

Symmetries and Fixed Points of Iterated Linear Optimization on the Elliptope

by

Heon Lee

Sc. B., Brown University, 2026

A Thesis submitted in partial fulfillment of the requirements for Honors
in the Department of Computer Science at Brown University



BROWN

Department of Computer Science

Brown University

Providence, Rhode Island

April 2026

Acknowledgements

Thank you, Professor Pedro Felzenszwalb, for your guidance over the past four years at Brown. I can never be more grateful for all that I've learned from you. Your willingness to explore every idea, regardless of the outcome, has shaped the way I approach challenges. I look forward to carrying the lessons from our work into my next chapter. I would also like to thank Professor Alice Paul for her time and for serving as the second reader for this thesis. Finally, thank you to my friends and family who were always there to support me. This journey would not have been possible without your constant support.

Contents

1	Introduction	1
1.1	Contributions	2
2	Iterated Linear Optimization	3
2.1	The Fixed Point Iteration Process	3
2.2	Symmetries and Equivariance	4
3	Iterated Linear Optimization on the Elliptope	7
3.1	Symmetries of the Elliptope	8
3.2	Fixed Points of the Elliptope	12
3.2.1	Definitions and Properties	12
3.2.2	Examples in Low Dimensions	13
3.3	Bijection for Irreducible Fixed Points	14
3.4	Iteration on Circulant Symmetric Matrices	19
3.4.1	Cyclic Graphs	22
4	Conclusions	24
A	Symmetries of the 3-elliptope	25
	Bibliography	29

Chapter 1

Introduction

Semidefinite programming (SDP) is an optimization technique for optimizing a linear function, subject to the constraint that an affine combination of symmetric matrices involving decision variables must be positive semidefinite (PSD) [12]. It generalizes linear programming (LP) and convex quadratic programming (QP) by incorporating constraints beyond linear inequalities. In LP, optimization is performed over a polyhedral feasible region defined by linear inequalities, and the variables are vectors constrained elementwise. In convex QP, optimization involves a quadratic objective function with linear constraints, but the feasible region remains restricted to vector decision variables.

SDP extends these frameworks by allowing matrix variables and PSD constraints on those matrices. Instead of elementwise scalar inequalities, SDP uses matrix inequalities, where the solution must lie in the cone of PSD matrices. This flexibility has led to its use across various fields, including control theory [13], where SDP is used in the context of linear matrix inequalities, and quantum mechanics [14], where it plays a role in quantum state estimation and error correction.

A significant strength of SDP is that it can be efficiently solved using interior point methods [10]. As a result, in many instances, by transforming discrete constraints into convex ones, we can relax many intractable combinatorial optimization problems into SDP problems for an approximate solution. A seminal application of SDP relaxation for combinatorial optimization problems was demonstrated by Goemans and Williamson [5] for the Max-Cut problem in 1994. The Max-Cut problem involves partitioning the vertices of a graph $G = (V, E)$ into two disjoint sets $(S, V \setminus S)$ to maximize the cut,

$$\delta(S) := \sum_{(u,v) \in E} |\mathbb{1}_{u \in S} - \mathbb{1}_{v \in S}|,$$

defined as the sum of the weights of edges crossing the partition. Goemans and Williamson introduced a 0.878-approximation algorithm by solving an SDP relaxation. They then used a randomized hyperplane rounding technique to obtain a feasible cut/integer solution from the relaxed solution. While the method is effective, this probabilistic method introduces randomness.

Recently, Felzenszwalb et al. [3] introduced a deterministic rounding method based on a fixed point iteration process. This process repeatedly applies a linear transformation T over a convex and compact domain. They show that limit points of the resulting sequence generated by iteratively applying T are fixed points of T . A special case of the map is when it is taken over the elliptope. In this context, these attractive fixed points correspond to its vertices, which represent the integer solutions to the Max-Cut problem. This deterministic approach offers an alternative to randomized rounding.

Building on their deterministic rounding method, Felzenszwalb et al. [4] applied the fixed-point iteration

process for a novel clustering algorithm via an SDP relaxation of the Max- k -Cut problem. They demonstrate that the vertices of the Max- k -Cut relaxation corresponds to partitions of the data into at most k clusters. Additionally, these vertices are again the attractive fixed points of the iterative process. They suggest, through experiments, that the fixed-point iteration for rounding the SDP relaxation yields better clustering results compared to randomized rounding.

Despite the experimental success (see e.g., [4]) of the fixed-point iteration method, an understanding of its dynamics remains limited. Beyond the general fact that limit points must be fixed points, much less is known about the necessary conditions for convergence, the geometry and organization of the fixed point set itself, and how the iteration is constrained by the structure of the feasible region. Alongside general guarantees for arbitrary compact convex domains, our main focus is the ellipsope, where we analyze how its geometry and symmetries organize the fixed points of the iteration.

In this thesis, we aim to address these gaps by studying the geometry of the ellipsope and how its symmetries interact with the iterative map T . Our main objectives are:

- To analyze the fixed points of the iterative transformation T on the ellipsope, with a focus on how symmetries influence these fixed points;
- To classify the symmetries of the ellipsope, drawing parallels to the symmetries of the cut polytope as studied by Deza et al. [2], and to understand the implications of these symmetries for the iterative process;
- To investigate the dynamics of iterated linear optimization.

1.1 Contributions

Our main contributions are as follows:

1. We introduce a related single-valued variant T' by selecting the minimizer of the L^2 norm within the image of T . We prove that the fixed points of T and T' coincide, simplifying the analysis of the iterative process;
2. We show that $T(T')$ is $\text{LinSym}(\Delta)$ -equivariant for the group of symmetries $\text{LinSym}(\Delta)$ (Proposition 2.2.3), and demonstrate that the property of being fixed by a symmetry $f \in \text{LinSym}(\Delta)$ is a loop invariant under iteration of $T(T')$;
3. We provide an alternative proof of the classification of the symmetries of the n -ellipsope for $n \neq 4$ (Theorem 3.1.7) by showing that its symmetry group is equivalent to that of the n -cut-polytope;
4. We establish a bijection between irreducible fixed points of different ranks (Proposition 3.3.8), demonstrating that the same symmetries fix them and are geometrically related to each other as antipodal points on the boundary of the ellipsope (Proposition 3.3.12);
5. Using Fourier analysis, we investigate circulant matrices and show that for these matrices, the iterative map T' converges to a fixed point in a single iteration (Proposition 3.4.4).

Chapter 2

Iterated Linear Optimization

2.1 The Fixed Point Iteration Process

In this section, we outline the fixed point iteration process as described by Felzenszwalb et al. [3] and introduce an alternative but related single-valued map.

We work in a compact and convex domain $\Delta \subset \mathcal{H}$ in some arbitrary Hilbert space $\mathcal{H}$ over $\mathbb{R}$. The map $T : \Delta \rightarrow \mathcal{P}(\Delta)$ is defined by maximizing the inner product with the input over Δ :

$$T(x) := \operatorname{argmax}_{y \in \Delta} \langle x, y \rangle.$$

Here, $T(x)$ is set-valued because multiple elements can maximize the inner product. For example, taking Δ to be the unit disk in $\mathbb{R}^2$, $T(O) = \Delta$ where O denotes the origin. Furthermore, it is nonempty since the inner product is a continuous function and the Extreme Value Theorem implies the attainment of the supremum in the compact domain Δ .

The fixed point iteration is then defined as

$$x_{i+1} \in T(x_i),$$

which generates a sequence of points

$$\{x_n\}_{n \in \mathbb{N}}.$$

Each step of the iteration, $x_{i+1} \in T(x_i)$, involves selecting an arbitrary element of $T(x_i)$. To remove the inherent ambiguity in selecting an element from $T(x_i)$, we propose selecting the minimizer of the norm induced by the inner product. To show that the approach is well-defined, we show that the norm minimizer must be unique.

We first demonstrate that $T(x)$ is a compact and convex set for any $x \in \Delta$.

Lemma 2.1.1. *$T(x)$ is compact and convex.*

Proof. We first show that $T(x)$ is convex. For any $y, z \in T(x)$, there exists $c \in \mathbb{R}$ such that $c = \langle x, y \rangle = \langle x, z \rangle$. For any $t \in [0, 1]$, convexity of Δ implies that $ty + (1 - t)z \in \Delta$. Furthermore,

$$\begin{aligned} \langle x, ty + (1 - t)z \rangle &= t\langle x, y \rangle + (1 - t)\langle x, z \rangle \\ &= tc + (1 - t)c \\ &= c. \end{aligned}$$

Hence, $ty + (1 - t)z \in T(x)$, proving $T(x)$ is convex.

We now show that $T(x)$ is compact. It suffices to show closedness, since a closed subset of a compact set is compact. To show that $T(x)$ is closed, let $\{y_i\}_{i \in \mathbb{N}}$ be a sequence in $T(x)$ converging to $y \in \mathcal{H}$. By compactness of Δ , $y \in \Delta$. Because $f_x(z) := \langle x, z \rangle$ is continuous, we know that the sequence $\{\langle x, y_i \rangle\}_{i \in \mathbb{N}}$ converges to $\langle x, y \rangle$. Given $\langle x, y_i \rangle = c$ for all i , it follows $\langle x, y \rangle = c$. Thus, $y \in T(x)$, showing that $T(x)$ is closed. $\square$

By the Hilbert Projection Theorem, we know that $T(x)$ has a unique norm minimizer. Therefore, we have the following well-defined map $T' : \Delta \rightarrow \Delta$ defined as follows:

$$T'(x) := \operatorname{argmin}_{y \in T(x)} \|y\|_2.$$

Our goal is to analyze the fixed points of T . As we show, the fixed points of T and T' are the same.

