Kubjas, Kaie; Robeva, Elina; Robinson, Richard Z.

Positive semidefinite rank and nested spectrahedra

Published in:
Linear and Multilinear Algebra

DOI:
10.1080/03081087.2017.1381664

Published: 01/01/2018

Document Version
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:
Kubjas, K., Robeva, E., & Robinson, R. Z. (2018). Positive semidefinite rank and nested spectrahedra. *Linear and Multilinear Algebra*, *66*(10), 1952-1974. https://doi.org/10.1080/03081087.2017.1381664

Positive semidefinite rank and nested spectrahedra

Kaie Kubjas[a], Elina Robeva[b] and Richard Z. Robinson[c]

[a]Department of Mathematics and Systems Analysis, Aalto University, PO Box 11100, FI-00076, Aalto Finland; [b]Department of Mathematics, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, USA; [c]Department of Mathematics, University of Washington, Box 354350, Seattle, WA 98195, USA

ARTICLE HISTORY
Compiled September 4, 2017

ABSTRACT
The set of matrices of given positive semidefinite rank is semialgebraic. In this paper we study the geometry of this set, and in small cases we describe its boundary. For general values of positive semidefinite rank we provide a conjecture for the description of this boundary. Our proof techniques are geometric in nature and rely on nesting spectrahedra between polytopes.

KEYWORDS
Positive semidefinite rank; algebraic boundaries; spectrahedra; polytopes

AMS CLASSIFICATION
15A23; 14P10

1. Introduction

Standard matrix factorization is used in a wide range of applications including statistics, optimization, and machine learning. To factor a given a matrix $M \in \mathbb{R}^{p\times q}$ of $\operatorname{rank}(M) = r$, we need to find size-$r$ vectors $a_1, ..., a_p, b_1, ..., b_q \in \mathbb{R}^r$ such that $M_{ij} = \langle a_i, b_j \rangle$.

Often times, however, the matrix at hand as well as the elements in the factorization are required to have certain positivity structure [5,10,11]. In statistical mixture models, for instance, we need to find a *nonnegative* factorization of the matrix at hand [3,9,17,26]. In other words, the vectors a_i and b_j need to be nonnegative. In the present article we study a more general type of factorization called positive semidefinite factorization. The vectors a_i and b_j in the decomposition are now replaced by $k \times k$ symmetric positive semidefinite matrices $A_i, B_j \in \mathcal{S}^k_+$, and k is the size of the positive semidefinite factorization of M. Here the space of symmetric $k \times k$ matrices is denoted by $\mathcal{S}^k$, the cone of $k \times k$ positive semidefinite matrices by $\mathcal{S}^k_+$, and the inner product on $\mathcal{S}^k$ is given by

$$\langle A, B\rangle = \operatorname{trace}(AB).$$

Definition 1.1. *Given a matrix $M \in \mathbb{R}^{p\times q}_{\geq 0}$ with nonnegative entries, a* positive semidefinite (psd) factorization *of size k is a collection of matrices $A_1, ..., A_p, B_1, ..., B_q \in \mathcal{S}^k_+$ such that $M_{ij} = \langle A_i, B_j\rangle$. The* positive semidefinite rank (psd rank) *of the matrix M is the smallest $k \in \mathbb{N}$ for which such a factorization exists. It is denoted by $\operatorname{rank}_{psd}(M)$.*

The nonnegativity constraint on the entries of M is natural here since for any two psd matrices $A, B \in \mathcal{S}_+^k$, it is always the case that $\langle A, B\rangle \geq 0$. To see this, write $A = UU^T, B = VV^T$ for some $U, V \in \mathbb{R}^{k\times k}$. Then, $\operatorname{trace}(AB) = \operatorname{trace}((V^TU)(V^TU)^T) \geq 0$ since $((V^TU)(V^TU)^T)$ is positive semidefinite. Thus, in order for M to have finite psd rank, its entries need to be nonnegative.

Given a polytope P, the smallest number k such that the polytope can be written as a projection of an affine slice of $\mathcal{S}_+^k$ is called the semidefinite extension complexity of P. This quantity is also equal to the psd rank of a slack matrix for the polytope P. This connection between positive semidefinite rank and semidefinite extension complexity is analogous to the connection between nonnegative rank and linear extension complexity, established in the seminal paper of Yannakakis [27]. This was the first paper in the line of work providing super-polynomial lower bounds on the linear and semidefinite extension complexities of families of polytopes [6,7,18,19,22]. The geometric aspects as well as many of the properties of psd rank have been studied in a number of recent articles [4,10–14].

In this paper we study the space $\mathcal{M}_{r,k}^{p\times q}$ (or $\mathcal{M}_{r,k}$ for short) of $p \times q$ nonnegative matrices of rank at most r and psd rank at most k. By Tarski-Seidenberg's Theorem [1, Theorem 2.76], this set is semialgebraic, i.e. it is defined by finitely many polynomial equations and inequalities, or it is a finite union of such sets. It lies inside the variety $\mathcal{V}_r^{p\times q}$ (or $\mathcal{V}_r$ for short) of $p \times q$ matrices of rank at most r. By the definition of psd rank, it is also a subset of $\mathbb{R}_{\geq 0}^{p\times q}$. We study the geometry of $\mathcal{M}_{r,k}^{p\times q}$, and in particular, we investigate the *boundary* $\partial\mathcal{M}_{r,k}^{p\times q}$ of $\mathcal{M}_{r,k}^{p\times q}$ as a subset of $\mathcal{V}_r^{p\times q}$.

A semialgebraic description of $\mathcal{M}_{r,k}^{p\times q}$ would allow one to check if a $p \times q$ matrix of rank r has psd rank at most k by verifying whether the matrix entries satisfy the equations and inequalities in the semialgebraic description. Finding a semialgebraic description is in general a difficult problem. Determining its boundaries is the first step towards achieving the final goal.

Definition 1.2. *The* topological boundary *of $\mathcal{M}_{r,k}^{p\times q}$, denoted by $\partial\mathcal{M}_{r,k}^{p\times q}$, is its boundary as a subset of $\mathcal{V}_r^{p\times q}$. In other words, it consists of all matrices $M \in \mathcal{V}_r^{p\times q}$ such that for every $\epsilon > 0$, the ball with radius ϵ and center M, denoted by $\mathcal{B}_\epsilon(M)$, satisfies the condition that $\mathcal{B}_\epsilon(M) \cap \mathcal{V}_r^{p\times q}$ intersects $\mathcal{M}_{r,k}^{p\times q}$ as well as its complement $\mathcal{V}_r^{p\times q}\setminus\mathcal{M}_{r,k}^{p\times q}$. The* algebraic boundary *of $\mathcal{M}_{r,k}^{p\times q}$, denoted by $\overline{\partial\mathcal{M}_{r,k}^{p\times q}}$ is the Zariski closure of $\partial\mathcal{M}_{r,k}^{p\times q}$ over $\mathbb{R}$.*

In Section 3, we completely describe $\partial\mathcal{M}_{3,2}^{p\times q}$, as well as $\overline{\partial\mathcal{M}_{3,2}^{p\times q}}$. More precisely, Corollary 3.7 shows that a matrix M lies on the boundary $\partial\mathcal{M}_{3,2}^{p\times q}$ if and only if in every psd factorization $M_{ij} = \langle A_i, B_j\rangle$, at least three of the matrices $A_1, \ldots, A_p$ and at least three of the matrices $B_1, \ldots, B_q$ have rank one.

In Sections 4 and 5, we study the general case $\partial\mathcal{M}_{r,k}^{p\times q}$. Conjecture 4.1 is an analogue of Corollary 3.7. It states that a matrix M lies on the boundary $\partial\mathcal{M}_{r,k}^{p\times q}$ if and only if in every psd factorization $M_{ij} = \langle A_i, B_j\rangle$, at least $k+1$ of the matrices $A_1, \ldots, A_p$ have rank one and at least $k+1$ of the matrices $B_1, \ldots, B_q$ have rank one. In Section 5.1, we give theoretical evidence supporting this conjecture in the simplest situation where $p = q = r = k+1$. In Section 5.2, we present computational examples. Our code is available at

`https://github.com/kaiekubjas/psd-rank`.

Our results are based on a geometric interpretation of psd rank, which is explained in Section 2. Given a nonnegative matrix M of rank r satisfying $M\mathbf{1} = \mathbf{1}$, we can associate to it nested polytopes $P \subseteq Q \subseteq \mathbb{R}^{r-1}$. Theorem 2.2, proved in [14], shows that M has psd rank at most k if and only if we can fit a projection of a slice of the cone of $k \times k$ positive semidefinite matrices $\mathcal{S}_+^k$ between P

and Q. When we restrict to the case when the rank of M is three, this result states that M has psd rank two if and only if we can nest an ellipse between the two nested polygons P and Q associated to M. In Theorem 3.6 we show that M lies on the boundary $\partial\mathcal{M}_{3,2}^{p\times q}$ if and only if every ellipse that nests between the two polygons P and Q, touches at least three vertices of P and at least three edges of Q. The statement of Conjecture 4.3 is analogous to the statement of Theorem 3.6 for the general case $\partial\mathcal{M}_{r,k}^{p\times q}$.

Acknowledgments

Part of this work was done while the first and second authors were visiting the Simons Institute for the Theory of Computing, UC Berkeley. We thank Kristian Ranestad and Bernd Sturmfels for very helpful discussions, Rekha Thomas for reading the first draft of the article and Sophia Sage Elia for making Figure 4.

2. Preliminaries

Many of the basic properties of psd rank have been studied in [4]. We give a brief overview of the results used in the present article.

2.1. *Bounds*

The psd rank of a matrix is bounded below by the inequality

$$\operatorname{rank}(M) \leq \binom{\operatorname{rank}_{\mathrm{psd}}(M)+1}{2}$$

since one can vectorize the symmetric matrices in a given psd factorization and consider the trace inner product as a dot product. On the other hand, the psd rank is upper bounded by the nonnegative rank

$$\operatorname{rank}_{psd}(M) \leq \operatorname{rank}_{+}(M)$$

since one can obtain a psd factorization from a nonnegative factorization by using diagonal matrices. The psd rank of M can be any integer satisfying these inequalities.

2.2. *Geometric description*

From nested polytopes to nonnegative matrices

We now describe the geometric interpretation of psd rank. Let $P \subseteq \mathbb{R}^{r-1}$ be a polytope and $Q \subseteq \mathbb{R}^{r-1}$ be a polyhedron such that $P \subseteq Q$. Assume that $P = \operatorname{conv}\{v_1, ..., v_p\}$ and Q is given by the inequality representation $Q = \{x \in \mathbb{R}^{r-1} : h_j^T x \leq z_j, j = 1, ..., q\}$, where $v_1, ..., , v_p, h_1, ..., h_q \in \mathbb{R}^{r-1}$ and $z_1, \ldots, z_q \in \mathbb{R}$. The *generalized slack matrix* of the pair P, Q, denoted by $S_{P,Q}$, is the $p \times q$ matrix whose (i, j)-th entry is $z_j - h_j^T v_i$.

