb{R}^{n \times n} : \begin{pmatrix} I & X \\ X^\top & I \end{pmatrix} \succeq 0 \right\}.$$

There exist symmetric matrices $A_{i,j} \in \mathbb{S}^{2^{n-1}}$ indexed by $(i,j) \in [n]^2$ such that

$$\text{conv}(\text{SO}(n)) = \left\{ X \in \mathbb{R}^{n \times n} : \sum_{i=1}^n \sum_{j=1}^n X_{i,j} A_{i,j} \preceq I_{2^{n-1}} \right\}.$$

In light of this fact, the convex relaxations we consider in the $O(n)$ setting can be directly solved with semidefinite programs (SDPs), assuming that \mathcal{C} is itself efficiently SDP-representable. In contrast, the convex relaxations we consider in the $\text{SO}(n)$ setting result in exponentially sized SDPs. For this reason, we also offer new algorithms for optimizing over $\text{conv}(\text{SO}(n))$ in the settings we consider that do not rely on the description of $\text{conv}(\text{SO}(n))$ given in Proposition 1.

1.1 Motivation

We first discuss *constrained* variants of Wahba’s problem. To set up Wahba’s problem, imagine that a satellite in space wants to determine its relative rotation (with respect to a reference rotation) given the observed direction of some number of far-away stars (or other objects).

Formally, we are given a set of (unit) vectors $v_1, \dots, v_k \in \mathbb{R}^3$, corresponding to the known directions of the k stars in the reference rotation, and (unit) vectors $u_1, \dots, u_k \in \mathbb{R}^3$, corresponding to the observed directions of the k stars in the satellite’s frame. Our goal is to find a rotation minimizing the observation error

$$\min_{X \in \text{SO}(3)} \sum_{i=1}^k \|Xu_i - v_i\|_2^2.$$

In [29], it was observed that this is equivalent to a linear optimization problem over $\text{SO}(3)$

$$\max_{X \in \text{SO}(3)} \left\langle \sum_{i=1}^k u_i v_i^\top, X \right\rangle,$$

and that this problem can be solved in turn using a single SVD computation. We will formalize this in Lemma 1.

Now, suppose we are given additional information about the true rotation X^* . We will incorporate this additional information as hard constraints into Wahba’s problem to get a constrained optimization problem over $\text{SO}(3)$.

For example, we may know that the true rotation X^* is within some angle, δ , of another rotation $X_0 \in \text{SO}(3)$. In this case, we would need to solve the problem¹

$$\max_{X \in \text{SO}(3)} \left\{ \left\langle \sum_{i=1}^k u_i v_i^\top, X \right\rangle : \langle X_0, X \rangle \geq 1 + 2 \cos(\delta) \right\}.$$

¹Any $X \in \text{SO}(3)$ can be thought of as a rotation of some angle δ around some axis, and this angle satisfies the equation $\text{tr}(X) = 1 + 2 \cos(\delta)$, which can be seen by considering the eigenvalues of X . Here, we say X is at an angle δ from Y if $Y^\top X$ is of angle δ .

	Hidden convexity	Algorithms
Diag. constraints	[18, Theorem 8]	Theorem 2
One constraint	Theorem 3 and Corollary 1	Theorem 4
SUT constraints	Theorem 5 and Corollaries 2 and 3	Theorem 6

Table 1: Summary of our hidden convexity results and algorithms for constrained optimization over $\text{SO}(n)$. We present accompanying results (Theorem 7) showing that our hidden convexity results are each “maximal” in certain senses (see Section 7).

Theorem 3 below implies that this problem has the same optimal value as its convex relaxation. We will additionally show that the optimum value of this problem can be efficiently computed using convex optimization techniques, even for $n \geq 3$.

As a second example, we may have a few high-fidelity observations, in which case we could introduce constraints for those observations:

$$\langle Xu_i, v_i \rangle = \langle u_i v_i^\top, X \rangle \geq \cos(\delta_i).$$

Theorem 5 below implies that feasibility problems with at most $n - 1$ such constraints and certain optimization problems over such constraints can be solved efficiently using convex optimization techniques.

High-dimensional settings, where $n \geq 3$, have found use in modeling nonlinear transformations on manifolds [25]. In this setting, the goal is to learn a nonlinear transformation mapping one manifold to another based on given point–point correspondences. Ovsjanikov et al. [25] suggest modeling this problem as that of finding an orthogonal (linear) transformation in the *space of functions* on the manifolds.² Note, this space of functions may be high dimensional. In this function space, point–point correspondence constraints or symmetry constraints can be naturally expressed as linear constraints on the linear transformation. Additional desirable properties of the nonlinear transformation can be encoded as an orthogonality constraint. Thus, this problem has a natural interpretation as an optimization problem of the form (1) over $\text{O}(n)$.

1.2 Statement of Results

Our main contributions show that certain linear images of $\text{SO}(n)$ are convex and that certain problems of the form (1) and its variants can be solved efficiently. We give an overview of these results in the order of the sections they appear in; see also Table 1. For our algorithms, we assume we are able to perform basic arithmetic operations (including taking square roots) exactly.

1.2.1 Feasibility problems on $\text{SO}(n)$ with diagonal constraints

A classical theorem of Horn, [18, Theorem 8], gives a first example of a hidden convexity result on $\text{SO}(n)$. Recall, the parity polytope is defined as

$$\text{PP}_n := \text{conv} \left\{ x \in \{\pm 1\}^n : \prod_{i=1}^n x_i = 1 \right\}.$$

²The function spaces are endowed with bases corresponding to (possibly a truncated set of) eigenfunctions of the Laplace–Beltrami operators.

Theorem 1. *Let $\text{diag} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^n$ map a square matrix to its diagonal elements, then $\text{diag}(\text{SO}(n))$ is convex (in fact, polyhedral) and its image is given by the parity polytope PP_n .*

In Appendix A, we give efficient separation and optimization oracles for PP_n . We give a definition of a separation oracle in Section 4. It follows that feasibility problems on $\text{SO}(n)$ with convex constraints on the diagonal elements, i.e., given convex $\mathcal{C} \subseteq \mathbb{R}^n$, decide the feasibility of

$$\{X \in \text{SO}(n) : \text{diag}(X) \in \mathcal{C}\}, \quad (2)$$

can be decided by testing the feasibility of $\text{PP}_n \cap \mathcal{C}$.

In Section 3, we complete this picture by showing how to construct an element of (2) given $d \in \text{PP}_n \cap \mathcal{C}$.

Theorem 2. *Given $d \in \text{PP}_n$, it is possible to construct $X \in \text{SO}(n)$ satisfying $\text{diag}(X) = d$ in time $O(n^2)$.*

1.2.2 Optimization on $\text{SO}(n)$ with one constraint

The main result of Section 4 is that the intersection of $\text{SO}(n)$ with any codimension-one affine space is connected.

Theorem 3. *Let $n \geq 3$, $A \in \mathbb{R}^{n \times n}$, and $c \in \mathbb{R}$. Then, the set $\text{SO}(n) \cap \{X \in \mathbb{R}^{n \times n} : \langle A, X \rangle = c\}$ is connected.*

A corollary of Theorem 3, which we prove in Section 4, is the following:

Corollary 1. *Let $n \geq 3$ and let $\pi : \text{SO}(n) \rightarrow \mathbb{R}^2$ be linear, then $\pi(\text{SO}(n))$ is convex.*

This follows as a set is convex if and only if its intersection with any one-dimensional affine subspace (i.e., a line) is connected. Theorem 3 may be of importance in its own right, as for some hyperplanes H , it may be the case that the set $\text{SO}(n) \cap H$ is a *connected* submanifold of $\text{SO}(n)$, in which case it may be possible to use Riemannian optimization methods to optimize functions on this set.

While these results imply that it is possible to use convex optimization techniques to solve problems of the form

$$\max_{X \in \text{SO}(n)} \{\langle A, X \rangle : \langle B, X \rangle \in [a, b]\} \quad (3)$$

(for example by replacing $\text{SO}(n)$ with $\text{conv}(\text{SO}(n))$), it is not clear how to do so *efficiently* when n is large. This is because Proposition 1 only guarantees an exponentially sized LMI representation for $\text{conv}(\text{SO}(n))$.

To address this issue, we give an efficient algorithm for problems of the form (3) based on running the ellipsoid algorithm on the two-dimensional image of $\text{SO}(n)$. It is noteworthy that because we use the ellipsoid algorithm in a constant-dimensional space, we do not face the infamously high running times of the ellipsoid method in high-dimensional spaces.

Theorem 4. *Let $n \geq 3$, $A, B \in \mathbb{R}^{n \times n}$ with $\|A\|_{\text{tr}} = \|B\|_{\text{tr}} = 1$. Here $\|\cdot\|_{\text{tr}}$ is the trace norm, defined formally in Section 2. Let X^* be the optimal solution to (3). We can compute $\langle A, X^* \rangle$ and $\langle B, X^* \rangle$ within an additive error of ϵ in time*

$$O\left(n^3 \log\left(\frac{1}{\epsilon}\right)^2\right).$$

Here, n^3 is the time complexity of computing the SVD of an $n \times n$ matrix.

Moreover, we will return $\alpha, \beta \in \mathbb{R}$ so that $|\alpha| + |\beta| = 1$ and

$$\langle \alpha A + \beta B, X^* \rangle + \epsilon \geq \max\{\langle \alpha A + \beta B, X \rangle : X \in \text{SO}(n)\}.$$

Remark 1. Let α, β denote the quantities returned in Theorem 4. While Theorem 4 does not directly return a minimizer of (3), we believe that any element of

$$\arg \max_{X \in \text{SO}(n)} \langle \alpha A + \beta B, X \rangle$$

should be a good approximation of a true minimizer under mild conditions. Such an element can be computed from $\alpha A + \beta B$ in the time of a single SVD decomposition. Analyzing this procedure is outside the scope of the current paper and we leave this question for future work.

1.2.3 Feasibility and optimization on $\text{SO}(n)$ with strictly upper triangular constraints

The last class of constraints we consider are constraints on the strictly upper triangular (SUT) entries of $X \in \text{SO}(n)$. Our main result in this direction shows that not only is the projection of $\text{SO}(n)$ onto its SUT entries convex, but furthermore, it is possible to optimize certain linear functions subject to convex constraints on the SUT entries using convex optimization.

Let $\pi_{\mathcal{T}}(X) = (X_{ij})_{i < j} \in \mathbb{R}^{\binom{n}{2}}$ denote the projection of X onto its SUT entries (i.e., those entries X_{ij} such that $i < j$). We will consider constraining the value of $\pi_{\mathcal{T}}(X)$ for $X \in \text{SO}(n)$ and then optimizing a linear function over this set. Let $A \in \mathbb{R}^{n \times n}$ and let \mathcal{C} be a nonempty closed convex subset of $\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$. Recall $\mathbb{B}_{\text{op}}(n)$ is the operator norm ball and is the same as $\text{conv}(\text{O}(n))$ by Proposition 1. We consider the problems

$$\sup_{X \in \text{SO}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}\} \tag{4}$$

$$\leq \sup_{X \in \text{O}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}\} \tag{5}$$

$$\leq \max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}\}. \tag{6}$$

The max in (6) will be justified in Section 5. Our main result on this topic is:

Theorem 5. *Let $A \in \mathbb{R}^{n \times n}$ be a diagonal matrix and let $\mathcal{C} \subseteq \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$ be a nonempty closed convex set. Then, equality holds between (5) and (6). If additionally $\det(A) \geq 0$, then equality holds between (4) to (6).*

A corollary of Theorem 5, which will be proved in Section 5, is the following:

Corollary 2. *It holds that $\pi_{\mathcal{T}}(\text{SO}(n)) = \pi_{\mathcal{T}}(\text{O}(n)) = \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$. In particular, all three sets are convex.*

We remark that any $n-1$ rank-one matrices $u_1 v_1^\top, \dots, u_{n-1} v_{n-1}^\top$ can be made strictly upper triangular simultaneously by left- and right-multiplying by $\text{SO}(n)$ matrices using Gram-Schmidt. In particular, optimization problems or feasibility problems with constraints on $\langle u_i v_i^\top, X \rangle$ are a special case of problems with SUT constraints (see Lemma 7). This holds too for any $n-1$ coordinate constraints. See Section 5.2 for a more detailed explanation.

Optimization over $\mathbb{B}_{\text{op}}(n)$ is tractable using a linearly sized SDP by Proposition 1, so the presentation of this theorem also can be turned into an efficient algorithm for performing such optimization whenever \mathcal{C} is itself efficiently SDP-representable.

We will further show strong structural results about the matrices in $\text{SO}(n)$ with fixed SUT entries. These structural results allow us to explicitly construct an optimizer of (4) given an optimizer of (6) under the assumptions of Theorem 5. They will additionally allow us to extend Theorem 5 to an approximation result for $\text{SO}(n)$ with $\det(A) < 0$. These structural results are summarized below and proven in parts throughout Section 6.

Theorem 6. *Let $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))) \subseteq \mathbb{R}^{\binom{n}{2}}$ and let $V_{\sigma} = \{X \in \text{O}(n) : \pi_{\mathcal{T}}(X) = \sigma\}$. The following assertions hold:*

1. $|V_{\sigma}| = 2^n$.
2. For each $i \in [n]$, there exist functions $\alpha_i(\sigma) < \beta_i(\sigma)$, so that $X_{i,i} \in \{\alpha_i, \beta_i\}$ for each $X \in V_{\sigma}$. We will suppress the dependence of α_i and β_i on σ for convenience of notation.
3. No two elements in V_{σ} have the same diagonal entries. That is, for each $d \in \{\alpha_1, \beta_1\} \times \{\alpha_2, \beta_2\} \times \cdots \times \{\alpha_n, \beta_n\}$, there is a unique $X \in V_{\sigma}$ so that $\text{diag}(X) = d$.
4. For each $i \in [n]$, α_i and β_i are continuous functions of σ . The function β_i is convex in σ , and the function α_i is concave in σ .
5. $X \in V_{\sigma}$ is in $\text{SO}(n)$ if and only if the number of i so that $X_{i,i} = \alpha_i$ is even.
6. Given $\rho \in \{-1, 1\}^n$, we can construct a matrix $X \in V_{\sigma}$ so that $X_{i,i} = \begin{cases} \alpha_i & \text{if } \rho_i = -1 \\ \beta_i & \text{if } \rho_i = 1 \end{cases}$ in time $O(n^3)$.