Proposition 2.1.2. *The fixed points of T and T' coincide.*

Proof. By definition, the set of fixed points of T' is a subset of the set of fixed points of T since $x = T'(x) \in T(x)$.

For the other inclusion, let $x \in T(x)$ be a fixed point of T . Let $y = T'(x)$. Because

$$\langle x, y \rangle = \langle x, x \rangle,$$

then

$$\begin{aligned} 0 &\leq \|y - x\|^2 \\ &= \|y\|^2 - 2\langle x, y \rangle + \|x\|^2 \\ &= \|y\|^2 - \|x\|^2, \end{aligned}$$

implying $\|x\|^2 \leq \|y\|^2$. Since y is the unique norm minimizer, we conclude $x = y$ as desired. $\square$

2.2 Symmetries and Equivariance

This section studies how symmetries of the domain Δ interact with iterating the maps T and T' . In particular, we show that T (T') preserves the set of fixed points of any symmetry. In other words, if we start our iteration at a point that is unchanged by a symmetry (a “fixed point” of the symmetry), then the iteration will remain within the set of such points. This constitutes a loop invariant. We will also show that symmetries preserve the set of fixed points of T (T'), generalizing to the notion of equivariant maps.

Define $\operatorname{Sym}(\Delta)$ as the group of symmetries of Δ , consisting of isometries of $\mathcal{H}$ that preserve Δ . The symmetries are affine by the Mazur-Ulam Theorem [8]. For our purposes, we focus on the subgroup $\operatorname{LinSym}(\Delta)$ of linear isometries that preserve Δ .

A point x is a fixed point of a symmetry f if $f(x) = x$. It is important to distinguish between these fixed points of the symmetry f and the fixed points of the iterative maps T and T' that we are ultimately interested in.

The linearity of f implies that the origin is a fixed point of f . This means that the norm is preserved since f is an isometry. This further implies that the inner product is preserved:

$$\langle x, y \rangle = \langle f(x), f(y) \rangle$$

for every $x, y \in V$.

We now show that the maps T and T' are $\operatorname{LinSym}(\Delta)$ -equivariant.

Definition 2.2.1 (G -set). Let G be a group. A G -set is a pair $(X, \cdot)$ consisting of a set X and a group action $\cdot : G \times X \rightarrow X$ such that for all $x \in X$ the following hold:

1. $e \cdot x = x$ where $e \in G$ is the identity element (Identity)
2. $(gh) \cdot x = g \cdot (h \cdot x)$ for all $g, h \in G$ (Compatibility)

Definition 2.2.2 (Equivariance). Let G be a group and $(X, \cdot_X), (Y, \cdot_Y)$ be G -sets. A function $f : X \rightarrow Y$ is said to be G -equivariant if for every $g \in G$ and every $x \in X$,

$$f(g \cdot_X x) = g \cdot_Y f(x).$$

To simplify notation, we may omit the subscripts and, where the context is clear, the group action symbol entirely, writing $f(gx) = gf(x)$.

Proposition 2.2.3. T and T' are $\text{LinSym}(\Delta)$ -equivariant.

Proof. Let $f \in \text{LinSym}(\Delta)$. As an isometry, f is invertible. It suffices to show $T = f^{-1} \circ T \circ f$ and $T' = f^{-1} \circ T' \circ f$.

We start by showing that

$$f \left(\underset{y \in \Delta}{\operatorname{argmax}} \langle x, y \rangle \right) = \underset{y \in \Delta}{\operatorname{argmax}} \langle x, f^{-1}(y) \rangle.$$

There exists a constant $c \in \mathbb{R}$ such that for any $z \in \underset{y \in \Delta}{\operatorname{argmax}} \langle x, y \rangle$, we have $\langle x, z \rangle = c$. Clearly $f(z)$ is a maximizer of $\langle x, f^{-1}(y) \rangle$, since $c = \langle x, z \rangle = \langle x, f^{-1}(f(z)) \rangle$. This implies that $f \left(\underset{y \in \Delta}{\operatorname{argmax}} \langle x, y \rangle \right) \subset \underset{y \in \Delta}{\operatorname{argmax}} \langle x, f^{-1}(y) \rangle$. Similarly, letting z be the maximizer of $\langle x, f^{-1}(y) \rangle$, we see that $f^{-1}(z)$ is the maximizer of $\langle x, y \rangle$, implying that $f \left(\underset{y \in \Delta}{\operatorname{argmax}} \langle x, y \rangle \right) \supset \underset{y \in \Delta}{\operatorname{argmax}} \langle x, f^{-1}(y) \rangle$.

We now observe that because inner product is preserved by f ,

$$\langle x, f^{-1}(y) \rangle = \langle f(x), f(f^{-1}(y)) \rangle = \langle f(x), y \rangle$$

for every $x, y \in \Delta$.

Thus, f and T commute as follows:

$$f(T(x)) = f \left(\underset{y \in \Delta}{\operatorname{argmax}} \langle x, y \rangle \right) = \underset{y \in \Delta}{\operatorname{argmax}} \langle x, f^{-1}(y) \rangle = \underset{y \in \Delta}{\operatorname{argmax}} \langle f(x), y \rangle = T(f(x)).$$

Since f preserves the norm, we know that $f(T'(x))$ is the element of smallest norm in the set $f(T(x))$. Similarly, $T'(f(x))$ is the element of smallest norm in $T(f(x))$. Then $f(T(x)) = T(f(x))$ implies that $f(T'(x)) = T'(f(x))$. Thus, T and T' are $\text{LinSym}(\Delta)$ -equivariant. $\square$

Given that the operators f and $T(T')$ commute, we can iteratively apply commutations between f and $T(T')$.

Corollary 2.2.4. $T^k = f^{-1} \circ T^k \circ f$ and $T'^k = f^{-1} \circ T'^k \circ f$ for any $k \in \mathbb{N}$.

As illustrated in Figure 2.1, the functions f and T^k commute, meaning $T^k(x) = f^{-1} \circ T^k \circ f(x)$ for any $k \in \mathbb{N}$. This implies that iterating from x is equivalent to iterating from $f(x)$ or $f^{-1}(x)$, followed by applying f^{-1} or f respectively.

$$\begin{array}{ccccccc} x_1 & \xrightarrow{T} & x_2 & \xrightarrow{T} & \cdots & \xrightarrow{T} & x_{k-1} & \xrightarrow{T} & x_k \\ f \downarrow & & f \downarrow & & & & f \downarrow & & f \downarrow \\ f(x_1) & \xrightarrow{T} & f(x_2) & \xrightarrow{T} & \cdots & \xrightarrow{T} & f(x_{k-1}) & \xrightarrow{T} & f(x_k) \end{array}$$

Figure 2.1: Commutative diagram illustrating the interaction between f and T^k .

Corollary 2.2.5. *If $f \in \text{LinSym}(\Delta)$ and $x \in \Delta$ is a fixed point of f , then $T'(x)$ is also a fixed point of f .*

Proof. This result follows directly from the commutativity of f and T' :

$$T'(x) = T'(f(x)) = f(T'(x)).$$

□

Because $T'(x) \in T(x)$, it follows that fixed points of f are also preserved under T . In other words, if $x = f(x)$, then there exists $y \in T(x)$ such that $y = f(y)$. Specifically, $y = T'(x) \in T(x)$.

The significance of Corollary 2.2.5 becomes apparent when we consider the iterative sequence generated by T' . If we start with a fixed point of f , x_0 , then the corollary implies $x_k = T'^k(x_0)$ remains a fixed point of f by induction.

Consequently, the property of “being a fixed point of f ” is maintained throughout the iterative process under T' , provided we begin at such a point. In the language of iterative algorithms, we can say that “the property of being a fixed point of f is a loop invariant for the iterative process under T' , when initialized at a fixed point of f .”

This loop invariant perspective is insightful because it tells us that if we initiate our iteration within the set of fixed points of f , the entire trajectory under repeated application of T' will be confined to this set. This simplifies the analysis of the iterative dynamics in certain contexts. We will see an application of this principle in Section 3.4.

The symmetries also preserve the attractiveness of fixed points of $T(T')$. We use the following definition of attractive fixed points as given in [3].

Definition 2.2.6. (Attractive fixed point) Let x be a fixed point of T . Then x is attractive if there exists $\epsilon > 0$ such that for any x_0 in an ϵ -neighborhood of x , x_0 iterates to x .

Definition 2.2.7. (Non-attractive fixed point) Let x be a fixed point of T . Then x is non-attractive if it is not an attractive fixed point.

Lemma 2.2.8. *If $f \in \text{LinSym}(\Delta)$, then x is a fixed point of $T(T')$ if and only if $f(x)$ is a fixed point of $T(T')$.*

Proof. It suffices to prove both statements in one direction because $f^{-1} \in \text{LinSym}(\Delta)$. By applying f to both sides, $f(x) = T'(x)$. Then by Proposition 2.2.3, $f(x) = T'(f(x))$, meaning $f(x)$ is a fixed point of $T(T')$. □

Corollary 2.2.9. *If $f \in \text{LinSym}(\Delta)$ and x is a fixed point of $T(T')$, then x and $f(x)$ share the same attractiveness.*

Proof. By Lemma 2.2.8, $f(x)$ is a fixed point of $T(T')$. To show that x and $f(x)$ share the same attractiveness, it suffices to show that the behavior of points in a small enough neighborhood of x and $f(x)$ are the same.

Let $N(x, \epsilon)$ be the intersection of an ϵ -ball around x and Δ . We know that $N(f(x), \epsilon) = f(N(x, \epsilon))$. Now, let $\{x_i\}$ be a sequence generated from $x_0 \in N(x, \epsilon)$ such that $x_{i+1} \in T(x_i)$. Then $\{y_i\} = \{f(x_i)\}$ is a possible sequence generated from $f(x_0) \in N(f(x), \epsilon)$ such that $y_{i+1} \in T(y_i)$. This is because $x_k \in T^k(x_0)$ if and only if $f(x_k) \in T^k(f(x_0))$ by Corollary 2.2.4.