Remark 2.1. *The generalized slack matrix depends on the representations of P and Q as the convex hull of finitely many points and as the intersection of finitely many half-spaces whereas the* *slack matrix depends only on P and Q. We will abuse the notation and write $S_{P,Q}$ for the generalized slack matrix as by the next result the* $\text{rank}_{psd}(S_{P,Q})$ *is independent of the representations of P and Q.*

Theorem 2.2 (Proposition 3.6 in [14]). *Let $P \subset \mathbb{R}^{r-1}$ be a polytope and $Q \subseteq \mathbb{R}^{r-1}$ a polyhedron such that $P \subseteq Q$. Then,* $\text{rank}_{psd}(S_{P,Q})$ *is the smallest integer k for which there exists an affine subspace L of $\mathcal{S}^k$ and a linear map π such that $P \subseteq \pi(L \cap \mathcal{S}^k_+) \subseteq Q$.*

A *spectrahedron* of size k is an affine slice of the cone $\mathcal{S}^k_+$ of $k \times k$ positive semidefinite matrices. A *spectrahedral shadow* of size k is a projection of a spectrahedron of size k. Therefore, Theorem 2.2 states that the matrix $S_{P,Q}$ has psd rank at most k if and only if one can fit a spectrahedral shadow of size k between P and Q.

Remark 2.3. *Given M, the polytopes P and Q are not unique, but the statement of Theorem 2.2 holds regardless of which pair P, Q such that $M = S_{P,Q}$, is chosen.*

From nonnegative matrices to nested polytopes

Given a $p \times q$ nonnegative matrix M, we can assume that it contains no zero rows as removing zero rows does not change its psd rank. Secondly, we may assume that $\mathbf{1}$ is contained in the column span of M as scaling its rows by positive scalars also keeps the psd rank fixed. Consider a rank-size factorization $M = AB$ with A having rows $A_i = (a_i^T, 1)$. Let

$$P = \text{conv}(a_1, \dots, a_p) \text{ and } Q = \{x \in \mathbb{R}^{r-1} : (x^T, 1)B \geq 0\}.$$

Then $P \subseteq Q$ and $S_{P,Q} = M$.

Without loss of generality, we may further assume that $M\mathbf{1} = \mathbf{1}$ by scaling the rows of M by its row sums. The following lemma shows that in this case we can choose P and Q to be bounded.

Lemma 2.4 (Lemma 4.1 in [4]). *Let $M \in \mathbb{R}^{p \times q}_{\geq 0}$ be a nonnegative matrix and assume that $M\mathbf{1} = \mathbf{1}$. Let* $\text{rank}(M) = r$. *Then, there exist polytopes $P, Q \subseteq \mathbb{R}^{r-1}$ such that $P \subseteq Q$ and M is the slack matrix of the pair P, Q.*

The geometry of $\mathcal{M}^{p \times q}_{r,k}$

A point $M \in \mathcal{M}^{p \times q}_{r,k}$ is an *interior point* of $\mathcal{M}^{p \times q}_{r,k}$ if there is an open ball $B_\epsilon(M) \subset \mathbb{R}^{p \times q}$ that satisfies $B_\epsilon(M) \cap \mathcal{V}^{p \times q}_r = B_\epsilon(M) \cap \mathcal{M}^{p \times q}_{r,k}$. By the following lemma, we can check whether a matrix lies in the interior or boundary of $\mathcal{M}^{p \times q}_{r,k}$ by checking this for its rescaling that satisfies $M\mathbf{1} = \mathbf{1}$.

Lemma 2.5. *A matrix $M \in \mathbb{R}^{p \times q}_{\geq 0}$ without zero rows lies in the interior of $\mathcal{M}^{p \times q}_{r,k}$ if and only if the matrix N, obtained from M by rescaling such that $N\mathbf{1} = \mathbf{1}$, lies in the interior of $\mathcal{M}^{p \times q}_{r,k} \cap \{P \in \mathbb{R}^{p \times q}_{\geq 0} : P\mathbf{1} = \mathbf{1}\}$ with respect to $\mathcal{V}^{p \times q}_r \cap \{P \in \mathbb{R}^{p \times q}_{\geq 0} : P\mathbf{1} = \mathbf{1}\}$.*

Proof. First assume that the rescaled matrix N lies in the interior of $\mathcal{M}^{p \times q}_{r,k} \cap R$, where $R = \{P \in \mathbb{R}^{p \times q}_{\geq 0} : P\mathbf{1} = \mathbf{1}\}$. Thus, there exists $\epsilon > 0$ such that $\mathcal{B}_\epsilon(N) \cap \mathcal{V}^{p \times q}_r \cap R \subseteq \mathcal{M}^{p \times q}_{r,k} \cap R$. Here we use the Frobenius norm on $\mathbb{R}^{p \times q}$. Let $\alpha_1, \dots, \alpha_p$ be the row sums of M, i.e. $M\mathbf{1} = \alpha$. Without loss of generality, assume that $0 < \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_p$. Then, consider the ball $\mathcal{B}_{\epsilon\alpha_1}(M)$. If a matrix $M' = M + A \in \mathcal{B}_{\epsilon\alpha_1}(M) \cap \mathcal{V}^{p \times q}_r$, then, after dividing the rows of M' by $\alpha_1, \dots, \alpha_p$

respectively, we obtain the matrix $N+B$, where B is the rescaled version of A. Since $\alpha_1 \leq \cdots \leq \alpha_p$, then $\|B\| \leq \frac{1}{\alpha_1}\|A\|$. Thus $N+B \in \mathcal{B}_\epsilon(N) \cap \mathcal{V}_r^{p\times q} \cap R \subseteq \mathcal{M}_{r,k}^{p\times q} \cap R$. Since rescaling of the rows by positive numbers does not change the rank or psd rank, we have $M' \in \mathcal{M}_{r,k}^{p\times q}$. Therefore, $\mathcal{B}_{\epsilon\alpha_1}(M) \cap \mathcal{V}_r^{p\times q} \subseteq \mathcal{M}_{r,k}^{p\times q}$, i.e. M is in the interior of $\mathcal{M}_{r,k}^{p\times q}$.

Now, assume that M lies in the interior of $\mathcal{M}_{r,k}^{p\times q}$. Then, there exists $\epsilon > 0$ such that $\mathcal{B}_\epsilon(M) \cap \mathcal{V}_r^{p\times q} \subseteq \mathcal{M}_{r,k}^{p\times q}$. Let $M\mathbf{1} = \alpha$, and assume that $0 < \alpha_1 \leq \alpha_1 \leq \cdots \leq \alpha_p$. Consider the ball $\mathcal{B}_{\epsilon/\alpha_p}(N)$. If $N' = N+B \in \mathcal{B}_{\epsilon/\alpha_p}(N) \cap \mathcal{V}_r^{p\times q} \cap R$, then after multiplying the rows of N' by $\alpha_1, \ldots, \alpha_p$ respectively we obtain the matrix $M' = M+A$, where A is the rescaled version of B, and $||A|| \leq \alpha_p ||B||$. Thus, $M' \in \mathcal{B}_\epsilon(M) \cap \mathcal{V}_r^{p\times q} \subseteq \mathcal{M}_{r,k}^{p\times q}$. Since rescaling of the rows by positive numbers does not change the rank or the psd rank, we have $N' \in \mathcal{M}_{r,k}^{p\times q}$. Thus, $\mathcal{B}_{\epsilon/\alpha_p}(N) \cap \mathcal{V}_r^{p\times q} \cap R \subseteq \mathcal{M}_{r,k}^{p\times q} \cap R$, so N lies in the interior of $\mathcal{M}_{r,k}^{p\times q} \cap R$. $\square$

Lemma 2.5 implies that if we want to study the topology of $\mathcal{M}_{r,k}^{p\times q}$ as a subset of $\mathcal{V}_r^{p\times q}$, we can restrict ourselves to the topology of the space $\mathcal{M}_{r,k}^{p\times q} \cap \{P \in R_{\geq 0}^{p\times q} : P\mathbf{1} = \mathbf{1}\}$ as a subset of $\mathcal{V}_r^{p\times q} \cap \{P \in R_{\geq 0}^{p\times q} : P\mathbf{1} = \mathbf{1}\}$, and Lemma 2.4 gives us a recipe for thinking of the elements of this space geometrically.

2.3. *Comparison with nonnegative rank*

Three different versions of nonnegative matrix factorizations appear in the literature: In [26] Vavasis considered the exact nonnegative factorization which asks whether a nonnegative matrix M has a nonnegative factorization of size equal to its rank. The geometric version of this question asks whether one can nest a simplex between the polytopes P and Q.

In [9] Gillis and Glineur defined restricted nonnegative rank as the minimum value r such that there exist $A \in \mathbb{R}_{\geq 0}^{p\times r}$ and $B \in \mathbb{R}_{\geq 0}^{r\times q}$ with $M = AB$ and $\text{rank}(A) = \text{rank}(M)$. The geometric interpretation of the restricted nonnegative rank asks for the minimal r such that there exist r points whose convex hull can be nested between P and Q.

The geometric version of the nonnegative rank factorization asks for the minimal r such that there exist r points whose convex hull can be nested between an $(r-1)$-dimensional polytope and a q-simplex. These polytopes are not P and Q as defined in this paper. See [3, Theorem 3.1] for details.

In the psd rank case there is no distinction between the psd rank and the restricted psd rank, because taking an intersection with a subspace does not change the size of a spectrahedral shadow while intersecting a polytope with a subspace can change the number of vertices. Conjecture 5.2 also suggests that there is no distinction between the spectrahedron and the spectrahedral shadow case which we can compare with simplices and polytopes in the nonnegative rank case, or equivalently the exact nonnegative matrix factorization and restricted nonnegative factorization case.

3. Matrices of rank three and psd rank two

In this section we study the set $\mathcal{M}_{3,2}^{p\times q}$ of $p\times q$ matrices of rank at most three and psd rank at most two. We completely characterize its topological and algebraic boundaries $\partial\mathcal{M}_{3,2}^{p\times q}$ and $\overline{\partial\mathcal{M}_{3,2}^{p\times q}}$.

Consider a matrix $M \in \mathbb{R}_{\geq 0}^{p\times q}$ of rank three. By Lemma 2.4, we get a 2-polytope P and a 2-polyhedron Q such that $P \subseteq Q \subset \mathbb{R}^2$. Theorem 2.2 now has the following simpler form.