1.2.4 Obstructions to Progress

Finally, Section 7 provides constructions showing that Theorem 1, Corollary 1, and Corollary 2 above are *maximal* in specific senses. We summarize the results here:

Theorem 7. *The following assertions hold:*

1. For any $n \geq 3$, the images of $\text{SO}(n)$ under linear maps to \mathbb{R}^2 are “maximally convex” in the following sense: There exists $\pi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^2$ so that $\pi(\text{SO}(n))$ is nonconvex.
2. The projection of $\text{SO}(n)$ onto its diagonal is “maximally convex” in the following sense: For $A \in \mathbb{R}^{n \times n}$, let $\pi(X) = (X_{11}, X_{22}, \dots, X_{nn}, \langle A, X \rangle)$. If A is not itself diagonal, then $\pi(\text{SO}(n))$ is not convex.
3. For any $n \geq 3$, the projection of $\text{SO}(n)$ onto its SUT entries is “maximally convex” in the following sense: If $\pi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ is any linear map with $\text{rank}(\pi) > \binom{n}{2}$, then $\pi(\text{SO}(n))$ is not convex.
4. The assumption $\det(C) \geq 0$ in Theorem 5 is necessary in the following sense: There exists $\sigma \in \mathbb{R}^{\binom{n}{2}}$ and a diagonal matrix C so that

$$\max_{X \in \text{SO}(n)} \{\langle C, X \rangle : \pi_{\mathcal{T}}(X) = \sigma\} < \max_{X \in \text{conv}(\text{SO}(n))} \{\langle C, X \rangle : \pi_{\mathcal{T}}(X) = \sigma\}$$

This theorem is proven in parts throughout Section 7.

1.3 Related literature

Hidden convexity results are scattered throughout the literature on optimization, numerical linear algebra, and matrix analysis. We recommend the following surveys/chapters for introductions to this subject [2, 26, 30]. Along these lines, our hidden convexity results and their subsequent applications in deriving convex SDP relaxations of nonconvex problems parallel Dines’ Theorem [9] and its application in deriving the S-lemma [12], a fundamental result in control theory and nonlinear optimization.

Our results extend existing hidden convexity results related to the (special) orthogonal group. Some of the earliest work in this line is [18, Theorem 8] stating that $\text{diag}(\text{SO}(n)) = \text{PP}_n$. There are other similar results concerning the convexity of the image of $\text{SO}(n)$ under various nonlinear maps, for example the famous Schur–Horn theorem [18]. Another paper along these lines is [11], which characterizes the possible projections of $\text{SO}(n)$ onto its rectangular submatrices. In particular, it is not hard to show using their results that the projection of $\text{SO}(n)$ onto a $k \times \ell$ rectangular submatrix is convex if and only if $k + \ell \leq n$. Our results extend [11] to nonrectangular coordinate patterns. For further work in this direction, see [16, 28].

Another important piece of related work is [27], which gives a description of $\text{conv}(\text{SO}(n))$ in terms of linear matrix inequalities (LMI). This LMI description is constructed using Lie group theory applied to $\text{SO}(n)$ and is related to the fact that the fundamental group of $\text{SO}(n)$ is $\mathbb{Z}/2\mathbb{Z}$. This fact will also be crucial in our proof of Theorem 3. Inspired by techniques from [27], we can view our hidden convexity results as new quadratic convexity results on the sphere in the spirit of Brickman’s Theorem [4]. Recall, Brickman’s Theorem states that for any $A, B \in \mathbb{S}^n$ and $n \geq 3$, that

$$\left\{ \begin{pmatrix} x^\top A x \\ x^\top B x \end{pmatrix} \in \mathbb{R}^2 : x \in \mathbf{S}^{n-1} \right\}$$

is convex. Here, \mathbf{S}^{n-1} is the sphere in \mathbb{R}^n . We elaborate on this connection in Appendix B.

There are many other examples in which people consider optimization over the special orthogonal group. These problems were implicitly studied for $\text{SO}(3)$ in [23], where they consider a formulation in terms of quadratic maps of quaternions. In another instance, [5] shows that certain standard semidefinite programming approaches to quadratic optimization problems applied to $\text{SO}(n)$ do not always produce the correct result. For this, they use the theory of nonnegative quadratic forms over real varieties developed in [3]. Some recent work of Gilman et al. [13] considers the exactness of SDP relaxations of quadratic optimization problems with variables in the Stiefel manifold $\{X \in \mathbb{R}^{n \times k} : X^\top X = I_k\}$ for some $k \leq n$. Note that when $k = n$, that this set is identical to $\text{O}(n)$. Gilman et al. [13] show that the natural SDP relaxation is exact for such problems when the operator defining the quadratic form is close enough to being diagonalizable.

2 Preliminaries

If $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear map, we denote by $\pi^* : \mathbb{R}^m \rightarrow \mathbb{R}^n$ the adjoint operator.

We will need to define a *maximal torus* in $\text{SO}(n)$. Fix some n for this section. Let $k = \lfloor \frac{n}{2} \rfloor$ and let

$R(\theta_1, \dots, \theta_k)$ denote the matrix in $\text{SO}(n)$ given by

$$R(\theta_1, \dots, \theta_k) := \begin{pmatrix} \cos(\theta_1) & \sin(\theta_1) & & & & \\ -\sin(\theta_1) & \cos(\theta_1) & & & & \\ & & \ddots & & & \\ & & & \cos(\theta_k) & \sin(\theta_k) & \\ & & & -\sin(\theta_k) & \cos(\theta_k) & \\ & & & & & \end{pmatrix} \quad (7)$$

if n is even, and

$$R(\theta_1, \dots, \theta_k) := \begin{pmatrix} \cos(\theta_1) & \sin(\theta_1) & & & & \\ -\sin(\theta_1) & \cos(\theta_1) & & & & \\ & & \ddots & & & \\ & & & \cos(\theta_k) & \sin(\theta_k) & \\ & & & -\sin(\theta_k) & \cos(\theta_k) & \\ & & & & & 1 \end{pmatrix} \quad (8)$$

if n is odd.

We define the maximal torus \mathbb{T} to be the set of matrices of the form $R(\theta_1, \dots, \theta_k)$ as the θ_i range over $[0, 2\pi)$. The following result is a special case of what is known as the Maximal Torus Theorem, but is a simple corollary of the real spectral theorem [19, Theorem 2.5.8] in our setting.

Theorem 8. *For any $X \in \text{SO}(n)$, there exists $U \in \text{SO}(n)$ so that $U^\top XU \in \mathbb{T}$. That is, $U^\top XU = R(\theta_1, \dots, \theta_k)$ for some $\theta_i \in [0, 2\pi)$.*

For $A \in \mathbb{R}^{n \times n}$, let $\|A\|_{\text{tr}}$ and $\|A\|_{\text{op}}$ denote the *trace norm* and *operator norm* of A . These are defined as the sum of the singular values of A and the maximum singular value of A respectively.

Define the *special trace* of $A \in \mathbb{R}^{n \times n}$ to be

$$\text{str}(A) := \max_{X \in \text{SO}(n)} \langle A, X \rangle.$$

This function is well-defined as $\text{SO}(n)$ is compact. Furthermore, $\text{str}(\cdot)$ is convex and 1-Lipschitz with respect to the trace norm. This holds because $\text{str}(\cdot)$ is defined as the pointwise maximum of linear functions which are individually 1-Lipschitz with respect to the trace norm. Finally, $\text{str}(\cdot)$ can be computed exactly using a single SVD:

Lemma 1 ([29, Problem 65-1]). *Given $A \in \mathbb{R}^{n \times n}$ with singular values $\sigma_1 \geq \dots \geq \sigma_n$,*

$$\text{str}(A) := \max_{X \in \text{SO}(n)} \langle A, X \rangle = \sum_{i=1}^{n-1} \sigma_i + \text{sign}(\det(A))\sigma_n.$$

We will also make use of the concept of *separation oracles* and the ellipsoid algorithm. If $C \subseteq \mathbb{R}^n$ is a compact convex set and $x \notin C$, then there is a hyperplane that separates x and C . This *separating hyperplane* is given by a nonzero vector $y \in \mathbb{R}^n$ so that $\langle y, x \rangle \geq \max\{\langle y, c \rangle : c \in C\}$. A ϵ -*weak separation oracle* for C is an oracle that on an input $x \in \mathbb{R}^n$, either correctly declares $x \in C + \mathbb{B}_\infty(0, \epsilon)$, or outputs $y \in \mathbb{R}^n$ so that y is a separating hyperplane between x and C . Here, $\mathbb{B}_\infty(a, r)$ is the ball of radius r in the L_∞ norm centered at a . The algorithmic equivalence between weak separation oracles and approximate optimization over convex sets is outlined in [15].

The ellipsoid algorithm as described in [14] provides the following guarantee for optimization in \mathbb{R}^2 .

Theorem 9. *Suppose we have access to an ϵ -weak separation oracle for closed compact $C \subseteq \mathbb{R}^2$, we are given a $R \in \mathbb{R}$ so that $C \subseteq \mathbb{B}_2(0, R)$ and C includes a ball of radius at least ϵ . There is an algorithm that optimizes a linear function with unit L_2 norm over C within an additive error of ϵ using at most $O(\log(\frac{R}{\epsilon}))$ calls to the weak separation oracle.*

3 Feasibility problems on $\text{SO}(n)$ with diagonal constraints

This section considers the feasibility problem

$$\{X \in \text{SO}(n) : \text{diag}(X) \in \mathcal{C}\}, \quad (9)$$

where $\mathcal{C} \subseteq \mathbb{R}^n$ is convex. We will assume that \mathcal{C} has an efficient separation oracle.

Horn [18, Theorem 8] shows that $\text{diag}(\text{SO}(n)) = \text{PP}_n$. As an immediate corollary, (9) is feasible if and only if $\text{PP}_n \cap \mathcal{C}$ is nonempty.

Appendix A shows how to efficiently separate from PP_n . Combined with a separation oracle for \mathcal{C} , we may then run an ellipsoid-style algorithm for deciding feasibility of $\text{PP}_n \cap \mathcal{C}$ (up to the usual errors). Supposing that $d \in \text{PP}_n \cap \mathcal{C}$ is found, it remains to see how to construct a witness $X \in \text{SO}(n)$ with $\text{diag}(X) = d$.

We will need the following description of PP_n given in [20, 22]:

$$\text{PP}_n = \{x \in [-1, 1]^n : \langle x, 1_n - 2 \cdot 1_S \rangle \leq n - 2, \quad \forall \text{ odd } S \subseteq [n]\}. \quad (10)$$

Here, 1_n is the all-ones vector, 1_S is the indicator vector of the set S and S is odd if $|S|$ is odd.

We will also need a constructive version of the Schur–Horn theorem. To state this, we need the notion of *majorization*. If $c, d \in \mathbb{R}^n$, let c^\uparrow and d^\uparrow be the results of sorting the entries of c and d in nondecreasing order. We say c majorizes d if $\sum_{i=1}^n c_i = \sum_{i=1}^n d_i$ and for each $\ell \in [n]$,

$$\sum_{i=1}^{\ell} c_i^\uparrow \leq \sum_{i=1}^{\ell} d_i^\uparrow.$$

The following lemma is essentially due to [6].

Lemma 2. *Given $c, d \in \mathbb{R}^n$ such that c majorizes d , it is possible to construct a sequence of matrices*

$$Q_1, \dots, Q_{n-1} \in \text{SO}(n)$$

in time $O(n \log n)$ satisfying

$$\text{diag} \left(\left(\prod_i^{n-1} Q_i \right)^\top \text{Diag}(c) \left(\prod_i^{n-1} Q_i \right) \right) = d.$$

Furthermore, each Q_i differs from the identity on only one principal 2×2 block, where it is a rotation matrix.

We are now ready to prove the following theorem.

Theorem 2. *Given $d \in \text{PP}_n$, it is possible to construct $X \in \text{SO}(n)$ satisfying $\text{diag}(X) = d$ in time $O(n^2)$.*

Proof. We focus on the case of even n for simplicity. The odd case follows analogously. Let $m := n/2$ and let $\theta_1, \dots, \theta_m \in [0, 2\pi)$ to be fixed later. Recall the definition of $R(\theta_1, \dots, \theta_k)$ from (7) and (8), and let $c := \text{diag}(R(\theta_1, \dots, \theta_m))$. If we can find $\theta_1, \dots, \theta_m$ so that c majorizes d , then we can apply Lemma 2 to produce the required element of $\text{SO}(n)$ with diagonal d . Equivalently, we may pick $c_1 = c_2, c_3 = c_4, \dots$ arbitrarily in $[-1, 1]$ and define $\theta_i = \arccos(c_{2i})$.

We will set c_i as follows: Let $t := \frac{1}{4}(n - \langle d, 1_n \rangle)$ and let $j - 1 = \lfloor t \rfloor$ be the integer part and $\delta := t - \lfloor t \rfloor$ be the fractional part of t . We set

$$c_1 = \dots = c_{2(j-1)} = -1, \quad c_{2j+1} = \dots = c_n = 1,$$

and the remaining elements we set as $c_{2j-1} = c_{2j} = 1 - 2\delta$. Note that $1 - 2\delta \in [-1, 1]$ since the fractional part $\delta \in [0, 1]$. Then, we have

$$\begin{aligned} \langle c, 1_n \rangle &= -2(j-1) + 2(1-2\delta) + 2(m-j) \\ &= \frac{-2}{4}(n - \langle d, 1_n \rangle - 4\delta) + 2(1-2\delta) + \frac{2}{4}(n + \langle d, 1_n \rangle - 4(1-\delta)) = \langle d, 1_n \rangle, \end{aligned}$$

where the second step was by our definition of c , and the third was by our choice of j .