Thus, all sequences starting in neighborhoods of x and $f(x)$ have the same behavior in terms of their distances to x and $f(x)$, respectively.

□

Chapter 3

Iterated Linear Optimization on the Elliptope

For the remainder of this thesis, we turn our attention to the fixed point iteration process applied to the Goemans-Williamson algorithm [5, 3]. To facilitate this analysis, we will focus our study on the elliptope. The elliptope is a geometric object of interest due to its role in SDP relaxations of combinatorial optimization problems (e.g., [5, 6]). Within this context of the Goemans-Williamson algorithm and the elliptope, we will further investigate the properties of fixed points and the dynamics of the fixed point iteration process.

The Max-Cut problem is a canonical example of SDP relaxation. Consider an undirected graph $G = (V, E)$ with n vertices and nonnegative edge weights $w_{ij} \geq 0$. The Max-Cut problem seeks a partition $S \subset V, V \setminus S$ such that the sum of the weights between the partitions is maximized. Formally,

$$\text{maximize } \frac{1}{4} \sum_{i,j} w_{ij} (1 - x_i x_j) \text{ subject to } x_i \in \{-1, +1\}. \quad (3.1)$$

The value of x_i correspond to the partitions, i.e., $S = \{i \in [n] : x_i = 1\}$ and $V \setminus S = \{i \in [n] : x_i = -1\}$.

Definition 3.0.1. (S_n) Let $S_n \subset \mathbb{R}^{n \times n}$ denote the set of $n \times n$ symmetric matrices.

The space S_n , equipped with the Frobenius inner product, forms an $\frac{n(n+1)}{2}$ -dimensional Hilbert space.

Definition 3.0.2. (Elliptope) The n -elliptope, denoted $\mathcal{L}_n$, is the set of $n \times n$ positive semidefinite (PSD) matrices with unit diagonal elements:

$$\mathcal{L}_n := \{X \in S_n : X \geq 0, X_{ii} = 1 \forall i \in \{1, \dots, n\}\}.$$

The elliptope is an $\binom{n}{2}$ -dimensional convex subset of S_n . The elements of the elliptope are called correlation matrices.

The Max-Cut problem is NP-hard. The Goemans-Williamson (GW) algorithm circumvents this difficulty by relaxing the discrete constraints to continuous ones. Instead of scalar variables $x_i \in \{-1, +1\}$, the algorithm associates a unit vector $v_i \in \mathbb{R}^n$ with each vertex $i \in V$. This transforms the integer quadratic program into a vector optimization problem:

$$\text{maximize } \frac{1}{4} \sum_{i,j} w_{ij} (1 - \langle v_i, v_j \rangle) \text{ subject to } v_i \in \mathbb{R}^n, |v_i| = 1. \quad (3.2)$$

There is a natural correspondence between the set of these unit vectors and the set of correlation matrices. Specifically, by defining the Gram matrix X such that $X_{ij} = \langle v_i, v_j \rangle$, we can equivalently formulate this problem as an

SDP over the elliptope, i.e.,

$$\text{maximize } \langle X, -W \rangle \text{ subject to } X \in \mathcal{L}_n. \quad (3.3)$$

The corresponding SDP can be solved to arbitrary precision in polynomial time. The final discrete partition is then obtained by randomly sampling a normal vector n uniformly on the unit sphere, and then partitioning by $S = \{i \in [n] : \langle n, v_i \rangle \geq 0\}$ and $V \setminus S$.

Felzenszwalb et al. [3] suggest an alternative rounding method using the fixed point iteration with T . This motivates the study of the dynamics of the iteration.

The remainder of this chapter is organized as follows. In Section 3.1, we classify the linear symmetries of the elliptope and discuss their relationship to the cut polytope. Section 3.2 defines and characterizes the fixed points of the elliptope. In Section 3.3, we establish a bijection between irreducible fixed points of different ranks. Finally, Section 3.4 investigates the dynamics of the iteration on circulant symmetric matrices.

3.1 Symmetries of the Elliptope

We begin by studying the symmetries of the elliptope through a closely related object, the cut polytope.

The cut polytope represents the convex hull of all cut vectors (or cut matrices) corresponding to all possible bipartitions (cuts) of a graph. Hence, it is the feasible region of the Max-Cut problem. Moreover, its extreme points are its vertices, corresponding to an incidence vector of a cut. The elliptope serves as a convex relaxation of the cut polytope and forms the feasible region for the SDP relaxation of the Max-Cut problem. Understanding the geometric and algebraic properties of the elliptope is essential for analyzing SDP relaxations and providing theoretical guarantees for related algorithms.

Deza et al. [2] conducted an extensive study on the symmetries of the cut polytope and its relatives, including the equicut polytope, k -cut polytope, and the Boolean quadric polytope.

Tropp [11] explored the geometric aspects of the elliptope in his study of the simplicial faces of the set of correlation matrices. He demonstrated that the elliptope approximates the cut polytope and the conditions under which sets of vertices of the elliptope generate simplicial faces. This further highlights how the elliptope approximates the cut polytope. Given these close relationships and the importance of symmetries in understanding geometric structures, it is natural to study the symmetries of the elliptope.

While the linear symmetries of the elliptope have been fully characterized by Li and Pierce [7] using operator theory, we present an alternative derivation. Our proof leverages the geometric relationship between the elliptope and the cut polytope. Unlike the operator-theoretic approach, our method relies on the symmetries of the underlying combinatorial polytope. Consequently, our proof is unable to resolve the specific case of $n = 4$, where the cut polytope possesses additional symmetries that do not extend to the elliptope. However, this approach offers a distinct geometric perspective that reinforces the structural connection between the relaxation and the original problem.

We first observe that the group of symmetries of $\mathcal{L}_n$ is finite.

Lemma 3.1.1. *$\text{LinSym}(\mathcal{L}_n)$ is finite.*

Proof. By Lemma 2.2.8, the linear symmetries of $\mathcal{L}_n$ map fixed points of T to fixed points of T . Additionally, linear symmetries map attractive fixed points to attractive fixed points by Corollary 2.2.9. The attractive fixed points are precisely the vertices of $\mathcal{L}_n$ (Theorem 20 in [3]), of which there are 2^{n-1} (Theorem 2.5 in [6]).

Since the symmetries are uniquely defined by how they permute the vertices, $|\text{LinSym}(\mathcal{L}_n)| \leq (2^{n-1})!$. $\square$

Elements of $\mathcal{L}_n$ can be interpreted as Gram matrices. Let $X \in \mathcal{L}_n$ and $M \in \mathbb{R}^{n \times k}$ such that $X = MM^T$. The n rows of M , denoted $m_1, \dots, m_n$, can be viewed as unit vectors in $\mathbb{R}^k$ such that $X_{ij} = \langle m_i, m_j \rangle$.

Symmetries of the elliptope must preserve the relative placement of these vectors. This includes the vertices, which correspond to a bipartition of $\{1, \dots, n\}$. Naturally, one can permute elements within each partition and swap any two elements between the partitions without altering the overall structure of the partitions. These operations are represented by conjugation with permutation and signature matrices, respectively.

We begin by showing that conjugations by signed permutation matrices induce linear symmetries of the elliptope. We then prove that these signed permutation matrices fully characterize all linear symmetries when $n \neq 4$.

Lemma 3.1.2. *Let A be a signed permutation matrix. Then the transformation defined by*

$$f(X) = AXA^T$$

for all $X \in \mathcal{L}_n$ is an element of $\text{LinSym}(\mathcal{L}_n)$.

Proof. Because f is a linear isometry if and only if the inner product is preserved, to show that $f \in \text{LinSym}(\mathcal{L}_n)$, it suffices to show that f preserves the inner product and maps $\mathcal{L}_n$ to itself. For any $X, Y \in \mathcal{L}_n$, we observe that

$$\begin{aligned} \langle f(X), f(Y) \rangle &= \text{tr}(f(X)^T f(Y)) \\ &= \text{tr} \left((AXA^T)^T (AYA^T) \right) \\ &= \text{tr} (X^T Y) \\ &= \langle X, Y \rangle. \end{aligned}$$

Thus, the inner product is preserved.

Now, consider a correlation matrix X . We want to show that $f(X)$ is also a correlation matrix, i.e., a PSD matrix with unit diagonal entries. Because X is PSD, for any $v \in \mathbb{R}^n$,

$$v^T X v \geq 0.$$

Observe that since $w = A^T v \in \mathbb{R}^n$,

$$v^T f(X) v = v^T A X A^T v = (A^T v)^T X (A^T v) = w^T X w \geq 0.$$

Thus, $f(X)$ is PSD.

Finally, conjugation by signature matrices leaves the diagonal entries of X , meaning $f(X)$ also has unit diagonal entries. Therefore, $f \in \text{LinSym}(\mathcal{L}_n)$. $\square$

The set of signed permutation matrices is the Coxeter group B_n . For any $A \in B_n$, there exists a signature matrix D and permutation matrix P such that $A = DP$. Given the correspondence between the groups B_n and $\text{LinSym}(\mathcal{L}_n)$, we define the homomorphism

$$\psi : B_n \rightarrow \text{LinSym}(\mathcal{L}_n).$$

To see that this is a homomorphism, it suffices to show that

$$\psi(AB) = \psi(A) \circ \psi(B)$$

for all $A, B \in B_n$. This follows from the fact that

$$\begin{aligned}\psi(AB)(X) &= (AB)X(AB)^T \\ &= A(BXB^T)A^T \\ &= A(\psi(B)(X))A^T \\ &= \psi(A) \circ \psi(B)(X)\end{aligned}$$

for every $X \in S_n$.