Corollary 3.1 (Proposition 4.1 in [14]). *Let M be a nonnegative matrix of rank three. Let $P \subseteq Q \subseteq \mathbb{R}^2$ be a polytope and a polyhedron for which $M = S_{P,Q}$. Then $\mathrm{rank}_{psd}(M) = 2$ if and only if there exists a half-conic (an ellipse, a parabola or a connected component of hyperbola) such that its convex hull C satisfies $P \subseteq C \subseteq Q$. In particular if Q is bounded, then $\mathrm{rank}_{psd}(M) = 2$ if and only if we can fit an ellipse between P and Q.*

If $M\mathbf{1} = \mathbf{1}$, then P and Q are bounded and the half-conic in Corollary 3.1 is an ellipse. We will call its convex hull an elliptical region. Using this geometric interpretation of psd rank two, we give a condition on when a matrix M lies in the interior of $\mathcal{M}_{3,2}^{p\times q}$.

Lemma 3.2. *Let $M \in \mathbb{R}^{p\times q}$ be such that $M\mathbf{1} = \mathbf{1}$ and $\mathrm{rank}(M) = r$. In a small neighborhood of M, there exists a continuous map $\mathcal{V}_r^{p\times q} \cap \{M \in \mathbb{R}^{p\times q} : M\mathbf{1} = \mathbf{1}\} \to \mathbb{R}^{p\times r} \times \mathbb{R}^{r\times q}, M \mapsto (A, B)$ such that $M = AB$ and the last column of A consists of ones.*

Proof. Let $\mathrm{rank}(M) = r$. Consider the rank-size factorization $M = AB$ where A consists of $r-1$ linearly independent columns of M and the column $\mathbf{1}$ such that $\mathbf{1}$ is not in the column span of the $r-1$ columns. Let M' be a submatrix of M given by r linearly independent rows and let A' be the submatrix of A given by the same r linearly independent rows. Then $B = (A')^{-1}M'$. Since the entries of A are also entries of M, we have written the entries of B in the entries of M. This map $M \mapsto (A, B)$ is continuous in the neighborhood of M where the set of linearly independent columns and rows used for constructing A' remain linearly independent. $\square$

Lemma 3.3. *Let M be a nonnegative matrix of rank three satisfying $M\mathbf{1} = \mathbf{1}$. Let $P \subseteq Q \subseteq \mathbb{R}^2$ be polytopes such that $M = S_{P,Q}$. Then M lies in the interior of $\mathcal{M}_{3,2}^{p\times q}$ if and only if there exists an elliptical region E such that $P \subset E \subset Q$ and the boundary of E does not contain any vertices of P.*

Proof. Assume first that M lies in the interior of $\mathcal{M}_{3,2}^{p\times q}$. By Corollary 3.1, there exists an elliptical region E such that $P \subseteq E \subseteq Q$. If the boundary of E does not contain any vertices of P, then we are done. Suppose that the boundary of E contains some vertices of P. We are going to find another elliptical region E' such that $P \subset E' \subset Q$ and the boundary of E' does not contain any vertices of P.

If an entry of M were 0, then M would be on the boundary of $\mathcal{M}_{3,2}^{p\times q}$, because $\mathcal{M}_{3,2}^{p\times q}$ is contained in the nonnegative orthant. Since M is in the interior of $\mathcal{M}_{3,2}^{p\times q}$, none of the entries of M are 0, so the boundary of the polygon Q does not contain any vertices of P. Moreover, there exists $\epsilon > 0$ such that $\mathcal{V}_3^{p\times q} \cap \mathcal{B}_\epsilon(M) \subset \mathcal{M}_{3,2}^{p\times q}$. Pick a point in the interior of the polygon P and consider the polygon tP obtained by a homothety centered at the selected point with some $t > 1$. Then, $P \subset tP \subseteq Q$ for a small enough $t > 1$, and P is strictly contained in tP. Now consider the generalized slack matrix of tP and Q and call it M_t. We can choose t close enough to 1 so that $M_t \in \mathcal{B}_\epsilon(M) \subseteq \mathcal{M}_{3,2}^{p\times q}$. Thus, M_t has psd rank at most two and there exists an elliptical region E' such that $tP \subset E' \subset Q$. Therefore $P \subset tP \subset E' \subset Q$ and the boundary of E' does not contain any vertices of P.

Now suppose that there exists an elliptical region E and polygons P and Q such that $P \subset E \subset Q$ and the boundary of E does not contain any vertices of P. It is possible to shrink E slightly by decreasing the sum of distances to the two focal points of the boundary ellipse by a small number such that the boundary ellipse of the new elliptical region E' does not touch any vertices of P and does not touch any edges of Q. By Lemma 3.2, for any matrix $M' \in \mathcal{B}_\epsilon(M) \cap \mathcal{V}_3^{p\times q} \cap \{M \in \mathbb{R}^{p\times q} : M\mathbf{1} = \mathbf{1}\}$ we obtain polyhedra that are small perturbations of P and Q and hence E' is nested between them. Therefore, $M' \in \mathcal{M}_{3,2}^{p\times q}$ and so $\mathcal{B}_\epsilon(M) \cap \mathcal{V}_3^{p\times q} \cap \{M \in \mathbb{R}^{p\times q} : M\mathbf{1} = \mathbf{1}\} = \mathcal{B}_\epsilon(M) \cap \mathcal{M}_{3,2}^{p\times q} \cap \{M \in \mathbb{R}^{p\times q} : M\mathbf{1} = \mathbf{1}\}$. By Lemma 2.5, the matrix M is in the interior of $\mathcal{M}_{3,2}^{p\times q}$. $\square$

We can now show how $\mathcal{M}_{3,2}^{p\times q}$ relates to the variety $\mathcal{V}_3^{p\times q}$.

Proposition 3.4. *The Zariski closure of $\mathcal{M}_{3,2}^{p\times q}$ over the real numbers is $\mathcal{V}_3^{p\times q}$.*

Proof. Suppose that there exists a ball $\mathcal{B} \subseteq \mathbb{R}^{p\times q}$ such that $\mathcal{B} \cap \mathcal{V}_3^{p\times q} \subseteq \mathcal{M}_{3,2}^{p\times q}$. This implies that the dimension of $\mathcal{M}_{3,2}^{p\times q}$ is equal to that of $\mathcal{V}_3^{p\times q}$, and since $\mathcal{M}_{3,2}^{p\times q} \subset \mathcal{V}_3^{p\times q}$ and $\mathcal{V}_3^{p\times q}$ is irreducible [2, Theorem 2.10], the Zariski closure of $\mathcal{M}_{3,2}^{p\times q}$ over the real numbers equals $\mathcal{V}_3^{p\times q}$.

We show how to find such a ball $\mathcal{B}$. By Lemmas 3.2 and 3.3, it would suffice to find nested polygons $P \subseteq Q \subseteq \mathbb{R}^2$ such that P has p vertices, Q has q edges and there exists an ellipse nested between them that does not touch the vertices of P. Such a configuration certainly exists, for example, we can consider a regular p-gon P centered at the origin with length 1 from the origin to any of its vertices, and a regular q-gon Q centered at the origin with length 5 from the origin to any of its edges. Then, we can fit a circle of radius 2 and center the origin between P and Q so that it does not touch the vertices of P. □

Remark 3.5. *The set of $p \times q$ matrices of psd rank at most k is connected as it is the image under the parametrization map of the connected set $(\mathcal{S}_+^k)^p \times (\mathcal{S}_+^k)^q$. If we also fix the rank, then it is not known if the corresponding set is connected.*

The following theorem is the main result of this section.

Theorem 3.6. *We describe the topological and algebraic boundaries of $\mathcal{M}_{3,2}^{p\times q}$.*

a. A matrix $M \in \mathcal{M}_{3,2}^{p\times q}$ satisfying $M\mathbf{1} = \mathbf{1}$ lies on the topological boundary $\partial\mathcal{M}_{3,2}^{p\times q}$ if and only if $M_{ij} = 0$ for some i, j, or each ellipse that fits between the polygons P and Q contains at least three vertices of the inner polygon P and is tangent to at least three edges of the outer polygon Q.

b. A matrix $M \in \overline{\mathcal{M}_{3,2}^{p\times q}} = \mathcal{V}_3^{p\times q}$ satisfying $M\mathbf{1} = \mathbf{1}$ lies on the algebraic boundary $\overline{\partial\mathcal{M}_{3,2}^{p\times q}}$ if and only if $M_{ij} = 0$ for some i, j or there exists an ellipse that contains at least three vertices of P and is tangent to at least three edges of Q.

c. The algebraic boundary of $\mathcal{M}_{3,2}^{p\times q}$ is the union of $\binom{p}{3}\binom{q}{3} + pq$ irreducible components. There are pq components each defined by an equation $M_{ij} = 0$ for $1 \le i \le p, 1 \le j \le q$, and there are $\binom{p}{3}\binom{q}{3}$ components each defined by one polynomial equation with 1035 terms homogeneous of degree 24 in the entries of M and homogeneous of degree 8 in each row and each column of a 3×3 submatrix of M. The equations for all of these components are considered inside the coordinate ring of the variety of $p \times q$ matrices of rank at most 3, which is defined by the 4×4 minors of M.

Proof. Let $\tilde{P}$ and $\tilde{Q}$ be the projective completions of $\text{cone}(P \times \{1\})$ and $\text{cone}(Q \times \{1\})$, i.e. the closures of images of $\text{cone}(P \times \{1\}) - \{0\}$ and $\text{cone}(Q \times \{1\}) - \{0\}$ under the map $\mathbb{R}^3 \to \mathbb{P}^2, (x, y, z) \mapsto [x : y : z]$. In [8], $\tilde{P}$ and $\tilde{Q}$ are called projective polyhedra. If P and Q are bounded, there is no need to take closure. Hence, in this case there is one-to-one correspondence between statements about incidence relations in the affine and projective case. In Section 2, we required A to have rows $A_i = (a_i^T, 1)$ and defined $P = \text{conv}(a_1, \dots, a_p)$. Similarly, the last row of B gave constant terms of inequalities defining Q. Thus $\text{cone}(P \times \{1\})$ is the cone over the rows of A and $\text{cone}(Q \times \{1\}) = \{x \in \mathbb{R}^3 : x^T B \ge 0\}$. This allows us to define $\tilde{P}$ and $\tilde{Q}$ for general M (even if $\mathbf{1}$ is not in the column span of M). Since on the real projective plane ellipses, parabolas and hyperbolas are equivalent, we will use the word "conic" instead of "ellipse". We will call a conic together with the contractible connected component of its complement, a disk.

(a) Only if: We show the contrapositive of the statement: If all the entries of M satisfying $M\mathbf{1} = \mathbf{1}$ are positive and there is a disk between $\tilde{P}$ and $\tilde{Q}$ whose boundary contains at most two vertices of $\tilde{P}$ or is tangent to at most two edges of $\tilde{Q}$, then M lies in the interior of $M_{3,2}^{p\times q}$.