Now we verify the majorization inequalities:

$$\begin{aligned} \forall k \leq 2(j-1) \quad \sum_{i=1}^k c_i &= -k \leq \sum_{i=1}^k d_i^\uparrow, \quad \text{and} \\ \forall k \geq 2j+1 \quad \sum_{i=k}^n c_i &= (n-k+1) \geq \sum_{i=k}^n d_i^\uparrow. \end{aligned}$$

Here, the last step in both inequalities hold because $d \in [-1, 1]^n$. Since $\langle c, 1_n \rangle = \langle d, 1_n \rangle$, the second set of inequalities are equivalent to

$$\forall k \geq 2j \quad \sum_{i=1}^k c_i = -k \leq \sum_{i=1}^k d_i^\uparrow.$$

We now verify the final inequality for index $k := 2j - 1$:

$$\sum_{i=1}^k c_i = -2(j-1) + (1-2\delta) = 1 - \frac{1}{2}(n - \langle c, 1_n \rangle) = \frac{1}{2}(\langle d, 1_n \rangle - (n-2)) \leq \sum_{i=1}^k d_i^\uparrow,$$

where the first step was by definition of c , in the second step we used that $j - 1 = t - \delta$, in the third step we used that $\langle c, 1_n \rangle = \langle d, 1_n \rangle$, and the final step was by the defining inequalities of PP_n given in Equation (10) for odd set $S = [k] = [2j - 1]$.

Setting $\cos(\theta_i) = c_{2i}$, we have $R(\theta_1, \dots, \theta_m) \in \text{SO}(n)$ with diagonal c majorizing d .

Now, apply Lemma 2 to get a matrix $U = \prod_{i=1}^{n-1} Q_i \in \text{SO}(n)$ such that

$$\text{diag}(U^\top \text{Diag}(c)U) = d.$$

For notational convenience, write R for $R(\theta_1, \dots, \theta_m)$. Then, $U^\top R U \in \text{SO}(n)$ satisfies

$$\begin{aligned} \text{diag}(U^\top R U) &= \text{diag}(U^\top \text{Diag}(c)U) + \text{diag}(U^\top (R - \text{Diag}(c))U) \\ &= \text{diag}(U^\top \text{Diag}(c)U) = d. \end{aligned}$$

Here, the second line follows as $R - \text{Diag}(c)$ is skew symmetric so that $U^\top (R - \text{Diag}(c))U$ must also be skew symmetric. In particular, $\text{diag}(U^\top (R - \text{Diag}(c))U) = 0$.

The time complexity follows from the fact that each Q_i differs from the identity only in a principal 2×2 block, so all $n - 1$ conjugations by Q_1, \dots, Q_{n-1} can be completed in $O(n^2)$ time. \blacksquare

4 Optimization on $\text{SO}(n)$ subject to one constraint

This section will discuss optimization of a linear function over $\text{SO}(n)$ subject to a single (possibly two-sided) linear constraint:

$$\max_{X \in \text{SO}(n)} \{ \langle A, X \rangle : \langle B, X \rangle \in [a, b] \}. \quad (11)$$

We will provide a proof that for problems of the form (11), the convex relaxation that replaces $\text{SO}(n)$ with $\text{conv}(\text{SO}(n))$ is exact. Moreover, we will give a practical algorithm for this problem that runs in roughly the same time as Wahba's problem without additional constraints, i.e., in the time to compute an SVD.

The technical core of these results lies in the following theorem:

Theorem 3. *Let $n \geq 3$, $A \in \mathbb{R}^{n \times n}$, and $c \in \mathbb{R}$. Then, the set $\text{SO}(n) \cap \{X \in \mathbb{R}^{n \times n} : \langle A, X \rangle = c\}$ is connected.*

This theorem implies the following fact using the observation that a subset of \mathbb{R}^2 is convex if and only if its intersection with every one-dimensional affine subspace is connected.

Corollary 1. *Let $n \geq 3$ and let $\pi : \text{SO}(n) \rightarrow \mathbb{R}^2$ be linear, then $\pi(\text{SO}(n))$ is convex.*

Proof. Let $x, y \in \pi(\text{SO}(n))$ denote a pair of distinct points. By definition, there exist $X, Y \in \text{SO}(n)$ so that $\pi(X) = x$ and $\pi(Y) = y$. Let $\ell \in \mathbb{R}^2 \setminus \{0\}$, $\alpha \in \mathbb{R}$ parameterize the affine line so that $\langle \ell, x \rangle = \langle \ell, y \rangle = \alpha$. Thus, $\pi(\text{SO}(n)) \cap \{z \in \mathbb{R}^2 : \langle \ell, z \rangle = \alpha\}$ is the linear image of the connected set

$$\text{SO}(n) \cap \{Z \in \text{SO}(n) : \langle \pi^*(\ell), Z \rangle = \alpha\}.$$

In particular, $\pi(\text{SO}(n)) \cap \{z \in \mathbb{R}^2 : \langle \ell, z \rangle = \alpha\}$ is connected so $[x, y]$ is contained in $\pi(\text{SO}(n))$. ■

As we have seen, this fact implies that any optimization problem of the form (11) can be solved as a convex optimization problem on $\text{conv}(\text{SO}(n))$. Thus, one could theoretically solve this problem with an exponentially sized SDP using Proposition 1. Alternatively, we give an algorithm which can successfully solve any such optimization problem in $O(n^3 \log^2(n))$ time.

In the case of $\text{SO}(3)$, Corollary 1 can be viewed as a corollary of Brickman's theorem [4], which states that the image of the unit sphere under a homogeneous quadratic map into \mathbb{R}^2 is always convex. This, together with the fact that $\text{SO}(3)$ is the image of a sphere under a quadratic map shows the result in that case (see also Appendix B). On the other hand, Corollary 1 does not follow directly from known quadratic convexity theorems for $n \geq 4$.

4.1 Topological preliminaries for the proof of Theorem 3

The proof of Theorem 3 will require some topological techniques, which we review here. We attempt to be as explicit as possible in our constructions and proofs to make them accessible to readers that are less familiar with such arguments. As a general reference for (algebraic) topology, we refer to [17].

As motivation, consider the (easy) problem of showing that the image of $\text{SO}(n)$ under a continuous map into \mathbb{R}^1 is connected (and hence convex). We can see this by appealing to two facts. First, $\text{SO}(n)$ is a connected set. Second, the image of a connected set under any continuous map is connected. Our proof of Theorem 3 will be spiritually similar and will require generalizations of

these two key ingredients. First, we will need to understand in what ways we can “move around” $\mathrm{SO}(n)$. This will be answered in Lemma 4, where we will build “loops” in $\mathrm{SO}(n)$ with nice properties. Second, we will need to understand how these loops behave under a continuous map into \mathbb{R}^2 . This will be answered in Lemma 3.

In order to formalize these statements, we will need some basic definitions from topology/homotopy theory that we present below. We encourage the reader to keep the following topological spaces in mind:

- $\mathrm{SO}(n) \subseteq \mathbb{R}^{n \times n}$ viewed as a subtopological space of $\mathbb{R}^{n \times n}$ with the standard topology.
- The punctured plane $\mathbb{R}^2 \setminus (0,0)$ viewed as a subtopological space of \mathbb{R}^2 with the standard topology.

Let X be a topological space with a designated base point $x \in X$. The fundamental group of X , denoted $\pi_1(X)$, is a group whose elements are (equivalence classes of) continuous functions $\gamma : [0, 1] \rightarrow X$ so that $\gamma(0) = \gamma(1) = x$. We will refer to such functions as *loops*. We will say that two loops γ_1 and γ_2 are equivalent if there exists a continuous function $T : [0, 1] \times [0, 1] \rightarrow X$ so that $T(0, t) = \gamma_1(t)$ and $T(1, t) = \gamma_2(t)$ for all $t \in [0, 1]$. We refer to such a T as a *homotopy*. Intuitively, two loops γ_1 and γ_2 are equivalent if γ_1 can be continuously deformed into γ_2 . This set of loops can be made into a group with the group operations being the concatenation of loops.

The identity element of X is represented by the constant loop given by $i(t) = x$ for all $t \in [0, 1]$. If X is path connected, then the fundamental group is independent of the choice of basepoint—this will be the case for all topological spaces we consider.

For example, the fundamental group of the punctured plane $\mathbb{R}^2 \setminus (0,0)$ is \mathbb{Z} . Explicitly, any loop in $\mathbb{R}^2 \setminus (0,0)$ can be identified with the number of times it winds anticlockwise around the origin. Less intuitively, we will also need the fact that for $n \geq 3$, the fundamental group of $\mathrm{SO}(n)$ is $\mathbb{Z}/2\mathbb{Z}$. An important consequence of this fact is that given any loop in $\mathrm{SO}(n)$ (for $n \geq 3$), simply concatenating this loop with itself results in the identity element in $\pi_1(\mathrm{SO}(n))$, which can then be continuously deformed to a point. This will be used in Lemma 3 below, for which Figure 2 shows a cartoon of the proof strategy.

This will be relevant to us in the following context:

Lemma 3. *Let $n \geq 3$. Suppose $f : \mathrm{SO}(n) \rightarrow \mathbb{R}^2$ is a continuous map and $\gamma : [0, 1] \rightarrow \mathrm{SO}(n)$ is a loop such that $(0,0)$ is not in the image $(f_*(\gamma))([0, 1])$. In this case, we may view $f_*(\gamma)$ as a loop in $\mathbb{R}^2 \setminus (0,0)$. If $f_*(\gamma)$ is not equivalent to the identity element in $\pi_1(\mathbb{R}^2 \setminus (0,0))$, then $(0,0) \in f(\mathrm{SO}(n))$.*

Proof. We may assume that γ is equivalent to the identity element in $\pi_1(\mathrm{SO}(n))$ by concatenating γ with itself. Indeed, this produces the identity element in $\pi_1(\mathrm{SO}(n))$ as $\pi_1(\mathrm{SO}(n)) = \mathbb{Z}/2\mathbb{Z}$. Additionally, if $f_*(\gamma)$ is not equivalent to the identity element in $\pi_1(\mathbb{R}^2 \setminus (0,0))$, then concatenating this loop with itself does not yield the identity element. ■

Let T be a homotopy between γ and the identity element i in $\pi_1(\mathrm{SO}(n))$, i.e., assume $T(0, t) = \gamma(t)$ and $T(1, t) = i(t)$. Now, consider the image $(f \circ T)([0, 1] \times [0, 1])$. We claim that $(0,0)$ is in this image. Indeed, supposing otherwise, we have constructed a homotopy between $f_*(\gamma)$ and the identity element in $\pi_1(\mathbb{R}^2 \setminus (0,0))$, a contradiction. ■

By adding a translation term to f , we can apply Lemma 3 with an arbitrary point $\beta \in \mathbb{R}^2$ in place of the origin. Figure 3 depicts a two-dimensional linear image of $\mathrm{SO}(3)$ and a loop γ in $\mathrm{SO}(3)$ and




Figure 2: A cartoon of the proof of Lemma 3. On the left, we depict a loop (solid line) in $\text{SO}(n)$, that is equivalent to the identity element, and its deformations (dashed lines) to a point. On the right is the image of the loop and its deformations under the map $f : \text{SO}(n) \rightarrow \mathbb{R}^2$. The origin in \mathbb{R}^2 is shown in blue. By assumption, $f_*(\gamma)$ “winds” around the origin a nonzero number of times. We will show that the image of the deformations must eventually “pass through” the origin. This shows that the origin is in the image of $f(\text{SO}(n))$.




Figure 3: The solid red area is the image of $\text{SO}(3)$ under a linear map into \mathbb{R}^2 . The blue ellipse is the image of a loop in $\text{SO}(3)$ under the same linear map. The (image of the) loop begins at some point on the blue ellipse and walks clockwise around the ellipse exactly once. Lemma 3 implies that every point in \mathbb{R}^2 contained within this blue ellipse must be contained in the image of $\text{SO}(3)$. The proof of Theorem 3 will follow a similar idea with a continuous but nonlinear function.

gives some intuition on how we will use Lemma 3 to prove Theorem 3.

4.2 Proof of Theorem 3

This subsection will contain a proof of Theorem 3. Let $H = \{X \in \mathbb{R}^{n \times n} : \langle A, X \rangle = c\}$. We will require the following construction.

Lemma 4. *Let $U, V \in H \cap \text{SO}(n)$. Then, there exists a continuous function $\gamma : [0, 1] \rightarrow \text{SO}(n)$ with the following properties:*

- $\gamma(0) = \gamma(1) = U$,
- $\gamma(\frac{1}{2}) = V$,
- either $\langle A, \gamma(t) \rangle = c$ for all $t \in (0, \frac{1}{2})$ or $\langle A, \gamma(t) \rangle > c$ for all $t \in (0, \frac{1}{2})$, and
- either $\langle A, \gamma(t) \rangle = c$ for all $t \in (\frac{1}{2}, 1)$ or $\langle A, \gamma(t) \rangle < c$ for all $t \in (\frac{1}{2}, 1)$.

Proof. We may write $U = R(\eta_1, \dots, \eta_k)$ for some $\eta_i \in [0, 2\pi)$ and $V = R(\theta_1, \dots, \theta_k)$ for some $\theta_i \in [0, 2\pi)$.

Let \mathbf{S}^1 denote the unit circle in \mathbb{R}^2 , thought of as the points $[0, 2\pi)$ where 0 and 2π are identified. Let \mathbb{T} denote the set of all matrices $R(\phi_1, \dots, \phi_k)$ as ϕ_1, \dots, ϕ_k range over \mathbf{S}^1 .