We now observe that ψ is not injective because for any $A \in B_n$, conjugation by A and $-A$ result in the same transformations. To make our homomorphism injective, we can quotient B_n by $\ker(\psi)$ and instead consider the homomorphism

$$\phi : B_n / \ker(\psi) \rightarrow \text{LinSym}(\mathcal{L}_n).$$

We claim that the kernel of ψ is $\langle -I \rangle$, the cyclic subgroup generated by $-I$.

Lemma 3.1.3. $\ker(\psi) = \langle -I \rangle$.

Proof. Suppose $A \in B_n$ such that $AXA^T = X$ for all $X \in \text{Sym}(n)$. In other words, $\psi(A)$ is the identity transformation. Clearly, if $A = \pm I$, then $\psi(A)$ is the identity transformation. Because A is a signed permutation matrix, there exists a signature matrix D and permutation matrix P such that $A = DP$.

Let J be the all-ones matrix. Notice that permuting any of the rows or columns of J leaves it unchanged, meaning $J = PJP^T$. Since $J = AJA^T$, then

$$J = AJA^T = D(PJP^T)D^T = DJD^T.$$

Suppose that $D \neq \pm I$. Then there exists i, j such that $D_{ii} = 1$ and $D_{jj} = -1$. This implies that $(DJD^T)_{ij} = -1$, meaning $J \neq DJD^T$. Thus, $D = \pm I$, implying that $A = \pm P$ for some permutation matrix P .

Now let X be the diagonal matrix such that $X_{ii} = i$. Because

$$X = AXA^T = (\pm P)X(\pm P^T) = PXP^T,$$

if some row and column i are permuted, then $(PXP^T)_{ii} \neq i$. Hence, it must be that $P = I$. Therefore, $\ker(\psi) = \langle -I \rangle$. □

By Lemma 3.1.3, it is immediately clear that the quotient map,

$$\phi : B_n / \langle -I \rangle \rightarrow \text{LinSym}(\mathcal{L}_n),$$

is an injective homomorphism. This implies that

$$B_n / \langle -I \rangle \cong \phi(B_n / \langle -I \rangle) \subset \text{LinSym}(\mathcal{L}_n).$$

As Tropp showed, the elliptope can be thought of as an approximate cut polytope sharing low-dimensional simplicial faces [11]. Perhaps, unsurprisingly then, the elliptope and cut polytope share the same symmetries. We first show that

$$\text{LinSym}(\mathcal{L}_n) \subset \text{LinSym}(\text{Cut}_n)$$

where $\text{Cut}_n \subset \mathbb{R}^{\binom{n}{2}}$ is the cut polytope.

Definition 3.1.4. (Cut polytope) Let $c \in \{\pm 1\}^n$ be an n -dimensional cut vector and cc^T be an n -dimensional cut matrix. The n -cut polytope is the convex hull of all possible cut matrices:

$$\text{Cut}_n := \text{conv}\{cc^T : c \in \{\pm 1\}^n\}.$$

Lemma 3.1.5. $\text{LinSym}(\mathcal{L}_n) \subset \text{LinSym}(\text{Cut}_n)$.

Proof. The vertices of the cut polytope and the elliptope are the same. Since the cut polytope is the convex hull of the vertices and the elliptope is a convex object containing the vertices, the cut polytope is a subset of the elliptope.

Let $f \in \text{LinSym}(\mathcal{L}_n)$. We know that $f(\mathcal{L}_n) = \mathcal{L}_n$. We want to show that $f(\text{Cut}_n) = \text{Cut}_n$. Let $\mathcal{V} = \{V_1, \dots, V_m\}$ be the vertices of Cut_n (and consequently $\mathcal{L}_n$) where $m = 2^{n-1}$. Any element $W \in \text{Cut}_n$ is a convex combination of V_i 's:

$$W = \sum_{i=1}^m \lambda_i V_i$$

where $\sum_{i=1}^m \lambda_i = 1$ and $\lambda_i \geq 0$. Since f is a linear transformation, it preserves convex combinations. This means that

$$f(W) = \sum_{i=1}^m \lambda_i f(V_i).$$

Since f maps vertices to vertices, we have that $f(W) \in \text{Cut}_n$ is a convex combination of the vertices. It follows that f must map Cut_n to itself. Hence, $\text{LinSym}(\mathcal{L}_n) \subset \text{LinSym}(\text{Cut}_n)$. □

By Lemma 3.1.3 and 3.1.5, we have that

$$B_n / \langle -I \rangle \cong \phi(B_n / \langle -I \rangle) \subset \text{LinSym}(\mathcal{L}_n) \subset \text{LinSym}(\text{Cut}_n).$$

The Coxeter group B_n has cardinality $2^n n!$, so the subgroup $B_n / \langle -I \rangle$ has $2^{n-1} n!$ elements. Notably, the set $\text{LinSym}(\text{Cut}_n)$ also contains $2^{n-1} n!$ elements when $n \neq 4$ ([2]).

Theorem 3.1.6. (Theorem 2.6. and 2.8. in [2])

1. For $n \neq 4$, $|\text{LinSym}(\text{Cut}_n)| = 2^{n-1} n!$.
2. For $n = 4$, $|\text{LinSym}(\text{Cut}_n)| = 2(4!)^2 > 2^3 4!$.

Therefore, when $n \neq 4$, we know the group of linear symmetries of the elliptope is isomorphic to $B_n / \langle -I \rangle$.

Theorem 3.1.7. For $n \neq 4$,

$$\text{LinSym}(\mathcal{L}_n) \cong B_n / \langle -I \rangle.$$

When $n = 4$, $\mathcal{L}_4$ has at least $|B_4 / \langle -I \rangle| = 2^3 4! = 192$ symmetries and at most $2(4!)^2 = 1152$ symmetries. As Li and Pierce [7] show, the isomorphism holds true for $n = 4$.

Theorem 3.1.8 (Theorem 1.1 of [7]). For $n \geq 2$,

$$\text{LinSym}(\mathcal{L}_n) \cong B_n / \langle -I \rangle.$$

Given that we have an explicit formula for the linear symmetries, we can establish the following results:

Observation 3.1.9. The linear symmetries of $\mathcal{L}_n$ preserve eigenvalues.

Proof. Let $f \in \text{LinSym}(\mathcal{L}_n)$ and $X \in \mathcal{L}_n$. There exists $A \in B_n$ such that $f(X) = AXA^T$ for all $X \in \mathcal{L}_n$. Also, there exists a diagonal matrix Λ and an orthogonal matrix U such that $X = U\Lambda U^T$. We observe that the diagonal entries of Λ are precisely the eigenvalues of X . Since

$$f(X) = AU\Lambda(AU)^T,$$

where $AU \in O(n)$, then the diagonal entries of Λ are also the eigenvalues of $f(X)$. Thus, f preserves eigenvalues. $\square$

Corollary 3.1.10. *Linear symmetries of $\mathcal{L}_n$ preserve rank.*

Proof. Because the eigenvalues are preserved, the rank is preserved. $\square$

3.2 Fixed Points of the Elliptope

3.2.1 Definitions and Properties

Define $\lambda_i(M)$ as the i th eigenvalue of M in descending order.

Definition 3.2.1. Define $\pi : \{X \in S_n : \text{diag}(X) = I\} \rightarrow \partial\mathcal{L}_n \cup \{I\}$ as:

$$\pi : M \mapsto \begin{cases} I + \frac{1}{1-\lambda_n(M)}(M - I) & \text{if } M \neq I, \\ I & \text{otherwise.} \end{cases}$$

One can check the well-definedness of the map. Certainly, the diagonal entries are all ones. To check that $\pi(M)$ is PSD, it suffices to observe that the eigenvalues of $\pi(M)$ are precisely $(1 + \frac{\lambda_i - 1}{1 - \lambda_n})$ with eigenvectors v where v is the eigenvector of λ_i as an eigenvalue of M . Since the eigenvalues are nonnegative, we conclude that $\pi(M) \in \mathcal{L}_n$.

We now check that $\pi(M) \in \partial\mathcal{L}_n$ for $M \neq I$. Recall that a correlation matrix M has full rank if and only if it is in the interior of the elliptope. Consequently, the condition $\lambda_n(M) = 0$ characterizes matrices M on the boundary of the elliptope. For $M \neq I$, the map π provides a projection of M onto the boundary of the elliptope. This projection is linear, occurring along the ray emanating from the centroid I of the elliptope and passing through M . Specifically, $\pi(M)$ is constructed to force the smallest eigenvalue of $\pi(M)$ to be zero, thereby guaranteeing that it lies on the boundary.

The fixed points of the n -elliptope can have arbitrary rank $k \in [n]$. When $k = n$, it is clear that the only fixed point is the identity matrix. It is unique because all other fixed points must lie on the boundary since, if $M \neq I$ is a fixed point in the interior, then $\pi(M) \neq M$ and

$$\langle \pi(M), M \rangle > \langle M, M \rangle.$$

The rank one fixed points are precisely the vertices [6].

Apart from the attractiveness of fixed points, we are interested in its reducibility. A reducible fixed point, up to a permutation of its rows and columns, can be viewed as a block diagonal matrix, where each block corresponds to a fixed point of an elliptope of a lower dimension. Therefore, understanding the irreducible fixed points provides the fundamental building blocks for characterizing all fixed points of the n -elliptope.

Definition 3.2.2. (Irreducible matrix) An $n \times n$ matrix M is irreducible if there does not exist a partition of $\{1, \dots, n\}$ into two nonempty sets A and B such that $M_{ij} = 0$ for every $i \in A$ and $j \in B$.