First, if there is a disk E between $\tilde{P}$ and $\tilde{Q}$ whose boundary touches neither of the polytopes, then M is in the interior of $\mathcal{M}_{3,2}^{p\times q}$ by Lemma 3.3. If at most two edges of $\tilde{Q}$ are tangent to the boundary of the disk E, then $\tilde{P} \subset E \subset \tilde{Q}$ can be transformed by a projective transformation such that the two tangent edges are $x = 0$ and $y = 0$ and that the points of tangency are $[0 : 1 : 1]$ and $[1 : 0 : 1]$. We denote the image of E by $\overline{E}$. The equation of the boundary of $\overline{E}$ has the form $ax^2 + bxy + cy^2 + dxz + eyz + fz^2 = 0$. We know that the only point that lies on the boundary of $\overline{E}$ with $x = 0$ is the point $[0 : 1 : 1]$ since the boundary of $\overline{E}$ touches the line $x = 0$ at $[0 : 1 : 1]$. If we plug in $x = 0$, we get

$$cy^2 + eyz + fz^2 = 0.$$

We may assume $c \geq 0$, hence we must have $cy^2+eyz+fz^2 = (y-z)^2$. Therefore, $c = 1, e = -2, f = 1$. Similarly, since the boundary of $\overline{E}$ touches the line $y = 0$ at $[1 : 0 : 1]$, when we plug in $y = 0$, we get that $ax^2 + dxz + fz^2 = (x - 1)^2$, so, $a = 1, d = -2, f = 1$. Thus, the boundary of $\overline{E}$ has the form

$$\{(x, y) : x^2 + bxy + y^2 - 2xz - 2yz + z^2 = 0\},$$

for some b. The conic is degenerate if and only if $b = 2$. Since the boundary of E is nondegenerate, also the boundary of $\overline{E}$ is nondegenerate. The double cone corresponding to $\overline{E}$ in $\mathbb{R}^3$ is defined by $x^2 + bxy + y^2 - 2xz - 2yz + z^2 \leq 0$. Since $x = 0$ and $y = 0$ are tangent to this double cone and touch it at the points $(0, 1, 1)$ and $(1, 0, 1)$, for all nonzero x and y we have $xy > 0$. This corresponds to $b < 2$. For a slightly smaller value of b, we obtain a slightly larger double cone. The disk $\overline{E}' \subseteq \mathbb{P}^2$ corresponding to this double cone contains $\overline{E}$ and touches $\overline{E}$ only at the points $[1 : 0 : 1]$ and $[0 : 1 : 1]$. Let E' be the preimage of $\overline{E}'$ under the projective transformation considered above. We have $\tilde{P} \subseteq E \subset E' \subseteq \tilde{Q}$ and the boundary of E' does not touch $\tilde{P}$. Thus, by Lemma 3.3, M lies in the interior of $\mathcal{M}_{3,2}^{p\times q}$. The case when the boundary of E goes through at most two vertices of $\tilde{P}$ follows by duality. An introduction to duality can be found in [21].

If: If an entry of M is zero, then M is on the topological boundary $\partial\mathcal{M}_{3,2}^{p\times q}$, because $\mathcal{M}_{3,2}^{p\times q} \subseteq \mathbb{R}_{\geq 0}^{p\times q}$. By Lemma 3.3, if $M \in \mathcal{M}_{3,2}^{p\times q}$ satisfying $M\mathbf{1} = \mathbf{1}$ lies in the interior, then there is a disk between $\tilde{P}$ and $\tilde{Q}$ that does not touch $\tilde{P}$. Thus, if the boundary of every disk nested between $\tilde{P}$ and $\tilde{Q}$ contains at least three vertices of $\tilde{P}$ and touches at least three edges of $\tilde{Q}$, then M lies on the boundary $\partial\mathcal{M}_{3,2}^{p\times q}$

$(b), (c)$ If $M \in \mathbb{R}^{p\times q}$ without nonnegativity constraints satisfies $M\mathbf{1} = \mathbf{1}$, then one can define polytopes P and Q as explained before Lemma 2.4. The difference is that $P \subseteq Q$ does not hold anymore, and we also might not have $\tilde{P} \subseteq \tilde{Q}$. Nevertheless, one can talk about vertices of $\tilde{P}$ and edges of $\tilde{Q}$. Hence given three points a, b, c in $\mathbb{P}^2$ and three lines d, e, f in $\mathbb{P}^2$, each given by three homogeneous coordinates, we seek the condition that there exists a conic X such that a, b, c lie on X and d, e, f are tangent to X.

Let $X = \begin{bmatrix} x_{11} & x_{12} & x_{13} \\ x_{12} & x_{22} & x_{23} \\ x_{13} & x_{23} & x_{33} \end{bmatrix}$ be the matrix of a conic. Then the corresponding conic goes through the points a, b, c if and only if

$$a^T X a = b^T X b = c^T X c = 0. \tag{3.1}$$

Similarly, the lines d, e, f are tangent to the conic if and only

$$d^T Y d = e^T Y e = f^T Y f = 0, \tag{3.2}$$

where $XY = I_3$. We seek to eliminate the variables X and Y.

Let $[a, b, c]$ denote the matrix whose columns are a, b, c. First we assume that $[a, b, c]$ is the 3×3-identity matrix. Then we proceed in two steps:

1) The equations (3.1) imply that x_{11}, x_{22}, x_{33} are zero. We make the corresponding replacements in equations (3.2).

2) We use [24, formula (4.5) on page 48] to get the resultant of three ternary quadrics to get a single polynomial in the entries of d, e, f.

Now we use invariant theory to obtain the desired polynomial in the general case. Let $g \in \mathrm{GL}_3(\mathbb{R})$. The conic X goes through the points a, b, c and touches the lines d, e, f if and only if the conic $g^{-T} X g^{-1}$ goes through the points ga, gb, gc and touches the lines $g^{-T}d, g^{-T}e, g^{-T}f$. Thus our desired polynomial belongs to the ring of invariants $\mathbb{R}[V^3 \oplus V^{*3}]^{\mathrm{GL}_3(\mathbb{R})}$ where $V = \mathbb{R}^3$ and the action of $\mathrm{GL}_3(\mathbb{R})$ on $V^3 \oplus V^{*3}$ is given by

$$g \cdot (a, b, c, d, e, f) := (ga, gb, gc, g^{-T}d, g^{-T}e, g^{-T}f).$$

The First Fundamental Theorem states that $\mathbb{R}[V^3 \oplus V^{*3}]^{\mathrm{GL}_3(\mathbb{R})}$ is generated by the bilinear functions $(i|j)$ on $V^3 \oplus V^{*3}$ defined by

$$(i|j) : (a, b, c, d, e, f) \mapsto ([a, b, c]^T [d, e, f])_{ij}.$$

For the FFT see for example [16, Chapter 2.1]. In the special case when $[a, b, c]$ is the 3×3 identity matrix, $(i|j)$ maps to the (i, j)-th entry of $[d, e, f]$. Hence to obtain the desired polynomial in the general case, we replace in the resultant obtained in the special case the entries of the matrix $[d, e, f]$ by the entries of the matrix $[a, b, c]^T [d, e, f]$.

`Maple` code for doing the steps in the previous paragraphs can be found at our website. This program outputs one polynomial of degree 1035 homogeneous of degree 8 in each of the rows and the columns of the matrix $\begin{bmatrix} - & a & - \\ - & b & - \\ - & c & - \end{bmatrix} \begin{bmatrix} | & | & | \\ d & e & f \\ | & | & | \end{bmatrix}$. By construction, if this homogeneous polynomial vanishes and the projective polyhedron $\tilde{P}$ with vertices a, b, c lies inside the projective polyhedron $\tilde{Q}$ with edges d, e, f and a, b, c, d, e, f are real, then there exists a conic nested between $\tilde{P}$ and $\tilde{Q}$ touching d, e, f and containing a, b, c. Therefore, the Zariski closure of the condition that the only possible conics that can fit between $\tilde{P}$ and $\tilde{Q}$ touch at least three edges of $\tilde{Q}$ and at least three vertices of $\tilde{P}$ is exactly that there exists a conic that touches at least three edges of $\tilde{Q}$ and at least three vertices of $\tilde{P}$. This proves (b).

To prove (c), let $M \in \mathcal{V}_3^{p \times q}$ be such that $M = AB$ and a, b, c are three of the rows of A and d, e, f are three of the columns of B. Then, the above-computed polynomial contains variables only from the entries of a 3×3 submatrix of M corresponding to these rows and columns. We can drop the assumption $M\mathbf{1} = \mathbf{1}$ here: Scaling a row of M by a constant corresponds to scaling the corresponding row of A by the same constant, which does not influence equations (3.1). For each

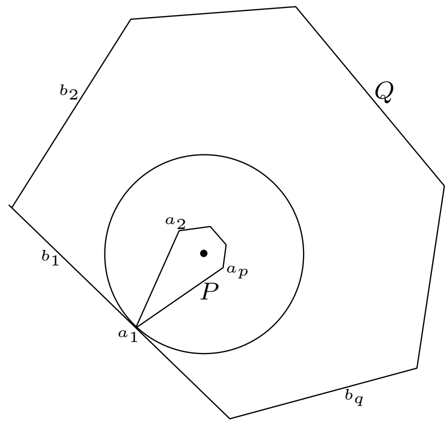

Figure 1.: A configuration with $M_{11} = 0$

three rows and three columns of M we have one such polynomial, so the algebraic boundary is given by the union over each three rows and three columns of M of the vanishing of the corresponding degree 24 polynomial with 1035 terms inside the variety defined by the 4×4 minors of M.

To finish the proof, we show that the components defined by $M_{ij} = 0$ are full dimensional components of the boundary of $\mathcal{M}^{p\times q}_{3,2}$. Without loss of generality, we will show it for $i = j = 1$. Pick a matrix $M \in \mathcal{M}^{p\times q}_{3,2}$ which corresponds to Figure 1.

Assume that the inner circle has radius r and the vertices $a_2, \ldots, a_p$ are within $r/10$ from the center of the circle, while the facets $b_2, \ldots, b_q$ are at least $10r$ away from the center. By Lemma 3.2, all matrices of rank 3 in a small ball around M can be represented by vectors $a'_1, \ldots, a'_p$ and $b'_1, \ldots, b'_q$ which lie in small balls around $a_1, \ldots, a_p$ and $b_1, \ldots, b_q$ respectively. Suppose that we consider a ball around M with small enough radius that the corresponding balls around $a_1, \ldots, a_p$ and $b_1, \ldots, b_q$ have radius at most $r/100$ respectively. Now, consider a point M' in this small ball which has $M'_{11} = 0$. We will show that M' has psd rank at most 2. Instead of considering the motion of all of $a'_1, \ldots, a'_p$ and $b'_1, \ldots, b'_q$, we can fix $a'_1 = a_1$ and the direction of b'_1 to the direction of b_1, and consider the relative motion of the remaining $a'_2, \ldots, a'_p$ and $b'_2, \ldots, b'_q$. We also fix the circle of radius r. Then, because of the specs that we have chosen the points $a'_2, \ldots, a'_p$ will remain in this circle, and then the facets $b'_2, \ldots, b'_q$ will remain outside of the circle. Therefore, M' has psd rank at most 2. This shows that the boundary component defined by $M_{11} = 0$ is a full dimensional boundary component. $\square$

Here is an algebraic version of Theorem 3.6.