Now, consider the expression $\langle A, R(\phi_1, \dots, \phi_k) \rangle$, which we can express as a sum over the blocks of $R(\phi_1, \dots, \phi_k)$:

$$\langle A, R(\phi_1, \dots, \phi_k) \rangle = \sum_{\ell=1}^k \left\langle \begin{pmatrix} A_{2i-1, 2i-1} & A_{2i-1, 2i} \\ A_{2i, 2i-1} & A_{2i, 2i} \end{pmatrix}, \begin{pmatrix} \cos(\phi_i) & \sin(\phi_i) \\ -\sin(\phi_i) & \cos(\phi_i) \end{pmatrix} \right\rangle + c_0,$$

where $c_0 = A_{nn}$ if n is odd and 0 otherwise. We may rewrite this as

$$\begin{aligned} \langle A, R(\phi_1, \dots, \phi_k) \rangle &= \sum_{\ell=1}^k \left\langle \begin{pmatrix} A_{2i-1, 2i-1} + A_{2i, 2i} \\ A_{2i-1, 2i} - A_{2i, 2i-1} \end{pmatrix}, \begin{pmatrix} \cos(\phi_i) \\ \sin(\phi_i) \end{pmatrix} \right\rangle + c_0 \\ &= \sum_{i=1}^k c_i \left\langle \begin{pmatrix} \cos(\hat{\phi}_i) \\ \sin(\hat{\phi}_i) \end{pmatrix}, \begin{pmatrix} \cos(\phi_i) \\ \sin(\phi_i) \end{pmatrix} \right\rangle + c_0 \\ &= \sum_{i=1}^k c_i \cos(\hat{\phi}_i - \phi_i) + c_0, \end{aligned}$$

where $c_i \geq 0$ and $\hat{\phi}_i \in \mathbf{S}^1$ are the polar coordinates for $\begin{pmatrix} A_{2i-1, 2i-1} + A_{2i, 2i} \\ A_{2i-1, 2i} - A_{2i, 2i-1} \end{pmatrix}$.

We will construct the loop $\gamma(t)$ by setting $\gamma(t) = R(\phi_1(t), \dots, \phi_k(t))$ where $\phi_i(t)$ will be defined shortly. Specifically, $\phi_i(t)$ will be a continuous function satisfying $\phi_i(0) = \phi_i(1) = \eta_i$ and $\phi_i(\frac{1}{2}) = \theta_i$. In particular, $\gamma(t)$ will be continuous and satisfy $\gamma(0) = \gamma(1) = U$ and $\gamma(1/2) = V$.

Now, fix some $i \in [k]$. We will parameterize $\phi_i(t) = \hat{\phi}_i - \delta(t)$. Note the requirement $\phi_i(0) = \phi_i(1) = \eta_i$ and $\phi_i(1/2) = \theta_i$ is equivalent to $\delta(0) = \delta(1) = \hat{\phi}_i - \eta_i$ and $\delta(1/2) = \hat{\phi}_i - \theta_i$. Note that $\cos(\hat{\phi}_i - \phi_i(t)) = \cos(\delta(t))$ is simply the “ x -coordinate” of $\delta(t)$.

Consider the locations of $\hat{\phi}_i - \eta_i$, $\hat{\phi}_i - \theta_i$ on the unit circle \mathbf{S}^1 . If $\hat{\phi}_i - \eta_i$ is on the bottom half of the unit circle, then let $\delta(3/4) = \pi$ and set $\delta(t)$ to move counterclockwise along the bottom half of the circle until $\delta(1) = \hat{\phi}_i - \eta_i$. Also set $\delta(0) = \hat{\phi}_i - \eta_i$ and set $\delta(t)$ to move counterclockwise along the bottom half of the circle until $\delta(1/4) = 0$. On the other hand, if $\hat{\phi}_i - \eta_i$ is on the top half of the unit circle, then let $\delta(3/4) = \pi$ and set $\delta(t)$ to move clockwise along the top half of the circle until $\delta(1) = \hat{\phi}_i - \eta_i$. Also set $\delta(0) = \hat{\phi}_i - \eta_i$ and set $\delta(t)$ to move clockwise along the top half of the circle until $\delta(1/4) = 0$. Note that in both cases, the function $\cos(\delta(t))$ is either constant or strictly increasing on each interval $[0, 1/4]$ and $[3/4, 1]$. Completely analogous ideas can be used to define $\delta(t)$ on the intervals $[1/4, 1/2]$ and $[1/2, 3/4]$ so that $\delta(1/2) = \hat{\phi}_i - \theta_i$ and $\cos(\delta(t))$ is either constant or strictly decreasing on each interval $[1/4, 1/2]$ and $[1/2, 3/4]$. (See Figure 4 for two examples of this construction.)

We then define $\gamma(t) := R(\phi_1(t), \phi_2(t), \dots, \phi_k(t))$. We see that this loop satisfies all of the properties desired based on properties of the individual $\phi_i(t)$. \blacksquare

We are now ready to prove Theorem 3.




Figure 4: Two examples of the construction of $\delta(t)$ in the proof of Lemma 4. The first row corresponds to $\hat{\phi}_i - \eta_i = \frac{\pi}{4}$ and $\hat{\phi}_i - \theta_i = -\frac{3\pi}{4}$. The second row corresponds to $\hat{\phi}_i - \eta_i = \frac{3\pi}{4}$ and $\hat{\phi}_i - \theta_i = \frac{\pi}{2}$.

Proof of Theorem 3. Suppose for the sake of contradiction that $H \cap \text{SO}(n)$ is not connected, which by definition means that there exist nonempty closed sets $\mathcal{U}, \mathcal{V} \subseteq H \cap \text{SO}(n)$ so that $\mathcal{U} \cap \mathcal{V} = \emptyset$, and $\mathcal{U} \cup \mathcal{V} = H \cap \text{SO}(n)$. As \mathcal{U} is closed, the distance function

$$\text{dist}_{\mathcal{U}}(X) := \min_{U \in \mathcal{U}} \|U - X\|_{\text{op}}$$

is well-defined. Let $\delta := \min_{V \in \mathcal{V}} \text{dist}_{\mathcal{U}}(V)$. As \mathcal{U} and \mathcal{V} are compact and disjoint, we have that $\delta > 0$. Define $f : \text{SO}(n) \rightarrow \mathbb{R}^2$ given by

$$f(X) = \begin{pmatrix} \langle A, X \rangle - c \\ \text{dist}_{\mathcal{U}}(X) - \delta/2 \end{pmatrix}$$

We claim that $(0, 0) \notin f(\text{SO}(n))$. Assume for contradiction that $(0, 0) \in f(\text{SO}(n))$, i.e., there exists an $X \in \text{SO}(n)$ so that $\langle A, X \rangle = c$ and $\text{dist}_{\mathcal{U}}(X) = \frac{\delta}{2}$. As $X \in H \cap \text{SO}(n)$, we have that either $X \in \mathcal{U}$ or $X \in \mathcal{V}$. In the first case, $\text{dist}_{\mathcal{U}}(X) = 0$, a contradiction. In the second case, we have that $\min_{V \in \mathcal{V}} \text{dist}_{\mathcal{U}}(V) \leq \text{dist}_{\mathcal{U}}(X) = \delta/2 < \delta = \min_{V \in \mathcal{V}} \text{dist}_{\mathcal{U}}(V)$, again a contradiction. Thus, $(0, 0) \notin f(\text{SO}(n))$.

Now, fix $U \in \mathcal{U}$ and $V \in \mathcal{V}$ and let γ be the loop constructed in Lemma 4. We claim that $\langle A, \gamma(t) \rangle > c$ for all $t \in (0, \frac{1}{2})$. Indeed, by Lemma 4, if this is not the case, then $\langle A, \gamma(t) \rangle = c$ for all $t \in [0, \frac{1}{2}]$ so that the path $\gamma(t)$ for $t \in [0, 1/2]$ connects U, V in $H \cap \text{SO}(n)$. This contradicts the fact that U and V are not connected in $H \cap \text{SO}(n)$.

We now verify that f and γ satisfy the assumptions of Lemma 3. To do so, we will exhibit a homotopy from $f_*(\gamma)$ to the loop $(\sin(2\pi t), -\cos(2\pi t))$. Let

$$T(s, t) := s f_*(\gamma)(t) + (1 - s)(\sin(2\pi t), -\cos(2\pi t)).$$

This is clearly a continuous function from $[0, 1] \times [0, 1] \rightarrow \mathbb{R}^2$. Thus, to verify that it is a valid homotopy in $\mathbb{R}^2 \setminus (0, 0)$ it remains to check that $T(s, t) \neq (0, 0)$ for any $(s, t) \in [0, 1] \times [0, 1]$. To

see this, note that for $t \in (0, \frac{1}{2})$, we have $\langle A, \gamma(t) \rangle > c$ so that $f(\gamma(t))_1 > 0$. Additionally, for all $t \in (0, \frac{1}{2})$, $\sin(2\pi t) > 0$. Thus, $T(s, t)_1 \neq 0$ for all $t \in (0, \frac{1}{2})$ and $s \in [0, 1]$. Similar arguments show that $T(s, t)_1 \neq 0$ for all $t \in (\frac{1}{2}, 1)$ and $s \in [0, 1]$, and that $T(s, t)_2 \neq 0$ for all $t \in \{0, \frac{1}{2}, 1\}$ and $s \in [0, 1]$.

Finally, Lemma 3 implies that $(0, 0) \in f(\text{SO}(n))$, a contradiction. We conclude that $H \cap \text{SO}(n)$ is connected. \blacksquare

4.3 Algorithms for One Constraint Optimization

Here, we aim to solve the optimization problem given in (3).

Theorem 4. *Let $n \geq 3$, $A, B \in \mathbb{R}^{n \times n}$ with $\|A\|_{\text{tr}} = \|B\|_{\text{tr}} = 1$. Here $\|\cdot\|_{\text{tr}}$ is the trace norm, defined formally in Section 2. Let X^* be the optimal solution to (3). We can compute $\langle A, X^* \rangle$ and $\langle B, X^* \rangle$ within an additive error of ϵ in time*

$$O\left(n^3 \log\left(\frac{1}{\epsilon}\right)^2\right).$$

Here, n^3 is the time complexity of computing the SVD of an $n \times n$ matrix.

Moreover, we will return $\alpha, \beta \in \mathbb{R}$ so that $|\alpha| + |\beta| = 1$ and

$$\langle \alpha A + \beta B, X^* \rangle + \epsilon \geq \max\{\langle \alpha A + \beta B, X \rangle : X \in \text{SO}(n)\}.$$

To state the algorithm we note that for any $A, B \in \mathbb{R}^{n \times n}$, (3) is equivalent to the following optimization problem

$$\max_{x \in \mathbb{R}^2} \{x_1 : x \in \pi(\text{SO}(n)), x_2 \in [a, b]\}.$$

Here, $\pi(\text{SO}(n))$ is the image of $\text{SO}(n)$ in \mathbb{R}^2 given by $\pi(X) = (\langle A, X \rangle, \langle B, X \rangle)$. By Corollary 1, we have that $\pi(\text{SO}(n))$ is convex. Therefore the above is a linear optimization problem over convex set $C := \pi(\text{SO}(n)) \cap \{x_2 \in [a, b]\}$, and we can apply standard methods from convex optimization to solve the above equivalent formulation of (3).

For this, we appeal to the ellipsoid algorithm, which we stated in Theorem 9. By Theorem 9, it suffices to provide a weak separation oracle for the set $C = \pi(\text{SO}(n)) \cap \{x_2 \in [a, b]\}$. As the second constraint is trivial to separate over, we focus on separating over the convex set $\pi(\text{SO}(n))$.

It will be useful to have a subroutine for minimizing $h(y)$ over the unit ℓ_1 -ball in \mathbb{R}_2 , where

$$h(y) := \text{str}(\pi^*(y)) - \langle y, x \rangle.$$

Lemma 5. *Given $x \in \mathbb{R}^2$ with $\|x\|_\infty \leq 1 + \epsilon$, we can construct \hat{y} with $\|\hat{y}\|_1 = 1$ so that*

$$h(\hat{y}) - \epsilon \leq \min_y \{h(y) : \|y\|_1 = 1\}$$

using at most $O\left(\log\left(\frac{1}{\epsilon}\right)\right)$ evaluations of h and additional computations.

Proof. We note that $\{y : \|y\|_1 = 1\}$ is a union of 4 line segments, so minimizing h on this set can be done by minimizing the following 4 univariate functions on $[0, 1]$:

$$g_{\sigma_1\sigma_2}(\alpha) = h(\sigma_1\alpha, \sigma_2(1-\alpha)) = \text{str}(\sigma_1\alpha A + \sigma_2(1-\alpha)B) - \sigma_1\alpha x_1 - \sigma_2(1-\alpha)x_2,$$

indexed by $\sigma_1, \sigma_2 \in \{\pm 1\}$. Each of the four functions $g_{\sigma_1\sigma_2}$ is a one-dimensional convex function with Lipschitz constant bounded by

$$\|A\|_{\text{tr}} + \|B\|_{\text{tr}} + \|x\|_1 \leq 4 + 2\epsilon.$$

For each $\sigma_1\sigma_2 \in \{\pm 1\}$, we may use golden section search [21] to find a $\hat{\alpha}_{\sigma_1\sigma_2} \in [0, 1]$ such that

$$g_{\sigma_1\sigma_2}(\hat{\alpha}_{\sigma_1\sigma_2}) \leq \min_{\alpha \in [0,1]} g_{\sigma_1\sigma_2}(\alpha) + \epsilon.$$

Each application of golden section search requires $O\left(\log\left(\frac{1}{\epsilon}\right)\right)$ evaluations of $g_{\sigma_1\sigma_2}$, or equivalently, evaluations of h . ■

Lemma 6. *Let $n \geq 3$ and $A, B \in \mathbb{R}^{n \times n}$ with $\|A\|_{\text{tr}} = \|B\|_{\text{tr}} = 1$. There is a weak separation oracle for the set $\pi(\text{SO}(n))$ that runs in time $O(n^3 \log(\frac{1}{\epsilon}))$.*

Proof. Suppose we are given $A, B \in \mathbb{R}^{n \times n}$ and $x \in \mathbb{R}^2$. If $\|x\|_\infty > 1 + \epsilon$, then in fact, $x \notin \pi(\text{SO}(n)) + \mathbb{B}_\infty(0, \epsilon)$ as, by Holder's inequality,

$$X \in \text{SO}(n) \implies \max\{|\langle A, X \rangle|, |\langle B, X \rangle|\} \leq \|X\|_{\text{op}} \max\{\|A\|_{\text{tr}}, \|B\|_{\text{tr}}\} \leq 1,$$

where the last step was by our assumption $\|A\|_{\text{tr}} = \|B\|_{\text{tr}} = 1$. Therefore, in this case, we may immediately terminate with one of $(\pm 1, 0)$ or $(0, \pm 1)$ as a separating hyperplane. For the remainder, we assume that $\|x\|_\infty \leq 1 + \epsilon$.