3.2.2 Examples in Low Dimensions

Example 3.2.3. (Fixed Points of $\mathcal{L}_3$ [3]) Elements of the 3-elliptope are correlation matrices in the form

$$\begin{pmatrix} 1 & x & y \\ x & 1 & z \\ y & z & 1 \end{pmatrix},$$

where $x, y, z \in [-1, 1]$. This set can be visualized as a subset of $\mathbb{R}^3$ through the coordinates (x, y, z) . The 3-elliptope is characterized by the inequality

$$x^2 + y^2 + z^2 - 2xyz - 1 \leq 0$$

for $x, y, z \in [-1, 1]$ [9]. Figure 3.1 illustrates $\mathcal{L}_3$ and its fixed points. These fixed points are categorized as follows:

Rank 1 Irreducible Fixed Points (Vertices): There are four such points:

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & -1 & 1 \\ -1 & 1 & -1 \\ 1 & -1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & -1 & -1 \\ -1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix}.$$

Rank 2 Irreducible Fixed Points: There are four such points:

$$\begin{pmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} \\ -\frac{1}{2} & -\frac{1}{2} & 1 \end{pmatrix}, \begin{pmatrix} 1 & \frac{1}{2} & -\frac{1}{2} \\ \frac{1}{2} & 1 & \frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} & 1 \end{pmatrix}, \begin{pmatrix} 1 & -\frac{1}{2} & \frac{1}{2} \\ -\frac{1}{2} & 1 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & 1 \end{pmatrix}, \begin{pmatrix} 1 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 1 & -\frac{1}{2} \\ \frac{1}{2} & -\frac{1}{2} & 1 \end{pmatrix}.$$

Rank 2 Reducible Fixed Points: There are six such points:

$$\begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -1 & 1 \end{pmatrix}.$$




Figure 3.1: 3-Elliptope Surface with Fixed Points. Vertices are red, rank 2 irreducible fixed points are blue, and rank 2 reducible fixed points are green.

The 24 symmetries of $\mathcal{L}_3$ specified in Theorem 3.1.7 are a mixture of rotations and reflections. The symmetries

can be decomposed into permutation and signature matrices. The six permutation matrices are

$$\begin{aligned} P_1 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, & P_2 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, & P_3 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \\ P_4 &= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, & P_5 &= \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, & P_6 &= \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \end{aligned}$$

and the four signature matrices are

$$\begin{aligned} S_1 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, & S_2 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ S_3 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, & S_4 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \end{aligned}$$

Then an arbitrary signed permutation matrix is of the form $S_i P_j$ for $i \in [4]$ and $j \in [6]$. In addition to $\mathcal{L}_n$ being a $\text{LinSym}(\mathcal{L}_n)$ -set, the set of (ir)reducible fixed points of rank k form a $\text{LinSym}(\mathcal{L}_n)$ -set. Since the number of fixed points of each type (vertices, rank 2 irreducible, rank 2 reducible) is fewer than the order of the symmetry group, and the action of the symmetry group is transitive on each of these sets of fixed points, it follows from the Orbit-Stabilizer Theorem that the stabilizer of each fixed point must be nontrivial. These stabilizers are detailed in Tables A.1, A.2, and A.3 for the vertices, rank 2 irreducible fixed points, and rank 2 reducible fixed points, respectively.

A natural question to consider is whether the set of fixed points of the elliptope is discrete. Felzenszwalb et al. [3] demonstrate this is not always the case by providing an uncountable family of rank 2 fixed points in $\mathcal{L}_4$.

Example 3.2.4. (Fixed Points of $\mathcal{L}_4$ [3]) In the 4-elliptope, the following matrix is a fixed point for any value of $c \in [-1, 1]$:

$$X(c) = \begin{pmatrix} 1 & -\sqrt{1-c^2} & 0 & c \\ -\sqrt{1-c^2} & 1 & -c & 0 \\ 0 & -c & 1 & -\sqrt{1-c^2} \\ c & 0 & -\sqrt{1-c^2} & 1 \end{pmatrix}.$$

Given this uncountable family, one may ask which nontrivial symmetries stabilize these fixed points, or if any such nontrivial symmetries exist. The answer is that nontrivial stabilizers do exist and depend on the specific value of c . For instance, if $c = 0.5$, one can check for symmetries that leave $X(0.5)$ unchanged under the action of $\text{LinSym}(\Delta)$. One such symmetry is:

$$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 \end{pmatrix}.$$

3.3 Bijection for Irreducible Fixed Points

In this section, we explore a bijection between the set of irreducible fixed points of rank k and rank $n - k$.

Felzenszwalb et al.[3] prove the following algebraic characterization for fixed points of the elliptope.

Theorem 3.3.1. (Theorem 15 in [3]) Let $X \in \mathcal{L}_n$. Then X is a fixed point of T if and only if there exists a diagonal matrix D such that

$$X^2 = DX.$$

Proposition 3.3.2. (Proposition 12 in [3]) If X is an irreducible fixed point of T , then there exists a constant $d \geq 1$ such that

$$X^2 = dX.$$

We remark that the converse is not true. For example,

$$X = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}$$

is a reducible fixed point of $\mathcal{L}_4$ satisfying $X^2 = 2X$

Corollary 3.3.3. If X is an irreducible fixed point of T , then

$$X^2 = \frac{n}{\text{rank}(X)}X.$$

Proof. Because X is symmetric, it can be diagonalized as

$$X = U\Lambda U^T$$

such that U is orthogonal and Λ is diagonal. By Proposition 3.3.2, we know that

$$U\Lambda^2U^T = X^2 = dX = dU\Lambda U^T,$$

which implies that

$$\Lambda^2 = d\Lambda.$$

Since each diagonal entry of Λ are eigenvalues of X , we know that

$$\lambda_i^2 = d\lambda_i$$

where λ_i is the i th eigenvalue of X . This implies that λ_i is either 0 or d . The trace of X is the sum of eigenvalues, meaning that

$$\text{tr}(X) = d \cdot \text{rank}(X).$$

Furthermore, X has unit diagonal entries, meaning its trace is n . Thus,

$$n = \text{tr}(X) = d \cdot \text{rank}(X),$$

implying that $d = \frac{n}{\text{rank}(X)}$ as desired. □

Definition 3.3.4. Let F_k be the set of irreducible fixed points of rank k :

$$F_k := \{X \in \mathcal{L}_n : X \in T(X) \text{ irreducible, rank}(X) = k\}.$$

Lemma 3.3.5. *Let $A \in \mathcal{L}_n$ be an irreducible matrix, and $\alpha, \beta \in \mathbb{R}$ such that $\alpha \neq 0$. Then*

$$B = \alpha A + \beta I$$

is irreducible.

Proof. First, we show that αA is irreducible. Suppose it was reducible. Then there exists a partition of $\{1, \dots, n\}$ to two nonempty sets S and T such that $(\alpha A)_{ij} = 0$ for every $i \in S$ and $j \in T$. However, since $\alpha \neq 0$, the entries $(\alpha A)_{ij}$ are zero if and only if the corresponding entries A_{ij} are zero. Thus, the same partition S and T would also satisfy $A_{ij} = 0$ for all $i \in S$ and $j \in T$, contradicting the irreducibility of A . Hence, αA is irreducible.

Now, let $A' = \alpha A$ be an irreducible matrix. It remains to show that $B = A' + \beta I$ is irreducible. Again, suppose B is reducible. Then there exists a partition of $\{1, \dots, n\}$ to two nonempty sets S and T such that $B_{ij} = 0$ for every $i \in S$ and $j \in T$. Since $i \neq j$, for any $i \in S$ and $j \in T$,

$$B_{ij} = (A' + \beta I)_{ij} = A'_{ij} + \beta I_{ij} = A'_{ij},$$

implying that B_{ij} is zero if and only if the corresponding A'_{ij} is zero. Thus, the same partition S and T would also satisfy $A'_{ij} = 0$ for all $i \in S$ and $j \in T$, which contradicts the irreducibility of A' . Therefore, B is irreducible.

Alternatively, it suffices to observe that irreducibility depends on the patterns of nonzeros in the off-diagonal. Scaling by a nonzero α does not change the pattern of nonzeros in the off-diagonal, and neither does changing the diagonal entries. □

Definition 3.3.6. Let $g : \partial\mathcal{L}_n \rightarrow \partial\mathcal{L}_n$ be defined as follows:

$$g : M \mapsto I - \frac{1}{\lambda_1 - 1}(M - I) = I - \frac{1}{\|M\|_2 - 1}(M - I),$$

where $\|\cdot\|_2$ denotes the spectral norm.

We first show that, indeed, $Y = g(X) \in \partial\mathcal{L}_n$ for $X \in \partial\mathcal{L}_n$. It suffices to show that Y is PSD with unit diagonal entries and that its smallest eigenvalue is zero. For any i , we know that

$$Y_{ii} = \left(I - \frac{1}{\|X\|_2 - 1}(X - I) \right)_{ii} = I_{ii} - \frac{1}{\|X\|_2 - 1}(X_{ii} - I_{ii}) = 1.$$

Thus, the diagonal entries of Y are one. For any eigenvalue λ of X , we know that any associated eigenvector v satisfies $Xv = \lambda v$. We show that v is an eigenvector of Y with eigenvalue $\mu = 1 - \frac{1}{\|X\|_2 - 1}(\lambda - 1)$, where μ is nonnegative because $\lambda \leq \|X\|_2$. Indeed, we see that

$$\begin{aligned} Yv &= \left(I - \frac{1}{\|X\|_2 - 1}(X - I) \right) v \\ &= v - \frac{1}{\|X\|_2 - 1}(\lambda v - v) \\ &= \left(1 - \frac{1}{\|X\|_2 - 1}(\lambda - 1) \right) v \\ &= \mu v. \end{aligned}$$

Because X is symmetric, it is diagonalizable, which implies that the algebraic and geometric multiplicities coincide. Thus, the eigenvalues of Y are precisely $\{1 - \frac{1}{\|X\|_2 - 1}(\lambda_i - 1)\}_{i \in [n]}$ counting multiplicity, implying that Y is PSD. Finally, the smallest eigenvalue of Y is $1 - \frac{1}{\|X\|_2 - 1}(\|X\|_2 - 1) = 0$. Thus, $Y \in \partial\mathcal{L}_n$.