Corollary 3.7. *A matrix $M \in \mathbb{R}^{p\times q}_{\geq 0}$ satisfying $M\mathbf{1} = \mathbf{1}$ lies on the boundary $\partial\mathcal{M}^{p\times q}_{3,2}$ if and only if for every size-2 psd factorization $M_{ij} = \langle A_i, B_j\rangle$, at least three of the matrices $A_1, \ldots, A_p \in \mathcal{S}^2_+$ have rank one and at least three of the matrices $B_1, \ldots, B_q \in \mathcal{S}^2_+$ have rank one.*

Proof. Suppose that $M \notin \partial\mathcal{M}^{p\times q}_{3,2}$. Let $P = \operatorname{cone}\{a_1, \ldots, a_p\}$ and $Q = \{x \in \mathbb{R}^{r-1} : \langle x, b_j\rangle \geq 0 \text{ for } j = 1, \ldots, q\}$ such that $M = S_{P,Q}$. By [14, Proposition 4.4] and Theorem 3.6, there exists an invertible linear map π such that $P \subseteq \pi(\mathcal{S}^2_+) \subseteq Q$ and the boundary of $\pi(\mathcal{S}^2_+)$ contains at most two rays of P or is tangent to at most two facets of Q.

The invertibility of π gives

$$\pi^{-1}(P) \subseteq \mathcal{S}^2_+ \subseteq \pi^{-1}(Q),$$

where $\pi^{-1}(P) = \text{cone}\{\pi^{-1}(a_1), \ldots, \pi^{-1}(a_p)\}$ and

$$\pi^{-1}(Q) = \{x \in L \cap \mathcal{S}^2 : \langle \pi(x), b_j \rangle \geq 0\} = \{x \in L \cap \mathcal{S}^2 : \langle x, \pi^T(b_j) \rangle \geq 0\}.$$

Thus $M = S_{\pi^{-1}(P), \pi^{-1}(Q)}$, since

$$M_{ij} = \langle a_i, b_j \rangle = \langle \pi(\pi^{-1}(a_i)), b_j \rangle = \langle \pi^{-1}(a_i), \pi^T(b_j) \rangle.$$

The inclusion $\pi^{-1}(P) \subseteq \mathcal{S}^2_+$ implies that $\pi^{-1}(a_1), \ldots, \pi^{-1}(a_p)$ are psd. Taking dual of the inclusion $\mathcal{S}^2_+ \subseteq \pi^{-1}(Q)$ gives that $\pi^T(b_1), \ldots, \pi^T(b_q)$ are psd. Since π is invertible, we know that either the boundary of $\mathcal{S}^2_+$ contains at most two rays of $\pi^{-1}(P)$ or is tangent to at most two facets of $\pi^{-1}(Q)$. Hence $\pi^{-1}(a_1), \ldots, \pi^{-1}(a_p), \pi^T(b_1), \ldots, \pi^T(b_q)$ gives a psd factorization of M with at most two of $\pi^{-1}(a_1), \ldots, \pi^{-1}(a_p)$ having rank one or at most two of $\pi^T(b_1), \ldots, \pi^T(b_q)$ having rank one.

Suppose that there exists a psd factorization of M, given by matrices $A_1, \ldots, A_p, B_1, \ldots, B_q \in \mathcal{S}^2_+$, such that at most two of the A_i have rank one. Consider $P = \text{cone}\{A_1, \ldots, A_p\}$ and $Q = \{x \in \mathcal{S}^2 : \langle x, B_j \rangle \geq 0, \forall j = 1, \ldots, q\}$. Then $P \subseteq \mathcal{S}^2_+ \subseteq Q$ and the boundary of $\mathcal{S}^2_+$ contains at most two rays of P. Using the inner product preserving bijection between $\mathcal{S}^2$ and $\mathbb{R}^3$, we can consider all objects in $\mathbb{R}^3$. In particular, the images of $A_1, \ldots, A_p, B_1, \ldots, B_q$ in $\mathbb{R}^3$ give a rank factorization of M. By Theorem 3.6 (a), we have $M \notin \partial \mathcal{M}^{p \times q}_{3,2}$. □

We now investigate the topological boundary more thoroughly.

Proposition 3.8. *Suppose $M \in \mathcal{M}^{p \times q}_{3,2}$ satisfying $M\mathbf{1} = \mathbf{1}$ is strictly positive. Then M lies on the topological boundary if and only if there exists a unique ellipse that fits between P and Q.*

Proof. A matrix in the relative interior of $\mathcal{M}_{3,2}$ will have multiple ellipses nested between P and Q: By the only if direction of the proof of Theorem 3.6 part (a), there exists an ellipse that is nested between P and Q such that the ellipse does not touch P. We can just take slight scalings of this ellipse to get multiple ellipses. This proves the "if" direction.

For the "only if" direction, suppose M lies on the topological boundary and E_0 and E_1 are two elliptical regions nested between P and Q. Let $E_{1/2}$ be the elliptical region determined by averaging the quadratics defining E_0 and E_1, i.e.

$$E_{1/2} = \{x : q_0(x) + q_1(x) \geq 0\} \text{ where } E_i = \{x : q_i(x) \geq 0\}.$$

It is straightforward to see that $E_{1/2}$ is nested between P and Q. Furthermore, if v is a vertex of P, then the boundary of $E_{1/2}$ passes through v if and only if both the boundaries of E_0 and E_1 pass through v. Similarly, if f is a facet of Q, then the boundary of $E_{1/2}$ is incident to f if and only if the boundaries of E_0 and E_1 are tangent to f at the same point. By Theorem 3.6, the boundary of $E_{1/2}$ must pass through three vertices of P and three edges of Q. Hence, there must exist six distinct points that both the boundaries of E_0 and E_1 pass through. No three of the six points are collinear, since the boundary ellipses of E_0 and E_1 pass through them. Since five distinct points in general position determine a unique conic, we must have that $E_0 = E_1$. □

Example 3.9. In the previous result, we examined the geometric configurations on the boundary of the semialgebraic set coming from strictly positive matrices. The simplest idea for such a matrix is to take two equilateral triangles and expand the inner one until we are on a boundary configuration as in Figure 2a.

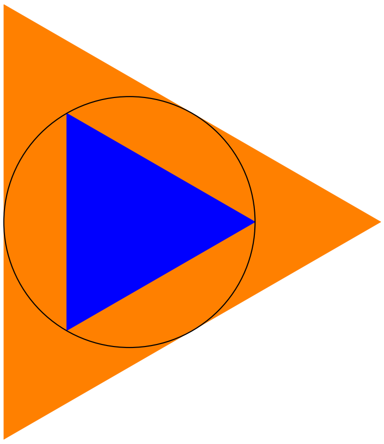

(a) Boundary configuration

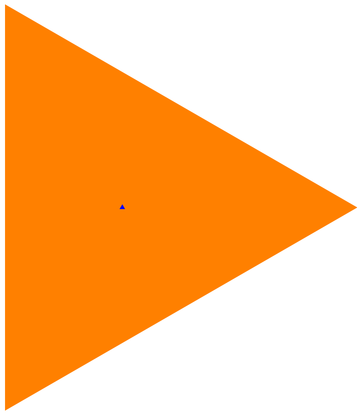

(b) Interior configuration which also lies on the algebraic boundary $\overline{\partial\mathcal{M}_{3,2}}$. Its slack matrix is a circulant matrix with $a \approx 1.05, b \approx 1.01, c \approx 1.03$

Figure 2.: Geometric configurations of matrices in $\mathcal{M}_{3,2}^{3\times 3}$

This configuration has the slack matrix

$$\frac{1}{6}\begin{bmatrix} 4 & 1 & 1 \\ 1 & 4 & 1 \\ 1 & 1 & 4 \end{bmatrix}. \tag{3.3}$$

The 1035 term boundary polynomial from Theorem 3.6 vanishes on this matrix, as we expect.

This matrix lies in the set of 3×3 circulant matrices which have the form

$$\begin{bmatrix} a & b & c \\ c & a & b \\ b & c & a \end{bmatrix}.$$

It was shown in [4, Example 2.7] that these matrices have psd rank at most two precisely when $a^2 + b^2 + c^2 - 2(ab + ac + bc) \leq 0$. As expected, whenever this polynomial vanishes, the 1035 term boundary polynomial vanishes as well. The matrix (3.3) is a regular point of the hypersurface defined by the boundary polynomial. Figure 2b shows an instance of parameters a, b, c such that the matrix is on the algebraic boundary but not on the topological boundary – the 1035 term boundary polynomial from Theorem 3.6 vanishes, but the matrix lies in the interior of $\mathcal{M}_{3,2}$.

We were interested in finding out if the 1035 term boundary polynomial could be used in an inequality to classify circulant matrices of psd rank at most two. The family of circulant matrices which have $c = 1$ and whose psd rank is at most two is depicted in Figure 3a. The boundary polynomial, shown in Figure 3b, takes both positive and negative values on the interior of the space. Figures 4a and 4b show the semialgebraic set and the boundary polynomial in the 3-dimensional space.