A nonzero vector $y \in \mathbb{R}^2$ defines a separating hyperplane between x and $\pi(\text{SO}(n))$ if and only if

$$\langle y, x \rangle \geq \max_{X \in \text{SO}(n)} \langle y, \pi(X) \rangle.$$

Recalling the definition of $\text{str}(\cdot)$ from Section 2, the expression on the right can be written as

$$\max_{X \in \text{SO}(n)} \langle y, \pi(X) \rangle = \max_{X \in \text{SO}(n)} \langle \pi^*(y), X \rangle = \text{str}(\pi^*(y)).$$

Thus, a nonzero $y \in \mathbb{R}^2$ defines a separating hyperplane if and only if $h(y) \leq 0$. Note that by Lemma 1, we can compute $h(y)$ for a given y using a single SVD. As h is 1-homogeneous, such a y exists if and only if one exists with $\|y\|_1 = 1$.

Now, we apply Lemma 5 to compute \hat{y} approximately minimizing $h(y)$ on the unit ℓ_1 ball.

If $h(\hat{y}) \leq 0$, then we may output \hat{y} as a separating hyperplane. For the remainder of the proof, suppose $h(\hat{y}) > 0$. By Lemma 5 and 1-homogeneity, $h(y) > -\epsilon$ for all $y \in \mathbb{B}_1(0, 1)$. We claim that $x \in \pi(\text{SO}(n)) + \mathbb{B}_\infty(0, \epsilon)$. If, to the contrary, $x \notin \pi(\text{SO}(n)) + \mathbb{B}_\infty(0, \epsilon)$, then by the separating hyperplane theorem, there would be some y so that

$$\langle y, x \rangle \geq \max\{\langle y, c + \delta \rangle : c \in \pi(\text{SO}(n)), \delta \in \mathbb{B}_\infty(0, \epsilon)\} = \max\{\langle y, c \rangle : c \in \pi(\text{SO}(n))\} + \epsilon \|y\|_1.$$

In particular, there would be some y with $\|y\|_1 = 1$ such that

$$h(y) = \text{str}(\pi^*(y)) - \langle y, x \rangle \leq -\epsilon,$$

which is a contradiction. ■

5 Optimization on $\text{SO}(n)$ and $\text{O}(n)$ with strict upper triangular constraints

In this section, we consider optimization problems with strict upper triangular (SUT) constraints over $\text{SO}(n)$ and $\text{O}(n)$: Let $A \in \mathbb{R}^{n \times n}$ and let \mathcal{C} be a nonempty closed convex subset of $\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$. We consider the problems (which we also stated in the introduction)

$$\sup_{X \in \text{SO}(n)} \{ \langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C} \} \quad (4)$$

$$\leq \sup_{X \in \text{O}(n)} \{ \langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C} \} \quad (5)$$

$$\leq \max_{X \in \mathbb{B}_{\text{op}}(n)} \{ \langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C} \}. \quad (6)$$

where the last inequality is because $\mathbb{B}_{\text{op}}(n)$ is the convex hull of $\text{O}(n)$. Here, we define the values of (4) and (5) to be $-\infty$ whenever they are infeasible. Since Equation (6) is always feasible (by our assumptions on \mathcal{C}), and $\mathbb{B}_{\text{op}}(n)$ is compact, the supremum is always achieved in this last case. Nonetheless, by compactness, (4) and (5) both achieve their maxima as long as they are feasible.

The following theorem is the main result of this section and shows that one or both of these inequalities hold at equality for certain choices of A .

Theorem 5. *Let $A \in \mathbb{R}^{n \times n}$ be a diagonal matrix and let $\mathcal{C} \subseteq \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$ be a nonempty closed convex set. Then, equality holds between (5) and (6). If additionally $\det(A) \geq 0$, then equality holds between (4) to (6).*

We will prove Theorem 5 in Section 5.1. As a byproduct of the proof, we will also see a numerical method for constructing optimizers of (4) or (5) from an optimizer of the convex program (6) by solving an SDP. In Section 5.2, we verify that our results in this section can be applied to problems with few coordinate constraints or rank-one constraints.

Before moving on, we note two hidden convexity properties implied by Theorem 5.

Corollary 2. *It holds that $\pi_{\mathcal{T}}(\text{SO}(n)) = \pi_{\mathcal{T}}(\text{O}(n)) = \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$. In particular, all three sets are convex.*

Proof. It is clear that $\pi_{\mathcal{T}}(\text{SO}(n)) \subseteq \pi_{\mathcal{T}}(\text{O}(n)) \subseteq \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$. Now, let $\sigma \in \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$ and set $A = I$ and $\mathcal{C} = \{\sigma\}$. Theorem 5 implies the feasibility of (4), i.e., $\sigma \in \pi_{\mathcal{T}}(\text{SO}(n))$, whence $\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)) \subseteq \pi_{\mathcal{T}}(\text{SO}(n))$. \blacksquare

Corollary 3. *Let $A \in \mathbb{R}^{n \times n}$ be a diagonal matrix, then*

$$\left\{ \left(\begin{array}{c} \langle A, X \rangle - \gamma \\ \pi_{\mathcal{T}}(X) \end{array} \right) : \begin{array}{c} X \in \text{O}(n) \\ \gamma \geq 0 \end{array} \right\} = \left\{ \left(\begin{array}{c} \langle A, X \rangle - \gamma \\ \pi_{\mathcal{T}}(X) \end{array} \right) : \begin{array}{c} X \in \mathbb{B}_{\text{op}}(n) \\ \gamma \geq 0 \end{array} \right\}.$$

If additionally $\det(A) \geq 0$, then

$$\left\{ \left(\begin{array}{c} \langle A, X \rangle - \gamma \\ \pi_{\mathcal{T}}(X) \end{array} \right) : \begin{array}{c} X \in \text{SO}(n) \\ \gamma \geq 0 \end{array} \right\} = \left\{ \left(\begin{array}{c} \langle A, X \rangle - \gamma \\ \pi_{\mathcal{T}}(X) \end{array} \right) : \begin{array}{c} X \in \mathbb{B}_{\text{op}}(n) \\ \gamma \geq 0 \end{array} \right\}.$$

Proof. We prove only the first claim as the second is proved analogously. For convenience, let \mathcal{L} and \mathcal{R} denote the left- and right-hand side sets in the first claim. As $O(n) \subseteq \mathbb{B}_{\text{op}}(n)$, we have that $\mathcal{L} \subseteq \mathcal{R}$. Now, suppose $(v, \sigma) \in \mathcal{R}$, so there exists $X \in \mathbb{B}_{\text{op}}(n)$ such that $\pi_{\mathcal{T}}(X) = \sigma$, $\langle A, X \rangle \geq v$. Let

$$v' = \max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) = \sigma\}.$$

By definition, $v' \geq v$. Next, by Theorem 5, there exists $X \in O(n)$ such that $\pi_{\mathcal{T}}(X) = \sigma$ and $\langle A, X \rangle = v'$. Then $(v', \sigma) \in \mathcal{L}$. As \mathcal{L} is closed under decreasing its first coordinate, $(v, \sigma) \in \mathcal{L}$. As $(v, \sigma) \in \mathcal{R}$ was arbitrary, we conclude $\mathcal{R} \subseteq \mathcal{L}$. \blacksquare

5.1 Proof of Theorem 5

We begin by proving the following special case of Theorem 5.

Proposition 2. *Let $A \in \mathbb{R}^{n \times n}$ be a diagonal matrix with $\det(A) \neq 0$ and let $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$. Then,*

$$\max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) = \sigma\} \tag{12}$$

has a unique optimizer \hat{X} . It holds that $\hat{X} \in O(n)$. If additionally $\det(A) > 0$, then $\hat{X} \in \text{SO}(n)$.

Proof. First, note that $\mathbb{B}_{\text{op}}(n)$ is full-dimensional in $\mathbb{R}^{n \times n}$ so its projection $\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$ is also full-dimensional. Thus, $\text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))) = \pi_{\mathcal{T}}(\text{int}(\mathbb{B}_{\text{op}}(n)))$ and (12) is strictly feasible. We deduce that strong duality and dual attainability hold in the following primal and dual problems

$$\begin{aligned} & \max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) = \sigma\} \\ &= \min_{Y \in \mathbb{R}^{n \times n}, \lambda \in \mathbb{R}^{\binom{n}{2}}} \{\langle \sigma, \lambda \rangle + \|Y\|_{\text{tr}} : Y + \pi_{\mathcal{T}}^*(\lambda) = A\}. \end{aligned}$$

Now, let $(\hat{Y}, \hat{\lambda})$ optimize the dual problem. Note that $\hat{Y} = A - \pi_{\mathcal{T}}^*(\hat{\lambda})$ is upper triangular with $\text{diag}(A)$ on its diagonal, so $\det(\hat{Y}) = \det(A) \neq 0$ by assumption, i.e., $\text{rank}(\hat{Y}) = n$.

Let $\hat{X} \in \mathbb{B}_{\text{op}}(n)$ be an arbitrary maximizer of (12), which exists by compactness of $\mathbb{B}_{\text{op}}(n)$. By strong duality,

$$\|\hat{Y}\|_{\text{tr}} + \langle \sigma, \hat{\lambda} \rangle = \langle A, \hat{X} \rangle = \langle \hat{Y} + \pi_{\mathcal{T}}^*(\hat{\lambda}), \hat{X} \rangle = \langle \hat{Y}, \hat{X} \rangle + \langle \sigma, \hat{\lambda} \rangle.$$

Thus, $\|\hat{Y}\|_{\text{tr}} = \langle \hat{Y}, \hat{X} \rangle$. Let $U\Sigma V^{\top} = \hat{Y}$ be an SVD of \hat{Y} . Then, $\text{tr}(\Sigma) = \|\hat{Y}\|_{\text{tr}} = \langle \hat{Y}, \hat{X} \rangle = \langle \Sigma, U^{\top} \hat{X} V \rangle$. Noting that $U^{\top} \hat{X} V \in \mathbb{B}_{\text{op}}(n)$ and that Σ has only positive diagonal entries, we deduce that $U^{\top} \hat{X} V = I$ so that $\hat{X} = UV^{\top} \in O(n)$. This proves the first claim.

Now, suppose $\det(A) > 0$. Then, $\det(\hat{Y}) = \det(A - \pi_{\mathcal{T}}^*(\hat{\lambda})) = \det(A) > 0$. Thus, $\det(\hat{X}) = \det(U) \det(V^{\top}) = \frac{\det(\hat{Y})}{\det(\Sigma)} > 0$. We conclude that $\hat{X} \in \text{SO}(n)$. \blacksquare

We may now prove Theorem 5 in full generality.

Proof of Theorem 5. Let \hat{X} be an optimizer of (6) and set $\sigma = \pi_{\mathcal{T}}(\hat{X})$. Now, let $\epsilon \in (0, 1]$ and define $\sigma_\epsilon := (1 - \epsilon)\sigma$. As $\mathbb{B}_{\text{op}}(n)$ is full-dimensional, we have that $\sigma_\epsilon \in \pi_{\mathcal{T}}(\text{int}(\mathbb{B}_{\text{op}}(n)))$. If $\det(A) \neq 0$, then define $A_\epsilon := A$. Otherwise, construct $A_\epsilon \in \mathbb{R}^{n \times n}$ by setting each zero diagonal entry of A to $\pm \frac{\epsilon}{n}$ in such a way that $\det(A_\epsilon) > 0$, which is possible because A is diagonal and its determinant is the product of its diagonal entries. Then, applying Proposition 2 with A_ϵ and σ_ϵ , there exists $X_\epsilon \in \text{O}(n)$ satisfying

$$\langle A, X_\epsilon \rangle \geq \langle A_\epsilon, X_\epsilon \rangle - \epsilon \geq (1 - \epsilon) \langle A_\epsilon, \hat{X} \rangle - \epsilon \geq (1 - \epsilon) (\langle A, \hat{X} \rangle - \epsilon) - \epsilon, \quad (13)$$

where the first and third steps are because $\|A - A_\epsilon\|_{\text{tr}} \leq \epsilon$ by our choice of A_ϵ and the fact that $\hat{X}, X_\epsilon \in \mathbb{B}_{\text{op}}(n)$, and the second step follows as $(1 - \epsilon)\hat{X}$ is feasible in

$$\max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A_\epsilon, X \rangle : \pi_{\mathcal{T}}(X) = \sigma_\epsilon\} = \langle A_\epsilon, X_\epsilon \rangle.$$

Next, consider a sequence $\{\epsilon_k\} \subseteq (0, 1]$ converging to zero and the corresponding sequence $\{X_{\epsilon_k}\} \subseteq \text{O}(n)$. As $\text{O}(n)$ is compact, $\{X_{\epsilon_k}\}$ has a subsequential limit $\tilde{X} \in \text{O}(n)$. By continuity, we have that $\langle A, \tilde{X} \rangle \geq \langle A, \hat{X} \rangle$ and $\pi_{\mathcal{T}}(\tilde{X}) = \sigma \in \mathcal{C}$. We deduce that equality holds between (5) and (6).

Finally, suppose $\det(A) \geq 0$ so that $\det(A_\epsilon) > 0$. Then, the sequence $\{X_{\epsilon_k}\}$ lies in $\text{SO}(n)$ so that the subsequential limit \tilde{X} may also be taken to live in $\text{SO}(n)$. ■

Remark 2. The proof of Theorem 5 suggests a numerical method for recovering an optimizer of (5) from an optimizer, \hat{X} , of (6). Let $\epsilon > 0$ be some small numerical parameter and let $A_\epsilon, \sigma_\epsilon$ be as defined in the proof of Theorem 5. Then, the unique maximizer of

$$\max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A_\epsilon, X \rangle : \pi_{\mathcal{T}}(X) = \sigma_\epsilon\}$$

is guaranteed to lie in $\text{O}(n)$ and is approximately optimal in the sense of (13) and approximately feasible as $\sigma_\epsilon \approx \sigma \in \mathcal{C}$.