We observe that applying g again yields the original eigenvalues associated with the same eigenvectors. Therefore, g is an involution. A geometric interpretation is that g maps boundary points to the antipodal points by mapping through the identity matrix to the opposite boundary. Therefore, g is an involution.

Proposition 3.3.7. g is an involution.

Proposition 3.3.8. g is a bijection between F_k and F_{n-k} for $k \in \{1, \dots, n-1\}$.

Proof. Let $X \in F_k$ and $Y = g(X)$. We first show that $Y \in \mathcal{L}_n$. It suffices to show that Y is PSD with unit diagonal entries. For any i , we know that

$$Y_{ii} = \left(I - \frac{1}{\|X\|_2^2 - 1} (X - I) \right)_{ii} = I_{ii} - \frac{1}{\|X\|_2^2 - 1} (X_{ii} - I_{ii}) = 1.$$

Thus, the diagonal entries of Y are one. Because X is an irreducible fixed point of rank k , it has k eigenvalues of $\frac{n}{k}$ and $n - k$ eigenvalues of 0. For any eigenvalue λ of X , we know that any associated eigenvector v satisfies $Xv = \lambda v$. We show that v is an eigenvector of Y with eigenvalue $\mu = \frac{n}{n-k} - \frac{k}{n-k}\lambda$. Indeed, we see that

$$\begin{aligned} Yv &= \left(I - \frac{1}{\|X\|_2^2 - 1} (X - I) \right) v \\ &= v - \frac{1}{\|X\|_2^2 - 1} (\lambda v - v) \\ &= v - \frac{1}{n/k - 1} (\lambda v - v) \\ &= v - \frac{k}{n - k} (\lambda v - v) \\ &= \left(\frac{n}{n - k} - \frac{k}{n - k} \lambda \right) v \\ &= \mu v. \end{aligned}$$

Because X is symmetric, it is diagonalizable, which implies that the algebraic and geometric multiplicities coincide. When $\lambda = 0$, $\mu = \frac{n}{n-k}$, and when $\lambda = \frac{n}{k}$, $\mu = 0$. Thus, Y has k eigenvalues of 0 and $n - k$ eigenvalues of $\frac{n}{n-k}$, implying that Y is PSD. Thus, $Y \in \mathcal{L}_n$.

We now show that Y is an irreducible fixed point of rank $n - k$. First, Y has rank $n - k$ because it has $n - k$ eigenvalues of $\frac{n}{n-k}$. To see that it is a fixed point, it suffices to show that $Y^2 = \frac{n}{n-k} Y$. We know that Y is diagonalizable by $Y = U \Lambda U^T$ with the diagonal entries of Λ being either 0 or $\frac{n}{n-k}$. This implies that $\Lambda^2 = \frac{n}{n-k} \Lambda$. Thus,

$$Y^2 = U \Lambda^2 U^T = \frac{n}{n-k} U \Lambda U^T = \frac{n}{n-k} Y.$$

By Lemma 3.3.5, Y is irreducible, and we conclude that $Y \in F_{n-k}$.

Finally, we observe that $g|_{F_k \cup F_{n-k}}$ is an involution because WLOG, for any $X \in F_k$ and $Y = g(X) \in F_{n-k}$,

$$\begin{aligned} g(g(X)) &= g(Y) \\ &= I - \frac{1}{\|Y\|_2^2 - 1} (Y - I) \\ &= I - \frac{1}{\|Y\|_2^2 - 1} \left(\left(I - \frac{1}{\|X\|_2^2 - 1} (X - I) \right) - I \right) \\ &= I + \frac{1}{\|Y\|_2^2 - 1} \left(\frac{1}{\|X\|_2^2 - 1} (X - I) \right) \\ &= I + \frac{1}{\frac{n}{n-k} - 1} \frac{1}{\frac{n}{k} - 1} (X - I) \\ &= I + (X - I) \\ &= X \end{aligned}$$

as desired. Therefore, g is a bijection between F_k and F_{n-k} . $\square$

We will now demonstrate that the bijection captures an important relationship between rank k and $n - k$ irreducible fixed points. The bijection preserves symmetries of the irreducible fixed points of T .

Lemma 3.3.9. *I is a fixed point of f for any $f \in \text{LinSym}(\mathcal{L}_n)$.*

Proof. I is a fixed point of T , since it satisfies $I^2 = dI$ for $d = 1$. By Lemma 2.2.8, $f(I)$ is a fixed point of T . Since f preserves norm and I is the only element of $\mathcal{L}_n$ with norm $\sqrt{n}$, we conclude that $f(I) = I$. $\square$

Theorem 3.3.10. *Let $X \in F_k$ and $Y = g(X)$. Let $f \in \text{LinSym}(\mathcal{L}_n)$ such that $X = f(X)$. Then $Y = f(Y)$.*

Proof. We conclude the proof through the linearity of f :

$$\begin{aligned} f(Y) &= f\left(I - \frac{1}{\|X\|_2^2 - 1}(X - I)\right) \\ &= f(I) - \frac{1}{\|X\|_2^2 - 1}(f(X) - f(I)) \\ &= I - \frac{1}{\|X\|_2^2 - 1}(X - I) \\ &= Y. \end{aligned}$$

$\square$

This bijection also has a geometric interpretation. For any fixed point $X \in F_k$, $Y = g(X)$ is its antipodal fixed point on the opposite side of the ellipsope's boundary. This means that the line connecting X and Y passes through the centroid of the ellipsope, which is the identity matrix I .

Since $\mathcal{L}_n$ is convex and compact, it has a well-defined and unique centroid in $\mathcal{L}_n$. The centroid is the fixed point of all symmetries, which is I for the ellipsope as shown in Lemma 3.3.9.

Every fixed point of T other than the identity matrix lies on the boundary of the ellipsope.

Lemma 3.3.11. *If $X \neq I$ is a fixed point of T , then it lies on the boundary of $\mathcal{L}_n$.*

Proof. Taking $\pi(X) \in \mathcal{L}_n$, we observe that

$$\langle \pi(X), X \rangle = \frac{1}{1 - \lambda_n(X)} \langle X, X \rangle + \left(1 - \frac{1}{1 - \lambda_n(X)}\right) \langle X, I \rangle.$$

The right-hand side is an affine combination of $\langle X, X \rangle$ and $\langle X, I \rangle$ because the smallest eigenvalue $\lambda_n(X)$ must lie in the interval $[0, 1]$. If $\lambda_n(X) > 1$, then the sum of the eigenvalues is larger than n , which cannot be true because $\text{tr}(X) = n$. Since $\langle X, X \rangle > \langle X, I \rangle$,

$$\frac{1}{1 - \lambda_n(X)} \langle X, X \rangle + \left(1 - \frac{1}{1 - \lambda_n(X)}\right) \langle X, I \rangle \geq \langle X, X \rangle.$$

Since X is a fixed point, we must have equality, which is achieved only if $\lambda_n(X) = 0$, implying X is on the boundary. $\square$

Lemma 3.3.11, in particular, implies that elements in F_k for $k \in \{1, \dots, n\}$ lie on the boundary of $\mathcal{L}_n$. Now, for any $X \in F_k$, consider the line going through X and I , the centroid of $\mathcal{L}_n$:

$$\ell = \overleftrightarrow{IX}.$$

The intersection of ℓ and $\partial\mathcal{L}_n$ contains two points since a straight line passing through an interior point of a convex set can intersect the boundary of the set at most twice—once entering and once exiting the set.

Let X and Y be the two points of intersection of ℓ and $\partial\mathcal{L}_n$. We claim that

$$Y = g(X).$$

Equivalently,

$$X = g(Y).$$

In other words, given X , a fixed point of T on the boundary, $g(X)$ is its antipodal point.

Proposition 3.3.12. *Let $X \in F_k$. Then $g(X) \in F_{n-k}$ is its antipodal point.*

Proof. Since ℓ and $\partial\mathcal{L}_n$ intersect at two points, one of the points of intersection is X , and the other is the point of intersection of $\overrightarrow{XI}$ and $\partial\mathcal{L}_n$. Hence, it suffices to show that $Y = g(X)$ is the point of this intersection.

First, we observe that Y is on the boundary of $\mathcal{L}_n$ by Lemma 3.3.11. Points on the open ray $\overrightarrow{XI}$ are in the form

$$I + \alpha(I - X)$$

for $\alpha > 0$. This is the form of Y where $\alpha = \frac{1}{\|X\|_2 - 1} > 0$. Since $\overrightarrow{XI}$ and $\partial\mathcal{L}_n$ have one point of intersection and Y is on $\overrightarrow{XI}$ and $\partial\mathcal{L}_n$, we conclude that Y is precisely the intersection of $\overrightarrow{XI}$ and $\partial\mathcal{L}_n$. $\square$

Example 3.3.13. (Bijection of Fixed Points in $\mathcal{L}_3$) Figure 3.1 illustrates how the bijection g maps irreducible fixed points to their antipodal counterparts. Here, the red and blue points map to each other by g .

3.4 Iteration on Circulant Symmetric Matrices

In Section 3.1, we asserted that understanding the symmetries of the elliptope would provide insights into the dynamics of the fixed point iteration process. To illustrate this, we now examine a specific type of symmetry: the n -cycle permutation, represented by the cyclic shift matrix P . For the remainder of the thesis, let f denote the symmetry operation defined as conjugation by P , i.e., $f(M) = PMP^T$. For the remainder of this section and in subsequent discussions, we will concentrate exclusively on circulant matrices, which are precisely those square matrices that remain invariant under the action of f (i.e., $f(C) = C$). Our goal is to analyze the behavior of these circulant matrices under the iterative map T .