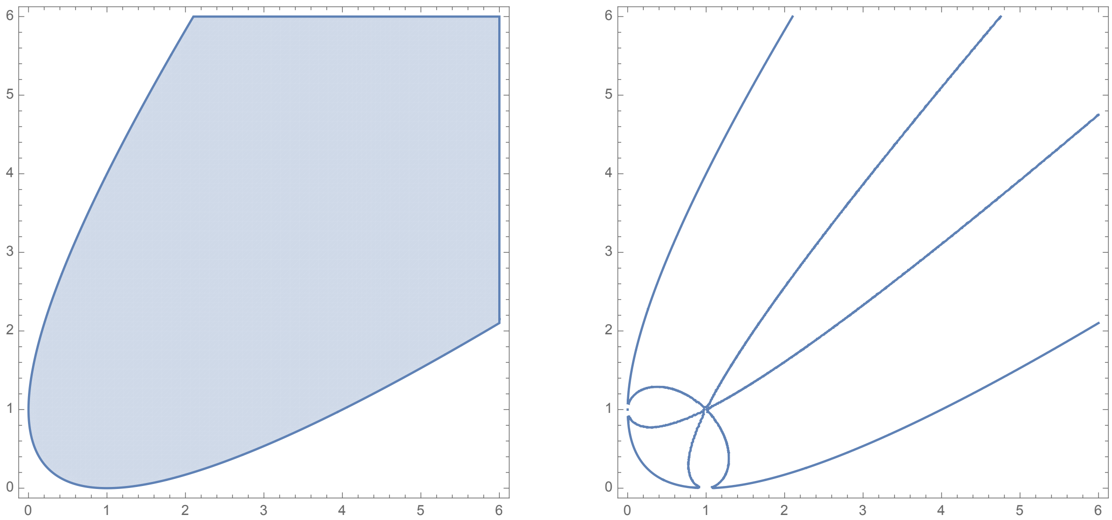

(a) Circulant matrices of psd rank at most 2 (b) The boundary polynomial

Figure 3.: 3×3 circulant matrices in $\mathbb{R}^2$

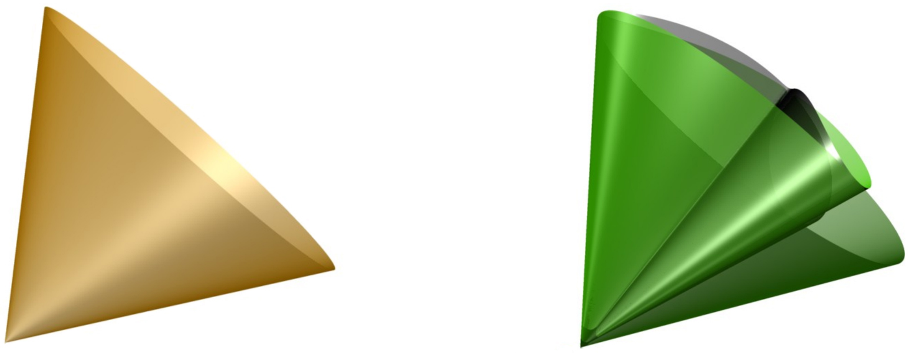

(a) Circulant matrices of psd rank at most 2 (b) The boundary polynomial

Figure 4.: 3×3 circulant matrices in $\mathbb{R}^3$

4. Matrices of higher psd rank

In Corollary 3.7, we showed that a matrix lies on the boundary $\partial \mathcal{M}_{3,2}^{p \times q}$ if and only if in every psd factorization $M_{ij} = \langle A_i, B_j \rangle$, at least three A_i's and at least three B_j's have rank one. In analogy with this result, we conjecture that a matrix lies on the boundary $\partial \mathcal{M}_{r,k}^{p \times q}$ if and only if in every psd factorization $M_{ij} = \langle A_i, B_j \rangle$, at least $k+1$ matrices A_i and at least $k+1$ matrices B_j have rank one.

Conjecture 4.1. *A matrix $M \in \mathbb{R}_{\geq 0}^{p \times q}$ satisfying $M\mathbf{1} = \mathbf{1}$ lies on the boundary $\partial \mathcal{M}_{r,k}^{p \times q}$ if and only if for every size-k psd factorization $M_{ij} = \langle A_i, B_j \rangle$, at least $k+1$ of the matrices $A_1, \ldots, A_p \in \mathcal{S}_+^2$ have rank one and at least $k+1$ of the matrices $B_1, \ldots, B_q \in \mathcal{S}_+^2$ have rank one.*

Let $M \in \mathbb{R}_{\geq 0}^{p \times q}$ be a full rank matrix, and let $P \subseteq Q \subseteq \mathbb{R}^{r-1}$ be nested polytopes such that

$M = S_{P,Q}$. By Theorem 2.2, the matrix M has psd rank at most k if and only if we can nest a spectrahedral shadow C of size k between P and Q. By definition, the spectrahedral shadow C is a linear projection of a spectrahedron $\tilde{C} = L \cap \mathcal{S}^k_+$ of size k.

Definition 4.2. *We say that a vector $v \in C$ lies in the* rank s locus *of C if there exists a $k \times k$ psd matrix in $\tilde{C}$ of rank s that projects onto v.*

The geometric version of the Conjecture 4.1 is:

Conjecture 4.3. *A matrix M is on the boundary $\partial\mathcal{M}^{p\times q}_{r,k}$ if and only if all spectrahedral shadows C of size k such that $P \subseteq C \subseteq Q$ contain $k+1$ vertices of P at rank one loci and touch $k+1$ facets of Q at rank $k-1$ loci.*

For $r = \binom{k+1}{2}$, one can show similarly to the proof of Corollary 3.7 that Conjectures 4.1 and 4.3 are equivalent. The case $r = \binom{k+1}{2}$ differs from other cases, because only for this case the linear map π is invertible.

The psd rank three and rank four setting corresponds to the geometric configuration where a 3-dimensional spectrahedral shadow of size three is nested between 3-dimensional polytopes. A detailed study of generic spectrahedral shadows can be found in [23].

Example 4.4. We now give an example of a geometric configuration as in Conjecture 4.3. We stipulate that the vertices of the interior polytope coincide with the nodes of the spectrahedron in Figure 5a and the facets of the outer polytope touch the boundary of this spectrahedron at rank two loci. In the dual picture, the vertices of the inner polytope lie on the rank one locus depicted in Figure 5b and the facets of the outer polytope contain the rank two locus of this spectrahedral shadow.

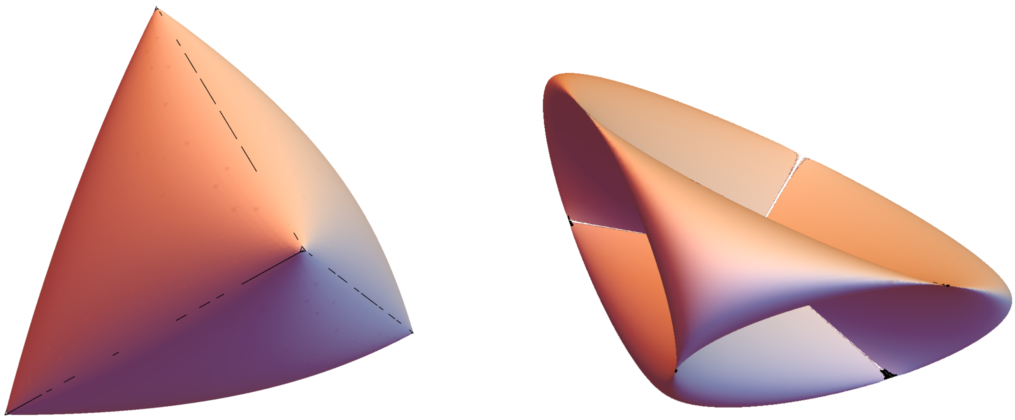

(a) Spectrahedron　　　　(b) Rank-one locus of the dual shadow

Figure 5.: 3-dimensional spectrahedral shadows

We end this section with a restatement of Conjecture 4.1 in a special case using Hadamard square roots.

Definition 4.5. *Given a nonnegative matrix M, let $\sqrt{M}$ denote a Hadamard square root of M obtained by replacing each entry in M by one of its two possible real square roots. The square root rank of a nonnegative matrix M, denoted as $\operatorname{rank}_{\sqrt{}}(M)$, is the minimum rank of a Hadamard square* *root of M.*

Lemma 4.6 (Lemma 2.4 in [13]). *The smallest k for which a nonnegative real matrix M admits a $\mathcal{S}_+^k$-factorization in which all factors are matrices of rank one is $k = \mathrm{rank}_{\surd}(M)$.*

Hence Conjecture 4.1 is equivalent to the statement that a matrix $M \in \mathcal{M}_{k+1,k}^{(k+1)\times(k+1)}$ lies on the boundary $\partial\mathcal{M}_{k+1,k}^{(k+1)\times(k+1)}$ if and only if its square root rank is at most k. We conclude this section with a conjecture which would lead to a semialgebraic description of $\mathcal{M}_{r,k}^{p\times q}$.

Conjecture 4.7. *Every matrix $M \in \mathcal{M}_{r,k}^{p\times q}$ has a psd factorization $M_{ij} = \langle A_i, B_j\rangle$, with at least k matrices A_i and $k-1$ matrices B_j being rank one, or at least $k-1$ matrices A_i and k matrices B_j being rank one.*

If this conjecture were true, there would be $\binom{p}{k}\binom{q}{k-1} + \binom{p}{k-1}\binom{q}{k}$ options for selecting the $2k-1$ rank-one matrices. For each such option we would be able to describe the semialgebraic set of all such matrices that have psd rank k.

5. Evidence towards Conjecture 4.1

In this section, we present partial evidence towards proving Conjecture 4.1 if $p = q = r = k+1$. Section 5.1 is theoretical in nature, while Section 5.2 exhibits computational results.

5.1. *Nested spectrahedra*

By Theorem 2.2 a matrix M for which $M\mathbf{1} = \mathbf{1}$ has psd rank k if and only if we can nest a spectrahedral shadow of size k between the polytopes P and Q corresponding to M. In the following lemma, we show that a $(k+1)\times(k+1)$ matrix M has psd rank k if and only if we can fit a spectrahedron of size k between P and Q. We show that if there is a spectrahedral shadow C nested between P and Q, then we can find a spectrahedron C' of the same size such that $P \subseteq C' \subseteq C \subseteq Q$.

Lemma 5.1. *Let $M \in \mathbb{R}_{\geq 0}^{(k+1)\times(k+1)}$ be a full-rank matrix such that $M\mathbf{1} = \mathbf{1}$. Then, M has psd rank at most k if and only if we can nest a spectrahedron of size k between the two polytopes P and Q corresponding to M.*

Proof. If we can fit a spectrahedron of size k between P and Q, then M has psd rank at most k.

Conversely, suppose that M has psd rank at most k. Then there exists a slice L of $\mathcal{S}_+^k$ and a linear map π such that $C = \pi(L \cap \mathcal{S}_+^k)$ lies between P and Q:

$$P \subseteq C \subseteq Q.$$

If π is a $1:1$ linear map, then, the image C is just a linear transformation of a spectrahedron, and is therefore a spectrahedron of the same size. So, assume that π is not $1:1$, i.e. it has nontrivial kernel.