Alternatively, we may shortcut solving two separate convex optimization problems by preemptively replacing A and \mathcal{C} with A_ϵ and \mathcal{C}_ϵ , in such a way that guarantees $\det(A_\epsilon) \neq 0$ and $\mathcal{C}_\epsilon \subseteq \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$. With these perturbed sets, Proposition 2 guarantees that any optimizer of

$$\max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A_\epsilon, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}_\epsilon\}$$

lies in $\text{O}(n)$. Analogous statements hold for the $\text{SO}(n)$ setting.

5.2 Applications to low-rank constraints

The following proposition shows that optimization and feasibility problems over $\text{SO}(n)$ with convex constraints on the values of $\langle B_i, X \rangle$, where $B_i \in \mathbb{R}^{n \times n}$ are low-rank matrices, can be seen as a special case of SUT constraints after a reparameterization of $\text{SO}(n)$.

Proposition 3. *Let $k \leq n - 1$. Fix $u_1, \dots, u_k \in \mathbb{R}^n$, and also fix $v_1, \dots, v_k \in \mathbb{R}^n$. For $i = 1, \dots, m$, let $B_i = \sum_{j=1}^k \beta_{ij} u_j v_j^\top$, for some $\beta_{i,j} \in \mathbb{R}$. Let $\mathcal{B}(X) := (\langle B_i, X \rangle)_{i \in [m]}$. Then $\mathcal{B}(\text{SO}(n))$ is convex.*

Theorem 5 then implies that we may decide feasibility of problems of the form $\mathcal{B}(\text{SO}(n)) \cap \mathcal{C}$ for compact convex \mathcal{C} via a simple SDP. As an example, Proposition 3 applies if $m \leq n - 1$ and B_i are each rank one as in the situation of Section 1.1.

In order to show this proposition, we will need a lemma.

Lemma 7. *Let $\{u_1, \dots, u_{n-1}\} \subseteq \mathbb{R}^n$ and $\{v_1, \dots, v_{n-1}\} \subseteq \mathbb{R}^n$. Then, there exists $U, V \in \text{SO}(n)$ such that for all $i \in [n-1]$, $v_i^\top (V^\top X U) u_i$ depends only on the strict upper triangular entries of X .*

Proof. It follows from the existence of QR decompositions that we may upper triangularize the $\{u_i\}$ with a special orthogonal matrix, i.e., there exists $U \in \text{SO}(n)$ such that $\text{supp}(U u_i) \subseteq [1, i]$ for each $i \in [n-1]$. Similarly, we may lower triangularize the $\{v_i\}$ with a special orthogonal matrix, i.e., there exists $V \in \text{SO}(n)$ such that $\text{supp}(V v_i) \subseteq [i+1, n]$. Then

$$\text{supp}(U u_i v_i^\top V^\top) \subseteq [1, i] \times [i+1, n] \text{ for } i \in [n-1].$$

Thus $(V v_i)^\top X (U u_i)$ depends only the strictly upper triangular entries of X . ■

We now show Proposition 3.

Proof of Proposition 3. Note that $\mathcal{B}(\text{SO}(n))$ is a linear image of the set

$$\left\{ \left(v_j^\top X u_j \right)_{j \in [k]} : X \in \text{SO}(n) \right\}. \quad (14)$$

Let $U, V \in \text{SO}(n)$ denote the matrices guaranteed by Lemma 7, then (14) is equivalent to

$$\left\{ \left(v_j^\top (V^\top X U) u_j \right)_{j \in [k]} : X \in \text{SO}(n) \right\} \quad (15)$$

by the fact that $\text{SO}(n) = V^\top \text{SO}(n) U$. Finally, by the assumed properties of U and V , we have that (15) is a linear image of $\pi_{\mathcal{T}}(\text{SO}(n))$ so that $\mathcal{B}(\text{SO}(n))$ is convex. ■

6 Explicit constructions for elements of $\text{SO}(n)$ with fixed strictly upper triangular entries

This section gives full characterizations and explicit constructions for $\pi_{\mathcal{T}}^{-1}(\sigma) \cap \text{SO}(n)$ and $\pi_{\mathcal{T}}^{-1}(\sigma) \cap \text{O}(n)$, where \mathcal{T} is the strictly upper triangular coordinates in $\mathbb{R}^{n \times n}$, for $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$. This will allow us to extend Theorem 5 to an approximation result in the remaining setting $\det(C) < 0$.

We overload notation below. Given $A \in \mathbb{S}_+^n$, let

$$\begin{aligned} \text{O}(A) &:= \{X \in \mathbb{R}^{n \times n} : X^\top X = A\} \\ \mathbb{B}_{\text{op}}(A) &:= \{X \in \mathbb{R}^{n \times n} : X^\top X \preceq A\}. \end{aligned}$$

Note that $\text{O}(I_n) = \text{O}(n)$ and $\mathbb{B}_{\text{op}}(I_n) = \mathbb{B}_{\text{op}}(n)$. Furthermore, if $A \in \mathbb{S}_{++}^n$, then

$$\text{int}(\mathbb{B}_{\text{op}}(A)) = \{X \in \mathbb{R}^{n \times n} : X^\top X \prec A\}$$

is full-dimensional. Thus, $\text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(A))) = \pi_{\mathcal{T}}(\text{int}(\mathbb{B}_{\text{op}}(A)))$.

We will require the following technical lemma.

Lemma 8. *Let $A \in \mathbb{S}_{+++}^n$ and $\tilde{U} \in \mathbb{R}^{n \times (n-1)}$. Suppose $\tilde{U}^\top \tilde{U} = A_{2,2}$, the bottom right $(n-1) \times (n-1)$ submatrix of A . Suppose also that the bottom $(n-1) \times (n-1)$ submatrix of \tilde{U} has full rank. Then, there exist exactly two choices of $u \in \mathbb{R}^n$ such that*

$$U = \begin{pmatrix} u & \tilde{U} \end{pmatrix} \in \text{O}(A)$$

and the two choices of u differ on their first coordinates. Furthermore, given $A_{2,2}^{-1}$, the two choices of u can be computed in $O(n^2)$ time.

Proof. Expanding the definition of U , we have that $U \in \mathcal{O}(A)$ if and only if

$$\begin{pmatrix} u^\top u & u^\top \tilde{U} \\ \tilde{U}^\top u & \tilde{U}^\top \tilde{U} \end{pmatrix} = \begin{pmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{pmatrix},$$

i.e., if and only if $\|u\|^2 = A_{1,1}$ and $\tilde{U}^\top u = A_{2,1}$.

We decompose \tilde{U}^\top as $\tilde{U}^\top = \begin{pmatrix} \hat{u} & \hat{U}^\top \end{pmatrix}$, where $\hat{u} \in \mathbb{R}^{n-1}$ and $\hat{U} \in \mathbb{R}^{(n-1) \times (n-1)}$. By assumption, \hat{U} is invertible. Thus, $\ker(\tilde{U}^\top)$ is one-dimensional and spanned by the vector

$$z := \begin{pmatrix} 1 \\ -\hat{U}^{-\top} \hat{u} \end{pmatrix}.$$

Next, note that $u_0 := \tilde{U} A_{2,2}^{-1} A_{2,1}$ satisfies $\tilde{U}^\top u_0 = A_{2,1}$. Thus, $u_0 + tz$ parameterizes the solutions of $\tilde{U}^\top u = A_{2,1}$.

Note that u_0 has squared norm

$$\|u_0\|^2 = A_{2,1}^\top A_{2,2}^{-1} (\tilde{U}^\top \tilde{U}) A_{2,2}^{-1} A_{2,1} = A_{1,2} A_{2,2}^{-1} A_{2,1} < A_{1,1},$$

where the last inequality follows by the Schur complement lemma and the assumption that $A \in \mathbb{S}_{++}^n$. We deduce that the quadratic equation $\|u_0 + tz\|^2 = A_{1,1}$ in t has exactly two solutions. In other words, there are exactly two choices of $u \in \mathbb{R}^n$ such that $U \in \mathcal{O}(A)$. Then, as $z_1 = 1$, we have that the two possible choices of u differ in their first coordinates.

We now turn to the time complexity. Note that $A_{2,2} = \tilde{U}^\top \tilde{U} = \hat{u} \hat{u}^\top + \hat{U}^\top \hat{U}$. Thus, $(\hat{U}^\top \hat{U})^{-1} = (A_{2,2} - \hat{u} \hat{u}^\top)^{-1}$. This quantity can be computed in $O(n^2)$ time given $A_{2,2}^{-1}$ using the Sherman–Morrison formula. Then, the quantity $-\hat{U}^{-\top} \hat{u}$ can be written as

$$\begin{aligned} -\hat{U}^{-\top} \hat{u} &= -\hat{U}^{-\top} \hat{U}^{-1} \hat{U} \hat{u} \\ &= -(A_{2,2} - \hat{u} \hat{u}^\top)^{-1} \hat{U} \hat{u}. \end{aligned}$$

We deduce that the quantities u_0 and z can both be computed in $O(n^2)$ time. Finally, computing the two choices of t can also be done within this time limit. \blacksquare

Remark 3. The output of the construction in Lemma 8 is continuous in \tilde{U} and A wherever it is defined. Formally, there are two continuous functions u_1 and u_2 from

$$\left\{ (\tilde{U}, A) \in \mathbb{R}^{n \times (n-1)} \times \mathbb{S}^n : \begin{array}{l} A \in \mathbb{S}_{++}^n \\ \tilde{U}^\top \tilde{U} = A_{2,2} \\ \hat{U} \text{ is invertible} \end{array} \right\}$$

to \mathbb{R}^n that track the two possible choices of u in Lemma 8. This follows from continuity of z , u_0 , and the coefficients in the quadratic equation $\|u_0 + tz\|^2 = A_{1,1}$ in the proof of Lemma 8.

The following proposition provides a parameterized construction for the entire set $\pi_{\mathcal{T}}^{-1}(\sigma) \cap \mathcal{O}(n)$.

Proposition 4. *Let $A \in \mathbb{S}_{++}^n$ and $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(A)))$. Then, $|\pi_{\mathcal{T}}^{-1}(\sigma) \cap \mathcal{O}(A)| = 2^n$. Furthermore, no two matrices in $\pi_{\mathcal{T}}^{-1}(\sigma) \cap \mathcal{O}(A)$ agree on all of their diagonal entries.*

Proof. We will induct on n . The claim is vacuously true for $n = 1$, thus assume $n \geq 2$.

Let $X \in \text{int}(\mathbb{B}_{\text{op}}(A))$ satisfy $\sigma = \pi_{\mathcal{T}}(X)$. Partition X and A as

$$X = \begin{pmatrix} \xi & x^\top \\ \bar{x} & X_{2,2} \end{pmatrix}, \quad A = \begin{pmatrix} \alpha & a^\top \\ a & A_{2,2} \end{pmatrix}.$$

As $X \in \text{int}(\mathbb{B}_{\text{op}}(A))$, we have that $A \succ X^\top X$ so that its bottom-right blocks are also ordered:

$$A_{2,2} \succ xx^\top + X_{2,2}^\top X_{2,2}.$$

Here, we have used the fact that the bottom-right block of $X^\top X$ is $xx^\top + X_{2,2}^\top X_{2,2}$.

Now, subtracting xx^\top from both sides, we have that $A_{2,2} - xx^\top \succ X_{2,2}^\top X_{2,2} \succeq 0$, i.e. $X_{2,2} \in \text{int}(\mathbb{B}_{\text{op}}(A_{2,2} - xx^\top))$ and $A_{2,2} - xx^\top$ is positive definite.

By induction, there exist exactly 2^{n-1} matrices $U_{2,2} \in O(A_{2,2} - xx^\top)$ matching the strictly upper triangular entries of $X_{2,2}$. For each of these choices, the matrix $U_{2,2}$ has rank $n - 1$. Define $\tilde{U} \in \mathbb{R}^{n \times (n-1)}$ as

$$\tilde{U} = \begin{pmatrix} x^\top \\ U_{2,2} \end{pmatrix}.$$

Note that $\tilde{U}^\top \tilde{U} = xx^\top + U_{2,2}^\top U_{2,2} = A_{2,2}$. By Lemma 8, for each choice of \tilde{U} , there are exactly two ways to append a column to the left of \tilde{U} to construct a matrix $U \in O(A)$. Furthermore, the two choices differ in their diagonal entry. Finally, we note that the strictly upper triangular entries of U match the strictly upper triangular entries of X . \blacksquare

For each $\rho \in \{\pm 1\}^n$, we may now define the map $X_\rho : \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))) \rightarrow O(n)$ to be the output of the above construction applied to $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$, where in the $k \times k$ submatrix we pick the larger (or smaller) possible diagonal entry if ρ_{n-k+1} is positive (or negative). For example, if $\sigma = 0 \in \mathbb{R}^{\binom{n}{2}}$, then $X_\rho(\sigma) = \text{Diag}(\rho) \in O(n)$. Inductively applying Remark 3, one may verify that $X_\rho(\sigma)$ is continuous as a function of $\sigma \in \text{int}(\mathbb{B}_{\text{op}}(n))$.

Lemma 9. *Given $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$ and $\rho \in \{\pm 1\}^n$, we can construct $X_\rho(\sigma)$ in time $O(n^3)$.*

Proof. We will apply the construction of Proposition 4 using Lemma 8 while recursively maintaining $A_{2,2}^{-1}$ in time $O(n^2)$. It is clear that we have access to $A_{2,2}^{-1}$ at the very top of the recursion as $A_{2,2}^{-1} = I_{n-1}^{-1} = I_{n-1}$. It remains to prove the following fact: Given $A \in \mathbb{S}^k$ and $x \in \mathbb{R}^k$ such that $A - xx^\top \succ 0$, it is possible to compute the inverse of the bottom-right $(k-1) \times (k-1)$ block of $A - xx^\top$ from the inverse of A in time $O(n^2)$.