Definition 3.4.1. Let $\mathcal{C}_n$ denote the set of all circulant matrices in the elliptope:

$$\mathcal{C}_n = \{M \in \mathcal{L}_n : M \text{ circulant}\}.$$

Elements of $\mathcal{C}_n$ are circulant and symmetric. They are thus in the form

$$\text{circulant}(m_0, m_1, m_2, \dots, m_2, m_1) = \begin{pmatrix} m_0 & m_1 & m_2 & \dots & m_2 & m_1 \\ m_1 & m_0 & m_1 & \dots & m_3 & m_2 \\ m_2 & m_1 & m_0 & \dots & m_4 & m_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ m_1 & m_2 & m_3 & \dots & m_1 & m_0 \end{pmatrix}.$$

For circulant matrices, we use the notation in which the lowercase letter denotes the first row of the matrix. For example, m denotes the first row of M and is given by $\begin{pmatrix} m_0 & m_1 & m_2 & \dots & m_2 & m_1 \end{pmatrix}$. We then write

$$M = \text{circulant}(m).$$

We begin by examining the eigenvalues of symmetric circulant matrices in the elliptope, followed by their application to the map T .

Definition 3.4.2. (Fourier matrix) The Fourier matrix F is the $n \times n$ unitary matrix whose k th column is given by

$$v_k = \frac{1}{\sqrt{n}} \begin{pmatrix} 1 & \omega_n^k & \omega_n^{2k} & \dots & \omega_n^{(n-1)k} \end{pmatrix}^T$$

where $\omega_n = e^{\frac{2\pi i}{n}}$ is the primitive n th root of unity.

The matrix M is diagonalized as

$$M = F^* \Lambda F$$

where $\Lambda = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{n-1})$ is the diagonal matrix of eigenvalues.

We know that a circulant matrix M is real and has eigenvalues λ_k if and only if

$$\lambda_k = \overline{\lambda_{n-k}}$$

for $k \in \{1, \dots, n-1\}$ (Chapter 2 Problem 15 in [1]). This implies the following lemma.

Lemma 3.4.3. Let $M = F^* \Lambda F$ with eigenvalues λ_k . Then $M \in \mathcal{C}_n$ if and only if

$$\lambda_k = \lambda_{n-k}$$

for $k \in \{1, \dots, n-1\}$, $\sum_{i=0}^{n-1} \lambda_i = n$, and the eigenvalues are nonnegative.

Proof. Let $M \in \mathcal{C}_n$. Because M is symmetric, it has real eigenvalues. Hence, M has eigenvalues λ_k if and only if

$$\lambda_k = \lambda_{n-k}$$

for $k \in \{1, \dots, n-1\}$.

If $M \in \mathcal{C}_n$, then it has trace n because it has unit diagonal entries. Thus, $\sum_{i=0}^{n-1} \lambda_i = n$. The eigenvalues are also nonnegative because M is PSD, completing one direction.

Now suppose M has nonnegative eigenvalues such that $\lambda_k = \lambda_{n-k}$ for $k \in \{1, \dots, n-1\}$ and $\sum_{i=0}^{n-1} \lambda_i = n$. Then M is a symmetric circulant matrix, as mentioned previously. Also, M is PSD because it has nonnegative eigenvalues. It now suffices to show that M has unit diagonal entries. We know that $\text{tr}(M) = \sum_{i=0}^{n-1} \lambda_i = n$. Since M is circulant, it has a constant diagonal, which implies that the diagonal entries of M are one. $\square$

In the context of fixed point iteration, circulant matrices are distinguished by their rapid convergence, reaching a fixed point of T in one iteration.

Proposition 3.4.4. Let $M \in \mathcal{C}_n$. Then $X = T'(M)$ is circulant and is a fixed point of T .

Proof. We know that $M = f(M)$ for $f \in \text{LinSym}(\mathcal{L}_n)$, where f is conjugation by the n -cycle permutation. By Lemma 2.2.5, $X = T'(M)$ is circulant.

It remains to show that X is a fixed point of T . Let

$$S := \underset{Y \in \mathcal{C}_n}{\text{argmax}} \langle M, Y \rangle = T(M) \cap \mathcal{C}_n.$$

Then we know that $X \in S$ and S is a nonempty subset of $\operatorname{argmax}_{Y \in \mathcal{L}_n} \langle M, Y \rangle$. Hence,

$$X = \operatorname{argmin}_{Y \in T(M)} \|Y\|_2 = \operatorname{argmin}_{Y \in S} \|Y\|_2.$$

For any circulant Y ,

$$\langle M, Y \rangle = n \langle m, y \rangle.$$

By Parseval's Theorem, we have that

$$\langle m, y \rangle = \frac{1}{n} \langle \mathcal{F}(m), \mathcal{F}(y) \rangle$$

where $\mathcal{F}$ denotes the discrete Fourier transform (DFT). Hence, we have the following equivalence

$$S = \operatorname{argmax}_{Y \in \mathcal{C}_n} \langle M, Y \rangle = \operatorname{argmax}_{Y \in \mathcal{C}_n} \langle \mathcal{F}(m), \mathcal{F}(y) \rangle = \operatorname{argmax}_{Y \in \mathcal{C}_n} \sum_{j=0}^{n-1} \lambda_j(M) \lambda_j(Y).$$

Let $\mathcal{J} = \operatorname{argmax}_{j \in \{0,1,\dots,n-1\}} \lambda_j(M)$ be the set of indices corresponding to the maximum eigenvalues of M . For any $j \notin \mathcal{J}$ and $Y \in S$, we show that

$$\lambda_j(Y) = 0.$$

Suppose not. Then there exists $j \notin \mathcal{J}$ such that $\lambda_j(Y) > 0$. Let $j' \in \mathcal{J}$. We can assume that $j, j' \neq 0$ since these separate cases can be handled similarly. The case when n is even and j or j' is $\frac{n}{2}$ should also be handled separately using the same argument. Consider $Y' := F^* \Lambda_{Y'} F$ where the eigenvalues of Y' are the same as Y except

$$\begin{aligned} \lambda_j(Y') &= 0, \\ \lambda_{n-j}(Y') &= 0, \\ \lambda_{j'}(Y') &= \lambda_j(Y) + \lambda_{j'}(Y), \\ \lambda_{n-j'}(Y') &= \lambda_{n-j}(Y) + \lambda_{n-j'}(Y). \end{aligned}$$

Clearly, $Y' \in \mathcal{C}_n$ by Lemma 3.4.3. Furthermore,

$$\sum_{j=0}^{n-1} \lambda_j(M) \lambda_j(Y') > \sum_{j=0}^{n-1} \lambda_j(M) \lambda_j(Y),$$

which contradicts $Y \in S$. Hence, $\lambda_j(Y) = 0$ for all $j \notin \mathcal{J}$.

Next, let $A := \{\alpha = (\alpha_0, \dots, \alpha_{n-1}) : \sum \alpha_i = 1, \alpha_i \geq 0 \forall i, \alpha_i = 0 \text{ if } i \notin \mathcal{J}\}$ be the set of convex coefficients such that for any $\alpha \in A$, $\alpha_i = 0$ if $i \notin \mathcal{J}$. For any $\alpha \in A$, let

$$Y_\alpha = F^* \operatorname{diag}(n\alpha) F.$$

Then $Y_\alpha \in \mathcal{C}_n$ because it is a real symmetric circulant matrix with trace n , meaning it has unit diagonal entries. Moreover, $Y_\alpha \in S$ because $\sum_{j=0}^{n-1} \lambda_{M,j} \lambda_{Y_\alpha,j}$ is maximized. This implies that

$$S = \{F^* \operatorname{diag}(n\alpha) F : \alpha \in A\}.$$

As a reminder, $X = \operatorname{argmin}_{Y \in S} \|Y\|_2^2$. This implies that $X = Y_\alpha$ for some $\alpha \in A$. We note that

$$\|Y_\alpha\|_2^2 = \operatorname{tr}(Y_\alpha^2) = \operatorname{tr}((F^* \operatorname{diag}(n\alpha) F)^2) = \operatorname{tr}(\operatorname{diag}(n\alpha)^2) = \sum_{i=0}^{n-1} (n\alpha_i)^2 = n^2 \sum_{i=0}^{n-1} \alpha_i^2.$$

Hence, it suffices to find $\alpha \in A$ such that $\sum_{i=0}^{n-1} \alpha_i^2$ is minimized. Because we know that $\sum_{i=0}^{n-1} \alpha_i = 1$ with the entries of α being nonnegative and $\alpha_i = 0$ if $i \notin \mathcal{J}$, then $\sum_{i=0}^{n-1} \alpha_i^2$ is minimized when $\alpha_i = \begin{cases} \frac{1}{|\mathcal{J}|} & \text{if } i \in \mathcal{J}, \\ 0 & \text{otherwise} \end{cases}$. This implies that the eigenvalues of $X = Y_\alpha = F^* \text{diag}(n\alpha)F$ are $\frac{1}{|\mathcal{J}|}$ if $i \in \mathcal{J}$ and 0 otherwise. Therefore,

$$X^2 = F^* \text{diag}(n\alpha)^2 F = F^* \left(\frac{n}{|\mathcal{J}|} \text{diag}(n\alpha) \right) F = \frac{n}{|\mathcal{J}|} X.$$

By Theorem 3.3.1, we conclude that X is a fixed point of T . $\square$

As anticipated from our discussion of symmetries in Section 3.1, Proposition 3.4.4 reveals a significant property of matrices invariant under cyclic permutations: their rapid convergence to fixed points of T (T') after a single application of T' . This connection between cyclic symmetry and fixed point behavior motivates us to explore specific instances where circulant matrices naturally arise in the context of the Max-Cut problem.