We can write

$$L \cap \mathcal{S}_+^k = \{(x_1,\ldots,x_s) \in \mathbb{R}^s : \sum_{i=1}^{s} x_i A_i + (1 - \sum_{i=1}^{s} x_i) A_{s+1} \succeq 0\}$$

for some $A_1, \ldots, A_{s+1} \in \mathcal{S}^k$. Let $u_1, \ldots, u_s$ be an orthonormal basis of $\mathbb{R}^s$ such that $\pi(u_i) = e_i$ for $i \in \{1, \ldots, r\}$ and $\ker(\pi) = \operatorname{span}(u_{k+1}, \ldots, u_s)$. Let U be the orthogonal matrix with columns $u_1, \ldots, u_s$. Consider new coordinates y such that $x = Uy$. We can write

$$L \cap \mathcal{S}_+^k = \{Uy \in \mathbb{R}^s : \sum_{i=1}^{s} y_i B_i + (1 - \sum_{i=1}^{s} y_i) B_{s+1} \succeq 0\},$$

where $B_1, \ldots, B_{s+1}$ are linear combinations of the A_i's. Then

$$C = \{(y_1, \ldots, y_k) \in \mathbb{R}^k : \exists y_{k+1}, \ldots, y_s \in \mathbb{R} \text{ s.t.} \sum_{i=1}^{s} y_i B_i + (1 - \sum_{i=1}^{s} y_i) B_{s+1} \succeq 0\}.$$

Since M is full rank, we can factor it as $M = AB$, where $A, B \in \mathbb{R}^{(k+1)\times(k+1)}$ and

$$A = \begin{pmatrix} 1 & 0 & \cdots & 0 & 1 \\ 0 & 1 & \cdots & 0 & 1 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & 1 & 1 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix}, \qquad\qquad B = A^{-1}M.$$

The inner polytope P comes from an affine slice of the conic hull of the rows of A. Let the slice be given by the last coordinate equal to 1. Then P is the standard simplex in $\mathbb{R}^k$, i.e.

$$P = \operatorname{conv}\{e_1, \ldots, e_k, 0\}.$$

Since $e_i \in P \subseteq C$ for $i \in \{1, \ldots, k\}$, then there exist $y_{k+1}^{(i)}, \ldots, y_s^{(i)} \in \mathbb{R}$ such that

$$D_i = B_i + \sum_{j=k+1}^{s} [y_j^{(i)} (B_j - B_{s+1})] \succeq 0.$$

Since $0 \in P \subseteq C$, then there exist $y_{k+1}^{(0)}, \ldots, y_s^{(0)} \in \mathbb{R}$ such that

$$D_{k+1} = B_{s+1} + \sum_{j=k+1}^{s} [y_j^{(0)} (B_j - B_{s+1})] \succeq 0.$$

Consider the spectrahedron

$$C' = \{(y_1, \ldots, y_k) : \sum_{i=1}^{k} y_i D_i + (1 - \sum_{i=1}^{k} y_i) D_{k+1} \succeq 0\}.$$

We have $e_i \in C'$ for $i \in \{1, \ldots, k\}$, since $D_i \succeq 0$. Also $0 \in C'$, since $D_{k+1} \succeq 0$. Thus $P \subseteq C'$.

Moreover, if $(y_1,\dots,y_k)\in C'$, then

$$0\preceq \sum_{i=1}^k y_iD_i+(1-\sum_{i=1}^k y_i)D_{k+1}=\sum_{i=1}^k y_i(B_i+\sum_{j=k+1}^s [y_j^{(i)}(B_j-B_{s+1})])$$

$$+(1-\sum_{i=1}^k y_i)(B_{s+1}+\sum_{j=k+1}^s [y_j^{(0)}(B_j-B_{s+1})])$$

$$=\sum_{i=1}^k y_iB_i+\sum_{j=k+1}^s (\sum_{i=1}^k y_iy_j^{(i)}-(1-\sum_{i=1}^k y_i)y_j^{(0)})B_j$$

$$+(1-\sum_{i=1}^k y_i-\sum_{j=k+1}^s (\sum_{i=1}^k y_iy_j^{(i)}-(1-\sum_{i=1}^k y_i)y_j^{(0)}))B_{s+1}.$$

Therefore $(y_1,\dots,y_k)\in C$ and $P\subseteq C'\subseteq C\subseteq Q$. $\square$

We conjecture that the statement of Lemma 5.1 holds for matrices of any size.

Conjecture 5.2. *Let $M\in\mathbb{R}^{p\times q}_{\geq 0}$ have rank $k+1$ and assume that $M\mathbf{1}=\mathbf{1}$. Then M has psd rank at most k if and only if we can nest a spectrahedron of size k between the two polytopes P and Q corresponding to M.*

We now turn our attention to matrices which lie on the boundary of the set of matrices of fixed size, rank, and psd rank. Our goal is to present partial evidence towards Conjecture 4.3. Suppose we have polytopes P and Q and a spectrahedron C such that $P\subseteq C\subseteq Q$. Further, assume that P has $k+1$ vertices. We show that if k of the $k+1$ vertices of the polytope P touch the spectrahedron C at rank-one loci, then we can find a smaller spectrahedron C' such that $P\subseteq C'\subseteq C\subseteq Q$. Moreover, since our construction of C' is analogous to the construction of nested ellipses in the proof of Theorem 3.6, we believe that, additionally, at least one of the facets of Q does not touch C'. This would mean that the matrix $S_{P,Q}$ does not lie on the boundary $\partial\mathcal{M}^{(k+1)\times(k+1)}_{k,k+1}$.

Lemma 5.3. *Let $P=conv(e_1,\dots,e_k,0)\subseteq\mathbb{R}^k$. Let C be a spectrahedron of size k such that $P\subseteq C$ and the vertices $e_1,\dots,e_k$ correspond to rank one matrices in C. Then there exists another spectrahedron C' of size k such that $P\subseteq C'\subseteq C$ with all $k+1$ vertices of P corresponding to rank one matrices in C'.*

Proof. The statement is trivial when $k=1$. We proceed by induction.

By the conditions in the statement of the lemma, we can assume that

$$C=\{(x_1,\dots,x_k)\in\mathbb{R}^k : x_1a_1a_1^T+x_2a_2a_2^T+\dots+x_ka_ka_k^T+(1-\sum_{i=1}^k x_i)B\succeq 0\},$$

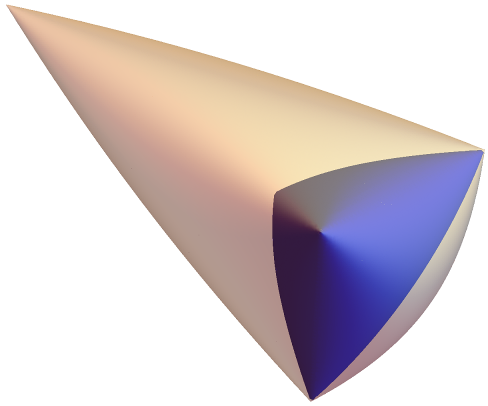

Figure 6.: The spectrahedra C (in light yellow) and C' (in blue) as in Lemma 5.3

where $a_1, \dots, a_k \in \mathbb{R}^k$ are vectors. We have $B \succeq 0$ since $0 \in C$.

Suppose first that $\dim(\mathrm{span}\{a_1, \dots, a_k\}) = \ell < k$. Let U be a change of coordinates that transforms $\mathrm{span}\{a_1, \dots, a_k\}$ into $\mathrm{span}\{e_1, \dots, e_l\}$. Denoting $a_i' = Ua_i$, we have

$$C = \{(x_1, \dots, x_k) \in \mathbb{R}^k : x_1 a_1'(a_1')^T + x_2 a_2'(a_2')^T + \cdots + x_k a_k'(a_k')^T + (1 - \sum_{i=1}^k x_i) UBU^T \succeq 0\},$$

where $B' = UBU^T$ is positive semidefinite. If $B'_{i,j} = 0$ for all $i, j \geq \ell + 1$, then, the statement reduces to the case of ℓ, which is true by induction. So suppose that $B'_{\ell+1,\ell+1} > 0$ (since $B' \succeq 0$). Choose a vector $d \in \mathbb{R}^k$ such that $d_{\ell+1} \neq 0$ and $dd^T \preceq B'$. Consider the spectrahedron

$$C' = \{(x_1, \dots, x_k) \in \mathbb{R}^k : x_1 a_1'(a_1')^T + x_2 a_2'(a_2')^T + \cdots + x_k a_k'(a_k')^T + (1 - \sum_{i=1}^k x_i) dd^T \succeq 0\}.$$

Clearly $e_1, \dots, e_k, 0 \in C'$. We will show that $C' \subseteq C$. Indeed, let $(x_1, \dots, x_k) \in C'$. Since $(a_i')_{\ell+1} = 0$ for $i \in \{1, \dots, k\}$, $d_{\ell+1} \neq 0$ and

$$x_1 a_1'(a_1')^T + x_2 a_2'(a_2')^T + \cdots + x_k a_k'(a_k')^T + (1 - \sum_{i=1}^k x_i) dd^T \succeq 0,$$

we have $(1 - \sum_{i=1}^k x_i) \geq 0$. But then

$$0 \preceq x_1 a_1'(a_1')^T + x_2 a_2'(a_2')^T + \cdots + x_k a_k'(a_k')^T + (1 - \sum_{i=1}^k x_i) dd^T$$

$$\preceq x_1 a_1'(a_1')^T + x_2 a_2'(a_2')^T + \cdots + x_k a_k'(a_k')^T + (1 - \sum_{i=1}^k x_i) B'$$

and therefore $C' \subseteq C$.

Now assume that $\dim(\text{span}\{a_1, \dots, a_k\}) = k$. Let U be an invertible transformation such that $Ua_i = e_i$. Then

$$C = \{(x_1, \dots, x_k) \in \mathbb{R}^k : x_1 e_1 e_1^T + x_2 e_2 e_2^T + \cdots + x_k e_k e_k^T + (1 - \sum_{i=1}^{k} x_i) UBU^T \succeq 0\},$$

where $B' = UBU^T$ is positive semidefinite. Let $d \in \mathbb{R}^k$ be such that $d_i = \sqrt{B'_{i,i}}$ and let $S \in \mathbb{R}^{k \times k}$ be such that

$$S_{i,j} = \begin{cases} \frac{B'_{i,j}}{\sqrt{B'_{i,i} B'_{j,j}}} & \text{if } B'_{i,i} B'_{j,j} \neq 0, \\ 1 & \text{if } B'_{i,i} B'_{j,j} = 0 \text{ and } i = j, \\ 0 & \text{if } B'_{i,i} B'_{j,j} = 0 \text{ and } i \neq j. \end{cases}$$

Since $B' \succeq 0$, also $S \succeq 0$, since it is obtained from B' by rescaling some rows and columns and by adding 1 on the diagonal in places that are 0 in B'. Let

$$C' = \{(x_1, \dots, x_k) \in \mathbb{R}^k : x_1 e_1 e_1^T + x_2 e_2 e_2^T + \cdots + x_k e_k e_k^T + (1 - \sum_{i=1}^{k} x_i) dd^T \succeq 0\}.$$

Then, clearly $e_1, \dots, e_k, 0 \in C'$. We will show that $C' \subseteq C$. Let $(x_1, \dots, x_k) \in C'$. Then

$$x_1 e_1 e_1^T + x_2 e_2 e_2^T + \cdots + x_k e_k e_k^T + (1 - \sum_{i=1}^{k} x_i) dd^T \succeq 0. \tag{5.1}$$

By the Schur Product Theorem, we know that the Hadamard product of two positive semidefinite matrices is positive semidefinite. Therefore, when we take the Hadamard product of the matrix (5.1) with S we get a positive semidefinite matrix. But that Hadamard product equals

$$x_1 e_1 e_1^T + x_2 e_2 e_2^T + \cdots + x_k e_k e_k^T + (1 - \sum_{i=1}^{k} x_i) B' \succeq 0,$$

and therefore $C' \subseteq C$.