Write

$$A - xx^\top = \begin{pmatrix} \alpha & a^\top \\ a & A_{2,2} \end{pmatrix}.$$

Note that $\alpha > 0$ by the assumption that $A - xx^\top \succ 0$. Then,

$$\begin{pmatrix} \alpha & \\ & A_{2,2} \end{pmatrix} = A - xx^\top - \begin{pmatrix} 0 & a^\top \\ a & 0 \end{pmatrix}.$$

Thus, we can compute $A_{2,2}^{-1}$ by computing the inverse of the expression on the right hand side and taking its bottom right block. This can be done via the Sherman–Morrison formula in time $O(n^2)$. Repeating once for each of the n entries results in $O(n^3)$ time in total. ■

By Proposition 4, $\text{diag}(\pi_{\mathcal{T}}^{-1}(\sigma) \cap O(n)) \subseteq \mathbb{R}^n$ is a set of 2^n distinct elements. The following result, due to Fiedler [11], allows us to deduce that the 2^n elements correspond to the vertices of a (scaled) hypercube.

Lemma 10 (Fiedler [11]). *Let $U \in O(n)$ and let $R \in \mathbb{R}^{a \times b}$ be a submatrix of U where $a + b > n$. Then, $\|R\|_{\text{op}} = 1$.*

Lemma 11. *Let $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$. There exist scalars $\alpha_i < \beta_i$ for $i \in [n]$ (independent of ρ) such that for all $\rho \in \{\pm 1\}^n$,*

$$X_{\rho}(\sigma)_{i,i} = \begin{cases} \alpha_i & \text{if } \rho_i = -1 \\ \beta_i & \text{if } \rho_i = 1 \end{cases}.$$

Proof. Fix $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$ and let $X \in \text{int}(\mathbb{B}_{\text{op}}(n))$ such that $\pi_{\mathcal{T}}(X) = \sigma$.

For $i \in [n]$, let \hat{R}_i denote the $i \times (n - i + 1)$ dimensional submatrix of X with bottom-left entry at coordinate (i, i) . Note that $\|\hat{R}_i\|_{\text{op}} \leq \|X\|_{\text{op}} < 1$ as $X \in \text{int}(\mathbb{B}_{\text{op}}(n))$. Let $R_i(s) \in \mathbb{R}^{i \times (n-i+1)}$ denote the matrix that replaces the bottom-left entry of \hat{R}_i with $s \in \mathbb{R}$. Then, $R_i(s)$ parameterizes a line that intersects the interior of the operator norm ball (which we will denote here by $\mathbb{B}_{\text{op}}(i, n - i + 1)$). As the operator norm ball is a compact convex body, there are exactly two choices of s , denoted $\alpha_i < \beta_i$, for which $\|R_i(s)\|_{\text{op}} = 1$.

Then, by Lemma 10, we deduce that $\text{diag}(\pi_{\mathcal{T}}^{-1}(\sigma) \cap O(n)) \subseteq \prod_i \{\alpha_i, \beta_i\}$. Combining with Proposition 4 completes the proof. ■

The following result states that the sign of $\det(X_{\rho}(\sigma))$ depends only on ρ .

Lemma 12. *Let $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$ and $\rho \in \{\pm 1\}^n$. Then $\det(X_{\rho}(\sigma)) = \prod_i \rho_i$.*

Proof. Fix $\rho \in \{\pm 1\}$ and $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$ and consider the continuous function $f(\alpha) := \det(X_{\rho}(\alpha\sigma))$ defined on $\alpha \in [0, 1]$. As $X_{\rho}(\alpha\sigma) \in O(n)$ for all $\alpha \in [0, 1]$, we have that f can only take on the values ± 1 . As f is also continuous, it must be constant so that $f(1) = f(0) = \det(X_{\rho}(0)) = \det(\text{Diag}(\rho)) = \prod_i \rho_i$, where we used $X_{\rho}(0) = \text{Diag}(\rho)$ as mentioned above. ■

6.1 Refinements of Theorem 5

The following theorem extends Theorem 5 to an approximation result in the remaining case to maximization over $\text{SO}(n)$ with SUT constraints and $\det(A) < 0$.

Theorem 10. *Let $A \in \mathbb{R}^{n \times n}$ be a diagonal matrix with $\det(A) < 0$ and let $\mathcal{C} \subseteq \pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n))$ be a nonempty closed convex set. Then (6) provides a $(1 - \frac{1}{n})$ -approximation of (4) in the following sense:*

$$\begin{aligned} & \max_{X \in \text{SO}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}\} \\ & \geq \left(1 - \frac{1}{n}\right) \max_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}\} + \frac{1}{n} \min_{X \in \mathbb{B}_{\text{op}}(n)} \{\langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C}\}. \end{aligned}$$

Proof. Let $\hat{\rho} = \text{sign}(\text{diag}(A))$ and $\hat{X} := \hat{X}_{\hat{\rho}} \in \mathbb{B}_{\text{op}}(n)$ maximize (6) with $\sigma = \pi_{\mathcal{T}}(\hat{X})$. We will only consider the case where $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$. The general case follows by continuity and compactness.

We will fix σ in the remainder of the proof and write X_{ρ} instead of $X_{\rho}(\sigma)$. Let (α_i, β_i) be the quantities furnished by Lemma 11 applied to σ . For $i \in [n]$, let $\rho^{(i)} \in \{\pm 1\}^n$ denote the vector that negates the i th entry of $\hat{\rho} = \text{sign}(\text{diag}(A))$, and by Lemma 12, for all $i \in [n]$, we have $X_{\rho^{(i)}} \in \text{SO}(n)$ and $\pi_{\mathcal{T}}(X_{\rho^{(i)}}) = \sigma \in \mathcal{C}$.

Then,

$$\begin{aligned} \max_{X \in \text{SO}(n)} \{ \langle A, X \rangle : \pi_{\mathcal{T}}(X) \in \mathcal{C} \} &\geq \max_{i \in [n]} \langle A, X_{\rho^{(i)}} \rangle \\ &= \langle A, \hat{X}_{\hat{\rho}} \rangle - \min_{i \in [n]} |A_{i,i}(\beta_i - \alpha_i)| \\ &\geq \langle A, \hat{X}_{\hat{\rho}} \rangle - \frac{1}{n} \left(\sum_{i=1}^n |A_{i,i}(\beta_i - \alpha_i)| \right) \\ &= \langle A, X_{\hat{\rho}} \rangle - \frac{1}{n} (\langle A, X_{\hat{\rho}} \rangle - \langle A, X_{-\hat{\rho}} \rangle) \\ &= \left(1 - \frac{1}{n} \right) \langle A, X_{\hat{\rho}} \rangle + \frac{1}{n} \langle A, X_{-\hat{\rho}} \rangle, \end{aligned}$$

where in the second line we used that $\text{diag}(X_{\rho^{(i)}})$ differs from $\text{diag}(\hat{X}) = \text{diag}(\hat{X}_{\hat{\rho}})$ only in the i -th entry, swapping α_i, β_i , and in the fourth line we used that $\hat{\rho} = \text{sign}(\text{diag}(A))$, so $\hat{X}_{\pm\hat{\rho}}$ satisfy $\pi_{\mathcal{T}}(X_{\pm\hat{\rho}}) = \sigma \in \mathcal{C}$ for $\sigma \in \text{int}(\pi_{\mathcal{T}}(\mathbb{B}_{\text{op}}(n)))$, and $\hat{X}_{\pm\hat{\rho}}$ are the maximizer and minimizer of the optimization problem Equation (5) over $\text{O}(n)$. ■

The approximation factor of $1 - \frac{1}{n}$ in the statement of Theorem 10 is optimal as can be seen by considering the case $A = \text{Diag}(1, 1, \dots, 1, -1)$ and $\mathcal{C} = \{\pi_{\mathcal{T}}(0_{n \times n})\}$. In this case, the optimal value of the optimization problem over $\text{SO}(n)$ is $n - 2$, whereas the maximum and minimum values over $\text{O}(n)$ are n and $-n$ respectively.

7 Obstructions to further generalization

This section collects a number of results showing that our hidden convexity results are essentially tight.

7.1 Maximality of Corollary 1

Recall that Corollary 1 shows any two dimensional linear image of $\text{SO}(n)$ is convex. The following result shows this is maximal in a specific sense.

Lemma 13. *For any $n \geq 3$, there exists a linear map $\pi : \text{SO}(n) \rightarrow \mathbb{R}^3$ so that $\pi(\text{SO}(n))$ is nonconvex.*

Proof. We define

$$\pi(X) = \left(X_{11}, X_{12}, \sum_{i=3}^n X_{ii} \right).$$

Let $H = \{x \in \mathbb{R}^3 : x_3 = n - 2\}$.

To see that $S = \pi(\text{SO}(n))$ is nonconvex, we show that $S \cap H$ is nonconvex. Consider a general $X \in \text{SO}(n)$. It holds that $\sum_{i=3}^n X_{ii} = n - 2$ if and only if $X_{ii} = 1$ for all $i > 2$. This occurs if and only if X is block diagonal, so that

$$X = \begin{pmatrix} \cos(\theta) & \sin(\theta) & & \\ -\sin(\theta) & \cos(\theta) & & \\ & & & \\ & & & I \end{pmatrix},$$

for some $\theta \in [0, 2\pi]$.

Now, if X has this form, then $\pi(X) = (\cos(\theta), \sin(\theta), n - 2)$, i.e.,

$$S \cap H = \{(\sin(\theta), \cos(\theta), n - 2) : \theta \in [0, 2\pi]\}.$$

This is clearly nonconvex, for example, $(0, 0, n - 2) \in \text{conv}(S \cap H) \setminus (S \cap H)$. ■

In fact, this construction gives us an example of a 2-constraint optimization problem over $\text{SO}(n)$ for which the $\text{conv}(\text{SO}(n))$ relaxation is not tight. Consider the following optimization problem:

$$\max_{X \in \text{SO}(n)} \left\{ \sum_{i=3}^n X_{ii} : \begin{array}{l} X_{1,1} = 0 \\ X_{1,2} = 0 \end{array} \right\}. \quad (16)$$

We have seen that it is not possible for a matrix in $\text{SO}(n)$ to attain a value of $n - 2$ in this problem, since any matrix in $\text{SO}(n)$ where $\sum_{i=3}^n X_{ii} = n - 2$ has the property that $X_{11}^2 + X_{12}^2 = 1$. However, $n - 2$ is attainable by a convex combination of matrices in $\text{SO}(n)$,

$$\frac{1}{2} \left(\begin{pmatrix} 1 & 0 \\ 0 & 1 \\ & & I \end{pmatrix} + \begin{pmatrix} -1 & 0 \\ 0 & -1 \\ & & I \end{pmatrix} \right).$$

Thus, the convex relaxation of (16) that replaces $\text{SO}(n)$ with $\text{conv}(\text{SO}(n))$ achieves value $n - 2$.

7.2 Maximality of [18, Theorem 8]

Horn's result [18, Theorem 8] shows that the diagonal projection of $\text{SO}(n)$ is a convex polytope. A slight modification of the construction from the previous subsection shows this is maximally convex in the following sense:

Lemma 14. *Let $A \in \mathbb{R}^{n \times n}$ be a nondiagonal matrix. Consider the linear map $\pi_A : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n+1}$, so that $\pi_A(X)_i = X_{ii}$ for $i = 1, \dots, n$, and $\pi_A(X)_{n+1} = \langle A, X \rangle$. Then $\pi_A(\text{SO}(n))$ is not convex.*

Proof. We first consider the case when A is not symmetric. By permuting coordinates, we may assume $A_{1,2} \neq A_{2,1}$. Define $H = \{x \in \mathbb{R}^{n+1} : x_i = 1 \text{ for } i = 3, \dots, n\}$ and consider $\pi(\text{SO}(n)) \cap H$.

We have seen that if $X \in \text{SO}(n)$ and $X_{ii} = 1$ for $i = 3, \dots, n$, then

$$X = \begin{pmatrix} \cos(\theta) & \sin(\theta) & & \\ -\sin(\theta) & \cos(\theta) & & \\ & & & \\ & & & I \end{pmatrix},$$

and therefore,

$$\langle A, X \rangle = (A_{11} + A_{22}) \cos(\theta) + (A_{1,2} - A_{2,1}) \sin(\theta) + \sum_{i=3}^n A_{ii}.$$

We will consider the linear map of $\pi(\text{SO}(n)) \cap H$ into \mathbb{R}^2 that sends $\pi(X)$ to

$$\left(X_{1,1}, \frac{\langle A, X \rangle - (A_{11} + A_{22})X_{1,1} - \sum_{i=3}^n A_{ii}}{A_{1,2} - A_{2,1}} \right) = (\cos(\theta), \sin(\theta)).$$

In other words, this linear map sends $\pi(\text{SO}(n)) \cap H$ to a circle. We conclude that $\pi(\text{SO}(n))$ is not convex.

Now, we consider the case when A is symmetric but not diagonal. We may permute coordinates to assume that $A_{1,2} = A_{2,1} \neq 0$. Let $D^{(i)}$ be a diagonal matrix where $D_{ii}^{(i)} = -1$ and for all $j \neq i$, $D_{jj}^{(i)} = 1$. Then, define $A' = D^{(1)}AD^{(2)}$, which is not symmetric. Note that

$$\pi_{A'}(X) = (X_{11}, X_{22}, X_{33}, \dots, X_{nn}, \langle A', X \rangle) = \tau(\pi_A(D^{(1)}XD^{(2)})),$$

where $\tau(x) = (-x_1, -x_2, x_3, \dots, x_n, x_{n+1})$. In particular, $\pi_A(\text{SO}(n))$ is convex if and only if $\pi_{A'}(\text{SO}(n))$ is convex, from which the claim follows. ■

7.3 Maximality of Corollary 2

Corollary 2 gives an example of an $m = \binom{n}{2}$ -dimensional linear projection of $\text{SO}(n)$ that is convex. The following lemma shows this is maximal in terms of dimension.