3.4.1 Cyclic Graphs

One important class of graphs where circulant matrices naturally emerge, particularly in the context of the Max-Cut problem, is the family of cycle graphs. The Laplacian matrix of a cycle graph possesses a circulant structure. In this section, we focus on characterizing the fixed points of T that the cycle graph iterates to.

Let L be the Laplacian of an n -cycle graph. The eigenvalues of the cycle graph are well known:

$$\mathcal{F}(\ell)_k = 4 \sin^2 \left(\frac{\pi k}{n} \right).$$

Lemma 3.4.5. *Let $X = T'(L)$. The DFT of x , the first row of X , is given by*

$$\mathcal{F}(x)_k = \begin{cases} \frac{n}{2} & \text{if } k = \frac{n-1}{2} \text{ or } \frac{n+1}{2} \\ 0 & \text{otherwise} \end{cases}$$

if n is odd and

$$\mathcal{F}(x)_k = \begin{cases} n & \text{if } k = \frac{n}{2} \\ 0 & \text{o.w.} \end{cases}$$

if n is even.

Proof. We know that

$$\mathcal{F}(\ell)_k = \lambda_{L,k} = 4 \sin^2 \left(\frac{\pi k}{n} \right).$$

The maximum of $4 \sin^2 \left(\frac{\pi x}{n} \right)$ is attained at $x = \frac{n}{2}$ where $x \in [0, n)$. Moreover, $4 \sin^2 \left(\frac{\pi x}{n} \right)$ is strictly increasing in $x \in [0, \frac{n}{2}]$ and strictly decreasing in $x \in [\frac{n}{2}, n)$. We note that

$$\mathcal{J} := \operatorname{argmax}_{k \in \{0, 1, \dots, n-1\}} \mathcal{F}(\ell)_k = \begin{cases} \left\{ \frac{n-1}{2}, \frac{n+1}{2} \right\} & \text{if } n \text{ is odd} \\ \left\{ \frac{n}{2} \right\} & \text{otherwise} \end{cases}.$$

In the proof of Proposition 3.4.4, we saw that

$$X = T'(L) = Y_\alpha = F^* \text{diag}(n\alpha)F$$

where $\alpha_i = \begin{cases} \frac{1}{|\mathcal{I}|} & \text{if } i \in \mathcal{I}, \\ 0 & \text{otherwise} \end{cases}$. Because

$$n\alpha = \mathcal{F}(x),$$

we conclude the lemma. □

Corollary 3.4.6. *Let $X = \text{circulant}(x) = T'(L)$. Then x is given by*

$$x_k = (-1)^k \cos\left(\frac{\pi k}{n}\right)$$

if n is odd and

$$x_k = (-1)^k$$

if n is even.

Proof. By Lemma 3.4.5, we know that

$$\mathcal{F}(x)_k = \begin{cases} \frac{n}{2} & \text{if } k = \frac{n-1}{2} \text{ or } \frac{n+1}{2} \\ 0 & \text{o.w.} \end{cases}$$

if n is odd and

$$\mathcal{F}(x)_k = \begin{cases} n & \text{if } k = \frac{n}{2} \\ 0 & \text{o.w.} \end{cases}$$

if n is even. By applying the inverse DFT on $\mathcal{F}(x)$, we obtain the desired form of x . □

Chapter 4

Conclusions

In this thesis, we investigated the geometry of the ellipsope and its influence on the dynamics of iterated linear optimization. We established that the iterative map T is equivariant under the linear symmetries of the domain, ensuring that the structural properties of the initialization are preserved throughout the process. Furthermore, we characterized the symmetries of the ellipsope when $n \neq 4$ via a geometric approach, identifying them with those of the cut polytope.

Building on these geometric insights, we analyzed the specific case of circulant matrices. We demonstrated that matrices invariant under cyclic permutation converge to a fixed point in a single iteration. While the general dynamics of the fixed-point iteration remain complex, our results highlight how symmetry constraints simplify the trajectory of the algorithm.

Appendix A

Symmetries of the 3-elliptope

This appendix records the full linear symmetry group of the 3-elliptope in an explicit, concrete form. Furthermore, we list which symmetry each irreducible fixed point is fixed by.

Recall that signed permutation matrices A can always be decomposed as $A = SP$ where S is a signature matrix and P is a permutation matrix. We know that $\text{LinSym}(\mathcal{L}_3)$ contains 24 elements, which are produced by a combination of four signature matrices and six permutation matrices.

The four signature matrices are:

$$\begin{aligned} S_1 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, & S_2 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ S_3 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, & S_4 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix} \end{aligned}$$

The six permutation matrices are:

$$\begin{aligned} P_1 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, & P_2 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \\ P_3 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, & P_4 &= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \\ P_5 &= \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, & P_6 &= \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \end{aligned}$$

Stabilizers of Vertices				
Symmetry	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & -1 & -1 \\ -1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & -1 & 1 \\ -1 & 1 & -1 \\ 1 & -1 & 1 \end{pmatrix}$
S_1P_1	O	O	O	O
S_2P_1				
S_3P_1				
S_4P_1				
S_1P_2	O	O		
S_2P_2				
S_3P_2				
S_4P_2			O	O
S_1P_3	O			O
S_2P_3				
S_3P_3		O	O	
S_4P_3				
S_1P_4	O			
S_2P_4		O		
S_3P_4			O	
S_4P_4				O
S_1P_5	O			
S_2P_5				O
S_3P_5		O		
S_4P_5			O	
S_1P_6	O		O	
S_2P_6		O		O
S_3P_6				
S_4P_6				

Table A.1: Stabilizers of Vertices

Stabilizers of Irreducible Fixed Points of Rank 2				
Symmetry	$\begin{pmatrix} 1 & -0.5 & -0.5 \\ -0.5 & 1 & -0.5 \\ -0.5 & -0.5 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & -0.5 \\ -0.5 & 0.5 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0.5 & -0.5 \\ 0.5 & 1 & 0.5 \\ -0.5 & 0.5 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & -0.5 & 0.5 \\ -0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{pmatrix}$
S_1P_1	O	O	O	O
S_2P_1				
S_3P_1				
S_4P_1				
S_1P_2	O	O		
S_2P_2				
S_3P_2				
S_4P_2			O	O
S_1P_3	O			O
S_2P_3				
S_3P_3		O	O	
S_4P_3				
S_1P_4	O			
S_2P_4		O		
S_3P_4			O	
S_4P_4				O
S_1P_5	O			
S_2P_5				O
S_3P_5		O		
S_4P_5			O	
S_1P_6	O		O	
S_2P_6		O		O
S_3P_6				
S_4P_6				

Table A.2: Stabilizers of Irreducible Fixed Points of Rank 2

Stabilizers of Reducible Fixed Points of Rank 2						
Symmetry	$\begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -1 & 1 \end{pmatrix}$
S_1P_1	O	O	O	O	O	O
S_2P_1			O	O		
S_3P_1	O	O				
S_4P_1					O	O
S_1P_2					O	O
S_2P_2						
S_3P_2						
S_4P_2					O	O
S_1P_3	O	O				
S_2P_3						
S_3P_3	O	O				
S_4P_3						
S_1P_4						
S_2P_4						
S_3P_4						
S_4P_4						
S_1P_5						
S_2P_5						
S_3P_5						
S_4P_5						
S_1P_6			O	O		
S_2P_6			O	O		
S_3P_6						
S_4P_6						

Table A.3: Stabilizers of Reducible Fixed Points of Rank 2

Bibliography

- [1] Philip J. Davis. *Circulant Matrices*. John Wiley & Sons, Inc., 1979.
- [2] Michel Deza, Viatcheslav P. Grishukhin, and Monique Laurent. The symmetries of the cut polytope and of some relatives. In P. Gritzmann and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics - The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 205–220. American Mathematical Society, 1991.
- [3] Pedro Felzenszwalb, Caroline Klivans, and Alice Paul. Iterated linear optimization. *Quarterly of Applied Mathematics*, 79(4):601–615, 2021.
- [4] Pedro Felzenszwalb, Caroline Klivans, and Alice Paul. Clustering with semidefinite programming and fixed point iteration. *Journal of Machine Learning Research*, 23(190):1–23, 2022.
- [5] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the Association for Computing Machinery*, 42(6):1115–1145, 1995.
- [6] Monique Laurent and Svatopluk Poljak. On a positive semidefinite relaxation of the cut polytope. *Linear Algebra and its Applications*, 223-224:439–461, 1995.
- [7] Chi-Kwong Li and Steve Pierce. Linear operators preserving correlation matrices. *Proceedings of the American Mathematical Society*, 131(1):55–63, 2002.
- [8] Stanisław Mazur and Stanisław Marcin Ulam. Sur les transformations isométriques d’espaces vectoriels normés. *Comptes Rendus de l’Académie des Sciences Paris*, 194:946–948, 1932.
- [9] Mateusz Michałek and Bernd Sturmfels. *Invitation to Nonlinear Algebra*. American Mathematical Society, 2021.
- [10] Yurii Nesterov and Arkadii Nemirovskii. *Interior Point Polynomial Methods in Convex Programming*. Society for Industrial and Applied Mathematics, 1994.
- [11] Joel A. Tropp. Simplicial faces of the set of correlation matrices. *Discrete & Computational Geometry*, 60:512–529, 2018.
- [12] Lieven Vandenbergh and Stephen Boyd. Semidefinite programming. *SIAM Review*, 38:49–95, 1996.
- [13] Lieven Vandenbergh and Stephen Boyd. Applications of semidefinite programming. *Applied Numerical Mathematics*, 29:283–299, 1999.
- [14] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018.