□

Let P and C be as in the statement of Lemma 5.3. Let $Q \subset \mathbb{R}^k$ be any polytope such that $P \subseteq C \subseteq Q$ and consider the slack matrix $S_{P,Q}$. The statement of Lemma 5.3 indicates that $S_{P,Q}$ does not lie on the boundary $\partial \mathcal{M}_{k+1,k}^{(k+1)\times(k+1)}$, because the new spectrahedron C' does not touch Q. As we saw in Section 3, in order for a matrix to lie on the boundary, the configuration $P \subseteq C \subseteq Q$ has to be very tight, and Lemma 5.3 shows that having k of the vertices of P lie in the rank one locus of C is not tight enough. Similarly, having k of the facets of Q touch C at rank $k-1$ loci will not be enough. This is why we believe that all $k+1$ vertices of P have to be in the rank one locus of C, and all $k+1$ of the facets of Q have to touch C at its rank $k-1$ locus.

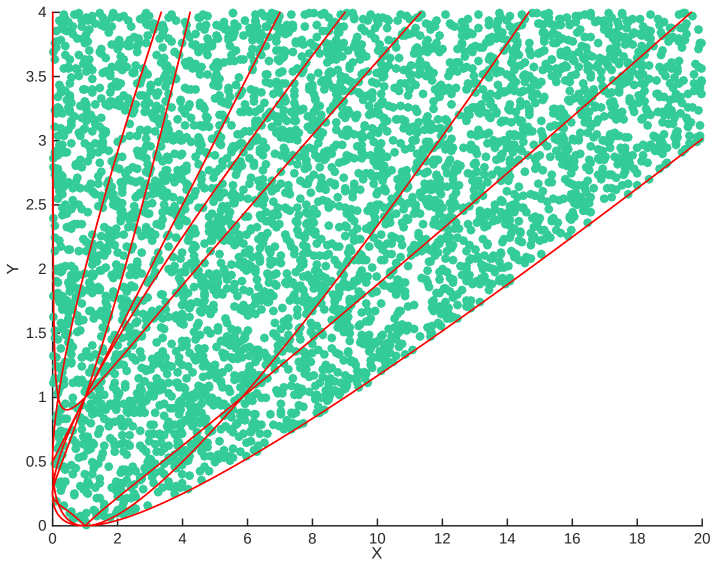

Figure 7.: A family of 4×4 circulant matrices of psd rank at most 3

5.2. *Computational evidence*

In this section we provide computational evidence for Conjecture 4.1 when $k > 2$.

Example 5.4. We consider the 2-dimensional family of 4×4 circulant matrices

$$\begin{bmatrix} a & b & 1 & b \\ b & a & b & 1 \\ 1 & b & a & b \\ b & 1 & b & a \end{bmatrix} \tag{5.2}$$

which is parametrized by a and b.

In Figure 7, the 4126 green dots correspond to randomly chosen matrices of the form (5.2) that have psd rank at most three. The psd rank is computed using the code provided by the authors of [25] adapted to the computation of psd rank [15, Section 5.6]. The red curves correspond to matrices of the form (5.2) that have a psd factorization by 3×3 rank one matrices. These curves are obtained by an elimination procedure in `Macaulay2`.

If the condition that $k+1$ matrices A_i and $k+1$ matrices B_j have rank one is equivalent to the matrix M being on the algebraic boundary $\overline{\partial \mathcal{M}^{p\times q}_{r,k}}$, then the set of matrices that have a psd factorization by such matrices should have codimension one inside the variety $\mathcal{V}^{p\times q}_r$ of $p \times q$ matrices of rank at most r. The dimension of $\mathcal{V}^{p\times q}_r$ is $pr + qr - r^2$. In the following example, we test several different assignments of ranks to each of the matrices A_i, B_j, and we mark those whose image has dimension $pr + qr - r^2 - 1$.

Example 5.5. Let $A_1, \ldots, A_p, B_1, \ldots, B_1 \in \mathcal{S}^k_+$ be symbolic matrices of ranks $r_1, \ldots, r_p, r'_1, \ldots, r'_q$. We construct a matrix M such that $M_{ij} = \langle A_i, B_j \rangle$. We vectorize the matrix M and compute its Jacobian J with respect to the entries of $A_1, \ldots, A_p, B_1, \ldots, B_q$. Finally we substitute the entries of $A_1, \ldots, A_p, B_1, \ldots, B_q$ by random nonnegative integers and compute the rank of J after this substitution. If $\mathrm{rank}(J) = pq - 1$, then the matrices that have a psd factorization by matrices of ranks $\{r_1, \ldots, r_p\}, \{r'_1, \ldots, r'_q\}$ give a candidate for a boundary component, assuming that the boundary components are only dependent on the ranks of the A_i's and the B_j's.

psd rank	p	q	ranks
3	4	4	{{1,1,1,1},{1,1,1,1}}
3	4	5	{{1,1,1,1},{1,1,1,1,2/3}}
3	4	6	{{1,1,1,1},{1,1,1,1,2/3,2/3}},{{1,1,1,2},{1,1,1,1,1,1}}
3	5	5	{{1,1,1,1,2/3},{1,1,1,1,2/3}}
3	5	6	{{1,1,1,1,2/3},{1,1,1,1,2/3,2/3}},{{1,1,1,2,3},{1,1,1,1,1,1}}
3	6	6	{{1,1,1,1,2/3,2/3},{1,1,1,1,2/3,2/3}},{{1,1,1,1,1,1},{1,1,1,2,3,3}}, {{1,1,1,1,1,1},{1,1,2,2,2,2}},{{1,1,1,1,1,2},{1,1,1,2,2,2}}

Table 1.: Ranks of matrices in the psd factorization of a psd rank three matrix that can potentially give boundary components

The possible candidates for $k = 3$ are summarized in Table 1. For all p, q the case where four matrices A_i and four matrices B_j have rank one and all other matrices have any rank greater than one are represented. These are the cases that appear in Conjecture 4.1. If any of the other candidates in Table 1 corresponded to a boundary component, then Conjecture 4.1 would be false.

If $k = 4$, $p = q = 10$, exactly five A_i and five B_j matrices have rank one and the rest of the matrices have rank two, then the Jacobian has rank 94. If the rest of the matrices in the psd factorization have rank three or four, then the Jacobian has rank 99 as expected. Hence if Conjecture 4.1 is true, then in general not every matrix on the boundary has a psd factorization with $k+1$ matrices A_i and $k+1$ matrices B_j having rank one, and rest of the matrices having rank two.

Example 5.6. Using the same strategy as in Example 5.5, we have checked that the Jacobian has the expected rank for $p = q = r = k + 1$ and $k < 10$.

References

[1] S. Basu, R. Pollack, and M.-F. Roy: *Algorithms in real algebraic geometry*, Springer (2005).

[2] W. Bruns and U. Vetter: *Determinantal rings*, Springer (1988).

[3] J. E. Cohen and U. G. Rothblum: *Nonnegative ranks, decompositions, and factorizations of nonnegative matrices*, Linear Algebra and its Applications 190 (1993), pages 149–168.

[4] H. Fawzi, J. Gouveia, P. A. Parrilo, R. Z. Robinson, and R. R. Thomas: *Positive semidefinite rank*, Mathematical Programming 153 (2015), pages 133–177.

[5] H. Fawzi and P. A. Parrilo: *Self-scaled bounds for atomic cone ranks: applications to nonnegative rank and cp-rank*, Mathematical Programming 158 (2016), pages 417–465.

[6] H. Fawzi, J. Saunderson, and P. A. Parrilo: *Equivariant semidefinite lifts and sum-of-squares hierarchies*, SIAM Journal on Optimization 25 (2015), pages 2212–2243.

[7] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf: *Linear vs. semidefinite extended formulations: exponential separation and strong lower bounds*, Proceedings of the forty-fourth annual ACM symposium on Theory of computing. ACM (2012).

[8] J. Gallier: *Notes on Convex Sets, Polytopes, Polyhedra, Combinatorial Topology, Voronoi Diagrams and Delaunay Triangulations*, available online at http://www.cis.upenn.edu/~jean/combtopol-n.pdf.

[9] N. Gillis and F. Glineur: *On the geometric interpretation of the nonnegative rank*, Linear Algebra and its Applications 437 (2012), pages 2685–2712.

[10] J. Gouveia, P. A. Parrilo, and R. R. Thomas: *Lifts of Convex Sets and Cone Factorizations*, Mathematics of Operations Research 38 (2013), pages 248–264.

[11] J. Gouveia, P. A. Parrilo, and R. R. Thomas: *Approximate cone factorizations and lifts of polytopes*, Mathematical Programming 151 (2015), pages 613–637.
[12] J. Gouveia, K. Pashkovich, R. Z. Robinson, and R. R. Thomas: *Four dimensional polytopes of minimum positive semidefinite rank*, arXiv:1506.00187.
[13] J. Gouveia, R. Z. Robinson, and R. R. Thomas: *Polytopes of minimum positive semidefinite rank*, Discrete & Computational Geometry 50 (2013), 679–699.
[14] J. Gouveia, R. Z. Robinson, and R. R. Thomas: *Worst-case results for positive semidefinite rank*, Mathematical Programming 153 (2015), pages 201–212.
[15] H. Klauck, T. Lee, D. O. Theis, and R. R. Thomas *Limitations of convex programming: lower bounds on extended formulations and factorization ranks (Dagstuhl Seminar 15082)*, Dagstuhl Reports 5 (2015), pages 2192–5283.
[16] H. Kraft and C. Procesi: *Classical Invariant Theory, a Primer*, available at https://math.unibas.ch/uploads/x4epersdb/files/primernew.pdf.
[17] K. Kubjas, E. Robeva, and B. Sturmfels: *Fixed Points of the EM algorithm and nonnegative rank boundaries*, Annals of Statistics 43 (2015), 422–461.
[18] J. R. Lee, P. Raghavendra, and D. Steurer: *Lower bounds on the size of semidefinite programming relaxations*, Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing. ACM (2015).
[19] J. R. Lee, P. Raghavendra, D. Steurer, and N. Tan: *On the power of symmetric LP and SDP relaxations*, in 2014 IEEE 29th Conference on Computational Complexity. IEEE (2014).
[20] K. Ranestad and B. Sturmfels: *On the convex hull of a space curve*, Advances in Geometry 12 (2012), pages 157–178.
[21] P. Rostalski and B. Sturmfels: *Dualities*, in Semidefinite Optimization and Convex Algebraic Geometry. SIAM (2013).
[22] T. Rothvoß: *The matching polytope has exponential extension complexity*, Proceedings of the 46th annual ACM symposium on theory of computing. ACM (2014).
[23] R. Sinn and B. Sturmfels: *Generic Spectrahedral Shadows*, SIAM Journal on Optimization 25 (2015), pages 1209–1220.
[24] B. Sturmfels: *Solving systems of polynomial equations*, American Mathematical Society (2002).
[25] A. Vandaele, N. Gillis, F. Glineur, and D. Tuyttens: *Heuristics for exact nonnegative matrix factorization*, Journal of Global Optimization (2015), pages 1–32.
[26] S. A. Vavasis: *On the complexity of nonnegative matrix factorization*, SIAM Journal on Optimization 20 (2009), pages 1364–1377.
[27] M. Yannakakis: *Expressing combinatorial optimization problems by Linear Programs*, Journal of Computer and System Sciences 43 (1991), 441–466.