Lemma 15. *Suppose $n \geq 3$ and $\pi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ satisfies $\text{rank}(\pi) > \binom{n}{2}$. Then, $\pi(\text{SO}(n))$ is non-convex.*

Proof. It suffices to show this statement in the case where $\text{rank}(\pi) = m$. Suppose that $\pi(\text{SO}(n)) = \pi(\text{conv}(\text{SO}(n)))$ is convex. A convex set has the property that it either contains an interior point or is contained in a proper affine subspace. As $\text{conv}(\text{SO}(n))$ is full-dimensional for all $n \geq 3$ and π has full rank, we deduce that $\pi(\text{conv}(\text{SO}(n)))$ cannot be contained in a proper affine subspace of \mathbb{R}^m .

Therefore, $\pi(\text{SO}(n))$ must have an interior point and in particular, has a positive measure.

This is a contradiction to Sard's lemma: $\text{SO}(n)$ is $\binom{n}{2}$ -dimensional and \mathbb{R}^m is m -dimensional, but Sard's lemma states that the image of a manifold of dimension less than m under a smooth map into \mathbb{R}^m must have measure zero. ■

7.4 Necessity of $\det(A) \geq 0$ in Theorem 5

We have shown that if A is a diagonal matrix with $\det(A) \geq 0$, then the optimization problem

$$\max_{X \in \text{SO}(n)} \{ \langle A, X \rangle : \pi_{\mathcal{T}}(X) = \sigma \}$$

agrees with the convex relaxation replacing $\text{SO}(n)$ with $\mathbb{B}_{\text{op}}(n)$ (see Theorem 5). The following numerical example shows that the assumption that $\det(A) \geq 0$ cannot be dropped in Theorem 5 even if we strengthen the convex relaxation by replacing $\text{SO}(n)$ with $\text{conv}(\text{SO}(n))$.

The following numerical example is computed using the cvxpy convex optimization package [8] and the description of $\text{conv}(\text{SO}(3))$ given in [27][Theorem 1.3]:

$$\max_{X \in \text{SO}(3)} \left\{ X_{1,1} + X_{2,2} - X_{3,3} : \begin{array}{l} X_{1,2} = 0.5 \\ X_{1,3} = 0.3 \\ X_{2,3} = 0.2 \end{array} \right\} = 0.921,$$

whereas

$$\max_{X \in \text{conv}(\text{SO}(3))} \left\{ X_{1,1} + X_{2,2} - X_{3,3} : \begin{array}{l} X_{1,2} = 0.5 \\ X_{1,3} = 0.3 \\ X_{2,3} = 0.2 \end{array} \right\} = 1.0.$$

8 Summary and open questions

In this paper, we proved new hidden convexity results inspired by solving constrained optimization problems over $\text{SO}(n)$ and $\text{O}(n)$. These results in turn show that specific structured instances of constrained optimization over $\text{SO}(n)$ and $\text{O}(n)$ can be efficiently solved via their convex relaxations. We close by posing some natural questions surrounding hidden convexity.

Convex coordinate projections. In general, it seems to be difficult to fully characterize the possible sets of coordinates $S \subseteq [n] \times [n]$ so that the projection of $\text{SO}(n)$ onto the coordinates in S is convex.

We will note some basic invariants of this question: clearly if $\pi_S(\text{SO}(n))$ is convex, then for all $T \subseteq S$, $\pi_T(\text{SO}(n))$ is convex. We say that S has a property *up to permutation* if there are permutations σ and ρ so that $\{(\sigma_i, \rho_j) : (i, j) \in S\}$ has this property. Similarly we say that S has a property *up to transposition* if either S or $\{(j, i) : (i, j) \in S\}$ has this property. Clearly, S has the property that $\pi_S(\text{SO}(n))$ is convex if and only if it has this property up to permutations and transposition.

Note also that by Lemma 10, that $\pi_S(\text{SO}(n))$ is not convex if S contains a rectangle of size $a \times b$ where $a + b > n$. Here, by rectangle we mean a subset of S of the form $A \times B \subseteq S$ where $A, B \subseteq [n]$ and $|A| = a$ and $|B| = b$.

Using this idea with additional casework (which we feel is ultimately uninformative) we can obtain the following characterization of the coordinate subsets of $[4] \times [4]$ so that $\pi_S(\text{SO}(4))$ is convex:

Lemma 16. *Let $S \subseteq [4] \times [4]$ be such that $\pi_S(\text{SO}(4))$ is convex. Then (up to permutations and transpositions), S is a subset of one of the following:*

- $T = \{(i, j) : i < j \in [4]\}$
- $D = \{(i, i) : i \in [4]\}$
- $F = \{(1, 1), (1, 2), (2, 3), (2, 4)\}$.

The structure of these examples suggest that there may be some rich combinatorial information hidden in the question of whether or not a given coordinate projection of $\text{SO}(n)$ is convex. In particular, we suspect the following: consider the decision problem **CONVEX** whose input is a set $S \subseteq [n] \times [n]$, and whose output is TRUE if $\pi_S(\text{SO}(n))$ is convex, and FALSE otherwise.

Conjecture 1. *The problem **CONVEX** is NP-hard.*

We will remark that it is not even clear if this problem is in NP, as there does not seem to be an efficient witness to the fact that $\pi_S(\text{SO}(n))$ is convex. We note that determining whether $S \subseteq [n] \times [n]$ is (up to permutations and transpositions) a subset of the upper triangular entries of $[n] \times [n]$ is NP-hard [10].

Small semidefinite representation of two-dimensional images. It is known that the smallest equivariant (see [27] for a definition) semidefinite representation of $\text{conv}(\text{SO}(n))$ is exponential in

size [27]. We leave open the question of whether $\pi(\text{SO}(n))$, where $\pi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^2$, may have a small (possibly linear-sized) semidefinite representation.

Hidden convexity of multiple copies of $\text{SO}(n)$. Finally, we also leave the study of convex images of direct products of $\text{SO}(n)$ to future work. Such results may be useful in applications such as cryo-EM [1], where the optimization problems contain multiple $\text{SO}(n)$ matrices.

Acknowledgments

Akshay Ramachandran was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme: AR from grant agreement no. 805241-QIP. Kevin Shu was supported by the ACO-ARC fellowship while conducting this work. Part of this work was completed while Alex L. Wang was supported by the Dutch Scientific Council (NWO) grant OCENW.GROOT.2019.015 (OPTIMAL).

References

- [1] A. S. Bandeira, Y. Chen, R. R. Lederman, and A. Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *Inverse Problems*, 36(6):064002, 2020.
- [2] A. Barvinok. *A course in convexity*, volume 54. Amer. Math. Soc, 2002.
- [3] G. Blekherman, G. Smith, and M. Velasco. Sums of squares and varieties of minimal degree. *J. Amer. Math. Soc.*, 29(3):893–913, 2016.
- [4] L. Brickman. On the field of values of a matrix. *Proc. Amer. Math. Soc.*, 12(1):61–66, 1961.
- [5] L. Brynte, V. Larsson, J. P. Iglesias, and C. Olsson and F. Kahl. On the tightness of semidefinite relaxations for rotation estimation. *J. Math. Imaging Vision*, pages 1–11, 2022.
- [6] N. Chan and K. Li. Diagonal elements and eigenvalues of a real symmetric matrix. *J. Math. Anal. Appl.*, 91(2):562–566, 1983.
- [7] Y. Chen, S. Huang, and R. Fitch. Active SLAM for mobile robots with area coverage and obstacle avoidance. *IEEE/ASME Trans. Mechatronics*, 25(3):1182–1192, 2020.
- [8] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *J. Mach. Learn. Res*, 17(1):2909–2913, 2016.
- [9] L. L. Dines. On the mapping of quadratic forms. *Bull. Amer. Math. Soc.*, 47(6):494–498, 1941.
- [10] G. Fertin, I. Rusu, and S. Vialette. Obtaining a triangular matrix by independent row-column permutations. In *Algorithms and Computation: 26th International Symposium, ISAAC 2015*, pages 165–175, 2015.
- [11] M. Fiedler. Suborthogonality and orthocentricity of matrices. *Linear Algebra Appl.*, 430(1): 296–307, 2009.
- [12] A. Fradkov and V. Yakubovich. The S-procedure and duality relations in nonconvex problems of quadratic programming. *Vestn. LGU, Ser. Mat., Mekh., Astron.*, 1:101–109, 1979.
- [13] K. Gilman, S. Burer, and L. Balzano. A semidefinite relaxation for sums of heterogeneous quadratics on the stiefel manifold. *arXiv preprint arXiv:2205.13653*, 2022.
- [14] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.
- [15] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.
- [16] V. Guillemin and R. Sjamaar. *Convexity Properties of Hamiltonian Group Actions*. AMS, 2005.
- [17] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [18] A. Horn. Doubly stochastic matrices and the diagonal of a rotation matrix. *Amer. J. Math.*, 76(3):620–630, 1954.
- [19] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 2012.

- [20] R. G. Jeroslow. On defining sets of vertices of the hypercube by linear inequalities. *Discrete Math.*, 11(2):119–124, 1975.
- [21] J. Kiefer. Sequential minimax search for a maximum. *Proc. Amer. Math. Soc.*, 4(3):502–506, 1953.
- [22] G. Lancia and P. Serafini. *The Parity Polytope*, pages 113–121. Springer International Publishing, 2018.
- [23] U. Lee and M. Mesbahi. Spacecraft reorientation in presence of attitude constraints via logarithmic barrier potentials. In *Proceedings of the 2011 American Control Conference*, pages 450–455. IEEE, 2011.
- [24] H. Maron, N. Dym, I. Kezurer, S. Kovalsky, and Y. Lipman. Point registration via efficient convex relaxation. *Trans. Graphics*, 2016.
- [25] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *Trans. Graphics*, 2012.
- [26] I. Pólik and T. Terlaky. A survey of the S-lemma. *SIAM Rev.*, 49(3):371–418, 2007.
- [27] J. Saunderson, P. Parrilo, and A. S. Willsky. Semidefinite descriptions of the convex hull of rotation matrices. *SIAM J. Optim.*, 25(3):1314–1343, 2015.
- [28] T-Y Tam. On the shape of numerical ranges associated with Lie groups. *Taiwanese J. Math.*, 5(3):497–506, 2001.
- [29] G. Wahba. A least squares estimate of satellite attitude. *SIAM Rev.*, 7(3):409–409, 1965.
- [30] Y. Xia. A survey of hidden convex optimization. *J. Oper. Res. Soc. China*, 8(1):1–28, 2020.

A Separation and optimization oracles for the parity polytope

It is possible to implement separation and optimization oracles for PP_n that run in $O(n \log n)$ time.

Separation. We will use the following description of PP_n given in [20, 22]:

$$PP_n := \{x \in [-1, 1]^n : \langle x, 1_n - 2 \cdot 1_S \rangle \leq n - 2, \quad \forall \text{ odd } S \subseteq [n]\}.$$

Here, we will say that $S \subseteq [n]$ is odd if $|S|$ is odd. Else, S is even. The set of constraints can be rewritten as

$$\min_{\text{odd } S \subseteq [n]} \langle x, 1_S \rangle \geq \frac{1}{2}(\langle x, 1_n \rangle - (n - 2)).$$

In $O(n \log n)$ time, we may sort the entries of x and compute the sums of all odd-length prefixes of the sorted vector. If every sum is at least $\frac{1}{2}(\langle x, 1_n \rangle - (n - 2))$, then $x \in PP_n$. Otherwise, we have found a separating hyperplane.

Optimization. We will use the vertex description

$$PP_n := \text{conv} \{1_n - 2 \cdot 1_S : \text{even } S \subseteq [n]\}.$$

Then, given $w \in \mathbb{R}^n$, we may optimize $\max_{x \in PP_n} \langle w, x \rangle$ by solving $\min_{\text{even } S \subseteq [n]} \langle w, 1_S \rangle$. We can construct a minimizer of the latter problem in $O(n \log n)$ time by sorting w and computing the even-length prefix sums of the sorted vector.

B Connections with quadratic convexity theorems

This appendix interprets hidden convexity results on $\text{SO}(n)$ as quadratic convexity results on the unit sphere.

A basic result in the Lie group theory of $\text{SO}(n)$ is the existence of a quadratic map $Q : \mathbb{R}^{2^{n-1}} \rightarrow \mathbb{R}^{n \times n}$ and a subset $\text{Spin}(n)$ of the unit sphere in $\mathbb{R}^{2^{n-1}}$ such that $Q(\text{Spin}(n)) = \text{SO}(n)$. This map is quadratic in the sense that there exists a collection of n^2 symmetric matrices $\{A_{ij}\} \subseteq \mathbb{S}^{2^{n-1}}$ indexed by $(i, j) \in [n]^2$ such that

$$(Q(x))_{i,j} = \langle x, A_{ij}x \rangle.$$

This result and its construction are explained in detail in [27, Appendix A]. It is additionally shown in [27, Theorem 1.1] that for any $Y \in \mathbb{R}^{n \times n}$,

$$\max_{X \in \text{SO}(n)} \langle Y, X \rangle = \max_{x \in \text{Spin}(n)} \langle Y, Q(x) \rangle = \max_{x \in \mathbf{S}^{2^{n-1}-1}} \langle Y, Q(x) \rangle. \quad (17)$$

Now, let $\pi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ be a linear function. Then,

$$\pi(\text{SO}(n)) = (\pi \circ Q)(\text{Spin}(n)) \subseteq (\pi \circ Q)(\mathbf{S}^{2^{n-1}-1}) \subseteq \text{conv}(\pi(\text{SO}(n))).$$

Here, the last inclusion follows by (17).

We deduce that if $\pi(\text{SO}(n))$ is convex, then equality holds throughout this chain and the image of the unit sphere $\mathbf{S}^{2^{n-1}-1}$ under the quadratic map $\pi \circ Q$ is convex. For example, combined with Corollary 2, we have that

$$\left\{ (\langle x, A_{ij}x \rangle)_{i < j} : x \in \mathbf{S}^{2^{n-1}-1} \right\} \subseteq \mathbb{R}^{\binom{n}{2}}$$

is convex. As another example, combined with [18, Theorem 8], we have that

$$\left\{ \begin{pmatrix} \langle x, A_{11}x \rangle \\ \vdots \\ \langle x, A_{nn}x \rangle \end{pmatrix} : x \in \mathbf{S}^{2^{n-1}-1} \right\} \subseteq \mathbb{R}^n$$

is convex (and equal to the polytope PP_n).